

RELATIONAL STRUCTURE OF COMPETITION IN
THE EXECUTIVE LABOR MARKET

Marko Coh

London Business School

Thesis submitted in partial fulfilment of the
Requirements for a Doctor of Philosophy degree at London Business School

Declaration

This thesis is the result of my studies at London Business School from September 2006 to May 2011. The work presented in the thesis is my own, unless otherwise stated in the text. Specifically, the research presented in Chapter 4 was done in collaboration with Professor Isabel Fernandez-Mateo at London Business School.

Signature:

Date:

EXECUTIVE SUMMARY

This thesis investigates when labor market participants pursue opportunities introduced to them by their labor market contacts. The thesis goes beyond the view of relationships in the labor market as a pathway to resources such as timely information about a vacancy or support in a new place of work, that in turn stimulate labor market participants to act on the opportunities. I argue that interactions also serve as a vehicle for transmission of non-resource flows, whose nature is conditioned by the roles that labor market participants and their contacts play in them. The thesis shows that the content of ties that emerge through time shapes which contact-channelled opportunities are pursued.

The empirical context of the thesis is the executive labor market, where executives and search firms interchangeably act as labor market participants and their contacts. I combine detailed statistical analyses of a unique dataset containing a history of interactions and search assignments at a global executive search firm with interviews among search consultants, candidates for executive-level jobs, and search firm clients. Executives represent an ideal population for a study of the role of tie content in the labor market as their position in the organizational hierarchies means they can, in their interactions with search consultants, play very different roles of candidates, clients, sources and general informants.

This thesis contains three separate studies that examine the progress of candidates for executive-level jobs to interviews and offers for vacancies mediated by the search firm. All the studies share a focus on the role of differences in the content of ties between executives and the search firm. The first study shows that prior ties built through interactions in which the search consultants played a boundary-spanning role facilitate the executives' willingness to subsequently participate in the recruitment process, while the ties in which the search consultants performed pure brokerage roles don't. The second study shows that the content of

candidate ties, specifically who rejected whom in prior searches, differentially affects female vs. male progress through the hiring pipeline. The third study demonstrates that the differences in tie content affect not only when labor market participants act on the information about opportunities, but that they also have identifiable economic implications.

The findings contribute to the sociological literature on labor markets by showing that the content of ties between labor market participants and their contacts developed through time shapes the outcomes in the labor market, and that when contacts present the participants with opportunities the latter may act on the basis of the content of the past non-resource flows in their interactions, rather than on the merit of the opportunity. The findings also contribute to the brokerage literature by demonstrating that when brokered parties and the brokerage organization have a past history of interactions, the brokered parties' past point of contact with the broker shapes the probability of participation in subsequent brokered market exchanges. Finally, the findings contribute to the literature on strategic human resource management by unpacking how the mobility of highly paid talent is set in motion, and in particular what role labor market intermediaries such as executive search firms play in it.

ACKNOWLEDGEMENTS

I have always been curious about why things work the way they do. As a child I have taken apart every toy car I got as a gift, as I wanted to know “what is inside”. If it was my natural curiosity that led me to pursue a PhD, it was the help and support of people at London Business School and in the other parts of my life that got me through to the end of it. Firstly I would like to express my gratitude to my supervisor Professor Isabel Fernandez-Mateo. Isabel taught the first course I took in my doctoral programme, and introduced me to the world of economic sociology. This early experience had a profound impact on the development of my research interests. Always having a sharp eye for the inconsistencies in my arguments and analyses, Isabel taught me how to do rigorous research, and thereby provided me with the best possible platform for my future work.

I am deeply grateful to the search consultants and managers at the search firm that provided the data for my thesis. They also gave their time and shared their insights to help me understand what the data means, and how the search process works. This thesis would not be possible without their generosity.

I would like to thank a number of the LBS faculty who helped me at different stages of the programme. Dr Lisa Cohen and Professor Freek Vermeulen acted as the members of my dissertation proposal committee, and helped me set the direction of my research. In the formative stages of my thesis Professor Phanish Puranam provided helpful feedback. Toward the end of the process I received helpful suggestions and advice from Dr Donal Crilly, Professor Gary Dushnitsky, Dr Ioannis Ioannou, Dr Brandon Lee, and Dr Lourdes Sosa. I am particularly grateful to Professor Costas Markides for all his support. I would also like to thank Dr Yiorgos Mylonadis who gave me the opportunity to obtain teaching experience in the executive education programmes.

I will always be indebted to my late mentor at the University of Ljubljana, Professor Miroslav Glas. Ever since I first came up with the idea of pursuing a PhD, he encouraged me to aim for the highest international level. I wouldn't be here without his visionary support and his belief in me.

My research and the overall experience have benefited from the numerous conversations with the past and present doctoral students in the Strategy & Entrepreneurship department at the LBS. In particular, I would like to thank: Ramon Lecuona, Eucman Lee, Dr Felipe Monteiro, Dr Mihaela Stan, and Dr Onal Vural. A very special thanks goes to Dr Susan Hill. In addition to being a thoughtful colleague, Susan has been a friend and a supporter one could only wish for. I am also grateful to Sharon Wilson, the department coordinator, for all the help with the administrative matters. Doctoral students in the other LBS departments, in particular Dr Ramin Baghai, Dr Marc Gabarro, Dr Oguzhan Karakas and Yvetta Simonyan, have been great companions on this journey.

My London friends outside of LBS were always there to share my ups and downs in the programme, usually with a good dose of humour and encouragement. I am particularly thankful to Vladimir Aleksic, Drasko Draskovic, Mirko Milojevic and Ruzica Simic. My time in London would not have been the same without them. Whenever I went back to Slovenia, my friends, especially Aleksandar Ivanovic and Gregor Trtnik, made me feel as if I had never left. That always filled me with energy for the challenges ahead.

It is hard to describe my gratitude to my family for all their love and support. My mom Zorica is one of the most successful women I know, and at the same time a wonderful mother. She is my role model, and her love and support for all my endeavours have been endless. My sister Dusica always made me aware of how proud she was of her little-big brother for taking on such a project. Her daughter Nika, my niece, has been my joy, and a welcome distraction from my work when I was back home in Slovenia. I am grateful to my

late dad Ivan, who imparted in me the desire to learn and to explore. He would have been proud I have taken in the footsteps I have.

Finally, my time in London gave me so much more than a PhD. It is here that I met my fiancée Vlatka. Vlatka is the most wonderful person I know, more than I could ever hope for in a life companion, and “moja ljube”. An academic herself, she understood exactly what I had been going through. With unfailing sensibility she always said and did just what I needed. She has been the best possible sounding board, and a limitless source of encouragement. With all my love and appreciation, I dedicate this thesis to her.

LIST OF TABLES

Table 2.1: Interviews.....	28
Table 2.2: Summary of individual, relationship and vacancy data.....	52
Table 3.1: Summary of descriptive statistics	83
Table 3.2: Linear probability models predicting who does not go to an interview with the search firm.....	88
Table 3.3: Linear probability models predicting who does not go to an interview with the hiring firm after being interviewed by the search firm.....	93
Table 3.4: Linear probability models predicting who is dropped by the search firm	97
Table 4.1: Summary of descriptive statistics	118
Table 4.2: OLS regression predicting natural log of the salary of the vacancy for which the candidate is considered in Step 0 of the hiring pipeline.....	120
Table 4.3: Linear probability models predicting who does not go to an interview with the search firm.....	124
Table 4.4: Linear probability models predicting who does not get an offer from the hiring firm.....	129
Table 5.1: Summary of descriptive statistics	143
Table 5.2: OLS regressions predicting the length of the search assignment ...	145
Table 5.3: OLS regressions predicting margin of the vacancy a candidate is allocated to by the search firm	149

LIST OF FIGURES

Figure 2.1: Triadic structure of the executive labor market.....	26
Figure 2.2: Hiring pipeline at Execo	49
Figure 3.1.1: Candidate tie	65
Figure 3.1.2: General tie.....	65
Figure 3.1.3: Client tie.....	65
Figure 3.1.4: Source tie	66
Figure 3.2: Tie content, brokerage and boundary spanning roles, and trust.....	71
Figure 3.3: Hiring pipeline at Execo (reduced sample)	78
Figure 4.1: Hiring pipeline at Execo	112

CONTENTS

EXECUTIVE SUMMARY	i
ACKNOWLEDGEMENTS	iii
LIST OF TABLES	vi
LIST OF FIGURES	vii
CHAPTER 1: INTRODUCTION	1
1.1 Social networks in labor markets	5
1.1.1 “Information-based” view of social networks in labor markets	6
1.1.2 “Resource-based” view of social networks in labor markets	7
1.1.3 “Socialization-based” view of social networks in labor markets	8
1.2 Brokerage	11
1.3 Objectives of the thesis.....	17
1.4 Outline of the thesis.....	22
CHAPTER 2: RESEARCH SETTING AND DATA DESCRIPTION.....	24
2.1 Research context: executive search.....	24
2.1.1 Search process.....	27
2.2 Quantitative data.....	40
2.2.1 Dependent variables.....	50
2.2.2 Individual data	53
2.2.3 Relationship data.....	54
2.2.4 Vacancy data.....	54
2.3 Data limitations	57
CHAPTER 3: EXECUTIVE PARTICIPATION IN THE COMPETITION FOR JOBS: THE ROLE OF CONTENT OF TIES WITH THE BROKER	60
3.1 Introduction	60
3.2 Brokerage and boundary spanning roles, trust development, and pursuit of labor market opportunities	67
3.4 The sample and variables	76
3.5 The analysis and results.....	82
3.5.1 Robustness of the results.....	100
3.6 Discussion and conclusions.....	103
CHAPTER 4: THE ROLE OF GENDER AND PRIOR RELATIONSHIPS IN THE EXECUTIVE SEARCH.....	108
4.1 Introduction	108
4.2 Prior relationships, gender and progress through the hiring pipeline	109

4.3 The sample and variables	111
4.4 The analysis and results.....	119
4.5 Discussion and conclusions.....	131
CHAPTER 5: VALUE CAPTURE IN EXECUTIVE SEARCH.....	136
5.1 Introduction	136
5.2 Trusting ties, ties based on rejections, and value capture	137
5.3 The sample and variables	139
5.4 The analysis and results.....	144
5.5 Discussion and conclusions.....	151
CHAPTER 6: CONCLUSIONS	153
6.1 Summary of findings	153
6.2 Discussion and contributions.....	155
6.3 Limitations.....	161
6.4 Future research	162
REFERENCES.....	165
APPENDIX.....	175

CHAPTER 1: INTRODUCTION

Ever since Mark Granovetter published his seminal article on the strength of weak ties in getting a job (1973), scholars in economic sociology and management have sought to contribute to our understanding of the functioning of the labor market (e.g. Fernandez, Castilla and Moore, 2000; Petersen, Saporta and Seidel, 2000; Castilla, 2005; Khurana, 2002). The main contribution of this line of work has been to demonstrate that relationships between parties in the labor market exert important influence on the processes and outcomes. For instance, job seekers may hear about job openings from their close contacts or more distant acquaintances and the degree of novelty of that information for the job seeker determines whether he/she ends up in that job (Granovetter, 1973; Yakubovich, 2005). Individuals referred to the organizations by their employees, or by labor market intermediaries, have been found to stand more chance of being interviewed and hired than isolated candidates (Petersen, Saporta and Seidel, 2000).

Within this line of work researchers have emphasised three different facets of the role that relationships play in the labor market. The first body of work could be termed “information-based”. A critical premise there is that variance in the labor market processes and outcomes is a direct consequence of variance in the attributes of information transmitted through the relationships (e.g. timeliness, Granovetter, 1973; depth, Petersen, Saporta and Seidel, 2000). The second body of work could be termed “resource-based”. The studies in that tradition argue that the benefits of relationships do not flow through the relationships themselves, but rather through leading to the individuals who possess valuable resources (Lin, Ensel and Vaughn (1981a, 1981b)). The third body of work might be termed “socialization-based”. The latter studies argue that contacts that refer individuals to the organizations may have the most impact once referrals are hired, as they can assist them with their socialization

into the organization and thereby facilitate the referral's performance and retention (Fernandez, Castilla and Moore, 2000; Castilla, 2005).

While the three streams diverge in their view of the specific relational mechanisms that govern the functioning of the labor market, they share an assumption that labor market participants act on the information and cues provided by their contacts. In other words, they assume that contacts differ in what they can provide to the labor market participants, and the latter act on the information or cues that are most relevant to their situation. In short, the causal mechanisms are focused on the resources provided by the relationships. These resources may be less social or more social in nature (an example of the first would be information about a vacancy, Granovetter, 1973, and of the second advice and support in the new workplace, Castilla, 2005).

The aforementioned assumption is somewhat simplistic as it ignores a number of other "currencies" that are flowing through relationships in the labor market, and may affect when labor market participants pursue opportunities. This is surprising as the broader field of economic sociology has made significant strides in the understanding of how currencies such as trust and mutual obligation shape inter-organizational exchanges (Uzzi, 1996, 1997, 1999). I suggest that this is in large part because prior labor market studies relied on data that distinguished between the existence or non-existence of a relationship (Fernandez, Castilla and Moore, 2000; Castilla, 2005), measured the strength of a relationship (Granovetter, 1973; Yakubovich, 2005), or captured the identity of the contacts (Petersen, Saporta and Seidel, 2000). They also only sampled those individuals who decided to pursue a job opportunity presented to them by a contact (e.g. only those who accepted to be referred to the hiring organizations, not also those who rejected). We therefore know very little on, first, how differences in the content of the relationships between labor market participants and their contacts affect labor market processes and outcomes, and second, when relationships with

contacts lead the individuals to pursue particular jobs, and when they lead organizations to pursue particular individuals.

Prior research has started addressing the topic of relationship content. Rich qualitative studies have examined the functioning of one segment of the labor market, the executive labor market (Khurana, 2002; Finlay and Coverdill, 2002). These studies have uncovered that the labor market intermediaries in the executive search firm sector refer individuals to hiring firms, and in the process develop a variety of relationships with both the individuals and the hiring firms. The authors describe the content that flows through these relationships, and how it affects the search firms' overall position in the market, as well as the behavior of the market participants. While these studies have provided plenty of insight into functioning of the labor market, there have been no quantitative empirical studies that would systematically examine how the content of the relationships between labor market participants and their labor market contacts shape which labor market opportunities the former pursue. A benefit of such a study would be that it would be able to control for a host of explanations, and thereby provide a sharper identification of the mechanisms shaping the labor market processes and outcomes.

In this thesis I conduct such a study. I am particularly focusing on the role of the non-resource content that flows through relationships (Podolny and Baron, 1997). My approach is focused on exploring which labor market opportunities the labor market participants pursue (i.e. which opportunity they decide to "compete" for). I hence position my study as an investigation of the relational structure of competition in the labor market. Building on the prior literature that addressed content of flows in the labor market I situate my study in the executive labor market, and study the role of the different kinds of prior relationships between job candidates and the search firm in the functioning of this segment of the market. I use a unique quantitative dataset obtained from electronic records of a global executive

search firm (“Execo” - a pseudonym). I combine my quantitative findings with insights from interviews with search consultants, hiring firm representatives, and executive candidates.

My study shows that the differences in content of prior relationships with the search firm indeed affect which job opportunities introduced by the search firm the executives pursue. Furthermore, these relationships also shape the search firm’s selection of candidates to interview and to introduce to the hiring firms. The two properties of prior relationships that have emerged as important in my study are first, the extent to which these relationships lead to the development of the job candidates' trust in the search firm, and second, the identity of the party that has previously rejected pursuit of a match (the search firm, the hiring firm or the candidate), as well as the candidate’s gender. These findings underline the importance of relationship flows other than resources in the functioning of the labor market.

In the remainder of this chapter I first review the main literature on social networks in labor markets. I organize the review around the facets of the relationships emphasized by the three groups of studies (i.e. information-based, resource-based and socialization-based). I indicate that the second and the third group of studies directly respond to the first one, and provide alternative explanations of the mechanisms governing the relationship between relationships and outcomes in the labor market. At the end of section 1.1 I highlight how my study contributes beyond the studies included in the review. As the search firm operates as a broker between candidates and the hiring firms, I also review the key literature on brokerage. At the end of section 1.2 I describe how my study, by focusing on the question of when labor market participants pursue opportunities, also contributes to the understanding of brokerage. In the end of the chapter (section 1.3) I summarize the objectives of the thesis, and provide its outline.

1.1 Social networks in labor markets

The literature exploring the role of social networks in labor markets goes back to 1973 when Mark Granovetter published the article “The strength of weak ties”. In this early formulation Granovetter argued that social networks are very important in the process of finding a job. He demonstrated that individuals are more likely to obtain jobs through their weak than through their strong ties (with the strength of the tie depending on the amount of time spent together by two people). Granovetter’s study and subsequent studies in that tradition can be included in the group of “information-based” studies as they emphasize the benefits of information transmitted through relationships. I present examples of such studies in section 1.1.1

Empirical support for the strength of weak ties hypothesis in subsequent studies has been mixed, and as an alternative Nan Lin and co-authors proposed the social resources theory (Lin, Ensel and Vaughn, 1981a). In section 1.1.2 I present the main tenets of this group of studies (that I term “resource-based”) and present examples of empirical studies.

Whilst there are literally hundreds of studies of the relationship between social networks and career success, less attention has been devoted to the relationship between social networks and the outcomes from the organizational perspective. While the particular social ties affect individuals’ probability of getting jobs or getting promoted, social ties of the organizational members also affect the organizational hiring outcomes (Fernandez, Castilla and Moore, 2000; Castilla, 2005). As these studies emphasise socialization as the key mechanism through which the labor market contact (a referrer in that case) affects the outcomes for the hired individuals, I term this group of studies “socialization-based”. I outline the conclusions of this group of studies in section 1.1.3.

1.1.1 “Information-based” view of social networks in labor markets

In this section I review the literature that explores how information transmitted through relationships affects labor market outcomes such as being interviewed and getting a job. The first theoretical approach is the weak tie theory introduced by Granovetter (1973). This approach focuses on the strength of ties used in the process of finding a job. Granovetter argued that the ties among members of a social clique are likely to be strong (frequent and emotionally intense). Information obtained through strong ties is likely to be redundant as the individual receiving the information is likely to have that piece of information already. The ties individuals have with members of other cliques are likely to be weak in the sense that they are less emotionally intense and less frequent. Information obtained through such ties is more likely to be unique. Granovetter then hypothesized that weak ties will be more likely than strong ties to channel timely information about job openings. He found evidence for his hypothesis. Later studies, however, found mixed support for the weak tie hypothesis.

Bridges and Villemez (1986) studied the relationship between the strength and income. They concluded that the effects that may appear to be caused by weak ties are actually effects of the increased productivity of workers who have successfully participated in the labor market for an extended period of time and have developed a wealth of contacts as a by-product. In another test of the strength of weak ties hypothesis Murray, Rankin and Magill (1981) examined reliance on weak ties for job-related information. They examined this question in the context of Canadian and US social and physical scientists and found that the reliance on weak ties significantly declined during the period between 1920 and 1978. They argued that these results suggest that although weak ties provide more unique information,

this information may be of little use in the actual job search as it may convey information about unobtainable or undesirable job openings.

A recent study provided robust evidence in support of the strength of the weak ties hypothesis. Yakubovich (2005) observed that all past research analyzed the relationship between the actor's network and the labor market outcomes. He re-cast the strength of the weak ties hypothesis into a proposition about the probability of getting a job as a function of the within-actor differences in tie strength. The author analyzed data on hires in one Russian city in 1998. He found that a worker's probability of obtaining a job through one of his/her weak ties was higher than the probability of obtaining it through one of his/her strong ties. With this study the author addressed the possibility that the apparent network effects are actually a result of unobserved individual productivity. He demonstrated that the strength of the tie does have an independent effect on the probability of getting a job.

Petersen, Saporta and Seidel (2000) provided a detailed case study of hiring in a mid-sized high-technology firm in California. These authors studied the entire hiring process and found that the referral method (hiring new workers on the basis of referrals from existing workers) had a strong impact on the outcome of all phases of the hiring process. Hires sourced through more informal channels were less likely to depart than hires sourced through other methods. The authors explained that their findings indicate that some hiring channels have superior information transmission properties: "for the social networks, it is unquestionably the quality of information they yield that counts." (p. 811).

1.1.2 "Resource-based" view of social networks in labor markets

The main part of this group of studies is the social resources approach, which emphasizes access to the resources held by the hierarchically superior alters (Lin, Ensel and

Vaughn, 1981a). The theory emerged as an alternative to the weak tie theory. The social resources theory focuses on the nature of resources in the network. It argues that it is not the weak ties per se that generate the expected outcomes, but the fact that such ties are more likely to lead to someone that possesses the resources beneficial for the ego's career advancement. Such an individual is considered a social resource. The contacts that provide access to information, access to resources and career sponsorship are examples of social resources. Lin, Ensel and Vaughn (1981a, 1981b) found that the alter's occupational prestige was positively related to the prestige of the job obtained by the ego. De Graaf and Flap (1988) conducted a cross-national comparison on the use of social ties in the job search process. They found that the relationship between access to the social resources and attainment may be affected by the institutional characteristics of labor markets. Marsden and Hurlbert (1988) found that the effects of social resource measures are outcome-specific in the sense that different social resources affect different outcomes, and that no social resource affects a broad range of outcomes.

Yakubovich and Lup (2006) studied an aspect of the organizational referral process. They proposed that a referral's chance of getting hired increased with the referrer's performance (and in that way the referrer features as a resource to the referral hire). Their sample was drawn from data on the online recruitment of sales agents in a virtual call centre. Their findings confirmed the hypothesis.

1.1.3 "Socialization-based" view of social networks in labor markets

This group of studies has emerged to examine the exact mechanisms that give rise to the benefits associated with referral hiring. The broader motivation of the authors was to depart from the previous information- and resource-focused accounts of social networks in

labor markets, and highlight the social dynamic between the provider and the recipient of information about the labor market opportunity (referrer and referral hire in the case of referral hiring). The authors examined both information-based and socialization-based mechanisms, and compared their relative explanatory power with respect to the benefits of referral hiring.

Fernandez, Castilla and Moore (2000) studied referral hiring to shed light on the notion of social capital. They suggested that through referral hiring organizations were acting as “social capitalists” who invested in the social connections of their employees to generate benefits in the form of better hiring outcomes. The authors addressed three explanations for why referral hiring would lead to better hiring outcomes than the hiring of isolated job seekers: richer pool of applicants, better match, and social enrichment. The authors found evidence supporting the first argument (richer pool of applicants). The mechanism underlying the richer pool of applicants from referrals was homophily – in that setting people referred people like themselves. The authors found no evidence in support of the better match argument. Referrals didn’t possess more information than non-referrals, and the employer didn’t seek additional information about referrals from the referrers. They found significant support for the third argument (social enrichment). Referral hires whose referrers stayed on the job turned over less than referrals whose referrers resigned. There was no difference in the turnover rates between the latter group and the non-referrals. In addition the authors found that the social enrichment process generated higher returns on investment in social capital of employees than the richer pool of applicants process. The authors concluded that the benefits to hiring through social networks also flow through previously unobserved socialization mechanisms.

Castilla (2005) analyzed the relationship between referral hiring and employee performance. His analysis showed that referrals initially performed better than non-referrals

but over time their performance converged. However, performance declined for referrals whose referrer left the organization. In addition the author found that performance, by itself, didn't increase turnover; performance growth, on the other hand, did influence turnover – referrals whose referrers left experienced a drop in performance, and were at a higher risk of leaving the organization. The author interpreted these findings as additional evidence of the impact of socialization processes on hiring outcomes.

The preceding literature review indicates that a common theme of past studies has been how relationships between labor market participants and their contacts shape the former's access to the valuable resources (e.g. timely information about a job opening, access to a top performer, channel for transmission of tacit knowledge about the idiosyncrasies of a workplace etc.). The focus has therefore been on the resources themselves, or the access to them. The studies, however, have not examined the content of the relationships in great depth. The strength of weak ties studies, for instance, have distinguished between strong and weak ties; however, the focus there was on the extent to which one or the other serve as a conduit for novel and timely information, rather than how they shape the disposition of the labor market participants to their contacts. This aspect appears fundamental though, as the content that flows through relationships may shape such properties of the relationship between two actors as trust and mutual obligation. These properties may in turn affect the willingness of the labor market participants to act on the information and cues provided by their contacts.

In my study, in contrast with the reviewed studies, I examine relationships with a multitude of ties with different content between the labor market participants and their contacts. In my setting the labor market participants are candidates for executive level jobs, and a search firm. The candidates and the search firm represent labor market contacts to each other. Also, in contrast with the extant studies I study sets of individuals that have been

simultaneously considered for an opportunity, and am able to distinguish who rejected the opportunity, as well as who was rejected for an opportunity by the hiring agent (the search firm). This allows me to examine how the different ties within a single relationship shape whether or not a labor market participant acts on information about an opportunity (in my case a vacancy, or availability of a job candidate).

1.2 Brokerage

The empirical context of my study is executive search. Executive search firms are a labor market intermediary, and like the other labor market intermediaries they broker the hiring process (i.e. mediate between individuals – candidates for jobs on one side, and the hiring firms on the other). They introduce job opportunities to executives, who decide to pursue them or to reject them. In this way the search firms correspond to my conceptualization of a labor market contact that provides information on labor market opportunities, which are then either pursued or not. To indicate what the brokerage role entails I provide a review of the brokerage literature below.

One of the commonly used definitions of brokerage is the one offered by Marsden (1982, p. 202). He defined brokerage as a process “by which intermediary actors facilitate transactions between other actors lacking access to or trust in one another.” Roots of interest in third parties can be traced back to Simmel (1950) who suggested that a triad is a fundamental unit of social structure, and that triadic social relations are qualitatively different from dyadic ones. Gould and Fernandez (1989) expanded on that notion; they suggested that any set of social actors can be divided into a set of mutually exclusive subgroups. Actors can then be brought back together in five different ways, which in turn produces five identifiable brokerage roles.

The first role is the co-ordinator. In this case, the actor that facilitates the exchange and the participants in the exchange are members of the same subgroup. The second role is the cosmopolitan or itinerant broker. In this role the broker belongs to one subgroup while the brokered parties belong to another subgroup. The third role is the gatekeeper. Gould and Fernandez explained this role using an example from politics: the broker plays the role of the gatekeeper for his or her party if he or she is asked to provide access to a certain member of his or her party by a member of a rivalling party. The fourth role is the representative. This role is analogous to the role of the gatekeeper; it occurs if a broker contacts a member of the rivalling party on behalf of a member of his or her party. The fifth role is the liaison. In this role the broker is an outsider to both the initiator of the relation and to the receiver of the relation. The actor's role is to establish a link between the two previously unconnected groups. These roles represent different non-overlapping types of brokerage relations. The individual broker, however, may perform any combination of the roles simultaneously.

In a subsequent empirical analysis (Fernandez and Gould, 1994) the same authors concluded that the brokerage position confers power. They highlighted that this relationship is not automatic and depends on the broker's identity and the alignment of the broker's actions with the type of role. For instance, in their setting (US health policy) brokers performing liaison and itinerant roles only obtained influence to the extent that they were uncommitted to specific agendas. The main contribution of Gould and Fernandez's theory was to provide formalized conceptualizations of different brokerage roles that guided much of the subsequent brokerage literature.

Another foundation of the brokerage literature is Burt's theory of structural holes (1992). In this seminal work Burt described structural holes - the absence of a relationship between two contacts - as a fundamental building block of competition. Actors spanning structural holes connect actors who would otherwise be disconnected. Structural-hole-

spanning corresponds to Gould and Fernandez's definition of the brokerage relation to the extent that it involves connecting two previously disconnected parties. Burt remarked that when there is a tension between parties connected by the actor spanning structural holes this actor may obtain control benefits. The notion of control benefits has roots in Simmel's role of *tertius gaudens*, the third one who benefits from the disunion of others (Burt, 1992, p. 31). This tension may exist before the broker's involvement, and the broker reaps benefits from revealing the tension or controlling the negotiations for resolving the conflict situation. Alternatively, a broker may create tension between parties, and then reap benefits from participating in the resolution of the conflict. Burt showed that there are also informational benefits to actors who span structural holes. Such actors have access to information from different parts of the social world. Such information is non-redundant and potentially valuable.

A substantial literature has documented that brokers obtain valuable outcomes due to their occupancy of the position spanning a structural hole. Burt (2000) reviewed the social structural literature and concluded that social capital is more a function of the occupancy of structural holes than of network closure. However, he also noted that there are several contingency factors that affect the extent to which this is the case. Burt (1997) also showed that for managers, the value of bridging structural holes is contingent on the number of peers. Podolny and Baron (1997), using a sample of the professional employees and managers in a high-tech company, demonstrated that having a large, sparse network of contacts is beneficial for the upward intra-organizational mobility. They also showed, however, that in such networks reliable performance evaluations are less likely to arise. Together, the findings of the latter two studies suggest that occupancy of the brokerage position is beneficial for promotion, yet occupants of such positions are harder to evaluate, which may temper the positive effect of the informational benefits of the position. Recently Xiao and Tsui (2007)

studied another contingency factor: organizational and national culture. They found that the occupancy of brokerage positions can be detrimental in highly collectivistic national cultures, and in organizations with high commitment cultures.

In a further specification of his argument, Burt (2004) provided compelling evidence that brokerage produces social capital (which he defines as an advantage that people have due to their position in the social structure, p. 351). He suggested that occupancy of the brokerage position affords visibility of the options otherwise unseen. Managers in his sample who bridged structural holes between groups were more likely to express ideas and have them evaluated as valuable. In a variant of this theme, Hargadon and Sutton (1997) provided an ethnographic study of a design firm IDEO that indicated that a brokerage position between the technology domains stimulates innovation. Obstfeld (2005) analyzed product development in an automotive manufacturer and concluded that the managers' "tertius iungens" orientation ("third one who joins") stimulates organizational innovation.

The sociological brokerage literature provided an important additional perspective for the study of market intermediation. Traditionally, the literature grounded in economics has seen market intermediaries as the neutral facilitators of transactions. Studies of labor markets suggested that the presence of an intermediary increases efficiency of the transactions (Bull, Ornati and Tedeschi, 1987; Yavas, 1994). However, economic sociologists found that brokers may importantly affect outcomes for the market participants. Bielby and Bielby (1999) used longitudinal data on writers for television and feature films and examined the extent to which agency representation affects writers' employment and earnings. They found that writers represented by the elite agencies experienced better career progression and earned more than equally accomplished writers represented by the non-elite agencies. The authors pointed out that the elite agencies were actively involved in the production process, performing functions beyond mere intermediation. Zuckerman (1999) found that the firms who weren't followed

by the analysts who followed their industry experienced a stock price discount. Zuckerman interpreted these findings as the evidence of the ability of market intermediaries to confer legitimacy on market participants, and thereby affect market outcomes. Khurana (2002) made a similar point. He argued that in the external CEO market, the intermediaries (the executive search firms) bridge institutional gaps that exist between buyers and sellers. In that way they enable transactions that might not otherwise take place. These studies provided evidence that market intermediaries do have a qualitative impact on the outcomes of transactions that they mediate. Subsequent research addressed the question of what shapes the impact of market intermediaries on market outcomes.

Brokers enjoy the advantage of visibility of options that would otherwise be unseen. That may help them come up with creative and innovative solutions (Burt, 2004). In the context of market intermediation this means that they may help brokered parties produce outcomes that these parties could not produce on their own (Hargadon and Sutton, 1997). However, intermediaries are also self-interested market actors that may deploy their social capital for their own benefit, and/or to the detriment of the brokered parties. Researchers have examined how contingencies affect brokers' deployment of social capital and how it affects brokered parties. Pollock, Porac and Wade (2004) introduced the concept of the network architect, and proposed that the transaction networks brokers build depend on how they resolve the trade-offs between short- and long-term considerations. The authors theorized that the resolution of these trade-offs depended on the broker's social resources, dependence on a particular market and the exogenous deal conditions. In an empirical study that built on these propositions Pollock (2004) examined when IPO underwriters deploy their social capital to the benefit of the seller or of the buyer. He showed that when the demand for shares in a particular IPO firm is high, the embeddedness of the underwriter with the institutional investors results in the lower underpricing. This means that the deal conditions will lead the

underwriter (acting as a broker) to adjust the allocation of shares to benefit the buyers with whom it has an established relationship.

Pollock's findings indicated that an important contingency affecting broker's behaviour is their relationship with brokered parties. Fernandez-Mateo (2007) provided further evidence of the importance of these relationships in inter-mediated markets. She studied the staffing sector and examined factors that affect the pricing of intermediaries' services. She showed that the relationship a broker has with the actors on one side of the market can importantly affect the outcomes for both these actors and the actors on the other side of the market. Specifically, she demonstrated that the broker's relationships with actors on one side of the market may constrain the prices that the broker can charge for its services. In her setting the staffing agency, acting as a broker, did not absorb the constraint imposed by the client firms, but rather transferred it onto the contractors. The findings in these studies underline that the relationships play an important role in brokerage in the market settings.

The preceding review of the brokerage literature reveals that this literature has generally been focused on the broker (either an individual, as in Burt, 1992; Burt, 1997; Podolny and Baron, 1997; Burt, 2004; Obstfeld, 2005; Xiao and Tsui, 2007; or an organization, as in Fernandez and Gould, 1994; Hargadon and Sutton, 1997; Khurana, 2002; Pollock, 2004). Only recently has the focus shifted towards the brokered parties. Extant evidence suggests that brokers may benefit at the expense of brokered parties (Fernandez-Mateo, 2007). This raises a fundamental, yet so far underexplored question: why do brokered parties agree to participate in brokerage? In this thesis I am addressing this question in the context of market brokerage. That is, I examine why it is that brokered parties agree to engage in a process that may end with a market exchange, and where the broker only invites

them to the process, while the eventual exchange takes place with a party represented by the broker.

Treatments of brokerage in economics emphasise that brokered parties may participate in brokerage because of the efficiency benefits that may accrue to them (Autor, 2009). These arguments are echoed by the social networks literature (Brass, 2009). However, observing decisions of brokered parties to participate in brokerage is hard, and, to the best of my knowledge, no field-based study has actually measured the decision to participate in brokerage, much less so participation in brokerage in a market context. Yet understanding how market participants, including labor market participants, come to pursue some opportunities but not others, is fundamental to our understanding of the functioning of markets, including the labor market.

In the first empirical part of the thesis (chapter 3) I am analyzing how different kinds of prior relationships between executives and the search firm affect the willingness of the executives to pursue an opportunity of a job in another organization. From the brokerage perspective, this amounts to studying when and why brokered parties agree to participate in brokerage. By addressing a gap in the labor market literature the thesis thereby simultaneously also addresses a gap in the brokerage literature.

1.3 Objectives of the thesis

Scholars interested in how relationships shape the functioning of the labor market have generated a significant body of literature that provides deep insights. The early studies have focused on understanding how information that flows through relationships affects labor market outcomes such as getting a job (Granovetter, 1973). These studies have focused on how the strength of ties between labor market participants (e.g. job seekers) and their contacts

who provide information about labor market opportunities, affect attainment in the labor market. A number of studies have been published within this stream, and it continues to be an active stream (e.g. Yakubovich, 2005). A “resource-based” stream has emerged as a reaction to mixed empirical evidence in the “information-based” stream and has de-emphasized the flow of information through ties, and emphasized the capacity of weak ties to lead to individuals who can provide valuable resources (Lin, Ensel and Vaughn, 1981a, 1981b). The third and most recent stream, the “socialization-based” stream (Fernandez, Castilla and Moore, 2000; Castilla, 2005), noted that previous research did not actually measure the flow of information or other resources through relationships, and hence could provide little conclusive evidence on the exact mechanisms that govern the generation of relationship-based benefits in the labor market. The studies in the latter stream demonstrated that the bulk of the benefits of relationship-based hiring may actually flow through post-hire socialization provided by the labor market contact.

The three streams described above share a focus on the resources that the relationships provide (either in the form of timely information, access to a well-resourced individual, or advice and support in the workplace). Their aim is to explain how these resources shape the functioning of the labor market, in particular the attainment of the individuals and the beneficial hiring outcomes for the organizations. There is very little research, though, on when labor market participants actually make decisions on the basis of the information and resources accessible through relationships. Similarly, the extant literature largely does not take into account variance in the content of relationships. As a consequence little is known on how the relationship content may affect the dynamics and the outcomes of the labor market processes where information about a labor market opportunity is obtained from a prior contact. For instance, one individual that is referred for a job may have a more or less trusting relationship with his/her referral than another individual who is also referred. If

the content of the relationships affects what decisions the labor market participants make when their contact presents them with a labor market opportunity, we need to develop a better understanding of how these decisions are made. I also consider economic implications of the differential impact of ties with different content. Prior research has explored economic consequences of relationships in the labor market (e.g. Fernandez, Castilla and Moore, 2000; Fernandez-Mateo, 2007; Bidwell and Fernandez-Mateo, 2010), but there the explanatory variables were either presence or absence of the relationship, or relationship duration. In contrast I use explanatory relationship variables that vary in their content. If ties with different content differentially affect the decisions of labor market participants, this can have implications for the distribution of economic benefits in the labor market which merit investigation.

In my study I begin to address these issues. I explore how different content of ties between candidates for executive level positions and an executive search firm affects the candidates' willingness to pursue particular job opportunities. I also examine how these ties affect the search firm's decisions with respect to candidate progress through the hiring process. Finally I assess the economic implications of the different kinds of ties, specifically their impact on the search firm's value capture in search assignments.

My argument is structured in the following way: first I identify four types of ties between executives and the search firm, based on the roles that the executive plays in interactions with the search firm. These roles are a candidate, a client, a source and a general contact. Building on the insights from the literature on trust, brokerage and boundary spanning I then develop a typology of these ties with respect to their capacity for transfer of information, and development of an executive's trust in the search firm. I then propose that among ties that have the same capacity for the transfer of relevant information, the tie that leads to the greatest development of trust will most stimulate an executive to pursue a job

opportunity introduced to him/her by the search firm (even as the economic incentive to do so declines). Since in this part of the study I am studying the candidates' decisions to pursue a job opportunity, and the candidates in that context are parties brokered by the search firm, I am therefore simultaneously examining why brokered parties agree to participate in brokerage. This part of the study thus endeavors to also contribute to the brokerage literature, and thereby to the broader literature in economic sociology and management.

I then proceed with the part of the study that is joint work with Professor Isabel Fernandez-Mateo. We take a closer look at the candidate tie. We first identify variance in the content of the candidate tie on the basis of how the candidate tie was developed. In executive search candidate ties get developed through rejections. That means that when an executive is a candidate for a job, unless he/she gets the job, he/she drops out of the process at some point. We distinguish the candidate ties that have been developed through the executives' rejections from the candidate ties that have been developed through the focal search firm's or the hiring firms' rejections. We do not make specific propositions, but rather examine whether the differences in the identity of the previously rejecting party matter for the candidates' progress in subsequent hiring processes (i.e. once the executive again becomes a candidate). We also explore whether there are any gender differences in the effect of prior rejections. This part of the study represents another way of examining how differences in the content of relationships between labor market participants and their contacts affect their decisions to pursue labor market opportunities.

In the final part of the study I move on to exploring the economic implications of ties with different content. Specifically, I examine how different kinds of ties affect the value that the search firm captures with each search assignment. By value capture I am referring to the portion of the overall created value that is kept by a given actor (in my case the portion of the overall value that is kept by the search firm). A property of my context that is important for

this question is that the search firm is paid on retainer. Retained executive search firms agree their fee with the hiring firms at the outset of the process, and the payment of the fee is not contingent on the successful completion of the process (i.e. someone being hired). I conceptualize the search firm's fee as the value that the search firm is expected to create with the search assignment (the rationale being that the search itself generates value for the hiring firm). The search firm keeps the whole value that is created, minus its costs in carrying out the search assignment. Value captured by the search firm is thus determined by the difference between the fee and its costs. While value captured in an assignment equals the profitability of an assignment (since it amounts to the difference between revenues and costs), I conceptualize it as value capture rather than profitability because positive difference between the fee and costs indicates that the search firm was able to negotiate with the hiring firm a compensation that is larger than would be merited by the search firm's costs alone. I take the perspective that to the extent that the search firm's costs can be estimated up-front, all else being equal, the difference between the fee and the costs indicates how much value the search firm was able to capture in the negotiations with the hiring firms. As the fee is a constant, and accrues to the search firm regardless of the outcome of the search, longer assignments and hence higher costs of a search mean less value is captured by the search firm. I do not make particular predictions about the relationship between tie content and value capture, but rather analyze how the results of this part of the study compare to the results of the first and the second part, and draw implications for our understanding of the role of relationships in the labor market.

1.4 Outline of the thesis

In chapter 2 I describe the key elements of the research context, executive search. I summarize how search firms obtain search assignments from the hiring firms, how the process evolves once a search firm starts a search assignment, how success is conceptualized in this context, and how the search consultants' roles shift in the course of the search assignment.

I then describe the unique dataset I am using in the thesis, and define the main variables. The dataset contains information on 47,431 individuals that were candidates for 998 executive-level vacancies during January 2005-May 2009. Although the firm is global, I use data from four of its UK offices, for reasons of completeness and comparability (more on that in chapter 2). One of the key aspects of the uniqueness of this dataset is that I can examine what happens when the roles in the relationships, and hence the content of ties, shift, while the identity of the actors remains constant. This is possible because I am studying a population whose income is within the top 1% of the income distribution in the UK. Members of this sector of society occupy positions from which they can act as clients to the search firms as well as sources, which represent roles that are very different to the candidate roles that executives have when they compete for jobs. This population therefore exhibits changes in roles, and as such allows a window into previously underexplored, yet theoretically important, phenomena shaped by differences in tie content.

Chapters 3, 4 and 5 constitute the empirical chapters of the thesis. In chapter 3 I address the question of whether and how prior general, client, source and candidate ties affect the executives' willingness to pursue a job opportunity. In chapter 4, which is a joint work with Professor Isabel Fernandez-Mateo, we study how prior candidate ties built on rejections of different parties (candidates and the search firm/the hiring firms) affect the candidates'

willingness to pursue a job opportunity, and the search firm's decisions to pursue particular candidates. In particular we explore how prior relationships differentially affect these decisions for men and women. In chapter 5, I analyze how the ties studied in chapters 3 and 4 affect the value that the search firm captures in the search process. In chapter 6 I provide a summary of the findings, describe the contributions, indicate the limitations and outline the directions for further research on social relationships in the labor market, in particular the very high end of it.

CHAPTER 2: RESEARCH SETTING AND DATA DESCRIPTION

2.1 Research context: executive search

The context of this thesis is executive search. Executive search is a process of identifying and recruiting individuals for the executive level positions. This process is managed by executive search firms, the professional service firms that are managing the process of filling in the executive positions. The search firms are hired and paid by the firms who are trying to hire executives from the outside. I analyze the brokerage process in the context of retained executive search. Retained executive search firms agree their fee with the hiring firms at the outset of the process, and the payment of the fee is not contingent on the successful completion of the process (i.e. someone being hired). Retainer agreement also means that the focal search firm is the only search firm authorized to approach candidates for the position (i.e. only one search firm at the time works on a given vacancy).

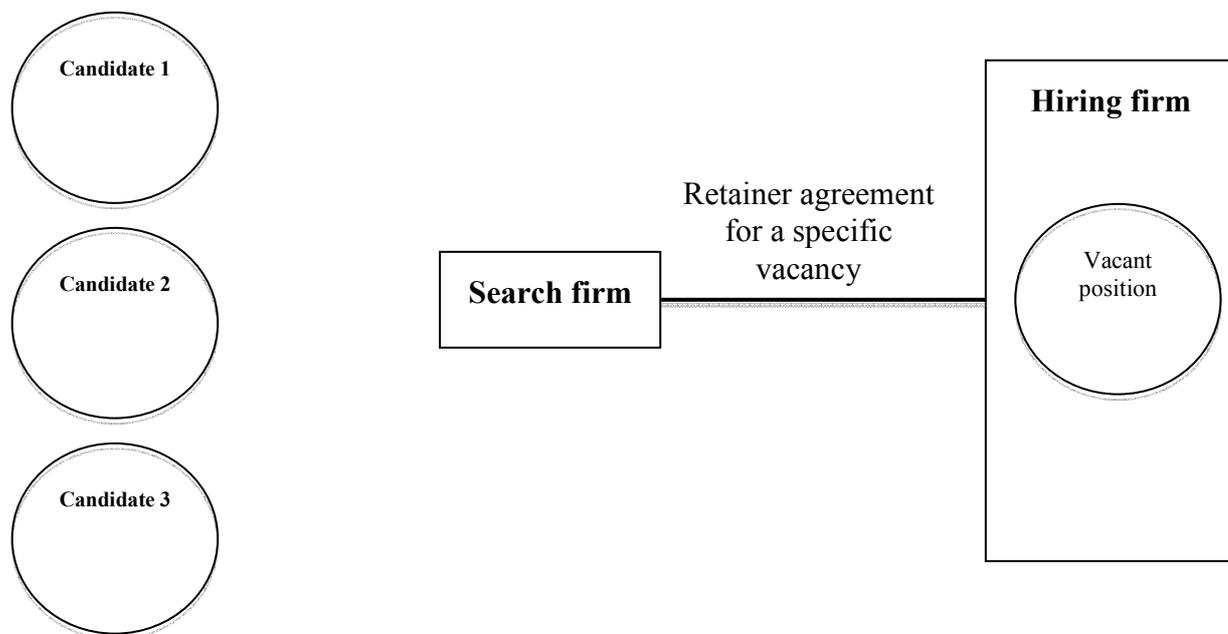
Executive search takes place in the executive labor market. One of the key properties of this market is that it is triadic in the sense that the hiring process is not a direct exchange between candidates for executive jobs and the hiring firms, but rather one that is mediated by the search firms. Mediation in this context means that the job candidates are interviewed and assessed by the search firm before they are interviewed by the hiring firms. In the US it has been estimated that executive search firms are involved in the recruitment of 35 percent of all managerial hires at salary levels higher than \$200,000 (McCool, 2008, p. 22).

Studies of this labor market (Khurana, 2002; Finlay and Coverdill, 2002) have often conceptualized its triadic structure as one where the search firm plays a role of the broker (while executives and the hiring firms are the brokered parties). Figure 2.1 presents the triadic structure of the executive labor market. Specifically, it depicts a situation where a hiring firm

has an executive-level vacancy, and has engaged an executive search firm who in turn has drawn up a list of executives, and thereby declared them candidates for the vacant position.

Figure 2.1: Triadic structure of the executive labor market

Consideration set



2.1.1 Search process

To gain an understanding of how the search firms go about including executives as candidates in the consideration sets for the vacancies, and an understanding of the process that follows until one of them is hired I conducted a series of interviews with the search consultants, clients (i.e. executives who have used services of the search firms in the past), and candidates (i.e. executives who were candidates in a recruitment process managed by a search firm). Specifically, I conducted 25 formal interviews. Interviews were conducted either in person or over the phone, and lasted between 45-90 minutes, with most of them lasting for 60 minutes. I took notes and recorded and then transcribed the interviews. The details of the interviewees are provided in table 2.1. In this table, in the column entitled “Format of the interview” I have marked with an asterisk the instances where I did not record the interview because the interviewee asked for it not to be recorded.

Table 2.1: Interviews

Interviewee number	Interviewee role	Gender	Type of firm	Industry	Date of the interview	Format of the interview
<i>Search consultants</i>						
1	Partner	M	Boutique	Investment banking	02/10/2008	In person*
2	Practice Leader	M	Global firm	Technology - board level searches	15/10/2008 14/01/2009	Phone In person*
3	Partner	F	Boutique	Real estate - board level searches	21/10/2008	Phone
4	Consultant	F	Global firm	Life sciences	22/10/2008	In person*
5	CEO	M	Global firm	Board level searches	27/10/2008	In person
6	Consultant	F	Global firm	Financial services	28/10/2008	Phone*
7	Author/consultant	F	Sole proprietor	General	04/11/2008	In person*
8	Managing director/founder	F	Boutique	Family business	10/11/2008	Phone*
9	Office leader	M	Global firm	Board level searches	11/11/2008	In person
10	CEO/founder	M	Boutique	Financial services - board level searches	14/11/2008	Phone
11	Consultant/researcher	F	Global firm	Board level searches	18/11/2008	In person*
12	Partner	M	Global firm (Execo)	Financial function - board level searches	07/01/2009	In person
13	Account manager	F	Global firm	General	28/01/2009	Phone
14	Head of research	F	Global firm (Execo)	General	09/02/2009 15/05/2009	In person In person
15	Database manager**	M	Global firm (Execo)	General	09/02/2009 15/05/2009 20/05/2009 22/09/2009	In person In person In person In person
16	Senior associate**	F	Global firm (Execo)	General	11/11/2009 16/06/2010	In person In person
<i>Candidates and clients</i>						
17	CEO*** (interviewed as both a candidate and a client)	M	Global firm	Software	05/11/2008	In person
18	HR director*** (interviewed as both a candidate and a client)	M	Global firm	Software	06/11/2008	In person
19	HR director (interviewed as a client)	M	Global firm	Insurance	16/12/2008	In person*

Notes:

* - did not record the interview as the interviewee asked not to be recorded

** - I followed up the in-person interviews over the phone and by e-mail

*** - interviewed as both a candidate and a client

I conducted 25 interviews with 19 individuals. All interviews but one were conducted between October 2008 and November 2009. 9 interviewees out of 19 were female. Out of 25 interviews, 22 were with search consultants. I define a search consultant broadly, and include in this category professionals working in executive search who have been actively conducting searches either at the time of the interview or in the past, as well as those professionals who have provided support services (such as research, database management etc.). Out of the 22 interviews, 9 have been conducted with 4 individuals from Execo (the search firm that provided the quantitative data for my study). These four individuals were, respectively: a partner (1 interview), head of research (2 interviews), database manager (4 interviews), and a senior associate (i.e. a consultant, 2 interviews). With the latter two individuals I conducted a number of follow-ups over the telephone and by e-mail. The primary goal of the interviews with search consultants from Execo was to understand what the various categories in its database mean, and how they are linked to the activities in the searches. I have thus incorporated the information obtained from those interviews in the definitions of variables and categories that I use in the quantitative analyses. The interviews with the senior associate (person number 16) were also useful for the illumination of the mechanisms that govern the relationships between the content of ties and the candidates' decision to pursue a job opportunity (in the empirical context, go for the search firm interview). I include quotes from those interviews in chapter 3.

The primary goal of the interviews with search consultants from firms other than Execo, though, was to obtain information on the search process in general. This was also the goal of the interviews with candidates (2 interviews) and clients (3 interviews; 2 interviews were with executives I interviewed also as candidates). An additional goal of the interviews with candidates and clients was to compare their view of the process with that of the search consultants.

I adopted several approaches to secure interviews. One approach was using London Business School's Portal. I sent an e-mail to a number of individuals who are identified in the School's directory as working in the executive search sector. Another approach was to ask London Business School Faculty to refer me to individuals in this sector. The interviews with the search consultants in firms other than Execo (interviewees 1-11 and 13), as well as the candidates/clients (interviewees 17 and 18) were arranged in these two ways. Upon conducting the initial interviews, I asked interviewees to refer me to their colleagues or other individuals who they felt could provide useful information for my research. This "snowballing" approach led me to interviews with the search consultants in Execo (interviewees 12, 14, 15 and 16). In November 2008, I attended the conference of Association of Executive Search Consultants (AESC), a global association representing the industry. There, I engaged in a number of informal discussions about the industry and the search process. In this way I secured the interview with interviewee 19.

Before starting the interview phase I researched the executive search industry in the UK. I also collected and analyzed information from the firms' web pages, and read articles and books written by the practitioners. At the same time I reviewed descriptive academic literature on executive search (Khurana, 2002; Finlay and Coverdill, 2002). These activities led me to an initial perception of the industry and the process. I developed my interview protocol on the basis of these initial perceptions. I opted for a semi-structured interview protocol. The topics I asked my interviewees about were the same, yet I often included additional questions, depending on the interviewee. When interviewing candidates and clients I adapted the questions accordingly.

I explored the following themes:

- search assignment sourcing
- candidate sourcing

- the search process
- definitions of success and measures of performance
- search consultants' roles

Below, I discuss the rationale for exploring these themes, and provide an overview of my key findings. Along with the findings I provide quotes from the interviewees. Where a quote is a short one, I embed it in the text. Where it is a longer one, I provide it in a separate paragraph.

Search assignment sourcing

In this theme, I focused on exploring the factors driving the decision to hire executive search firms, and the decision to hire a particular consultant/firm. The first aspect is relevant for the expectations of the hiring client, as these may importantly shape the process. The second aspect is important for understanding the factors of differentiation, which may then translate into the factors of competitive advantage in that setting. With respect to the first point, the consensus emerged that search consultants tend to be engaged for three reasons (any combination of them): first, the hiring firm does not have the resources to conduct the search on its own; second, search consultants have greater knowledge of, and access to, the candidates; third, the hiring firm is intent on appointing an internal candidate, but wants to benchmark that candidate against external candidates.

Once the hiring firm (i.e. the firm with the executive-level vacancy) makes a decision to hire an executive search firm, the question is which consultant to engage. Again there was a consensus that there are three broad factors that tend to affect that decision, and their relevance varies with respect to the level of the market (i.e. CEO, CFO etc. vs. middle-management positions). The first factor is individual networks. The choice of a particular

consultant hinges overwhelmingly on the relationships of the hiring client (the individual). This client will tend to choose a consultant with whom he or she has an established relationship. Since the executive search process is fundamentally a process occurring between individuals, the key component in the choice is inter-individual trust. The second factor is firm reputation. This factor contributes the insurance; should the placement go wrong the hiring client has a plausible defence in that he or she chose a reputable search firm which should guarantee a positive outcome. The third factor is the individual consultant's track record in a given market, sector or function. Since the executive labor market is characterized by a relatively high level of mobility and consequentially high levels of change, the recency of knowledge about the market plays an important role. Experience may exhibit a curvilinear relationship with attractiveness of a particular consultant. In the words of a partner at a boutique search firm who specializes in board-level searches in the real-estate industry (interviewee 3): "you want someone who is still thoughtful and keen, but at the same time has got experience". Search consultants agreed that it is difficult to say what affects the relative importance of the different factors. There did seem to be a consensus, though, that at higher levels of the executive market (e.g. board level) the overwhelming factor is the individual relationships between the hiring committee and the search consultants. At the lower levels, firm reputation and the consultants' knowledge of a particular market/function may play more of a role.

An interesting component of the decision to hire a particular firm/consultant is the extent to which the selection is made on the basis of a process. When a hiring firm does not have a preferred supplier of search services, or wants to compare the existing provider to other providers, it may opt for a "pitching" exercise. The pitching exercise is a competition between search providers. When a firm has a vacancy it may invite a number of search providers to "pitch" the way in which they would conduct the process to fill that vacancy.

The chosen provider then signs a retainer agreement, which means that this provider is the sole agent in charge of conducting the search process on behalf of the hiring firm.

Candidate sourcing

One of the distinguishing characteristics of the executive labor market, relative to other labor markets, is that individuals usually do not appear as sellers in the market, but rather have to be “converted” into sellers by the search consultants. The latter justify their conversion activities by claiming that they are “opportunity introducers”, who try to increase efficiency of the executive labor market by introducing the right opportunities to the right individuals. CEO of a global search firm (interviewee 5) commented: “sensible people want to keep their eyes open to what the opportunities are”. Search consultants, therefore, scan the executive market landscape to uncover “latent candidates”, individuals who may be interested in departing their current position should an appropriate alternative be introduced to them.

In terms of the technology of candidate sourcing, consultants use a number of approaches. They assemble an initial list of candidates on the basis of the information in their databases, which includes information on candidates in the previous searches. Another approach is “market mapping”, which means identifying individuals who are, at a given point in time, performing the role for which the consultant was retained. To map the market, consultants use both publicly available information and information they obtain from their networks of contacts in the industry. After this identification phase, consultants engage in an interaction with the identified candidates. This can take the form of telephone conversations, conducted by either the consultants themselves or in-house researchers (this generally depends on the type of the firm – in global firms there would be more division of labor, while in boutiques most activities, including the initial screening, tend to be performed by the

consultants themselves). Candidates who express interest may then be formally interviewed by the search firm. After the interviews the search firm draws up a “short-list” of candidates to be presented to the hiring firm. This aspect means that search firms act as an important filter in the labor market as it is they who ultimately decide who among the large number of potentially suitable candidates will be among the few introduced to the hiring firms.

The search process

As the managing director of the London office of a global search firm (interviewee 9) put it: “The search process starts with a question (a business problem), and it ends with an answer (a person).” This quote highlights a belief held by the search consultants that recruitment at the executive level is less about filling a vacancy and more about finding a person that has the capacity to solve a particular business problem of the hiring organization. This has important implications for considerations of success in executive search, to which I turn later. Here, I describe the stages in executive search.

The first stage in search is the definition of the business problem that hiring a new executive might solve. A job specification is drafted on the basis of this business problem. Generally, clients draft the first version, which is then co-developed with the search consultants. The extent of the consultants’ involvement in the development of the job specification varies widely. This involvement has consistently been emphasised by interviewees as a factor that both shapes the candidate pool and the outcomes of the search process. A practice leader for a global search firm (interviewee 2) summarized it in the following way: “the more time and effort you invest in that early stage of the process the more likely you are to get a successful outcome at the end”. A HR manager for a software firm that regularly works with the search consultants (interviewee 18) agreed that the

definition of the brief is important with respect to the likely outcomes: “I don’t think consultants push back on the brief enough”. To this he added that he wants a “brutally honest discussion about the brief”.

The second stage is the sourcing of candidates. While the identification of potential candidates is a relatively straightforward task, the conversion of these individuals into actual candidates requires a persistent and thoughtful activity on the part of the search consultants. The number of interested candidates is to some extent a result of the search consultants’ sourcing activities. It also rests on the consultants’ ability to understand who will “connect with” a certain job description.

Once candidates express interest, they are interviewed and assessed by the search consultants, and often by the assessment specialists. The assessment also includes taking references, which essentially involves obtaining opinion on the candidates by the referral sources nominated and not nominated by the candidates. The outcome of this part of the process is a short-list of candidates who are then introduced to the hiring firm. This part of the process rests heavily on the consultants’ judgement and their ability to screen different candidates.

The assessment phase is followed by the placement phase, where short-listed candidates are introduced to the hiring committee in the hiring firm, who then interview the candidates and (normally) make offers. The consultants’ involvement at this stage reduces somewhat, as the interviewing and the assessment are conducted by the hiring committee. However, once a preferred candidate is identified, the consultants’ involvement intensifies. At that point the consultant attempts to create conditions for a successful closure of the process. This can involve a number of activities, such as organizing informal events that involve the members of the hiring committee and the candidate, along with their significant others. This is supposed to increase the “feel-good” factor of the process.

In terms of the process itself, consultants at that stage try to reinforce the logic of the match-up between the preferred candidate and the hiring firm. Once the offer is made, consultants work on both sides to iron out a mutually acceptable compensation package. They also try to put the candidate off from accepting a counter-offer which may be issued by his or her current employer. On the issue of compensation negotiations, two opposing views emerged: the HR directors' view, which contends that a consultant's involvement in the compensation negotiations creates confusion about whose side the search consultants are on (as they should from the start to the finish of the process act as a representative of the client); and the candidates' view, which emphasises two advantages that the consultants' involvement brings: knowledge of the competitive compensation, and the capacity to diffuse potential tensions.

A CEO of a global software firm (interviewee 17) described his view of the first point from the perspective of a candidate:

"I'm going in as their favourite candidate to see my prospective boss to talk about pay, terms, conditions, the usual. Headhunter working as intermediary. And I am quite comfortable; I haven't yet been offered a job, and he is going to say I'm giving you this for 30p a year and I think oh that was a bit disappointing. Now what do we do? Do I walk out? So I ask her can you help and she says absolutely. And then we parted harmoniously, and she was on the phone saying they will make you an offer, and would you consider this number (obviously higher than 30p). She was able to go through not prolonged and acrimonious negotiation but she got me to move it up, and she has incentive to move it up because that increases her amount she takes, so it's useful to have her in the loop. But when she said I think this is probably the highest we can push it I was ready to listen to her. Because she's been around the block and she knows what good looks like. And she's got a good idea of leveraging of roles etc. Etc. So she was able to bring back that conversation quickly and harmoniously to a conclusion and then get it to a point where we actually signed the contract. Which is ultimately where she wants to get to."

Continuing with elaboration of the second point, he added:

"It's a pretty sensitive thing. The stakes have got higher and higher because you've both invested in this relationship and you're trying to get to a point where they are recruiting "I want you to feel really positive about the company, what we're going to do, how exciting it is going to be, the great team that you'll be joining". And then we kind of break all the paths when we get down to how much you're worth. And uh, that's a big different conversation to have. So it is really good to be able to, someone who's recruited, to be able to do all of that."

While the goal of the hiring firms is to place someone in a position, not all commenced searches end with a placement. Search consultants estimated that the rate of searches that are discontinued is around 15-20% (and the HR director of a global insurance firm (interviewee 19) argued that this is quite a large number). The search consultants cited two common reasons: first, mobility of the client (the hiring client leaves the position and a new person comes in), and second, occurrence of major events unanticipated at the onset of searches (e.g. firm takeover). Another often-cited factor was the consultants' inability or unwillingness to resolve differences of opinions within the hiring team. Practice leader of a global search firm (interviewee 2) described this situation in the following way:

“... one client really likes one candidate but another client really likes another candidate, not just on chemistry which of course can happen anyway, but because they are actually looking for different characteristics, or are at different stages in their career, etc. ...”.

According to several search consultants once a placement has been successfully made they tend to keep in touch with the candidate to make sure he/she is settling in well. Many large firms offer an on-boarding service, however, this service is generally not a part of the search process, and is procured as a separate service. There was some indication, though, that more methodological follow up may have some positive consequences, such as better perception on the part of the client of the professionalism of the process. As a founder of a boutique consultancy (interviewee 10) remarked, “it is one definition of professionalism if you institutionalize the follow-up”.

Definitions of success and measures of performance

In terms of defining what constitutes success in this setting, search consultants overwhelmingly emphasised that a successful search is one that results in a sustainable

placement, i.e. a placement where the candidate stays in the job for as long as was expected. Early departure, by extension, would therefore imply an unsuccessful search. The sustainability of placements is a measure of success for search consultants (as it is likely to affect repeat business), but it may also be a measure of success for clients, since early departures may involve additional costs. Several consultants mentioned that, as a rule, when the tenure of the placed executive is longer than three years it might be considered a success. Practice leader of a global search firm (interviewee 2), however, cautioned that the desired tenure may vary widely across sectors, and may be contingent on the particular situation.

A metric mentioned by the consultants and the managers of search firms was client satisfaction, which is generally measured with a survey (this metric was also mentioned by Execo's head of research (interviewee 14)). Consultants emphasised that there may be differences in client satisfaction with respect to the process and the outcome. Clients may be satisfied with the process, but not satisfied with the outcome. This may be the case when the preferred candidate exits the process at a late stage, or decides to accept a counter-offer. Clients may also be satisfied with the outcome, but not with the process. This may occur when a preferred candidate accepts the offer, yet this is not a consequence of the consultants' activities. Different combinations of scores on client satisfaction dimensions may have different implications for repeat business. The performance dimension on which consultants' opinions converged was the speed of the search. Consultants thought that speed is a desired component, yet may not be the best indicator of whether a search was successful. They thought that a focus on speed may lead to compromises in professionalism during the process, which may in turn have negative long-term consequences (such as reducing the sustainability of placements). This view was shared by the HR directors. HR director of a global software firm (interviewee 18) remarked that "you can't compromise on time. But, if they cannot do it

in 14 weeks, they will never be able to do it. If they do have a good network, they can do it quickly and well.”

Search consultants' roles

As implied in the preceding text, search consultants importantly influence outcomes of searches, and potentially the processes that take place after the search (such as subsequent executive turnover). Still, their impact is of the influencing rather than determining kind. A founder of a boutique search firm (interviewee 10) succinctly summarized this point: “you can’t make things happen, you can’t make someone take the job”. Partner at Execo (interviewee 12) provided a similar remark “all we can do is reinforce the logic, and tell them where the role could take them to in their career”. A consultants’ role is also introduction of objectivity into the process. Even though they are hired and paid by the clients they see themselves as a third party who needs to intervene when clients and candidates want to do something against their better judgement. The role of the objective third party also brings transparency into the process. This part of the role includes soliciting candidates’ and clients’ private thoughts, and relaying them to the opposite party in a way that facilitates the process. While in the actual process their role is one of an influencer, in the general business arena their position may be seen as far more powerful. As a group they are seen as “well connected”, as mentioned by the CEO of a global software firm (interviewee 17). The exact nature of their “connectedness” does not seem to be readily apparent or understood. This may increase their power in the social system, as evidenced by the quote of the same interviewee (17): “you don’t want to leave any scorched earth between you and one of these captains of the network”.

2.2 Quantitative data

As I describe in chapter 1, research on the role of relationships in the labor market lacks a systematic quantitative study that would examine how the content of the relationships between labor market participants and their labor market contacts shapes which labor market opportunities the former pursue. Executive search firms continually interact with executives, and, as the information I obtained from the interviews indicates the same executive may play a different role at different points in time (e.g. he/she may be a client of the search firm at time $t=1$, and a candidate for a vacancy at another firm at time $t=2$), this leads to the development of ties with different content within a single relationship between an executive and the search firm. Executive search therefore represents a suitable context to study the overall question.

The ultimate objective of my interviewing, hence, was to obtain access to quantitative data that would allow me to conduct a study that would address the gap outlined above. As I interviewed a partner of Execo he indicated that the firm collects data on all of its interactions with the executives, and that it would be willing to provide access to it. The agreement of data access was reached in February 2009, and after having signed the confidentiality agreement I obtained data from the electronic records of Execo in May 2009¹.

An important property of data from a large global firm is that the firm does searches in a number of industries, and across a range of jobs/positions. Such a data source provides two benefits: first, in executive search the majority of candidates that are put on the initial lists are never contacted; a large firm contacts enough individuals to generate a sample that is suitable for statistical analysis. Second, since data come from a number of industries the results are not affected by the specifics of any particular industry. I asked Execo's search

¹ The firm's condition to provide data was that it remains anonymous, so I unfortunately cannot provide more detail about it.

consultants to assess to what extent their processes correspond to the general description of the search process I previously developed on the basis of interviews with search consultants from other firms. They confirmed that the process that Execo employs to conduct searches corresponds to the process employed by the retained executive search firms in general.

Even though the firm conducts searches globally, for this thesis I use data from four of its UK offices. The firm's data manager and the head of research suggested that this data is the most complete and the most likely to be reliable. The dataset I use in the thesis contains information on 998 search assignments that the search firm conducted between January 2005 and May 2009.

Execo's database contains data for 1,053 vacancies for the period of observation. Execo's consultants explained in the interviews that even though Execo constructed a consideration set for all 1,053 vacancies, some of them didn't start as Execo eventually was either not retained for the assignment, or the assignment was taken away from it because the hiring firm discontinued the search, or the search was given to another search firm. Unfortunately the database did not include variables that would allow me to reliably distinguish between these scenarios, so I assumed that if Execo did not interview anyone in the consideration set, or if there was a high proportion of candidates where no contact with them was made the assignment, in all likelihood, didn't start. As I am interested in studying the progress of candidates through the hiring pipeline, I dropped those vacancies where this process cannot be observed. Specifically, I dropped those vacancies where the outcome for all observations was unknown, where Execo did not interview anyone and where the proportion of unknown observations was more than 75% and no one was not interested (i.e. Execo probably didn't speak to anyone in the consideration set). I also dropped those vacancies where more than 90% of the observations had an unknown outcome, somebody was not interested (i.e. Execo did speak to them), and the vacancy was coded as "lost" in the

database (i.e. the search assignment was taken away from Execo at some point). As a consequence 55 vacancies (1,526 observations) were dropped.

47,431 individuals were included in the consideration sets for the 998 assignments. Since some individuals were considered for more than one vacancy, the dataset consists of 64,505 observations. The relationship data is based on the contacts between executives and the search firm recorded in the search firm's database. These contacts span the period between January 2001 and May 2009 (as the search firm has been collecting the relationship data since January 2001). The assignment level data and most of the individual level data spans the period January 2005-May 2009 as the search firm started organizing data by assignments in January 2005. The data is collected for the period up to May 2009 as in that month the search firm provided me the snapshot of its electronic records. The period of observation for the analyses in the thesis is therefore January 2005 – May 2009.

The first step in the executive search process is the construction of the consideration set. A consideration set is an initial list of candidates for a vacancy. There are two ways in which candidates enter the consideration set: by being included on the list by the search firm, and by responding to an advertisement posted by the search firm. In the first case search consultants and researchers use information in the database, and information obtained through their contacts, to identify individuals whose qualifications correspond to the qualifications required by the vacancy. 80.25% of observations in the dataset correspond to the candidates who have been put on the list by the search firm. 19.75% of observations correspond to the candidates who entered the initial lists by responding to an advertisement.

Once the consideration set is constructed, the search firm contacts some individuals on this list. These contacts are normally made over the phone, and/or in person. The initial contacts constitute transition to the next step in the hiring pipeline: the search firm interview. The purpose of these contacts is to identify candidates who are qualified and willing to go for

an interview, as well as to persuade candidates that the search firm wants to put forward to go for an interview. A quote from an Execo senior associate (interviewee 16) underlines that the firm tries to persuade the candidates who most closely correspond to the hiring firm requirements to go for a search firm interview:

“We try hardest to get the best candidates, always. And it may be that those candidates that we think tick all those boxes might not want to leave their company, but we will try hard to convince them.”

There are three possible outcomes when moving from step 0 (the consideration set) to the step 1 (the search firm interview): the candidate rejects going for the search firm interview, the search firm declares the candidate unsuitable (before or after having spoken to him/her), and the candidate goes for the search firm interview. 11,188 candidates in the dataset (17.34% of the total number of observations in that step, 64,505) have rejected going for the search firm interview, 34,999 (54.25%) were declared unsuitable by the search firm, and 7,972 (12.35%) went for the search firm interview. For 10,346 candidates in the consideration set (16.03%) the outcome in step 1 was unknown. I explain how I treated the unknown observations below.

The next step after the search firm interview is the hiring firm interview (step 2). There are three possible outcomes in the transition from step 1 to step 2, and they mirror the outcomes of transition from step 0 to step 1. Out of 7,972 candidates that made it to step 1 (search firm interview) 695 candidates (8.71%) rejected going for the hiring firm interview, 2,842 (35.65%) were declared unsuitable by the search firm, and 3,978 (49.89%) went for the interview with the hiring firm. For 457 observations (5.73%) the outcome of going from step 1 to step 2 was unknown.

Step 3 is the job offer. 476 candidates (11.96% of 3,978 observations where candidates were interviewed by the hiring firm) did not want to consider a job offer, 2,500

candidates (62.84%) were declared unsuitable by the hiring firm, while 570 candidates (14.32%) received a job offer. For 432 candidates (10.85%) the outcome of going from step 2 to step 3 was unknown. Final step (step 4) is candidate placement into the hiring firm. In the transition from step 3 (a job offer) to step 4 (candidate placement) there are two possible outcomes: the candidate rejects the offer, or the candidate accepts the offer (is placed with the hiring firm). 75 candidates (13.16% of the 570 observations where candidates received an offer) rejected the offer, while 479 (84.03%) accepted it. For 16 candidates (2.80%) the outcome was unknown.

I coded the observations as “unknown”² for a particular step if either of the following occurred: the candidate had more than one observation in the step coded 1 (i.e. as if the candidate was for instance both not suitable and not interested in that step); or that he/she was coded as having progressed to the step (e.g. was included in the consideration set; or was interviewed by Execo), but then had all outcomes for the step and all steps beyond coded as 0 (i.e. as if none of the outcomes materialized). In short, the unknown observations are those where it is unclear what (if anything) happened with the candidate in a given step, and beyond. Since the proportion of the unknown observations, especially for step 1, was large, I first assessed the reasons for which an outcome is coded as unknown.

The first scenario, which applied to a very small proportion of all unknown observations, can be explained in a straightforward manner as a mistake made when Execo

² I conducted a separate analysis of the characteristics of the vacancies and individual characteristics of the candidates for whom the outcome was unknown in going from step 0 to step 1 (10,346). Vacancies where for more than 75% of candidates in the consideration set the outcome of that step is unknown are high salary vacancies (paying on average around £187,000, compared to the £157,000 paid by vacancies with less than 75% of candidates with the unknown outcomes). Those vacancies are also characterized by a stronger prior relationship between Execo and the hiring firm (2.6 prior search assignments for those vacancies, compared to 1.4 for vacancies with less than 75% of the observations resulting in the unknown outcome). Neither of these two differences was statistically significant, though. In terms of individual characteristics, the observations with the unknown outcome are associated with more missing individual data (significant difference in the chi-square tests). Among those candidates for whom I have individual data, those with the unknown outcomes are older, less likely to have an MBA, and have higher salaries when put in the consideration set. The proportion of women is the same in the unknown and the known group, and there is no difference in tenure on the current job. The individuals in the unknown group are therefore both more qualified (on the basis of the higher salary), and less (on the basis of the MBA), so it does not seem that the outcome for them is unknown due to their “quality”.

was entering information into the database, since a candidate cannot be both not interested and not suitable at the same time (also he/she cannot both drop out of the process because of rejection, and then be interviewed). The second scenario is somewhat more complex. It can be partially explained by the proportion of candidates that get interviewed out of the candidates included in the consideration set. The average number of candidates in the consideration set in the dataset is 64. The average number of candidates for a vacancy interviewed by the search firm is only 8, though (12.5% of the consideration set). If the search firm is, on average, looking to generate a set of 8 candidates that it interviews, and it stops having preliminary conversations with candidates once it has generated this set, a large number of candidates are bound not to get contacted. That should not automatically mean that the non-contacted candidates were deemed not suitable by the search firm; they may not have been “processed” simply because of the small ratio of interviewed candidates to the considered candidates. Execo unfortunately does not use a specific code that would allow me to distinguish among the reasons for why there was no information on the candidate once he/she was included in the dataset. To refrain from imposing any assumptions of the reasons for no information I coded these observations as having an unknown outcome.

I considered several ways in which these observations could be incorporated in the analysis. A procedure that would correct for selection bias was not feasible in this case because selection into the group of unknown outcomes, and selection into the outcomes captured by the key dependent variables (not interested and not suitable; see an explanation of these variables later in the chapter) takes place at the same time. That is, in this context it is not that first some observations take on the unknown outcome and some don't, and afterwards the outcomes of interest are observed among those with the “known” outcomes.

Not being able to apply a standard procedure that would correct for a potential selection bias, I conducted separate analyses to verify whether the occurrence of the unknown

observations was significantly predicted by some attribute of the candidates, their relationships with Execo, or characteristics of the vacancy. The purpose of these analyses was to identify whether selection into the unknown category is significantly predicted by some of those variables. I hence ran logit models with the four unknown variables (unknown in step 1, step 2, step 3 and step 4) as the dependent variables. I included as the independent variables individual attribute variables, relationship variables and vacancy variables that I describe later in this chapter.

I first ran the logit models for the overall sample (N=64,505). This is also the sample that is used for the analyses in Chapter 4. Results indicated that for female candidates and candidates who applied for a vacancy the outcome in step 1 of the process is less likely to be unknown. It was more likely to be unknown if the candidates had more non-candidate contacts with Execo. Among the variables of interest in this thesis (that I describe in detail later in the chapter) the outcome was less likely to be unknown when candidates were rejected for more jobs in the past. Having applied for the vacancy also significantly predicted the unknown outcome in steps 2 and 3 (lowered the probability). In contrast with step 1, a candidate was less rather than more likely to be unknown in step 4 when he/she had more non-candidate contacts with Execo.

I repeated the analyses for the samples that I use in Chapter 5 (N=51,764 - only candidates included in the consideration set by the search firm; the rationale provided in Chapter 5), and Chapter 3 (N=6,308 - candidates for whom I had the full set of individual characteristics; more information about this sample in section 2.2.2). In the sample used in Chapter 5 the outcome for step 1 was more likely to be unknown when the salary was higher, and less likely when the assignment was advertised (which is consistent with previous results obtained with t-tests). Among the variables of interest, the outcome was less likely to be unknown when candidates rejected more, or were rejected more in the past. It was more

likely if the candidates had prior client ties with the search firm. I didn't conduct the analyses for the unknown outcomes in steps 2, 3 and 4 for this sample as in this chapter I only analyze outcomes in step 1.

In the sample used in Chapter 3 the outcome for step 1 was less likely to be unknown for older candidates, and for those that applied for vacancies. Among the variables of interest in the chapter, the unknown outcome was more likely for candidates who had a prior client or general tie with the search firm. In step 2 the unknown outcome was more likely when the candidates had a prior source tie, and if they applied for the vacancies. In step 3 the outcome was less likely to be unknown for women, and for candidates with an MBA. None of the variables had a significant relationship with the unknown outcome in step 4.

Since I had no information that would allow me to code the unknown outcomes as either not interested or not suitable (i.e. rejection by the candidate or rejection by the search firm), I conducted robustness checks by conducting the analyses in the thesis assuming first that all unknown outcomes are in fact outcomes where the candidate rejected the search firm or the hiring firm, and second, that they are in fact outcomes where the search firm or the hiring firm rejected the candidate. These assumptions are extreme, and predict unlikely scenarios, as in reality it is most likely that some of the unknown outcomes were in fact candidate rejections, and some the search firm/hiring firm rejections.

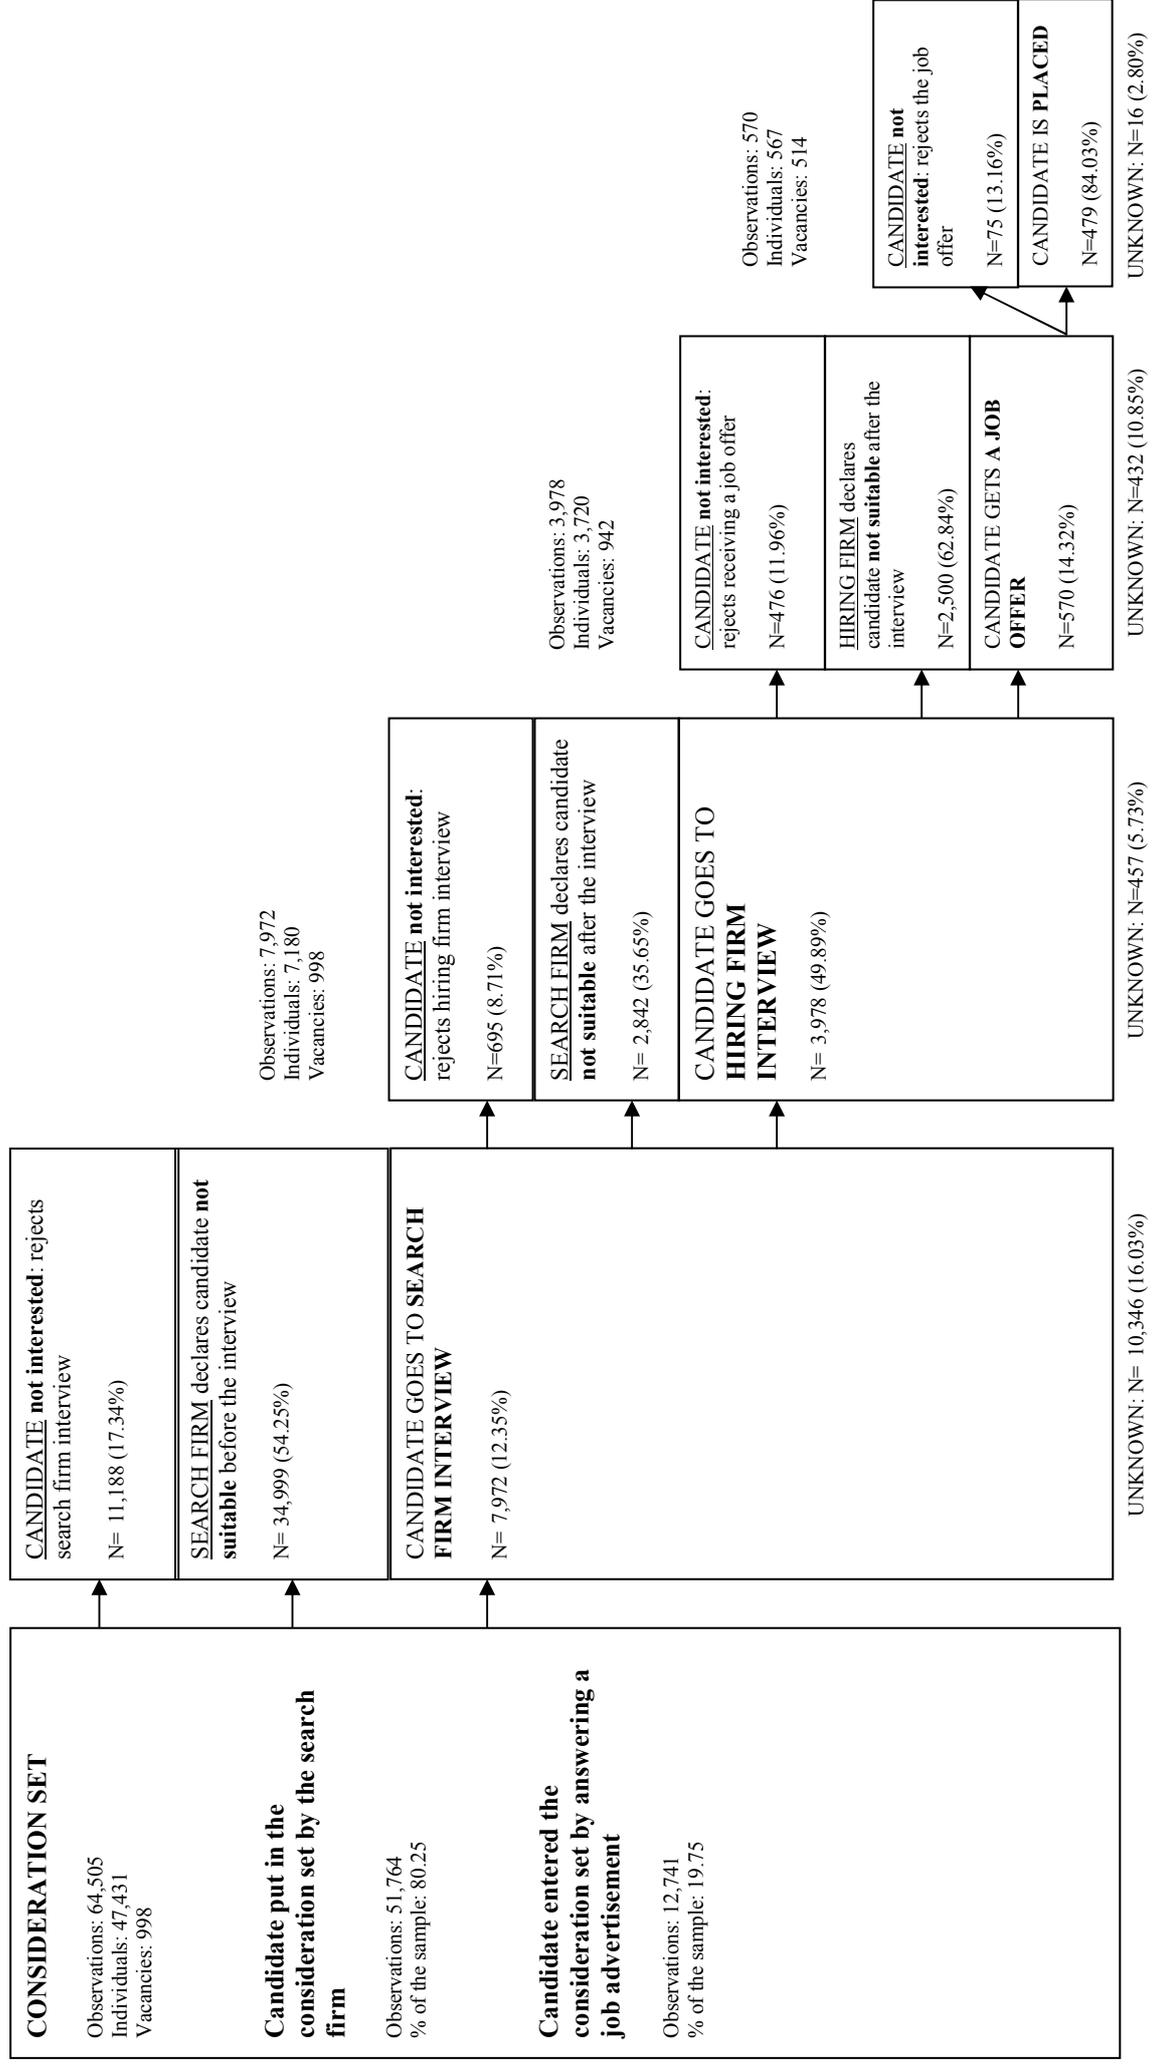
I first conducted the analyses with the assumption that all unknown outcomes were in fact candidate rejections. The results were unaffected, apart from the interaction between salary ratio and general tie in Model 4 in Table 3.2 in chapter 3, which became insignificant (see chapter 3 for more detail). I then conducted the analyses with the assumption that all unknown outcomes were in fact search firm/hiring firm rejections. No substantive results changed. Most results therefore remained robust to different assumptions. The result mentioned above did change, however only when the assumption of unknown outcomes

actually being candidate rejections was employed. I have no information that would suggest that this assumption should be preferred to the assumption that unknown outcomes were actually search firm/hiring firm rejections of candidates.

Since I had no strong reason to adopt one or the other assumption, I decided to adopt a conservative approach and not impose any assumptions about the unknown observations, but rather treat them as missing data. I hence conducted the analyses in the thesis using only observations for which the outcome was "known" (i.e. either not interested, not suitable or interview/offer/accepted offer). However, I still used the information about the unknown outcomes to create a past relationship variable (the number of times a candidate had an unknown outcome in the past), and used it in the analyses in chapters 4 and 5 as a control.

Figure 2.2 below schematically represents the hiring pipeline at Execo, and thereby summarizes the discussion in this section (2.2).

Figure 2.2: Hiring pipeline at Execo



2.2.1 Dependent variables

The datasets so far have enabled the researchers to observe who progresses through the hiring pipeline. But since the progress is an outcome jointly produced by the individuals and the hiring agents, supply and demand could not be fully disentangled. An advantage of this dataset is that it allows the researcher to unambiguously determine whose decision it was that a candidate does not progress through the hiring pipeline. It thereby enables a separate study of who does not progress through the hiring pipeline because he/she pulled out, and a study of who does not progress because he/she was dropped by the hiring agent (in this case the search firm or the hiring firm). Specifically, it allows the precise identification of how prior relationships shape the decision making of actors on both the supply and the demand side.

Binary variables “candidate not interested”, and “candidate declared not suitable”, constitute the main dependent variables in this thesis. Specifically, in the analyses I model the following outcomes:

- the candidate rejects the search firm interview (as opposed to goes for the search firm interview)
- the candidate is declared unsuitable by the search firm (as opposed to goes for the search firm interview)
- the candidate rejects the hiring firm interview (as opposed to goes for the hiring firm interview)
- the candidate is declared unsuitable by the search firm after the search firm interview (as opposed to goes for the hiring firm interview)
- the candidate rejects receiving a job offer (as opposed to gets a job offer)
- the candidate is declared unsuitable by the hiring firm (as opposed to gets a job offer)

I use three other dependent variables in this thesis: natural log of the vacancy salary, margin (the ratio between the search firm's fee and the vacancy salary), and length of the search assignment. I define these in more detail in the relevant chapters of the thesis.

Table 2.2 presents the summary of the individual, relationship and vacancy data.

Table 2.2: Summary of individual, relationship and vacancy data⁽¹⁾

Individual data (N=6,308)				Relationship data (N=64,505)			
Variable		Mean	Std. Dev.	Variable		Mean	Std. Dev.
Salary at the time of inclusion in the consideration set (in £)		137,812	67,066	Number of jobs considered		.765	1.449
Salary over the vacancy salary		.343	.474	Number of jobs: candidate not interested		.256	.737
Age		43.643	5.532	Number of jobs: candidate not suitable		.362	.878
MBA		.204	.403	Number of jobs: candidate placed		.023	.162
Job tenure		3.414	2.715	Number of contacts in other roles		2.427	6.324
Vacancy data (N=998)							
Variable						Mean	Std. Dev.
Maximum salary for the position (in £)						161,418	96,764
Number of the search firm's prior assignments for the hiring firm						1.433	2.579
Job industry (%)				Job function (%)			
Media	3.80	Infrastructure	6.31	Financial services professional	10.52	Marketing director	6.81
IT	7.41	Government	9.11	Consultant	1.40	Sales director	3.20
Pharmaceuticals	7.91	Health	0.70	Board member	4.30	Divisional sales director	2.20
Agriculture	0.10	Education	0.30	CFO	6.71	Legal and governance	6.31
Finance	33.66	NGO	1.20	Divisional finance director	3.10	Non-executive director	1.30
Professions	9.21	Other	3.70	CEO	7.71	Pharmaceuticals scientist	1.40
Leisure	5.91	Missing	0.00	Divisional managing director	13.42	Management – other	20.84
Engineering and manufacturing	12.12			Operations director	2.60	Other	15.03
Retail	14.82			Government	3.80	Missing	0.30
Energy	2.10			HR director	4.40		

(1) A given vacancy may correspond to more than one industry and more than one job function, hence the sum of percentages exceeds 100.

2.2.2 Individual data

I have information on four individual characteristics of candidates: their salary at the time Execo included them in the consideration set, age, having an MBA, and the tenure in their current job. Execo's database contains information on all four of these characteristics for only a small portion of the sample (6,308 observations, 9.77% of the dataset)³. For those 6,308 observations, the average salary of candidates in the consideration set was £137,000, putting them well into the top 1% of the income distribution in the United Kingdom (in 2008 the cut-off point for the top 1 percent was £118,027, according to the UK Office for National Statistics). For 34% of the group with all individual characteristics the salary they had at the time they were put in the consideration set exceeded the maximum salary the hiring firm was willing to pay. 20% of candidates had an MBA, and were on average 43 years old. Their

³ The maximum vacancy salary is larger for the portion of the sample with the missing individual data (£156,000 vs. £148,000). People in the reduced sample (portion of the sample with the individual characteristics) have been considered for more jobs (almost by a factor of 3). They have also have gone to more Execo interviews, and have been interviewed more by the hiring firms (by the factor of four). The latter two results could explain why there is more data for the group in the reduced sample - these people have been more successful in the past, i.e. have gone to more interviews, hence Execo had both more opportunity and more need to process the data from their CVs. Execo's consultants commented that the researchers and the consultants may not have the CV of the person when he/she is put in the consideration set. This may be so because people whose CV Execo does not have may be recommended for inclusion by Execo's sources and other contacts. Even if Execo does have a CV, it is possible that the individual data used in this thesis - salary, age, MBA and job tenure - is not stated the CV, and needs to be collected separately later on. Candidates in the reduced sample have also had more non-candidate contacts with Execo. They have also been placed more often in the past and have responded to an advertisement less. There is no difference in the number of prior search assignments that Execo did before vacancies in the large and the reduced sample. There are fewer women in the reduced sample. In terms of assignment industry, in the large sample there are more assignments in the finance industry (35% vs. 25%). There is also a larger proportion of observations with job function financial services professional (13% vs. 5%). Conversely, proportion of CFO and HR jobs is greater in the reduced sample (5% vs. 14%, and 4% vs. 9%, respectively). In sum, the reduced sample appears biased towards candidates that are more connected to Execo as they have been successful in progressing through the hiring pipeline in the past. They are more connected to Execo also in the other roles. When asked about the reasons for the lack of data on education and employment history the consultants at Execo explained that the education and employment history data is collected on the first or second contact with the candidates, i.e. the point at which they send in their CV. Many executives end up in the database, yet never send their CV to the search firm. Another reason for the missing data is that inputting the data in the database is a researchers' job, and according to the search consultants these are not always fully diligent. The search consultants, though, did not feel that there is some systematic bias in inputting the education and employment history data. They also explained that age is missing in a number of cases since during the observation period legislation was passed that prohibited discrimination on the basis of age. The focal search firm, like many others, decided to stop recording executives' date of birth, hence there is missing data on that variable for many observations.

average job tenure was 3.41 years. Within this group women represented 13.30% (less than in the whole sample - 16.53%).

2.2.3 Relationship data

While the dataset contains individual characteristics for only for a small portion of the sample, it includes the relational data for all observations. The key part of the relational data is data on the relationship between the executives and the search firm. When an executive is included in the consideration set, I can observe how many jobs he/she has been considered for by Execo in the past, how many jobs he/she has rejected (at any stage of the process), how many jobs he/she was rejected for (by Execo or the hiring firms), whether he/she was placed by Execo in the past, and how many contacts in the non-candidate roles the person had with Execo. The non-candidate roles include the general role, the role of a client and the role of a provider of references (I describe these in more detail in chapter 3). Executives have on average been considered for 0.76 jobs, have rejected 0.25, were rejected for 0.36, and were placed in 0.02 (in the remainder of jobs they were considered for the outcome was unknown). They had 2.42 contacts in the non-candidate roles.

2.2.4 Vacancy data

At the vacancy level, the data included in this thesis is data on the vacancy salary, the relationship between Execo and the hiring firm, data on the industry of the vacancy, and job function of the vacancy. On average the maximum vacancy salary is £161,000, meaning that the vacancies studied in this thesis are at the very high end of the labor market. I measured the relationship between Execo and the hiring firm using the number of prior search assignments Execo conducted for the hiring firm in the observation period. The average was

1.43 searches. In terms of the search assignment industry, over a third of the assignments were in the finance industry (33.36%), followed by retail (14.82%) and engineering and manufacturing (12.12%). The most represented job functions were management-other (20.84%)⁴, divisional managing director (13.42%), and financial services professional (10.50%). In chapter 5 I also use data on the assignment length, Execo's margin (fee/vacancy salary), number of candidates in the consideration set, and Execo being the preferred supplier to the hiring firm. I provide the descriptive statistics for these variables in that chapter.

To analyze the relationship between the independent and dependent variables I used the linear probability model regression. This is a regression method that uses the assumptions of the regular OLS methodology to model a binary dependent variable. In particular, I exploited a feature of the dataset that a given candidate may have been a candidate for several vacancies, and conducted the analysis with linear probability models with the individual fixed effects (i.e. candidate fixed effects – I use the two terms interchangeably in the thesis). This approach allowed me to fully control for the unobserved individual heterogeneity in stable characteristics which might affect the outcomes captured by the dependent variables.

An advantage of the linear probability model is that it allows a relatively simple and intuitive interpretation of the coefficients (Waguespack and Sorenson, 2010). This is important in my context as several of my theoretical arguments require an analysis of interaction terms, which are not very intuitively interpreted in non-regression models. Nevertheless, the linear probability model may be problematic, in particular in two ways: first, the predicted probabilities of an outcome can be greater than 1, and less than 0 (which is inconsistent with the probability theory); second, it assumes that the effects of the independent variables on the dependent variable are linear (while in fact they may depend on the values of both). For these reasons I adopted the following approach: I ran the models

⁴ Execo's consultants explained that the code management-other corresponds broadly to the middle management positions.

using both linear probability regressions and logit regressions (which more appropriately model the relationship between the independent and dependent variables). I then compared the results that they produce. In all regressions the main results were the same for both kinds of methods. I hence presented the results for the linear probability models in the tables, for ease of presentation and interpretation.

For analyses in chapters 3 and 4 I constructed a panel with candidate-by-vacancy as a unit of analysis. That is, the rows in the panel had the following structure: candidate 1 – vacancy 1, candidate 1 – vacancy 2, candidate 2 – vacancy 1, candidate 3 – vacancy 1, candidate 3 – vacancy 2 etc. The number of the vacancy is the number of the vacancy for the candidate (i.e. vacancy 1 for candidate 1 was the first vacancy for which he/she was included in the consideration set during the period of observation, vacancy 2 for candidate 1 was the second vacancy for which he/she was included in the consideration set etc.). The time elapsed between vacancies for which the same candidate was considered differed across candidates.

In the analysis presented in chapter 5 I used as the unit of analysis hiring firm-by-vacancy. This was possible because the search firm has conducted more than one search assignment for some firms during the period of observation. The rows in the panel used in chapter 5 had the following structure: hiring firm 1 – vacancy 1, hiring firm 1 – vacancy 2, hiring firm 2 – vacancy 1, hiring firm 3 – vacancy 1, hiring firm 3 – vacancy 2 etc. In contrast with the panel in chapters 3 and 4 in the panel in chapter 5 the number of the vacancy is the number for the hiring firm (i.e. vacancy 1 for hiring firm 1 was the hiring firm's first vacancy that the search firm conducted a search assignment for, vacancy 2 for hiring firm 1 was the second vacancy the search firm conducted an assignment for etc.). Time elapsed between vacancies for which the search firm conducted an assignment differed across hiring firms.

2.3 Data limitations

The data described in the previous sections has a number of advantages. To the best of my knowledge, it is the first dataset that allows the researcher to study how prior relationships affect future decisions of the actors both on the supply and the demand side of the labor market. Since many executives were included in the consideration sets for more than one vacancy, the data allows for individual fixed effects estimations. Also, it is the first time that an executive search firm has given researchers access to data at such a level of detail. The latter means that this is one of the rare instances in which we can study the high-end labor market at a great level of detail. Nevertheless, the data has limitations. I outline four of them below.

The first limitation is that I am modelling individual decision making only on the basis of the executives' ties with the search firm. Executives are likely to have many other relationships, and they are likely to shape the portfolio of options they have available at any point in time. Since I don't have the data on these other relationships, I don't know what other options executives had when they were approached by Execo. I try to deal with this limitation indirectly by using models that control for individual fixed effects which control for time-invariant unobserved heterogeneity, an element of which may be a propensity to develop relationships. This limitation nevertheless puts boundaries on the conclusions I can make with this data.

The second limitation is that I don't have reliable information on the individuals' employment for the entire sample. Specifically, for the portion of the sample I don't have salary data for (about 90%), I don't know how the vacancy compares to their current employment, and, in extreme, whether they are in employment at the time (i.e. I don't know whether there is no salary information as the candidates had no job at the time). I therefore

assume that the candidates are in employment, as they otherwise wouldn't be attached to the executive-level vacancies. This assumption appears reasonable for those candidates that are attached to the vacancies by the search firm, however may be less sustainable for those candidates that have applied in response to an advertisement (19.75% of the sample). I try to address this limitation by using a control variable that takes on a value of 1 if a person applied for the vacancy. In chapter 3, where the theoretical arguments are about job searches of employed individuals, I only look at the portion of the sample for which I know that the individuals were employed when they were included in the consideration set for a vacancy (as the database contained information on their salary at that time).

The third limitation is that I cannot observe the extent to which a person is a good match for a given vacancy. Ideally I would have access to the search briefs that describe the profile the hiring firm is looking for. I would then use this information to see how an individual's profile corresponds to the hiring firm's requirements. I have attempted to approximate the match by including among the controls the ratio between the person's salary at the time when he/she is put in the consideration set, and the vacancy salary. If the salary accurately captures a person's human capital, and if the vacancy salary is closely correlated with the human capital that the hiring firm is trying to acquire, this ratio would be an approximation of the amount of the required human capital an individual possesses. However, given that this is only an approximation, and in addition one that I could only use for a small part of the sample, the limitation remains.

The fourth limitation is that I have little information on the hiring firm (apart from the industry in which it operates). I also have little information on the firm a candidate is working for when he/she is attached to a vacancy. The search firm's condition for sharing the data was that the hiring firms, as the search firm's clients, remain anonymous. While this is a clear limitation, a helpful aspect of executive search in that respect is that the search firm divulges

the information about the hiring firm sequentially, and consequently the candidates know little about the hiring firm in the early stages of the process. As I focus on the early stages of the process in this thesis (i.e. the stages when the candidates get to know little about the hiring firm) this limitation is somewhat lessened, but is still present.

The fifth limitation is that, relative to the whole dataset, only a small number of candidates are made an offer and hired into the hiring firm. While the dataset consists of 64,505 observations, only 570 candidates were made an offer, and 479 were placed (i.e. hired by the hiring firm). This means that in 570 out of 998 vacancies (57%) an offer was made, and in 479 out of 998 (47%) an offer was accepted. Because of the small sample sizes at the stage of the offer and the placement the candidate fixed effects models did not converge. I hence cannot engage with the literature that examines why individuals get hired, or why they move from one organization to the other. I thus make no conclusions about such outcomes.

CHAPTER 3: EXECUTIVE PARTICIPATION IN THE COMPETITION FOR JOBS: THE ROLE OF CONTENT OF TIES WITH THE BROKER

3.1 Introduction

In this chapter I start analyzing how differences in content that flows through ties between executives and the search firm affect executives' willingness to subsequently pursue vacancies introduced to them by the search firm. That is, I examine how ties developed at time $t=1$ shape the executives' willingness to be interviewed for a vacancy that is introduced to them by the search firm at time $t=2$. This addresses the overall question of when relationships with the contacts lead individuals to pursue job opportunities in the following way: the executive search firm is a contact with whom an executive has a history of interactions, and who introduces an opportunity (a job) to the executive. The executive may decide to act on the information about the opportunity (by going for a job interview), or let it pass (by not going for a job interview). I propose that the content of the interactions that an executive has had with the search firm prior to being informed about a vacancy will affect whether or not the executive will decide to act and go for an interview for the focal vacancy.

Extant literature has demonstrated that prior history of interactions shapes the likelihood (Gulati, 1995) and parameters of future exchanges (Uzzi, 1999). That stream, that focuses on embeddedness of actors in close market or inter-organizational relationships, views interactions primarily as a platform for exchange of information as well as favors. It then adopts a premise that actors can take the established stock of the exchanged information and favors into account when considering further exchanges with the partners with whom they have exchanged in the past (Uzzi, 1996, 1997). However, one might argue that interactions also provide an opportunity for observation of the behavior of partners. Behavior

is, to an extent, dependent on the role that a given actor plays in a particular interaction (Shapiro 1987). Role may be conceptualized as "a set of recurrent behaviors, appropriate to a particular position in a social system." (Polzer, 1995a, p. 495). Roles define what behavior is appropriate for the actor enacting the role. They are defined by the expectations of role senders, i.e. other members of the role set with whom the role incumbent interacts. Specifically, role expectations are "beliefs and attitudes held by members of the role set what behaviors are appropriate for the person in the role." (Polzer, 1995b, p.496).

Prior research has shown that the extent to which a role incumbent is autonomous in his/her role (i.e. has discretion over his/her behavior and decisions while enacting the role) shapes the extent of trust that those that interact with him/her will place in him/her (Perrone, Zaheer and McEvily, 2003). Trust refers to "the willingness to accept vulnerability based on positive expectations about another's intentions or behaviors." (McEvily, Perrone and Zaheer, 2003, p. 92). Trust is therefore an expectation held by those that interact with an actor that the actor can be relied on to fulfill obligations, will behave in a predictable manner, and will act and negotiate fairly when the probability of opportunism is present (Zaheer, McEvily and Perrone, 1998, p. 143). An important element of this conceptualization is that trust is seen as an expectation rather than a conviction, and thereby reflects an uncertain anticipation of the referent's future behavior (Zaheer, McEvily and Perrone, 1998, p. 143).

Actors have been shown to trust the role incumbent less when the roles constrain their behavior more (Perrone, Zaheer and McEvily, 2003). Greater constraint means that role incumbents have less freedom in balancing the diverse and sometimes conflicting expectations from the members of their role set. To illustrate this point consider the case of the purchasing managers described by Perrone, Zaheer and McEvily (2003). An organization employing the purchasing managers has certain expectations about how they should deal with their opposite numbers, supplying managers of the supplier organizations. These expectations

are likely to be different from the expectations of the supplying managers. Purchasing managers need to balance these expectations to reach agreements that are beneficial for their employing organizations and acceptable to the supplying managers. The purchasing managers' employing organizations may give them more or less autonomy in dealing with the supplying managers. Hence, some behavior and decisions of the purchasing managers may be prescribed by the organization, while some may be based on the purchasing managers' personal judgement. The supplying managers, though, generally cannot tell which behavior and decisions are stipulated by the organization. Purchasing managers may make use of this uncertainty, and for instance refuse to give concessions to the supplying managers, citing organizational constraints as the reason. The supplying managers therefore can accurately assess the purchasing managers' underlying motivations only from the part of the latter's behavior and decisions that are not prescribed by their organizations. The implication is that only the organizationally unconstrained part of the purchasing managers' behavior and decisions can provide accurate cues for the supplying managers' formation of trust in the purchasing managers. Role autonomy, by determining the extent of the purchasing managers' freedom in devising their actions and behaviors, shapes supplying managers' trust development (Perrone, Zaheer and McEvily, 2003).

While a given role, such as that of the purchasing manager described above, may be characterized by different degrees of role autonomy in different organizations, different roles in the same organization are also generally characterized by a different degree of role autonomy. It is quite clear, for instance, that a CEO has greater role autonomy than a sales clerk. When external parties interact with the incumbents of different roles in an organization, the autonomy of these roles plays a part in the trust that the external parties put in the role incumbent. This may in turn affect the trust that the external parties place in the overall organization. As Zaheer, McEvily and Perrone (1998) explain, it is conceptually coherent to

view trust as being placed in an individual, *or* a group of individuals, such as a group comprising an organization. These authors also provide evidence that interpersonal and inter-organizational trust are positively correlated. That is, if individuals trust particular individuals from the other organization, they will tend to trust other individuals from that organization. This process may be underpinned by the mechanism of transferability of trust, whereby an individual transfers trust in one member of a group to another one, even if he/she has no prior experience with the second member. This kind of trust is characteristic-based trust (Zucker, 1986), and may apply to transfer of trust to members of an external group (i.e. a group of which one is not a member). That is, an individual who is external to an organization, and trusts an individual in that organization, may get to trust another member of the organization if the trust transfer takes place. Trust transfer is driven by perceived similarity among members of the collective, and to the extent that this similarity is apparent to the individuals, they may formulate perceptions of the whole collectives (such as organizations) on the basis of their histories of interactions with the individual members of the collectives (McEvily, Zaheer and Perrone, 2003).

In the context of executive search an executive may play four kinds of roles, and hence four kinds of ties may develop between an executive and the search firm. These are: candidate tie, general tie, client tie and source tie. Candidate tie develops when an executive plays the role of a candidate for a position for which the search firm is recruiting. General tie develops when an executive and the search firm engage in a basic social contact. Such social contact may be a meeting in which the executive and the search consultants exchange information about the state of the industry or the market, yet the search firm does not have a particular vacancy for the executive at the time and is not conducting a search for which the executive would be a client. Client tie develops when an executive, in interactions with the search firm, plays the role of a representative of an organization the search firm is working

with (or is trying to work with). Source ties develop when an executive, in interactions with the search firm, plays the role of an informant to the search firm. In such interactions the executive provides a reference for candidates the search firm is considering at the time. Figures 3.1.1-4 demonstrate how the four kinds of ties form.

Figure 3.1.1: Candidate tie

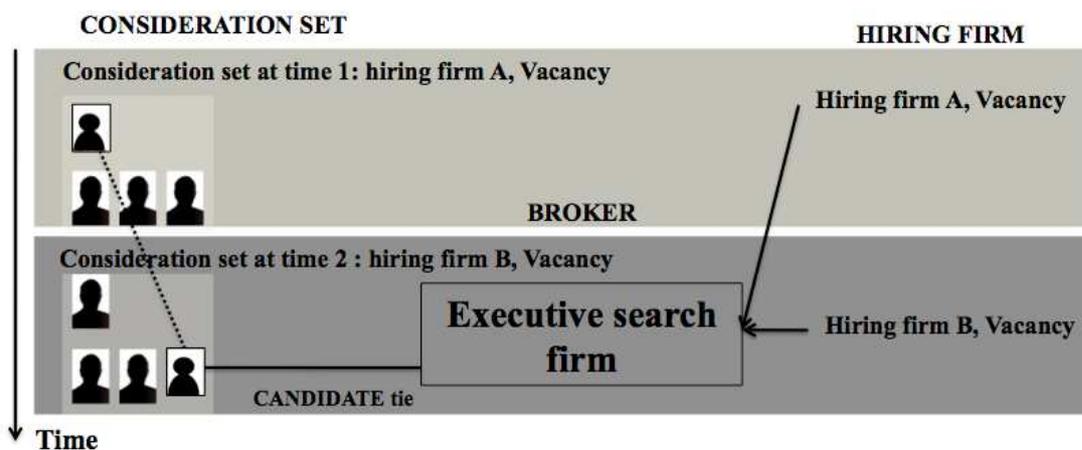


Figure 3.1.2: General tie

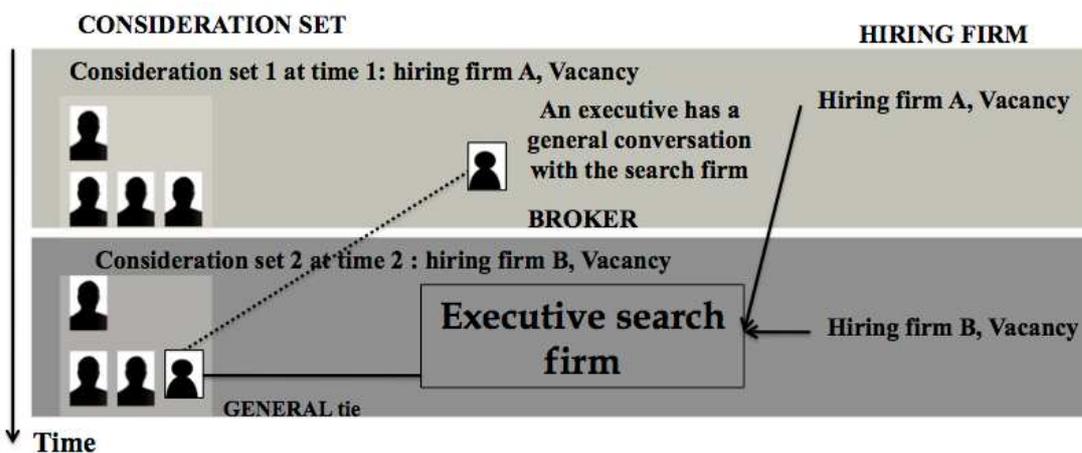


Figure 3.1.3: Client tie

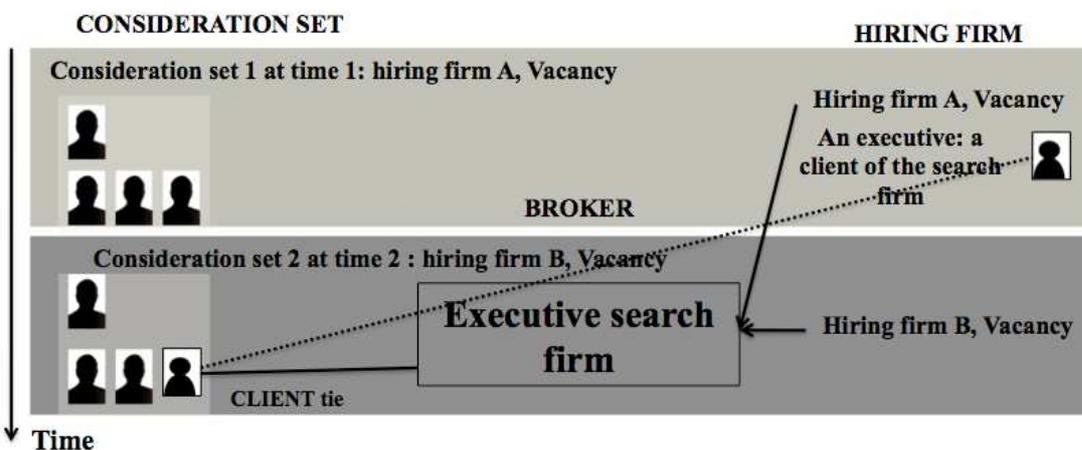
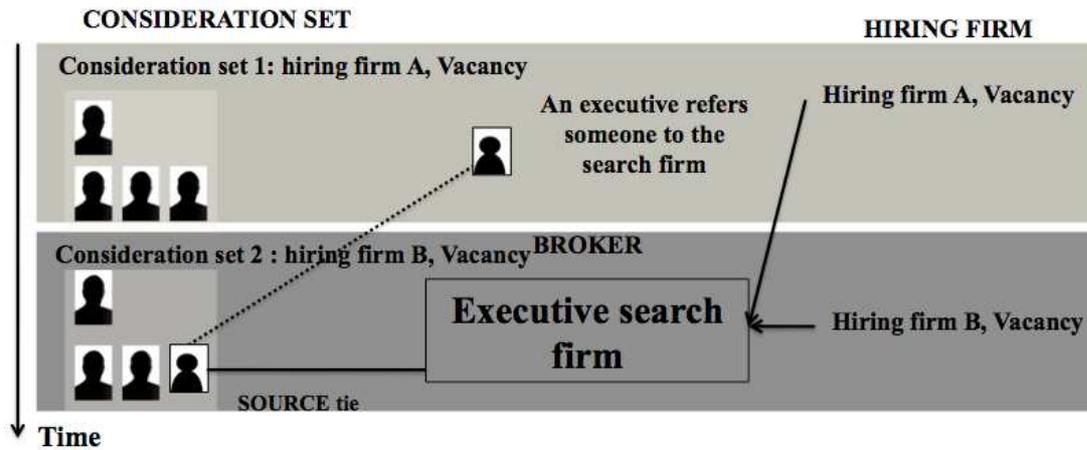


Figure 3.1.4: Source tie



Candidate tie forms when an executive is a candidate for vacancy at one hiring firm (e.g. hiring firm A), and is then included in a consideration set as a candidate for a different vacancy at time $t=2$ (e.g. vacancy at hiring firm B). When the executive is a candidate for a vacancy at time $t=2$, he/she has a candidate tie to the search firm, due to him/her being a candidate in a search conducted by the search firm at time $t=1$. The mechanics of the general, client and source tie are similar. When an executive is a candidate for a vacancy at time $t=2$, he/she has a general, client or source tie if he/she at time $t=1$ had a general conversation with the search consultants from the focal search firm (general tie), was a client of the search firm, or the search firm attempted to turn him/her into a client (client tie), or referred someone to the search firm, i.e. provided a reference (source tie).

The four ties differ in the content of interactions. This is so because both role of the executive, as well as roles of the representatives of the search firm (i.e. search consultants) differ. Specifically, I argue that the search consultants' roles differ in the extent to which they are pure brokerage roles, or boundary-spanning roles. I further argue that this affects the executives' development of trust in the search firm, and subsequently has an impact on the executives' decisions to be interviewed for vacancies.

3.2 Brokerage and boundary spanning roles, trust development, and pursuit of labor market opportunities

A search firm operates as a broker between candidates for executive-level vacancies and the hiring firms. In order to perform their brokerage function, the search firms need to first obtain clients (the hiring firms). To do so, search consultants need to interact with current and potential clients. To be able to successfully match candidates and hiring firms search consultants need to have accurate and timely information about candidates, which they obtain through references. They also need to have a deep insight into the dynamics of the industry sectors in which they operate. To obtain both kinds of resources, search consultants need to interact with executives who can provide them. I described the former kinds of interactions as source contacts, and the latter as general contacts. In contacts with executives search consultants play different roles, and I argue that the types of roles search consultants play when they interact with executives shape the degree of trust executives place in search consultants.

Prior research demonstrated that brokerage actors may take on different types of roles at different points in time and in different situations to accomplish their goals. In a single project, they may shift from a role of a strategic actor that extracts advantage from their position, to a relational expert that connects others to foster creativity and innovation (Long Lingo and O'Mahony, 2010). Search firms, as brokerage organizations, also shift their consultants between different roles to accomplish the firms' objective of successful matching of executive candidates and the hiring firms. The search consultants' roles I describe above could be classified into two distinct types: pure brokerage roles and boundary-spanning roles. The distinction between the two types of role in brokerage organizations is subtle, yet crucial from the perspective of trust development. Specifically, brokers can span boundaries, but not

all boundary spanners broker (Fleming and Waguespack, 2007). Boundary spanners are individuals at the boundary of organizations, who identify, translate and relay information within and across firms (Tushman, 1977). They are individuals that represent organizations, and engage with external parties on the organizations' behalf (Perrone, Zaheer and McEvily, 2003; Broschak, 2004). They can also be individuals that reside at the boundary of several domains, such as at the intersection of technological areas (Fleming and Waguespack, 2007). When an individual, working for a brokerage organization, performs a boundary-spanning role, he/she does not have a particular mandate by his/her organization, but rather interacts with the external parties in the capacity of an individual with a lot of information, often of a diverse kind. The latter is a natural consequence of working for a brokerage organization that resides at the boundary of several domains. An example of such a role in executive search is a general contact. When a search consultant interacts with executives as a general contact, he/she simply exchanges information about the industry, the market trends etc. Importantly, he/she does not have a particular organizational mandate to secure an upfront specified benefit for the search firm from the conversation.

Individuals performing pure brokerage roles in brokerage organizations, on the other hand, are boundary spanners who are simultaneously occupying a position from which they are actively attempting to obtain benefits from at least one of the brokered parties. They are placed in that position by their organizations, who give them a mandate to obtain specific information, resources, or other benefits from the brokered parties. Examples of such roles in executive search are candidate, client and source contacts with executives. In all these contacts executives encounter search consultants as they are performing a pure brokerage role. In the candidate contact search consultants broker a recruitment process between candidates and the hiring firm. The goal that the search firm sets for search consultants is to secure the best candidates for the hiring firms. In the client contact the search consultants broker the

relationship between the executive and his/her employer on one side, and the search firm on the other. The mandate there is also clear: obtain business from the executive's employer, or ensure that the current search is seen through to the end. In the source contact the search consultants broker the relationship between the executive, acting as a source, and the search firm. The mandate that a search consultant has in that case is to obtain the reference.

In all of the three roles where a search consultant is performing a pure brokerage role the consultant's role autonomy is lower than in the one role where a consultant is performing a boundary spanning role (general contact). Prior research suggests that lower role autonomy results in actors placing lower trust in the role incumbent. Greater role autonomy, in contrast, results in more trust being placed in the role incumbent (Perrone, Zaheer and McEvily, 2003). In executive search that would mean that general contacts, characterized by more role autonomy than either candidate, client, or source contacts, should lead to a greater development of trust of an executive in the search consultant.

When an executive subsequently becomes a candidate, he/she is likely to transfer trust previously developed through a general tie with the search consultant onto the search consultant(s) that works with him/her when he/she becomes a candidate. This is to be expected because of the role that similarity plays in trust transfer (Williams, 2001). Namely, search consultants are familiar with the content of prior interactions of their colleagues with a given executive, and tend to evoke previous conversations, especially if they resulted in the development of a positive disposition on the part of the executive. A quote from an Execo senior associate (interviewee 16) highlights that search consultants are well aware of the content of an executive's prior interactions with the search firm, and try to use the accumulated information in their favor when they work with executives as candidates.

" What the consultants will be mindful of is how they are going to sell particular job to a candidate. Researcher will prepare the consultant for an interview, give them information on

the jobs candidates previously rejected, along with the reasons, and the consultant can use this information in the interview to sell other points, to bring the candidate around."

As the consultants evoke the previous interactions the executive had with the search firm, the executive will be able to see a similarity between the consultants, which would facilitate trust transfer. Naturally, if the search consultant that the executive interacts with when he/she becomes a candidate is the same consultant that he/she interacted with in other roles in the past, previously accumulated trust in the consultant simply extends into the new situation (Gulati, 1995).

Figure 3.2 below summarizes the preceding discussion by tabulating the different ties that form in executive search, different roles that are associated with different ties, the degree of role autonomy associated with the roles, and the degree of trust that is expected to develop on the basis of role autonomy associated with different roles.

Figure 3.2: Tie content, brokerage and boundary spanning roles, and trust

Tie	Role	Consultant brokers between	Role autonomy	Trust development
Candidate	Pure brokerage	Candidates and hiring firms	Low	Low
General	Boundary spanning	No one	High	High
Client	Pure brokerage	Executive's employing organization and the search firm	Low	Low
Source	Pure brokerage	Executive and the search firm	Low	Low

Trust developed through prior interactions may have a bearing on the probability of subsequent exchanges (Gulati, 1995). Trust becomes a particularly important criterion in deciding a course of action when the actors face uncertainty about the future actions of their partners (Zaheer, McEvily and Perrone, 1998). When actors trust their prospective partners, they are more likely to take a "leap of faith" and expect that the latter will not act in an opportunistic manner (Zaheer and Venkatraman, 1995; McEvily, Perrone and Zaheer, 2003). When such an expectation is in place, an actor will more probably decide to enter into the process of exchange with a particular partner.

Executive search is characterized by a high degree of uncertainty of the outcome for any particular candidate, as the number of candidates in the consideration set tends to be large (the average in the overall dataset in the thesis is 64 candidates per vacancy). Candidates have little information about the parameters of the competition (such as the number and the quality of the competing candidates), and have to rely on the information provided by the search consultants (who in turn are fully informed). If the candidate has trust in the search firm, he/she may presume that the search firm will make an approach only when a candidate has a

reasonable probability of a successful outcome⁵. Everything else being equal, candidates should be less likely to reject a search firm interview invitation when they have more pre-existing trust in the search consultants from the focal search firm. As I described, more trust is likely to develop in the general ties. However, this does not yet tell us whether trust developed in candidate, client and source ties is merely lower than in the general tie, or is actually close to zero or even negative.

To be able to rank the four kinds of ties on the trust dimension, I asked Execo's consultants to compare and contrast the ties with respect to their potential effect on trust that executives develop in the search firm. Execo's senior associate (interviewee 16) provided the following comments:

Candidate tie:

“People who are familiar with being candidates, it is a much more level relationship [than if they are clients], a more trusting relationship. They are entrusting you to advise them on their career. However, once a candidate has been on three shortlists, four shortlists, and hasn't been picked, the consultant starts to think that there is something wrong with them, or they are interviewing badly, and they tend to get dropped. And you can imagine them (candidates) going into a meeting with a client: the person who has been approached once is fresh, is confident, says what they want to say to the client, has not been tarnished by four, five, six other recruitment processes. People do lose confidence, and the clients pick up on it.”

General tie:

“Generals, if you've met them and invested some time in them, the headhunter may not place any particular value on that, but the candidate certainly does: wasn't this headhunter kind to have a chat with me, they must be a good person, I will trust them.”

Client tie:

⁵ Participation in searches may be costly. Unsuccessful participation may be seen as a failed job search, and interpreted as revealing lower quality than the quality implied by the employee's current salary (Lazear, 1986). In addition, the search firm may be more reluctant to put previously unsuccessful candidates forward as candidates in the subsequent searches. Still, an executive might engage in a search to obtain bargaining leverage with his or her current employer (Boswell, Boudreau and Dunford, 2004). However, employers may be selective in the counter-offers (Barron, Berger and Black, 2006), making searching through a search firm to obtain leverage a relatively risky strategy.

“You’re selling to clients because you want their money. So the balance of power in the relationship is a client is up here, and the headhunter is down low, sort of going please, give me your money. [When they are approached as candidates] clients think “I know what your process is because I bought it. You’ve explained it to me, so I don’t trust you at all. I know exactly what is going on and I’ll make an objective decision based on what I think, and I don’t care what you think.”

Source tie:

“You need information from them. So you need to sell why they have to give you something that’s not really in their interest. They might call in the favour one day, but at the moment it is not particularly helpful. The sources probably feel sort of I don’t owe you anything headhunter, you haven’t invested time in me before so I’ll make an objective decision based on what I think about the market.”

The quotes suggest that the client and source ties are associated with little, if any development of executives' trust in the search firm. It therefore appears we may predict that when a candidate has these two kinds of tie prior to becoming a candidate, he/she will be no more or less likely to reject the invitation to go for a search firm interview. Candidate ties, though, are less clear. On the one hand candidate ties may stimulate trust development because they indicate that an executive was in the past content with being advised on his/her career by the search firm. On the other hand, though, candidate ties also indicate that the candidate has in the past rejected vacancies, or was rejected by the search firm or its clients, which may result in a reluctance to participate in future recruitment processes. That, in turn, appears incompatible with a development of trust (this point is addressed in more detail in chapter 4). Since the data doesn't allow me to identify the point at which the trust-facilitating aspects of candidate ties are over-weighted by the trust diminishing ones, I do not theorize about the impact of the candidate tie on the probability of rejection of the search firm interviews, but rather use it as a control. General ties, on the other hand, have an unambiguously linear relationship with trust development - more general contacts should therefore lead to more development of trust in the search firm.

Before developing this conclusion into a hypothesis we need to take into account that the search firm does not exert the same effort on persuading candidates to go for the search firm interview, but rather expands more effort on the most desired candidates. As per the quote from an Execo senior associate (interviewee 16) in chapter 2, Execo always tries to get the best candidates to the interview, regardless of their willingness to participate in the brokered recruitment process. Candidate desirability and the desirability of vacancies for candidates are likely to be inversely related: more desirable candidates will find the vacancies less desirable (as they are more in demand), and vice versa. The implication is that the impact of ties on the executives' decisions to participate in the search firm interviews will not be the same for all candidates in all vacancies, but will rather depend on the relative desirability of the candidates in particular vacancies. In other words, if the search firm exerts more effort on persuading the most desirable candidates, then the level of trust in the decision making matters only for them. The less desirable candidates will be less likely to be persuaded to go for the interviews in the first place, and hence their level of trust in the search firm will be inconsequential. I therefore focus the hypothesis on the moderating effect of the trust developed on the basis of ties, i.e. on how trust moderates the relationship between the relative desirability of a vacancy for a candidate, and the probability of him/her rejecting the search firm interview. If trust increases the desirable candidates' expectations that the search firm will act in good faith in the recruitment process, such candidates will be less likely to reject the search firm interviews.

The preceding discussion indicates that a measure for “the best candidates” needs to be taken into account. As described in chapter 2, I unfortunately do not have the data on how an individual candidate compares to the requirements stated in the vacancy brief. However, I use a proxy to get at that measure. The proxy is the ratio of the candidate’s current salary to the salary of the vacancy.

Human capital literature on executive compensation (Milkovich and Newman, 2005; Harris and Helfat, 1997; Castanias and Helfat, 2001; Sturman, Walsh and Chermie, 2008) contends that of all forms of executive compensation base salary most closely reflects the executive's human capital. The maximum vacancy salary is specified by the hiring firm ex-ante, and therefore closely corresponds to the human capital that it is trying to acquire. The search firm informs the candidates of this amount at the very beginning of the process. Salary held by an executive could be seen as the valuation of his/her human capital by his/her current employer⁶. Ratio between the salary held by an executive, and the maximum salary offered by the hiring firm is therefore a useful indicator of the extent to which an executive's human capital corresponds to the human capital sought by the hiring firm. The closer this ratio is to 1, the more of the required human capital a given executive possesses. When this ratio is above 1, the executive is likely to be overqualified for the vacant position.

Human capital literature would predict that as the salary ratio increases executives, as the suppliers of human capital, will be less likely to participate in the brokered market exchange (i.e. more likely to reject the search firm interview). In other words, the willingness to change the status quo will be lower when an executive has less to gain, in monetary terms, from changing his/her job. As argued above, though, previously developed trust in the search firm should negatively moderate this relationship (i.e. reduce the probability of rejection). General ties, but not client and source ties, should therefore reduce the probability of rejection of the search firm interview.

⁶ Salary reflects an employer's valuation of the employee's total human capital (i.e. the sum of general and firm-specific). When employees change jobs, their firm-specific human capital is lost, and therefore the salary offered by the hiring firms rewards only the employees' general human capital. To account for this discrepancy between the salary at the current employer and the vacancy salary in the empirical analysis I control for the firm-specific component of the salary (using tenure in the employer as a proxy).

Hypothesis 1: As salary ratio between the present salary and the vacancy salary increases, the candidates will be less likely to reject participation in the brokered recruitment process (reject the search firm interview) if they have a prior general tie, but not prior client or source tie, with the search firm (the broker).

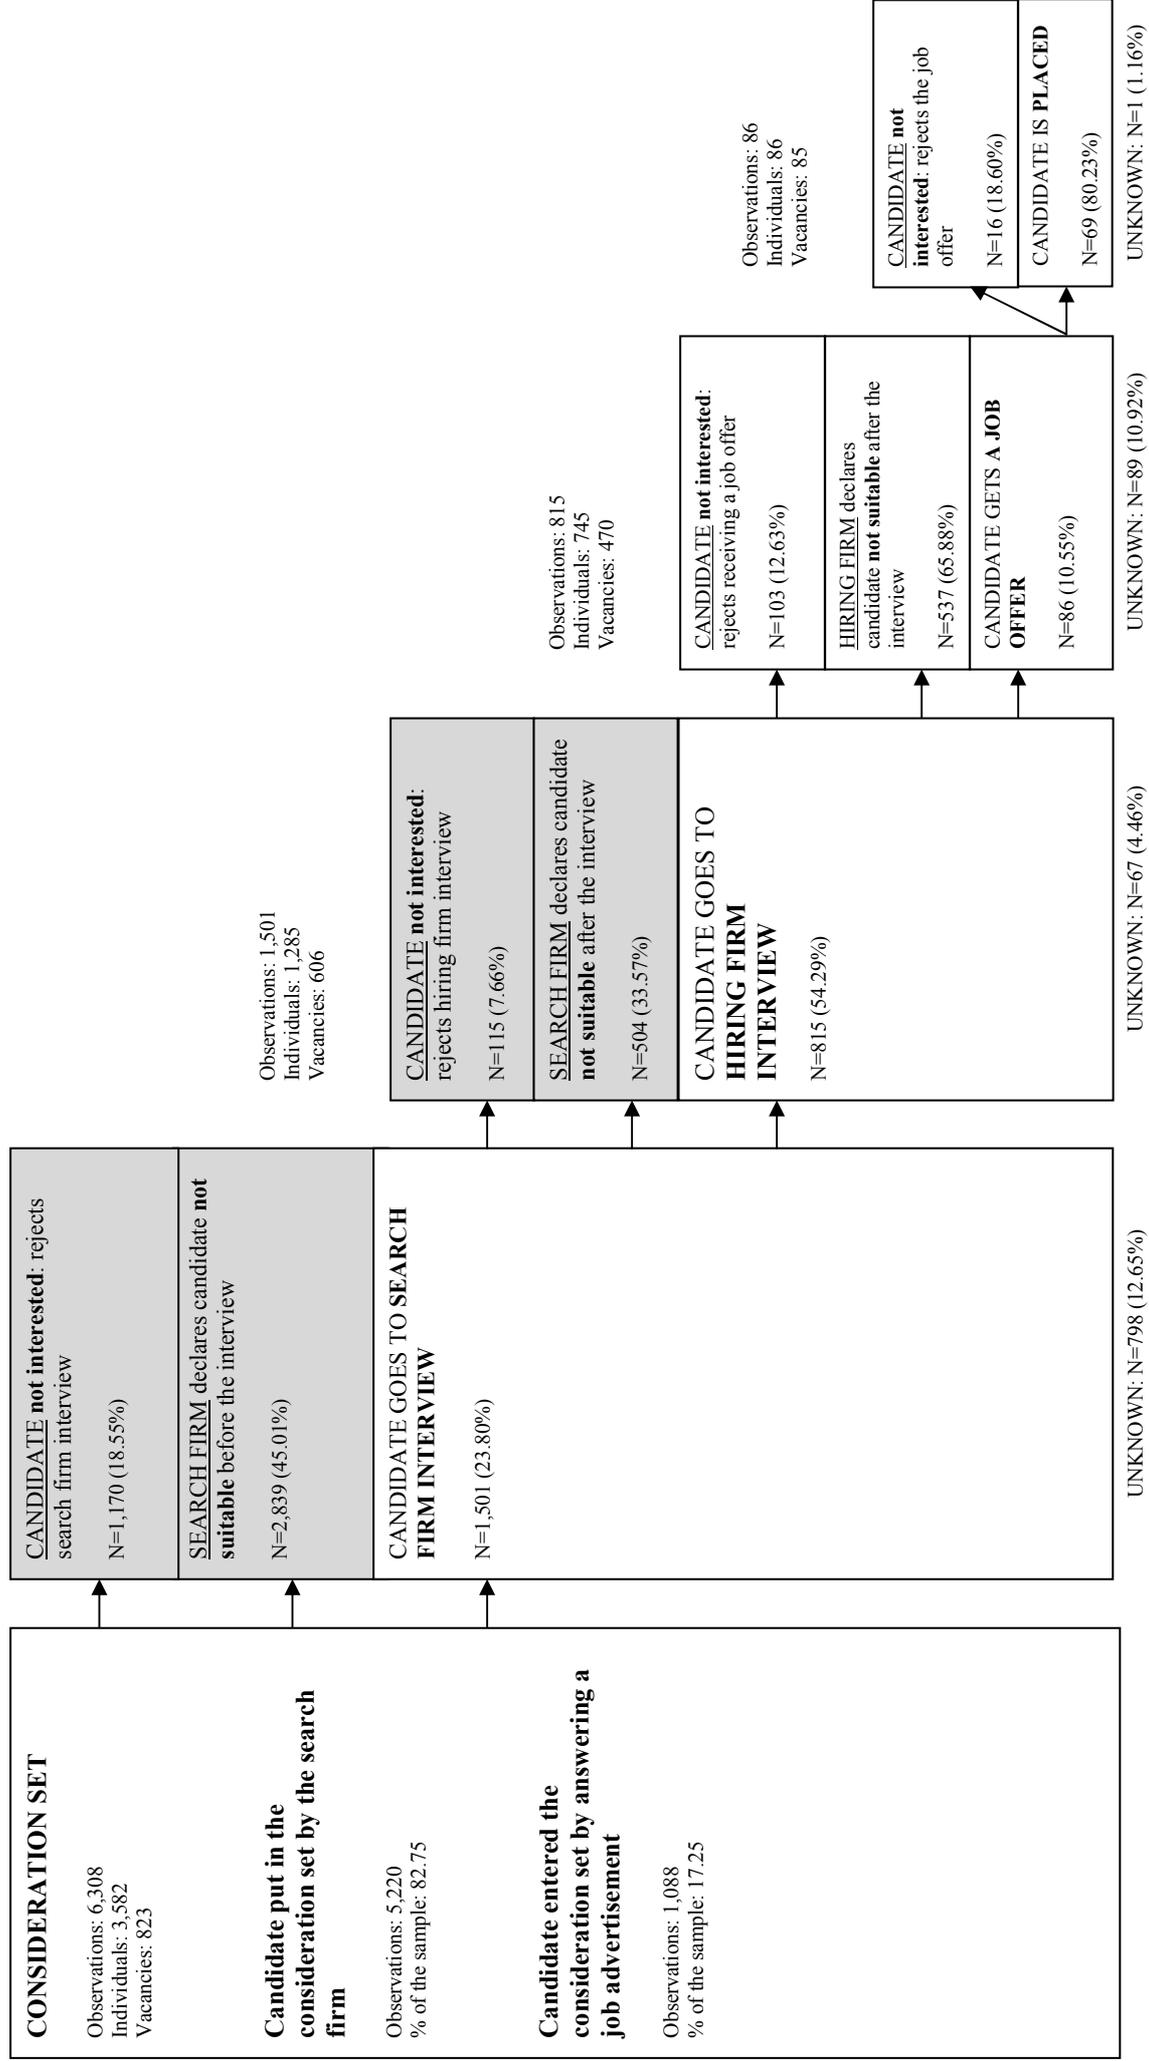
3.4 The sample and variables

In this chapter I am using a sample of 3,582 individuals that were included on the consideration sets for 823 vacancies. The sample consists of 6,308 observations. Only a relatively small proportion of the dataset was useful for the analysis in this chapter since the majority of entries include little or no information on the individual characteristics. I describe the reasons for the missing data in chapter 2. In that footnote I also describe the differences between the reduced sample (the sample with the full array of individual characteristics), and the rest of the sample.

To test the hypothesis advanced in the previous section I am studying the probability that the candidate will reject going for the search firm interview, as opposed to him/her going for the search firm interview. Once a candidate goes for a search firm interview, he/she might still pull out of the process before going for the hiring firm interview. Specifically, I model the probability that the candidate, having made it to the search firm interview, rejects going for the hiring firm interview (as opposed to him/her going to the hiring firm interview). These are the two outcomes of interest. In addition, to better understand the mechanisms, I also model the probability that the search firm declares a candidate not suitable for the vacancy before and after the search firm interview (as opposed to them going for the search firm interview/hiring firm interview). Figure 3.3 represents the hiring pipeline at Execo for the reduced sample (the sample with the full array of the individual data), which is used for the

analysis in this chapter. Shaded areas represent the outcomes that I am studying. The two main outcomes of interest are the two at the top (marked candidate not interested). Even though the outcomes of interest pertain to the decisions made by the individuals, to better understand the search firm's decisions I also study the two outcomes on the search firm's side. I marked them with a grey shade (labelled search firm declares candidate not suitable).

Figure 3.3: Hiring pipeline at Execo (reduced sample)



Variables used in the analysis are described below. Table 3.1 provides a summary of descriptive statistics for the variables used in the analysis.

Variables

Dependent variables⁷

- *A candidate rejects the search firm interview*: this is a dummy variable that takes on value 1 if a candidate rejects the search firm interview, and 0 if he/she goes for the search firm interview. When a candidate rejects the search firm interview, he/she effectively rejects participation in the brokered recruitment process.

- *A candidate rejects the hiring firm interview*: this is a dummy variable that takes on value 1 if a candidate rejects the hiring firm interview, and 0 if he/she goes for the hiring firm interview. When a candidate rejects the hiring firm interview, he/she effectively rejects participation in the brokered recruitment process (after having initially agreed and having been interviewed by the search firm).

Independent variable

- *General tie*: this variable measures the number of *general* contacts that the person has had with the search firm prior to being included on the list of candidates for a vacancy. General contact means that the person has met with the search firm to discuss relevant events in the person's industry, the trends, etc.

⁷ In this chapter I am not hypothesising about the outcomes on the search firm side (i.e. the search firm dropping the candidates from the hiring pipeline before and after the search firm interview). I hence do not include these two outcomes in the list of the dependent variables. As I do study them to understand the mechanisms, I reiterate what they mean in the analysis section, though.

Control variables

- *Salary ratio*: this variable is a ratio between a candidate's salary at the time when he/she was included in the consideration set, and the maximum vacancy salary that the hiring firm is willing to pay to the successful candidate. This variable takes on values larger than 0.
- *Salary at the time of the contact over the vacancy salary*: this is a dummy variable that takes on value 1 if the candidate's salary when he/she is included in the consideration set exceeds the vacancy salary.
- *Gender*: search consultants, when they make contacts with individuals, enter their gender in the database. I use this information to assign gender to potential candidates. This is a dummy variable that takes on value 1 if the candidate is female.
- *Age*: Age is measured at the time when a candidate is included in a consideration set. I calculated this variable by subtracting the date of the candidate's birth from the date they were included in the consideration set.
- *MBA*: since education may affect both the person's suitability for the job, and his/her preferences and ability, I included MBA degree as an indicator of the managerial education of the individual. This is a dummy variable that takes on value 1 if a person has an MBA degree.
- *Job tenure*: this variable measures the number of years that the person has held the job they were in at the time when he/she was included in the consideration set.
- *Candidate previously placed by the search firm*: this variable is a dummy variable and takes on value 1 if the person has been placed by Execo in the past.
- *Candidate answered an advertisement*: this variable is a dummy variable and takes on value 1 if the person entered the consideration set by responding to an advertisement.
- *Prior search firm relationship with the hiring firm*: research in the staffing sector has shown that prior relationships between labor market intermediaries and hiring firms may affect

individual-level outcomes (Fernandez-Mateo, 2007; Bidwell and Fernandez-Mateo, 2010). In the executive search setting long-term relationships between clients and search firms are relatively common (sometimes these relationships are formalized with a preferred supplier agreement). This variable measures the number of search assignments that the search firm has conducted for the hiring firm between January 2005 and the start date of the focal assignment.

- *Candidate tie*: I used contact logs in the search firm's database to construct this variable. I counted the number of vacancies the person was considered for by Execo between January 2001 and the date he/she was put in the consideration set for the current vacancy.

- *Client tie*: this variable measures the number of *client* contacts that the person has had with the search firm prior to being included in the consideration set for the current vacancy. Client contact means that the person has been a representative of a hiring firm for some vacancy in the past, or the search firm has held a meeting with the executive through which it attempted to obtain new assignments from the executive's employer.

- *Source tie*: this variable measures the number of *source* contacts that the person has had with the search firm prior to being included in the consideration set for the current vacancy. Source contact means that the search firm contacted the person to request a reference on a particular individual, or to obtain information on potential candidates for the current search firm's assignments.

- *Assignment industry controls*: for every assignment the search firm codes the industries in which the hiring firm (or a part of the hiring firm) operates. I include dummies for 16 industries (plus another dummy for cases where this information is missing): media, IT, pharmaceuticals, agriculture, finance, professions, leisure, engineering and manufacturing, retail, energy, infrastructure, government, health, education, NGO and other. I did not treat one of the industries as a reference category, but rather included all industry dummies in the regression models, and let Stata automatically define one of them as a reference, and exclude

it from the analysis. I adopted the same approach with job function dummies and year dummies.

- *Job function controls*: for every assignment the search firm codes the job function(s) that the hired person will occupy. I include dummies for 18 job functions (plus another dummy for cases where this information is missing): financial services professional, consultant, board member, chief financial officer (CFO), divisional finance director, chief executive officer (CEO), divisional managing director, operations director, government, human resources director (HR director), marketing director, sales director, divisional sales director, legal and governance, non-executive director, pharmaceuticals scientist, management – other, and other.

- *Year controls*: to control for the potential exogenous effects on the executive labor market (economic cycles, corporate performance etc.) I include year dummies.

3.5 The analysis and results

Table 3.1 presents a summary of the descriptive statistics.

Table 3.1: Summary of descriptive statistics^{(1), (2)}

Individual level variables (N=6,308); correlation with salary ratio													
Variable	Mean	Std. Dev.	Cor(p-value)	Variable	Mean	Std. Dev.	Cor(p-value)						
Salary ratio (1)	.971	.421	/	Candidate answered an advertisement (1=yes) (8)	.172	.377	.041 (0.001)						
Salary at the time of the contact over the vacancy salary (1=yes) (2)	.343	.474	.689 (0.000)	Number of the search firm's prior assignments for the hiring firm (9)	1.524	3.023	-.022 (0.071)						
Gender (1=female) (3)	.133	.340	-.031 (0.012)	Number of jobs considered (10)	1.830	2.063	.039 (0.001)						
Age (4)	43.643	5.532	.165 (0.000)	Number of general contacts (11)	1.049	1.697	.174 (0.000)						
MBA (1=yes) (5)	.204	.403	.020 (0.107)	Number of client contacts (12)	1.993	5.473	.196 (0.000)						
Job tenure (in years) (6)	3.414	2.715	-.062 (0.000)	Number of source contacts (13)	1.795	3.015	.120 (0.000)						
Candidate previously placed by the search firm (1=yes) (7)	.060	.237	.031 (0.012)										
Job industry (%); correlation with salary ratio							Job function (%); correlation with salary ratio						
		%	Cor(p-value)				%	Cor(p-value)					
Media		3.89	.015 (0.224)		Financial services professional		8.51	-.060 (0.000)					
IT		8.14	-.020 (0.108)		Consultant		1.46	-.049 (0.000)					
Pharmaceuticals		5.59	.000 (0.956)		Board member		4.98	-.060 (0.000)					
Agriculture		0.12	-.003 (0.776)		CFO		7.78	.018 (0.147)					
Finance		32.20	-.046 (0.000)		Divisional finance director		3.40	.070 (0.000)					
Professions		8.02	-.011 (0.382)		CEO		7.78	.002 (0.867)					
Leisure		6.44	.005 (0.641)		Divisional managing director		14.46	.053 (0.000)					
Engineering and manufacturing		13.49	-.010 (0.397)		Operations director		3.04	-.000 (0.945)					
Retail		16.65	.004 (0.706)		Government		4.01	.095 (0.000)					
Energy		2.31	-.011 (0.350)		HR director		5.35	.010 (0.427)					
Infrastructure		7.05	.009 (0.434)		Marketing director		6.93	.000 (0.973)					
Government		9.60	.119 (0.000)		Sales director		3.52	-.022 (0.071)					
Health		0.61	-.009 (0.443)		Divisional sales director		1.70	-.003 (0.783)					
Education		0.24	-.018 (0.144)		Legal and governance		3.16	-.041 (0.001)					
NGO		1.22	-.007 (0.545)		Non-executive director		1.22	.044 (0.000)					
Other		4.50	-.009 (0.443)		Pharmaceuticals scientist		0.61	-.003 (0.759)					
Missing		0.00	/		Management – other		22.11	-.042 (0.000)					
					Other		15.31	-.043 (0.000)					
					Missing		0.36	-.005 (0.685)					
Correlations (p-value)	1	2	3	4	5	6	7	8	9	10	11	12	13
1	1.000												
2	0.689 (0.000)	1.000											
3	-0.031 (0.013)	-0.029 (0.021)	1.000										
4	0.166 (0.000)	0.137 (0.000)	-0.107 (0.000)	1.000									
5	0.020 (0.112)	0.011 (0.397)	-0.045 (0.000)	-0.013 (0.312)	1.000								
6	-0.062 (0.000)	-0.075 (0.000)	-0.050 (0.000)	0.158 (0.000)	-0.034 (0.007)	1.000							
7	0.031 (0.013)	0.028 (0.027)	0.022 (0.077)	0.053 (0.000)	-0.001 (0.966)	-0.072 (0.000)	1.000						
8	0.041 (0.001)	0.040 (0.001)	-0.071 (0.000)	0.056 (0.000)	0.080 (0.000)	0.035 (0.005)	-0.011 (0.372)	1.000					
9	-0.023 (0.071)	-0.014 (0.261)	0.026 (0.043)	-0.003 (0.783)	-0.001 (0.948)	-0.004 (0.762)	-0.023 (0.069)	-0.014 (0.266)	1.000				

10	0.039 (0.002)	0.049 (0.000)	0.058 (0.000)	-0.005 (0.677)	-0.007 (0.562)	-0.097 (0.000)	0.147 (0.000)	-0.134 (0.000)	0.057 (0.000)	1.000			
11	0.174 (0.000)	0.147 (0.000)	0.030 (0.017)	0.107 (0.000)	0.014 (0.267)	-0.145 (0.000)	0.133 (0.000)	-0.114 (0.000)	-0.030 (0.016)	0.282 (0.000)	1.000		
12	0.196 (0.000)	0.126 (0.000)	0.119 (0.000)	0.159 (0.000)	-0.002 (0.882)	-0.033 (0.009)	0.152 (0.000)	-0.107 (0.000)	-0.029 (0.019)	0.187 (0.000)	0.375 (0.000)	1.000	
13	0.121 (0.000)	0.117 (0.000)	0.040 (0.001)	0.148 (0.000)	0.008 (0.550)	-0.105 (0.000)	0.191 (0.000)	-0.170 (0.000)	0.021 (0.103)	0.350 (0.000)	0.412 (0.000)	0.330 (0.000)	1.000

(1) Candidate statistics are calculated at the “candidate by vacancy” level. Statistics for the number of the search firm’s prior assignments for the hiring firm, calculated at the level of the vacancy, are very similar. Statistics for job industry and job function calculated at the level of the vacancy. Correlation with the salary ratio calculated at the “candidate by vacancy” level. A given vacancy may correspond to several industries and job functions.

(2) Numbers in brackets next to the names of individual level variables denote the number with which the variable is represented in the correlation matrix.

The average salary ratio in the sample (N=6,308) is .971, which means that on average candidates' salaries at the time of being included in a consideration set are 2.9% below the maximum vacancy salary. 34% of candidates in the sample have salaries that are above the vacancy salary. For candidates whose salary does not exceed the vacancy salary, the average salary ratio is .761, meaning that on average the current salary for that part of the sample falls short of the maximum vacancy salary by 23.9%. 13% of the sample is women (less than in the whole sample, 16.53%). The candidates are on average 43.4 years old, and 20% of them have a MBA. They have been in their present job for 3.4 years. 6% of the candidates were previously placed by the search firm, and 17% got into the consideration sets by responding to an advertisement. They have, on average, been considered for 1.8 jobs in the past, and had just over one general contact. They had 2 client contacts, and 1.7 source contacts. On average Execo conducted 1.5 search assignments for the hiring firms before starting the assignment for the current vacancy.

Almost a third of the vacancies were in the finance industry (32.20%). The second most represented industry was retail (16.65%), followed by engineering and manufacturing (13.49%). Other industries represented less than 10% of the vacancies. The most represented job function was management-other (which roughly corresponds to the middle management positions), with 22.11%. The second was other executive jobs (idiosyncratic to the hiring firms), with 15.31%, and the third divisional managing director (14.46%). Both CEO and CFO jobs represent 7.78% of the sample.

Table 3.1 indicates that, as expected, ties are correlated. The correlation coefficients between the four kinds of ties (candidate, general, client and source) are positive and statistically significant. They range from 0.197 (correlation between the candidate and client ties) to 0.412 (general and source ties). In addition, the ties are significantly correlated with other individual-level variables. All four of them have a positive correlation with whether the

candidate was previously placed (the coefficients range from 0.133 for the general tie, to 0.191 for the source tie). The non-candidate ties are also positively correlated with salary ratio (the coefficients range from 0.121 for the source tie, to 0.196 for the client tie). General tie is also positively and significantly correlated with the variable indicating that the candidate's salary at the time of the contact was over the vacancy salary ($r=0.147$), and negatively with tenure in the candidate's present job ($r=-0.145$).

For the analysis I constructed a panel where the unit of analysis was each “candidate by vacancy” pair. I then conducted a series of linear probability models as the dependent variables are binary and the interpretation of the coefficients in such models is easier than in logit models (for a similar use of linear probability models see Waguespack and Sorenson, 2010). One of the critical considerations in any study seeking to understand the impact of relationships is the extent to which the observed effects are a consequence of the unobserved heterogeneity between actors, as opposed to the genuine impact of the relationships themselves. Specifically, different actors may have different relationships due to their individual characteristics, and thereby the effect of these relationships may actually be a side effect of the characteristics of the individuals. I am able to control for such endogeneity as the advantage of the data I am using in this thesis is that some individuals have been considered for several vacancies. I hence conducted a series of linear probability models that take into account individual fixed effects. I also took into account that the standard errors are not likely to be independent across observations for the same individual, and therefore clustered them by individual. I conducted all analyses in Stata 11, using the command `xtreg, vce(robust) fe8,9,10`.

⁸ In Stata 11 `vce(robust)` in the fixed effects specifications is equivalent to `vce(cluster panelvar)`. That means that by specifying `vce(robust)` in the fixed effects specifications one is effectively controlling for heteroskedasticity of the error terms (as cluster option by definition includes also the robust option), as well as for the non-independence of the errors across the cluster in the panel (in my case, an individual).

⁹ When conducting an analysis of the binary dependent variables with linear probability models, an appropriate robustness check is to conduct logistic models as well. While this holds when looking at the cross-section

To test the hypothesis I adopt the following empirical strategy: I test whether the general tie has the predicted impact on the reduction of the high salary ratio candidates' rate of rejection of the search firm interviews. The analysis is therefore focused on the moderating effect of the general tie. It is important to consider moderating as opposed to the main effect of the general tie as the tie in the hypothesis implies different effects on the rejections at the different levels of salary ratio.

Table 3.2 presents the results. The dependent variable in all models is candidate rejects the search firm interview. In all models the sample used for testing the hypothesis includes those observations that were known (i.e. either Execo coded all three possible outcomes for that step as 0, or more than one outcome was coded as 1), and that were not unsuitable (i.e. the search firm didn't drop these candidates). The models therefore evaluate the probability that a candidate rejects the search firm interview, as opposed to goes for the interview.

specifications, it is more complex in the fixed effects specifications. Namely, when doing logistic regressions with the fixed effects the observations where there is no variance on the dependent variable across the cluster in the panel get dropped. In my case this means that the observations where a given individual has always rejected the search firm interview, or never rejected the search firm interview, get dropped. Linear probability models, however, keep all observations in the fixed effects specifications. Due to the resulting differences in the sample sizes the linear probability models and logistic models are not comparable in the fixed effects specifications in the same way that they are in the cross-sectional specifications. I therefore do not compare the results of the linear probability models with fixed effects to the logistic models with fixed effects in the text. None of the substantive results change with the logit specification though.

¹⁰ As described in chapter 2, the three possible outcomes when a candidate is included in the consideration set are: the candidate rejects going for the search firm interview, the candidate is dropped by the search firm, or the candidate goes for the search firm interview. Since the three outcomes are determined simultaneously, the most appropriate model to analyze their probability would be a multinomial logistic regression. The issue with the multinomial regression, particularly relevant in my case, is that it does not enable a specification with the fixed effects. Since being able to run the models with fixed effects is one of the key advantages of my dataset, running the multinomial regressions models would not take advantage of this feature. I therefore focus on linear probability models with fixed effects, and conduct estimations for the rejection of the search firm interview by holding the search firm's rejection at 0 in order to compare the candidate's rejection of the search firm interview to the candidate going for the interview. I employ the same approach for all steps (I use the hiring firm rejections in place of the search firm rejections for the progression from step 2 to step 3). Nevertheless, I conducted all linear probability models in cross-section, and compared the results to the results of the multinomial regressions. The results were similar, and led to the same conclusions. I am not reporting the results for the cross-section models and the multinomial models as they only partially control for unobserved individual heterogeneity, but can provide them on request.

Table 3.2: Linear probability models predicting who does not go to an interview with the search firm

Variable	SUPPLY (candidate not interested)			
	Model 1	Model 2	Model 3 (individual fixed effects)	Model 4 (individual fixed effects)
Female (1=woman)	0.0508+	0.0505+	/	/
	[0.028]	[0.028]		
Age	-0.0057**	-0.0056**	0.0034	0.0020
	[0.002]	[0.002]	[0.072]	[0.072]
MBA (1=yes)	0.0309	0.0308	/	/
	[0.024]	[0.024]		
Job tenure	0.0094**	0.0093**	0.0108	0.0107
	[0.004]	[0.004]	[0.020]	[0.020]
Candidate previously placed by the search firm (1=yes)	-0.1021*	-0.1043*	0.3477	0.3255
	[0.041]	[0.041]	[0.249]	[0.245]
Candidate answered an advertisement (1=yes)	-0.4123**	-0.4127**	-0.3645**	-0.3663**
	[0.024]	[0.024]	[0.063]	[0.063]
Number of the search firm's prior assignments for the hiring firm	-0.0014	-0.0014	-0.0021	-0.0019
	[0.004]	[0.004]	[0.007]	[0.007]
Number of jobs considered	0.0137*	0.0135*	-0.0205	-0.0214
	[0.006]	[0.006]	[0.019]	[0.018]
Number of general contacts	0.0043	0.0117	-0.0338	0.0879
	[0.007]	[0.019]	[0.044]	[0.067]
Number of client contacts	0.0012	0.0040	0.0170*	0.0220+
	[0.002]	[0.004]	[0.008]	[0.012]
Number of source contacts	0.0008	-0.0044	0.0090	-0.0137
	[0.004]	[0.010]	[0.017]	[0.040]
Salary at the time of contact over the vacancy salary (1=yes)	0.0758**	0.0687*	0.0599	0.0545
	[0.028]	[0.029]	[0.068]	[0.068]
Salary ratio	0.1038**	0.1221**	0.3572**	0.5003**
	[0.037]	[0.047]	[0.115]	[0.129]
Salary ratio * general tie		-0.0073		-0.1422*
		[0.017]		[0.057]
Salary ratio * client tie		-0.0024		-0.0064
		[0.003]		[0.011]
Salary ratio * source tie		0.0054		0.0268
		[0.010]		[0.032]
Constant	0.4005**	0.3824*	-0.2299	-0.2920
	[0.154]	[0.155]	[3.293]	[3.271]
N observations	2,671	2,671	2,671	2,671
N of candidates	1,989	1,989	1,989	1,989
R-squared	0.131	0.131	0.124	0.133

+ p<0.10, * p<0.05, ** p<0.01

All tests are two-tailed.

Robust standard errors in parentheses in all models are clustered by individual.

All models include the following controls: assignment industry, job function, year dummies.

R-squared reported for models 1 and 2 are the R-squared for the OLS regression. R-squared reported for models 3 and 4 are within R-squared.

Models 1 and 2 are cross-sectional, and are presented for comparison with the fixed effects models 3 and 4 (the models testing hypothesis 1). Model 1 includes the control variables. In the cross-section older candidates, those that have been previously placed by the search firm, and those that have answered an advertisement, are less likely to reject the interview invitation. Candidates who have been in their current job longer, and who have been considered for more vacancies by Execo in the past, are more likely to reject. As expected, the coefficient on the salary ratio is positive, meaning that as the candidates' current salary approaches the vacancy salary, the candidates are more likely to reject the interview invitation. In addition, candidates whose salary exceeds the vacancy salary are more likely to reject the invitation than candidates whose salary is below the vacancy salary. The latter two results are expected, and in line with the predictions of the human capital literature. Model 2 includes the interaction of salary ratio with the three types of non-candidate ties, including the main independent variable (the general tie). All interaction coefficients are insignificant, meaning that in the cross section the general tie does not have an impact on the probability of the interview rejection.

Models 3 and 4 provide a specification with individual fixed effects. Model 3 includes the controls. Candidates who have responded to an advertisement are less likely to reject the interview invitation. Client ties increase the probability of interview rejections. This is consistent with the preceding description of the client relationship between the search firm and the candidates. Client relationships are characterized by a search consultant's pure brokerage role, creating a context that does not stimulate the executive's development of trust in the search firm (and might even lead to the development of negative trust, i.e. distrust). As expected, salary ratio is positive. Model 4 provides a test of hypothesis 1. It includes the interaction of salary ratio with the general tie. The coefficient on the interaction term is

negative and significant¹¹. This means that while the candidates with high salary ratio are more likely to reject the interview invitations, the probability of rejection will be lower when the candidates have a general tie with the search firm¹². The within-individual specification controls for unobserved time-invariant heterogeneity among the individuals. That is, since the unobserved individual characteristics that might affect the patterns of the individuals' relationships are controlled for, it can be claimed with more confidence that the observed significant effect of the general tie is not caused by the differences in individual characteristics.

The question still remains, though, why we observe a significant interaction between salary ratio and the general tie in individual fixed effects specifications, but not in the cross-sectional specification. One potential explanation is the accumulation of trust through repeated general contacts. The cross-sectional specification estimates the effects taking into account both the individuals who have not been considered for vacancies in the past, and the individuals who have. The individual fixed effects specification, though, estimates the effects on the basis of the individuals who have been considered for at least one vacancy before the current vacancy. Those individuals who have not been considered for vacancies in the past have a significantly lower number of general contacts than those who have been considered

¹¹ The result is the same if I use the logged values of the relationship variables.

¹² An issue with linear probability models is that they may generate the predicted probabilities that are less than 0, and more than 1. I examined the results of model 4, and found that for 81 observations (3.03% of the analyzed sample), the predicted probability of the interview rejection was less than 0, while for 86 observations (3.21%) the predicted probability was more than 1. These results call for an examination of how well the model predicts probability of the outcome coded as 1. I used the procedure described in Herron (1999), which yields percentage of the correctly predicted observations by the model. I first calculated the predicted probabilities that the outcome for each observation in the analyzed sample equals 1. I then generated a variable that took on a value of 1 if the predicted probability of this value being 1 was equal to or more than 0.5 (an arbitrary cut-off point suggested in the article), and 0 if less than 0.5. I then compared the generated values to the actual values for the observations. For model 4 58.89 observations were predicted correctly. This value appeared relatively low, so I calculated the percentage of correctly predicted observations by a cross-sectional linear probability model, and a cross-sectional logistic model for the same sample. The percentage of correctly predicted observations for the first was 65.48%, and for the second 65.74%. The logistic model is fully accurate as it takes into account that the predicted probabilities have to be between 0 and 1. Given that the percentage of correctly predicted observations for the linear probability model with fixed effects (Model 4 in Table 3.2) was lower by 7 percentage points we may conclude that this model may not be accurate enough. As a robustness check I therefore ran the xtlogit version of Model 4. While the procedure estimated the coefficients for only 634 observations, the results were the same as in the linear probability model with fixed effects, suggesting that the results of the latter are accurate.

for at least one vacancy (0.507 vs. 1.283, $p=0.000$). To the extent that more general interactions lead to a greater accumulation of trust, the effect of the general tie and the associated effect of trust will be stronger within-individual, than across individuals. The results in table 3.2 are consistent with that interpretation, and thereby provide evidence in support of hypothesis 1.

Prior relationships are also conduits for transfer of private information (Uzzi, 1997; Podolny, 2001). General ties may serve as the channel for transmission of the private information on the candidates' motivations, abilities etc. This information might in turn result in a better matching of candidates with general ties. That possibility needs to be addressed as the finding that candidates with prior general ties reject the interviews less might simply mean that they are better matched to the vacancies because of previously transmitted private information. To examine this possibility I analyze the candidates' decision making when going from step 1 to step 2, i.e. from the search firm interview to the hiring firm interview. At the search firm interview the candidates get more information about the vacancy and the hiring firm, while the search firm gets more specific information about the candidate. The quote below is a description of the search firm interview by an Execo's senior associate (interviewee 16):

“Our interview is hour and a half, for some jobs which are more technical and we use something called competency based interviewing, so we'll have a very specific list of competencies for the job, some of which will be technical, some will be stylistic. And it's quite well documented way to interview, but it's forensic. We will ask very specific questions designed to draw out people's experience. Asking them about how they would respond to certain situations and use examples from the past of when they've dealt with similar situations. So it's very basic, but the way someone's reacted to a situation in the past is a good indicator of how they might respond to something in the future, and you can also determine perhaps what they've learned, whether they are really going to fit in with the client's organization.”

The quote highlights that the decision to proceed after the search firm interview will rest to a large extent on the degree to which an individual is a good match for the vacancy. If the general ties are associated with a better match, then they should facilitate the candidates' progress from the search firm interview to the hiring firm interview. Since better-matched candidates should be less likely to reject the invitations to go to the hiring firm interview, the general tie should reduce the hiring firm interview rejection rates. In addition, better-matched candidates should be less likely to be dropped by the search firm therefore the general tie should reduce the search firm's rejection of the candidates after the search firm interview.

Table 3.3 provides models that analyze the candidates' decision making once they have made it to the search firm interview. As in table 3.2, models 1 and 2 are cross-sectional, and used for comparison with models 3 and 4 (that use individual fixed effects). In all models the sample used for the analysis includes those observations that were not unknown (i.e. Execo coded all three possible outcomes in the step as 0, or more than one outcome was coded as 1), and that were not unsuitable (i.e. the search firm didn't drop these candidates). The models therefore evaluate the probability that a candidate rejects the hiring firm interview, as opposed to goes for the interview.

Table 3.3: Linear probability models predicting who does not go to an interview with the hiring firm after being interviewed by the search firm

Variable	SUPPLY (candidate not interested)			
	Model 1	Model 2	Model 3 (individual fixed effects)	Model 4 (individual fixed effects)
Female (1=woman)	0.0461 [0.034]	0.0438 [0.034]	/	/
Age	-0.0067** [0.002]	-0.0064** [0.002]	0.1506 [0.153]	0.1403 [0.145]
MBA (1=yes)	-0.0012 [0.028]	-0.0033 [0.028]	/	/
Job tenure	-0.0036 [0.004]	-0.0040 [0.004]	0.0131 [0.037]	0.0042 [0.045]
Candidate previously placed by the search firm (1=yes)	-0.0801* [0.039]	-0.0836* [0.039]	-0.4092 [0.559]	-0.3815 [0.560]
Candidate answered an advertisement (1=yes)	-0.0014 [0.037]	-0.0005 [0.037]	-0.0704 [0.151]	-0.0238 [0.139]
Number of the search firm's prior assignments for the hiring firm	0.0059 [0.005]	0.0059 [0.005]	0.0268 [0.024]	0.0216 [0.023]
Number of jobs considered	0.0056 [0.007]	0.0058 [0.007]	-0.0480 [0.034]	-0.0414 [0.032]
Number of general contacts	0.0037 [0.009]	0.0275 [0.022]	-0.1184 [0.098]	-0.6015** [0.209]
Number of client contacts	-0.0000 [0.002]	0.0049 [0.005]	0.0608* [0.029]	0.1227 [0.097]
Number of source contacts	0.0008 [0.004]	-0.0164 [0.013]	-0.0120 [0.041]	-0.1910* [0.092]
Salary at the time of contact over the vacancy salary (1=yes)	-0.0444 [0.038]	-0.0490 [0.037]	0.3342* [0.134]	0.2730* [0.117]
Salary ratio	0.1235* [0.061]	0.1349* [0.066]	-0.4476 [0.293]	-0.9709** [0.329]
Salary ratio * general tie		-0.0239 [0.019]		0.4391* [0.178]
Salary ratio * client tie		-0.0046 [0.004]		-0.0382 [0.092]
Salary ratio * source tie		0.0178 [0.014]		0.1956* [0.093]
Constant	0.1237 [0.133]	0.1111 [0.139]	-5.7936 [6.484]	-4.5857 [6.088]
N observations	930	930	930	930
N of candidates	840	840	840	840
R-squared	0.087	0.090	0.529	0.585

+ p<0.10, * p<0.05, ** p<0.01

All tests are two-tailed.

Robust standard errors in parentheses are clustered by individual.

All models include the following controls: assignment industry, job function, year dummies.

R-squared reported for models 1 and 2 are the R-squared for the OLS regression. R-squared reported for models 3 and 4 are within R-squared.

Model 1 includes the controls. As in the previous step, in the cross-section the older candidates and candidates that have been placed before by Execo are less likely to reject the hiring firm interview. At the same time, and as in the previous step, the candidates with higher salary ratio are more likely to reject the hiring firm interview. The general tie is not significant in this model. Model 2 includes the interactions with the general tie, as well as the interactions with the two other non-candidate ties (client and source). All interactions, including the general tie, are insignificant, lending no support to the possibility that general ties are associated with a better match of candidates¹³.

Models 3 and 4 are specifications with individual fixed effects¹⁴. Model 3 includes the controls. Within-individual, candidates with a prior client tie are more likely to reject the hiring firm interview. This is consistent with the result for the prior step. It appears that the low trust, or distrust, in the search firm's handling of the candidates that develops in the client relationships exerts a strong influence on the candidates' decision making. High salary ratio stimulates the rejection of the hiring firm interviews, but only when the executive's salary exceeds the vacancy salary. When the salary ratio is below 1, higher salary ratio leads to a lower probability of the hiring firm interview rejection (the effect is only marginally significant though). Model 4 includes the interaction of salary ratio with the general tie. The interaction is positive, while the main effect of the salary ratio becomes negative. This means that as the salary ratio increases, the candidates are less likely to reject the hiring firm interviews. A potential explanation could be that after the search firm interview, when

¹³ No one can reject a hiring firm interview without being interviewed by the search firm in the first place. I ran the models in table 3.3 also by correcting for sample selection, using the procedure described in Hamilton and Nickerson (2003) (i.e. by including the inverse Mills ratio calculated using the selection equation as a control variable in the second stage equation). I used a dummy variable, that takes on value 1 if the vacancy was advertised, as the instrument. When correcting for sample selection, the salary ratio was insignificant in both model 1 and model 2, however the inverse Mills ratio was also insignificant in both models, suggesting that the sample selection is not significantly affecting the results.

¹⁴ I considered running the individual fixed effects models with the specification that would correct for the sample selection in a similar way as I did for the cross-sectional models. However, this would result in inconsistent estimates, as described by Wooldridge in the standard econometrics textbook (Wooldridge, 2002, p. 582). I therefore estimated only standard individual fixed effects models.

candidates obtain more information about the vacancy and the hiring firm, the better matched candidates will be less likely to reject the hiring firm interview. To the extent that the salary ratio, as long as it is below 1, captures how well a candidate's human capital corresponds to the human capital required by the vacancy, this result is consistent with the idea of the lower rejection rates of the better matched candidates. The positive and significant coefficient on the interaction between salary ratio and the general tie, though, means that the high salary ratio candidates with a general tie are more likely to reject the hiring firm interviews. This result is in contrast with the idea that the candidates with the general tie are better matched to the vacancies, since if they were, they should be less likely to reject going for the hiring firm interview. We can therefore not ascribe the result in the previous step to the general ties leading to a better matching of the candidates.

The result in the model 4 is somewhat puzzling as the effect of the general tie on the decision making of the candidates after the search firm interview is the opposite of the effect prior to the search firm interview. One potential explanation might be that there is something in the profiles of candidates with prior general ties that makes these individuals harder to allocate to vacancies for which they would be a good match. As explained by Execo's consultants, the firm prefers to contact executives as clients and sources, as each of these contacts may have an immediate benefit for the search firm. The fact that the firm develops a general tie with a particular individual, as opposed to one of the other two ties, means that at the time the search firm "did not know what to do with that particular person" (but felt that there might be some benefit to be gained from this person in the future). The result in model 4 suggests that this ambiguity may adversely affect the quality of the match when such person becomes a candidate¹⁵. The quality of the match becomes more apparent after the

¹⁵ As I am using an individual fixed effects specification, the impact of the time-invariant characteristics of individuals is fully controlled for. It would seem that a characteristic such as "ambiguity of the profile" may be a time-invariant characteristic of the individual. However, the extent of the ambiguity of the profile at a point in time is relative to the search firm's portfolio of vacancies. The latter changes in time. That is, the profile of a

search firm interview, and at that stage the candidates who are not well matched to the vacancies may select themselves out. To the extent that the profile ambiguity is a characteristic of the general ties, these may be associated with a lower quality of the match, and in turn lead to a greater likelihood of the interview rejection¹⁶.

To assess the plausibility of this interpretation I conducted an analysis of the search firm's decisions. Table 3.4 provides models that analyze the search firm's decision making before and after the search firm interview. All models in the table use the individual fixed effects specification. In all models the sample used for the analysis includes those observations that were not unknown (i.e. Execo coded all three possible outcomes in the step as 0, or more than one outcome was coded as 1), and that were not non-interested (i.e. the candidates did not reject the search firm interview). The models therefore estimate the probability that a candidate is dropped before the search firm interview (as opposed to goes for the interview), and that a candidate is dropped after the search firm interview (as opposed to goes for the hiring firm interview).

given individual may at one point be ambiguous as the search firm does not at the time have vacancies that correspond to his/her profile, and at another point be relatively unambiguous.

¹⁶ This conclusion is based on the comparison of the rejection rates of the candidates with general ties, to the candidates without general ties. It seems somewhat counter-intuitive that the candidates with some prior interactions would be less well-matched to the vacancies than candidates with whom the search firm has had no prior interaction. However, it is important to note that the comparison group in model 4 includes both candidates who had no general contacts with Execo, but had some client and source contacts, and candidates who had none of the three types of contacts. I hence tested whether the results in model 4 are due to candidates having/not having a prior relationship with Execo, or due to having a particular kind of relationship. I re-ran model 4 using a three-way interaction between salary ratio, general tie and a dummy that took on value 1 if the candidate had no prior client or source contact with Execo. The interaction was insignificant, suggesting that regardless of a candidate's prior relationship in other roles, general ties appear to be associated with a lower quality of match. This would be consistent with the explanation that the search firm, at some point in the past, did not quite know "what to do with a person", developed a general tie with him/her as opposed to a client or a source tie, and subsequently this ambiguity of the person's profile carried over also to the matching of this person to the vacancies.

Table 3.4: Linear probability models predicting who is dropped by the search firm

Variable	DEMAND (candidate classified as not suitable by the search firm BEFORE the interview)		DEMAND (candidate classified as not suitable by the search firm AFTER the interview)	
	Model 1 (individual fixed effects)	Model 2 (individual fixed effects)	Model 3 (individual fixed effects)	Model 4 (individual fixed effects)
Female (1=woman)	/	/	/	/
Age	-0.0600 [0.043]	-0.0610 [0.043]	-0.1530 [0.159]	-0.1898 [0.167]
MBA (1=yes)	/	/	/	/
Job tenure	0.0122 [0.015]	0.0118 [0.015]	-0.0243 [0.029]	-0.0320 [0.030]
Candidate previously placed by the search firm (1=yes)	0.3155** [0.118]	0.2980* [0.125]	-0.6857 [0.417]	-0.9404* [0.433]
Candidate answered an advertisement (1=yes)	0.1598** [0.039]	0.1613** [0.039]	0.0589 [0.126]	0.1014 [0.131]
Number of the search firm's prior assignments for the hiring firm	-0.0017 [0.005]	-0.0016 [0.005]	-0.0272+ [0.016]	-0.0242 [0.016]
Number of jobs considered	0.0441** [0.013]	0.0444** [0.013]	-0.0145 [0.036]	0.0065 [0.036]
Number of general contacts	-0.0335 [0.022]	-0.0539 [0.036]	-0.1837* [0.085]	-0.5309** [0.138]
Number of client contacts	0.0008 [0.010]	0.0004 [0.014]	0.0160 [0.023]	0.0783 [0.057]
Number of source contacts	-0.0044 [0.014]	-0.0006 [0.019]	0.0797* [0.039]	0.1825+ [0.108]
Salary at the time of contact over the vacancy salary (1=yes)	-0.0017 [0.040]	-0.0001 [0.040]	0.0458 [0.146]	0.0857 [0.137]
Salary ratio	0.1005 [0.065]	0.0721 [0.080]	-0.6985** [0.268]	-0.9743** [0.286]
Salary ratio * general tie		0.0207 [0.029]		0.3828** [0.100]
Salary ratio * client tie		0.0007 [0.010]		-0.0576 [0.046]
Salary ratio * source tie		-0.0037 [0.015]		-0.1085 [0.092]
Constant	3.1174 [1.966]	3.1947 [1.969]	7.9114 [6.821]	9.8274 [7.146]
N observations	4,340	4,340	1,319	1,319
N of candidates	2,895	2,895	1,152	1,152
R-squared	0.062	0.063	0.336	0.365

+ p<0.10, * p<0.05, ** p<0.01

All tests are two-tailed.

Robust standard errors in parentheses are clustered by individual.

All models include the following controls: assignment industry, job function, year dummies.

Reported R-squared for all models is within R-squared.

Models 1 and 2 analyze the search firm's decisions before the search firm interview. Model 1 includes controls. Candidates that have been previously placed are more likely to be dropped by the search firm, and so are candidates who got into the consideration set by responding to an advertisement. The first result appears counter-intuitive, but it may be explained by the search firm's off-limits policy, whereby candidates who have recently been placed are not to be approached again for a specific period of time (3-5 years). Candidates that have been considered for several jobs in the past are also more likely to be dropped. Execo's consultants explained that the search firm is reluctant to put forward "serial candidates" because their numerous candidacies indicate that they have not been successful in the previous searches. Model 2 includes the interaction of salary ratio with the general tie. The interaction is insignificant (while it was significant for the same stage on the candidate's side).

Models 3 and 4 analyze the search firm's decisions after the search firm interview. Model 3 includes controls. The salary ratio at that stage is negative. Keeping the vacancy salary constant, the search firm is less likely to drop candidates with high salaries. Model 4 includes the interaction of salary ratio with the general tie. The coefficient on the interaction term is positive and significant. This means that as the salary ratio goes up, the search firm is less likely to drop candidates after the search firm interview, but it is more likely to do so if they have a general tie with the search firm. This result is consistent with the idea that the general tie may characterize a candidate with an ambiguous profile, which in turn reduces the quality of the candidate-vacancy match for these candidates. To the extent that the search firm's decision to send the candidates for the hiring firm interviews are based on the quality of the match, as revealed at the search firm interview, this aspect of the general tie could explain the higher dropping rates of candidates with whom the search firm has a general tie¹⁷.

¹⁷ Another possibility could be that the search firm is dropping the candidates with a high salary ratio and the general ties because these candidates are more often dropped by the hiring firms. I unfortunately don't have a

The results for the step after the search firm interview suggest that the effect of the general tie indeed depends on the uncertainty of the quality of the candidate-vacancy match. After the search firm interview, the candidates have more information about the vacancy and the hiring firm, and can better assess their match. They can rely on factual information, and hence appear to rely less on the trust that they have in the search firm. Before the search firm interview, when the uncertainty about the match is greater, and the candidates have little factual information, they seem to rely on their trust in the search firm developed through prior interactions.

The quotes by Execo's senior associate (interviewee 16) on general, client and source ties I provided earlier in the chapter highlight the differences in the extent to which the three types of tie are conducive to the candidate's development of trust in the search firm. The description of the general tie emphasises that the trust development process is likely to be asymmetric, i.e. that while the individual might develop trust in the search firm, the search firm is not likely to see a candidate it has a general tie with in some different way. This is consistent with the results in model 4 in table 3.2 and model 2 in table 3.4. The quotes also suggested a sharp distinction between the general tie and the client and source ties in terms of potential for the development of a candidate's trust. This difference is borne out by the results,

large enough sample at the hiring firm interview stage to conduct an individual fixed effects analysis that would test for this possibility (Stata reports an insufficient observations error). Nevertheless, I ran a cross-sectional model with robust standard errors clustered across the individual, corrected for the selection. The latter was necessary as the candidates can only be dropped by the hiring firms if they make it to the hiring firm interview, so selection into the hiring firm interview needs to be corrected for. I used as the instrument the existence of a prior relationship between Execo and the hiring firm (a dummy variable that takes on value 1 if Execo has conducted at least one search assignment for the hiring firm in the past, and 0 if not). I then adopted the procedure suggested by Hamilton and Nickerson (2003). I first ran a probit model predicting the probability that a candidate makes it to the hiring firm interview (once he/she made it to the search firm interview). I then calculated the inverse Mills ratio as the ratio of the probability density function/cumulative distribution function. I then included the calculated Mills ratio as a control in the linear probability model predicting the likelihood of the hiring firm dropping a candidate. The coefficient on the interaction between salary ratio and the general tie was positive and significant. The hiring firms are therefore more likely to drop candidates with a high salary ratio when they have a prior general tie. I then examined the probability that the high salary ratio candidates with a general tie are more likely to reject the offers (once extended to them). I adopted a similar procedure that corrects for the selection into the job offer, and modelled the probability that a candidate rejects an offer (using the number of people that have responded to an advertisement in the assignment as the instrument). I found no significant effect of the general tie (also when I ran a non-corrected model). These results provide further (albeit only indicative) evidence that the candidates with general ties may be less well-matched to the vacancies, as they are more likely to be dropped by both the search firm and the hiring firm.

as the client and source ties have an insignificant effect on a candidate's decision making before the search firm interview. The quote for the client tie indeed suggests that when an executive is in a client role, he/she might actually develop a certain level of distrust in how the search firm deals with individuals when they are candidates. This observation is consistent with the findings in model 3 in table 3.2, and model 3 in table 3.3 (the positive relationship between the client tie and the candidates' rejection of the search firm/hiring firm interviews).

3.5.1 Robustness of the results

The proposition in this chapter is that trust in a relationship is positively correlated with role autonomy. This, in turn, implies that it is not mere familiarity of executives with the search firm that produces the effects in this setting, but rather trust that facilitates the search firm's persuasion. While the results are consistent with this assumption, an opposite assumption could also be valid: familiarity between executives and the search firm may override the effects of trust and drive the outcomes in the manner predicted by the embeddedness theory (Uzzi, 1996, 1997). Research in the embeddedness tradition generally does not separate out ties between two actors and explore the effects of each kind of tie (as I have done in this chapter). Rather, it aggregates ties and explores the effect of the overall embeddedness between actors on the economic outcomes (e.g. Uzzi, 1999; Beckman and Haunschild, 2002). It generally assumes greater embeddedness will facilitate economic exchanges. In the context of executive search the embeddedness logic would be the following: the more embedded the relationship between an executive and the search firm, the more likely it is that the broker's persuasion efforts will be successful.

I tested this logic in two ways. First, I added up the number of contacts in all three types of non-candidate interactions (general, client and source) and tested the impact of this

cumulative variable on the probability of rejection of the search firm interview. Second, I considered the extent of multiplexity of ties as the extent of embeddedness of ties (Uzzi, 1999). I created a variable that took on values according to the number of different kinds of non-candidate ties between an executive and the search firm (0-3; 0 if no prior ties, 3 if an executive had a prior general, client and source tie with the search firm).

Replication of model 4 in table 3.2 with the interaction between salary ratio and the cumulative tie (in place of the interaction of salary ratio with the general tie) produced an insignificant interaction effect ($p=0.151$). This result suggests that a *cumulative* non-candidate tie between an executive and the search firm does not facilitate the broker's persuasion efforts. Cumulative tie has a different profile of familiarity and trust than the general tie. In addition to the general tie, the cumulative tie includes client and source ties, which involve lower search consultants' role autonomy, and hence lower trust development. In fact, quotes by Execo's senior associate (interviewee 16) suggested that client and source ties may even stimulate the development of distrust, which would mean that the cumulative tie might be characterized with little trust, which would in turn explain the insignificant coefficient.

The result for the cumulative tie also suggests that it is unlikely that the results are driven by mutual obligation. It is often argued that more embedded ties involve more mutual obligations between actors, as more favors get exchanged in longer and closer relationships (Uzzi, 1996, 1997). According to the results in the paragraph above mutual obligation, to the extent that it is present in more embedded relationships in this setting, is not the causal mechanism for the results. Also, theoretically speaking, it does not seem that mutual obligation should play a major role in this setting, for two reasons: first, I found significant effect of the general tie, and general tie is one where no specific favors are exchanged, and hence no future obligations are generated. A quote on the general tie does suggest that

executives may want to trust to the search consultants and talk to them again because they (the search consultants) took the time to brief them on the market dynamics, but this hardly constitutes an obligation. Second, obligations are generated in client and source ties, and even then the obligation is of the search firm to the candidate, so the candidate does not have obligations to the search firm.

Another mechanism that may be present in more embedded relationships is peer pressure. In this context the peers of executives could be other executives. I don't have information on the decisions of other executives the focal executives may have relationships with (or are in some other way relevant for the decision making of the focal executives), hence cannot rule out this possibility, which thereby remains an area for further research. If peers of the focal executives would be the search consultants though, then it seems unlikely that peer pressure would be driving the results as my interviews indicated that search consultants, in general, do not put much pressure on candidates to go for a particular interview, or accept a particular offer (recall quotes by a founder of a boutique search firm (interviewee 10), and a partner at Execo (interviewee 12) in chapter 2).

I also replicated model 4 with the multiplex tie in place of the general tie. The coefficient on the interaction term was negative and marginally significant ($p=0.052$). Since the significance of the coefficient was not within 5%, it was not strong enough to claim that the breadth of the relationship has a significant effect on the participation in brokerage (once we control for the economic incentives). The result for the multiplex variable reinforces the finding obtained using the cumulative tie variable: in this context it is not familiarity per se, but trust that develops through particular kinds of interactions that facilitates the brokered parties' participation in brokerage.¹⁸

¹⁸ A potential explanation for this result could be that the search firm exerts less effort on persuading high salary ratio candidates with whom it has a more embedded prior tie. Although anecdotal evidence from the interviews suggests this is not the case, I ran the analysis in which I estimated the probability of the search firm's rejection of candidates with whom it has a prior embedded tie (greater rates of rejection would indicate lower propensity

3.6 Discussion and conclusions

In this chapter I analyze how prior non-candidate ties between executives and the search firm affect the likelihood that the executives invited to a search firm interview for a vacancy reject the invitation. I have conceptualized the rejection of the search firm interview as the refusal to participate in the brokerage process. I found support for the hypothesis that the prior ties that lead to the development of the candidates' trust in the search firm reduce the executives' reluctance to go for a search firm interview. Since I tested the hypothesis in the context of the labor market, I addressed the possibility that these particular ties are associated with the better matching of the candidates to the vacancies, and that the observed relationship may be caused by the better match rather than the development of trust. I found a reasonable amount of evidence that this is likely not to be the case. I also considered the possibility that the results may be caused by the familiarity associated with all long-term relationships, rather than trust that only develops in a particular kind of tie in my setting. I found no evidence for the familiarity-only explanation.

The study presented in this chapter sheds new light on the role of executive search firms in the high-end managerial labor market. Executive search firms have traditionally been seen as a classical intermediary whose main function is matching (e.g. Cappelli, 2008). Qualitative analyses of the setting have discussed the other functions these organizations are performing, such as buffering between candidates and hiring firms (Khurana, 2002; Finlay and Coverdill, 2002). There has been little work though that would explore how executive search firms set the mobility of executives in motion (but see Hamori, 2010). More specifically, little is known about when executives accept the search firms' invitations to

for persuasion of these candidates). Both main effects of the cumulative tie and multiplex tie, and the interaction terms salary ratio*cumulative tie and salary ratio*multiplex tie were insignificant, lending no support for the lower effort explanation.

actively engage in the pursuit of a new job. This chapter points to the importance of prior interactions between executives and the search firm. Since executives and search firms regularly interact, the chances are that when a new vacancy comes up, and the search firm introduces that vacancy to the executive, the two already have a history of prior interactions. The findings in this chapter demonstrate that the content of these interactions is important, as it can increase or decrease the probability that a given executive will agree to pursue a particular vacancy. Since executives' prior relationships with the search firms shape who competes for vacant positions, search firms play an important role in executives' careers and hiring firms' access to pools of talent.

The findings also have implications for our understanding of how prior interactions between market actors shape market processes. One of the pillars of the embedded ties literature is the notion that the more types of tie there are between two actors, the more they are embedded in that relationship (Uzzi, 1999; Beckman and Haunschild, 2002). Greater embeddedness is associated with a greater familiarity and greater trust, as well as greater levels of mutual obligation and reciprocity (Blau, 1964). This additive assumption appears less sustainable when we consider the trust profiles of the ties within relationships between the same market actors. Some ties may lead to a development of high levels of trust, while some may lead to low levels, or even distrust. More interactions, or a greater breadth of interactions, and hence more embeddedness, may therefore not necessarily result in more trust and, by implication, more exchanges. A fruitful area for future research would be an identification of the conditions under which this is more likely to hold. Nevertheless, this basic result departs in an intriguing way from the extant literature.

The findings also have implications for brokerage theory. The original conceptualization of the theory of structural holes (Burt, 1992) provides little guidance with respect to the question of when brokerage will actually take place. That literature is focused

on the broker, and hence little is known about how the broker secures the participation of the brokered parties. Treatments of brokerage in economics emphasise that the brokered parties may participate in brokerage because of the efficiency benefits that may accrue to them (Autor, 2009). These arguments are echoed by the social networks literature (Brass, 2009). No previous study, to my knowledge, has actually measured the decision to participate in brokerage, though. In this chapter I am addressing this gap by using the decision to participate in brokerage as a dependent variable. The study in this chapter thereby provides an empirically grounded analysis of how brokerage process is set in motion.

Evidence in the present study suggests that the executives' trust may facilitate market brokerage. This result is consistent with studies outside of market context that have demonstrated that brokerage takes place when the brokered parties believe that the broker acts in their interest. In the organizational context a broker's intentions to bring parties together facilitates innovation (Obstfeld, 2005). In the public sphere brokers who are tasked with the role of a liaison need to be seen as impartial by the brokered parties in order to obtain influence and actually broker the proceedings (Fernandez and Gould, 1994; Padgett and Ansell, 1993). The present study provides empirical evidence that also in the market settings trusting ties can set the brokerage process in motion.

The study in this chapter demonstrates that brokerage organizations that operate in market settings have a complex set of relationships with the brokered parties. In particular, brokered parties can have different points of contact with the brokerage organizations. The differences in the attributes of these points of contact may importantly affect how the exchanges between the brokerage organizations and the brokered parties unfold over time. Some points of contact involve pure brokerage roles, and when brokered parties encounter the brokerage organizations through these points of contact, future exchanges may be less likely. Other points of contact, though, involve roles closer in character to boundary-spanning

roles, and if brokered parties encounter the brokerage organizations through these points of contact, future exchanges may be more likely. In short, brokerage organizations are not monolithic, and the point of contact with the brokered parties matters for the latter's subsequent participation in brokerage.

While I study the theoretical arguments in the context of executive search, there are other brokered markets where brokered parties may interact or exchange with the brokerage organizations through different points of contact. In bond markets, for instance, a firm may issue bonds through an investment bank, and in that way operate as a supplier (supplying bonds to the market). The same firm may, at another point in time, and through the same investment bank, invest excess cash in the bonds of other firms. In the latter scenario the firm is a buyer. A similar dynamic is present in the private equity market, where private equity firms both buy and sell stakes in firms, and may do so through the same investment banks. Relationships developed on one side of the market may influence subsequent exchanges on the other side. Studies of the relational mechanisms in investment banking and private equity have provided important theoretical contributions to the understanding of the interface between markets and organizations (Podolny, 2001; Rider, 2009). Studies of prior ties in these and other settings, that would explicitly take into account the shifts in roles, could uncover new mechanisms shaping important market outcomes.

This chapter analyzed how the content of ties in executive search affects the executive candidates' pursuit of vacancies introduced to them by the search firms. The next chapter focuses on one of the ties, candidate tie. Its point of departure is the observation that candidate ties develop through prior rejections. It analyzes how prior rejections of progress through the hiring pipeline affect progress in the subsequent searches. Rejections by both the candidates and the search firm/the hiring firms are considered. Specifically, the chapter focuses on whether these rejections have different consequences for women than for men. In

this way it examines how the content of ties in combination with actor attributes affects the likelihood that labor market actors pursue opportunities introduced to them by their labor market contacts.

CHAPTER 4: THE ROLE OF GENDER AND PRIOR RELATIONSHIPS IN THE EXECUTIVE SEARCH¹⁹

4.1 Introduction

Previous chapter has demonstrated that the content that flows through relationships can importantly affect the labor market participants' pursuit of opportunities. This chapter examines how the content of relationships may interact with the attributes of labor market participants, and thereby shape the functioning of the labor market. One of the attributes that received the most attention of labor market researchers is gender (e.g Petersen and Morgan, 1995; Morgan, 1998; Fernandez-Mateo, 2009). Researchers have also explored how relationships and gender interplay to shape labor market processes and outcomes (Fernandez and Sosa, 2005). Gender may importantly affect who pursues what labor market opportunities, particularly in the high end of the labor market, as the preferences for highly demanding jobs may differ for women and men (see Bertrand, Goldin, and Katz, 2009).

Little is known, though, on how gender and relationships in the labor market jointly affect which vacancies women and men pursue, and which ones they are prevented from accessing. In particular we know little about the role that labor market intermediaries, such as search firms, and the relationships they have with female and male job candidates, play in who even competes for high paying jobs. The focus of this chapter is to shed some light on this question, by examining how prior candidate ties affect the salary of the vacancies that female and male executive candidates are considered for, and exploring how they affect the candidates' progress through the hiring pipeline described in chapter 2. Specifically, we focus

¹⁹ The research presented in this chapter was done in collaboration with Professor Isabel Fernandez-Mateo at London Business School.

on the differences in how the candidate ties with a particular candidate have developed, i.e. through candidate rejections of previous vacancies, or through search firm/hiring firm rejections of the candidate. We don't develop particular hypotheses, but rather empirically analyse these relationships, and discuss at the end of the chapter how the findings inform our understanding of the labor market participants' pursuit of labor market opportunities.

In order to understand the role of executive search firms in the recruitment process, we need to consider two crucial aspects of how these organizations work. First, for any given search there are two actors on the demand side, the search firm and the hiring firm, who make decisions about which candidates pass through to each step of the hiring pipeline (candidates make decisions on the supply side). These two hiring actors may have different incentives and thus play different roles at each stage (see Fernandez-Mateo and King, 2010). Second, searches are often not independent. Both individuals and hiring firms may deal with the same search firm repeatedly, as they either use its services to fill jobs or are contacted as potential candidates for positions. There is some evidence that the relationships that labor market intermediaries maintain with candidates have significant effects on workers' wages (Bidwell and Fernandez-Mateo, 2010). Whether these relationships also determine who is more likely to be considered and hired for jobs is an open empirical question.

4.2 Prior relationships, gender and progress through the hiring pipeline

Although, to the best of our knowledge, no other study has systematically examined how gender affects how prior relationships between candidates and search firms may affect hiring outcomes, relationships have been shown to play a role in recruitment processes more generally. As indicated by the literature review in chapter 2 there is a considerable amount of evidence on the role of social ties in helping individuals find jobs.

We build on this work and examine whether relationships between candidates and executive search firms may differentially affect women's chances in the hiring process. Some authors have speculated that the fact that the labor market for high-level managers is brokered by a relatively small number of search firms may put women at a disadvantage (Brett and Stroh, 1997). The argument is that women are not as well connected to these brokers, which are staffed mostly by men and may operate under principles of social similarity – thus affecting which candidates they present to their clients (see also Bielby and Bielby, 1999; Finlay and Coverdill, 2002). Although these claims have not been formally examined, they implicitly assume that women have either weaker ties to executive search firms or the ties they do have do not generate as many returns as men's. This is not a far-fetched assumption, since prior literature has shown that women's patterns of relationships may differ from men's (Ibarra, 1997) and that, even when they do not differ, women may benefit less from relationships than men do (Ibarra, 1992; Burt, 1998). The specific ways in which these gendered relational processes may affect women's access to high-end jobs brokered by executive search firms have not been examined so far.

In this chapter we start filling in this gap. A particular advantage of the dataset we use²⁰ is that it allows us to distinguish who rejected whom in the current search, and who did it in the past searches. This is particularly relevant for understanding women's access to high-end positions. It is sometimes claimed that women may opt out of competition for these jobs given the heavy demands they impose on individuals (see Bertrand, Goldin, and Katz, 2009). Yet, unless we observe each actor's decisions at each stage of the process, it is difficult to draw such an implication.

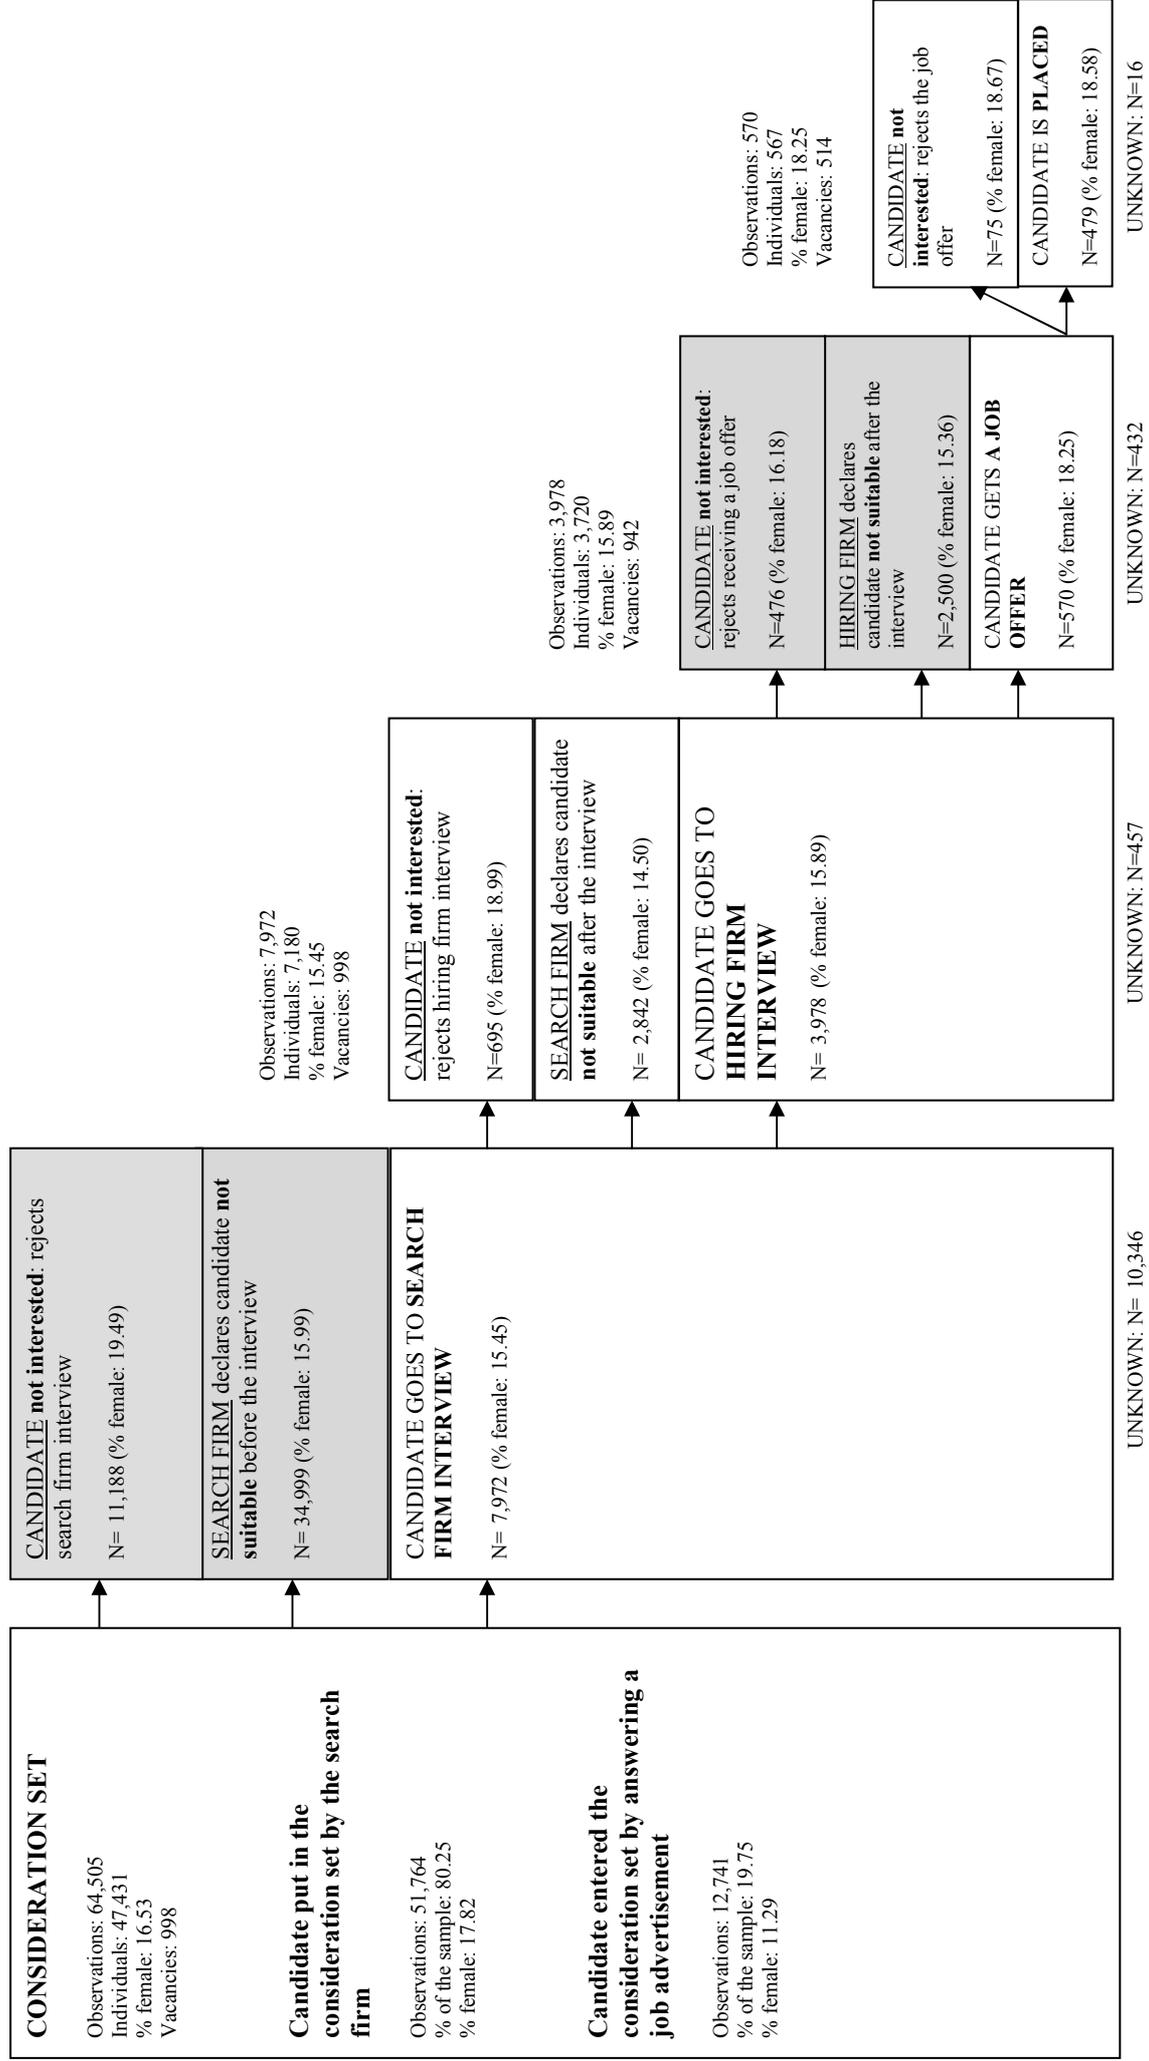
²⁰ Chapter 2 provides a detailed description of the dataset.

4.3 The sample and variables

For this chapter we are using the complete dataset (N=64,505). This is in contrast to chapter 3, where only those observations for whom the full array of individual controls was available were included in the sample (N=6,308). An advantage of using the whole dataset is that it gives greater statistical power. A disadvantage is that we cannot control for economic incentives of candidates, which, as demonstrated by chapter 3, are relevant for their decisions to reject vacancies. We reasoned that given the purpose of the analysis in this chapter the advantage of using the whole dataset out-weighs the disadvantage. Namely, the intent in this chapter is to understand both the decisions of candidates, and the search firm and the hiring firms. For the latter two parties economic incentives of the candidates are likely to be less relevant in their decisions over whom to reject than other factors that we explore. We therefore felt that the price of not being able to control for economic incentives was a price worth paying for being able to obtain much greater statistical power.

The dataset is described in detail in chapter 2. Figure 2.2 in chapter 2 provides the numbers of candidates in different steps, along with the number of candidates that were interviewed, number of candidates that rejected interviews/offers, and number of candidates that were dropped by the search firm/the hiring firm. Figure 4.1 below provides information on the percentage of women in each of the categories.

Figure 4.1: Hiring pipeline at Execo



In the overall dataset women represent 16.53% of the observations. The percentage of women among the candidates that were included in the consideration sets by the search firm is slightly larger, 17.82%. The percentage of women among candidates who have answered an advertisement is much smaller though, 11.29%. The percentage of women that have rejected the search firm interview is 19.49%, larger than the overall proportion of women in the dataset, indicating the possibility that women are more likely to reject the search firm interviews (we examine this possibility below). The percentage of women that are rejected at that early stage is lower than the overall percentage, indicating the possibility that the search firm is less likely to drop women. These proportions hold also in step 2 (i.e. after the search firm interview). As a consequence, the percentage of women that is interviewed by the search firm and the hiring firm is lower than the proportion of women in the overall dataset. In contrast, this percentage is higher at the offer stage. The proportion of women that receive an offer, after having been interviewed by the hiring firm, is 18.25%, almost 2 percentage points higher than the percentage in the overall dataset. These numbers suggest that in this setting women are dropping from the hiring process because they are pulling out (as opposed to being rejected by the search firm or the hiring firms). To explore this possibility further, we move on to the multivariate analysis.

Our statistical analysis uses each “candidate by job” pair as unit of observation. That is, we build a panel of candidates and the jobs they were ever considered for. Panel structure is described in detail in chapter 2. We then study the determinants of candidates being put in the consideration set for a given job (using OLS models), as well as their probability of moving through the several steps of the process in Figure 4.1 (using linear probability models with binary dependent variables for each outcome of interest). The regressions are thus run at the individual level of analysis, and they include robust errors clustered by individual, and individual fixed effects.

We also considered using the event history analysis approach (EHA). EHA is a methodology used to study events, where an event refers to a change from one state to another (Castilla, 2007). EHA is used to analyze longitudinal data for a sample of individuals or units during a period of time when a series of events may occur. The focus is to understand the main determinants of the occurrence and timing of events. Given that our intent was to understand the determinants of rejections, and that candidates are always at risk of further candidacies and further rejections, our data would lend itself well to the EHA. A rejection is essentially a transition from a state of being a candidate to state of being a non-interested or not-suitable candidate. EHA could be used to assess the effect of independent variables, including the effect of prior rejections, on the hazard of these transitions.

However, any analysis of decisions needs to take into account the unobserved heterogeneity between individuals. This is commonly done with the individual fixed effects models. EHA for repeated events with fixed effects would be a viable option for such analysis, as it allows estimation of different baseline hazards for each stratum, e.g. individual (Allison, 1996; Allison, 2009). The stratifying approach has for instance been previously used to study product introductions after acquisitions (Puranam, Singh and Zollo, 2006) and patent granting (Reitzig and Puranam, 2009). In the former study the authors designed the study in such a way that first post-acquisition product introductions have a unique hazard that is different from the second product introduction, which is in turn different from the third product introduction etc. A similar approach was used in the latter study where patents from different technology areas were considered to have different baseline hazards.

The stratifying proportional hazard technique, though, is sensitive to the degree of censoring in the data (Allison, 1996). Specifically, when during the period of observation there are few transitions between the states, and consequently censoring is significant (i.e. few transition events occur during the observation window), the estimates of coefficients for

variables that describe the preceding event history tend to be severely biased (Allison, 1996). More specifically, the fixed effects proportional hazard estimator tends to find misleadingly large negative effects of the number of prior events (Allison, 1996, p. 221). The dataset used in this thesis is characterized by severe censoring, as individuals on average experience a small number of events (i.e. transitions from being a candidate for one vacancy to being a not-interested candidate/not-suitable candidate/placed candidate, to a candidate for another vacancy). For instance, 57% of candidates in the sample have not experienced a transition into being a candidate again, after having been a candidate for the first time during the observation period. The average number of rejections is hence also very low (see the descriptives below). Given these data properties, there is a possibility that the effects of the number of prior rejections would be biased towards large negative values. Due to the possibility of biased estimates we decided not to use EHA, but rather use a linear probability methodology with individual fixed effects. Due to the shortcomings of the linear probability methodology (described in chapter 2), we also ran the logit regressions for all our models with the binary dependent variable. The main results held with both approaches.

We next describe the main variables used in the analysis. The dependent variables are shaded in grey in Figure 4.1.

Variables

Dependent variables

- *Maximum salary for the position*: the maximum salary for the position is the maximum base salary that the hiring firm is willing to pay the successful candidate. This salary is set at the outset of the search. As the distribution of this variable is skewed, we use the natural log of the absolute value in the analysis.

- *A candidate rejects the search firm interview*: this is a dummy variable that takes on value 1 if a candidate rejects the search firm interview, and 0 if he/she goes for the search firm interview.
- *A candidate rejects receiving a job offer*: this is a dummy variable that takes on value 1 if a candidate rejects receiving a job offer once he/she has been interviewed by the hiring firm, and 0 if he/she receives an offer.
- *The search firm declares a candidate not suitable before search firm interview*: this is a dummy variable that takes on value 1 if the search firm drops a candidate before the search firm interview, and 0 if he/she goes for the search firm interview.
- *The hiring firm declares the candidate not suitable after the interview*: this is a dummy variable that takes on value 1 if a hiring firm does not extend a job offer to a candidate after interviewing him/her, and 0 if he/she receives an offer.

Independent variables

- *Number of jobs that the candidate rejected (“not interested”)*: This variable counts the number of jobs for which the candidate was considered, yet he or she rejected to move forward with the hiring process. The rejection may have occurred at any step in the pipeline (i.e. prior to the search firm interview, after the search firm interview, after the hiring firm interview, or after having received an offer).
- *Number of jobs for which the candidate was rejected (“not suitable”)*: This variable measures the number of positions for which the candidate was considered but was classified as not suitable at any step of the process.
- *Gender*

Control variables

We used the following controls:

- *Number of contacts in other roles*: this variable measures the number of meetings between Execo and the candidate in which the candidate played a non-candidate role (i.e. a client, a general source of information on the industry, or a provider of referrals).

- *Number of jobs considered - the outcome unknown*: we counted the number of vacancies for which the outcome was unknown at any stage of the process. We coded the outcome as unknown if the person was included in the consideration set for a vacancy, yet Execo coded all subsequent categories as 0 (i.e. stating that the candidate was not not-interested, not not-suitable, not interviewed by Execo, not not-interested after the Execo interview etc.). We also coded the outcome as unknown if we had evidence that the person made it to an Execo interview, yet had all categories for the step after the Execo interview coded as 0 (likewise for the hiring firm interview, and the offer). Finally, we coded the outcome as unknown if the person had more than one category in a step coded as 1 (e.g. as if the person was both not interested and not suitable before the search firm interview in the same vacancy, not interested and not suitable after the search firm interview etc.).

In addition we used the control variables that were also used for analyses in chapter 3 (and are described there): *Candidate previously placed by the search firm*, *Candidate answered an advertisement*, *Prior search firm relationship with the hiring firm*, *Number of jobs considered*, *Assignment industry controls*, *Job function controls* and *Year controls*.

Table 4.1 below summarizes the descriptive statistics.

Table 4.1: Summary of descriptive statistics⁽¹⁾

Individual level variables, by gender (N=64,505)							
Variable	Men N=53,841	Women N=10,664	p-value (t-test)	Variable	Men N=53,841	Women N=10,664	p-value (t- test)
Maximum salary for the position	156,589	148,777	0.000	Number of jobs considered	.770	.738	0.033
Candidate previously placed by the search firm (1=yes)	.022	.018	0.027	Number of jobs: candidate not interested	.249	.290	0.000
Candidate answered an advertisement (1=yes)	.209	.134	0.000	Number of jobs: candidate not suitable	.373	.304	0.000
Number of the search firm's prior assignments for the hiring firm	1.469	1.656	0.000	Number of jobs: the outcome unknown	.062	.061	0.789
Number of contacts in other roles	2.386	2.637	0.000				
Job industry (%)				Job function (%)			
	Men N=53,841	Women N=10,664	p-value (chi square)		Men N=53,841	Women N=10,664	p-value (chi square)
Media	2.64	3.23	0.000	Financial services professional	12.91	10.45	0.000
IT	6.76	6.77	0.976	Consultant	2.17	1.65	0.000
Pharmaceuticals	4.22	8.18	0.000	Board member	4.44	3.23	0.000
Agriculture	0.09	0.13	0.249	CFO	7.28	3.23	0.000
Finance	34.73	30.79	0.000	Divisional finance director	4.82	2.25	0.000
Professions	9.99	8.72	0.000	CEO	7.07	3.46	0.000
Leisure	7.39	6.96	0.120	Divisional managing director	15.49	8.28	0.000
Engineering and manufacturing	12.86	8.14	0.000	Operations director	3.08	1.84	0.000
Retail	15.37	17.83	0.000	Government	3.33	3.77	0.019
Energy	1.73	2.82	0.000	HR director	3.21	14.45	0.000
Infrastructure	8.96	7.05	0.000	Marketing director	5.94	7.16	0.000
Government	8.87	9.18	0.305	Sales director	3.33	2.34	0.000
Health	0.70	0.71	0.954	Divisional sales director	2.08	1.13	0.000
Education	0.32	0.17	0.012	Legal and governance	3.85	7.21	0.000
NGO	1.47	1.83	0.005	Non-executive director	1.15	0.91	0.032
Other	4.40	4.18	0.294	Pharmaceuticals scientist	0.91	.0135	0.000
Missing	0.00	0.00	/	Management – other	20.06	27.41	0.000
				Other	13.45	11.48	0.000
				Missing	0.26	0.09	0.001

(1) Candidate statistics are calculated at the “candidate by vacancy” level. Statistics for maximum salary for the position, and number of the search firm’s prior assignments for the hiring firm, calculated at the level of the vacancy, are very similar. Statistics for job industry and job function calculated at the “candidate by vacancy” level. A given vacancy may correspond to several industries and job functions.

4.4 The analysis and results

We start our exploration of gender differences in this setting by asking a very simple question: do men and women compete for different types of jobs? (Step 0 in Figure 4.1). The descriptive statistics in Table 4.1 indicate that there are few differences in male and female probability of being considered for jobs in different industries. Women are more likely to be found in media, pharmaceuticals, retail and energy, while men are more prevalent in finance, professions, engineering and manufacturing. Differences in job functions are more noticeable, with women less prevalent in finance jobs (investment banking, CFO, etc.) and top management positions (CEO, Divisional Managing Director, etc.), and much more prevalent in Human Resources. We also explored whether more male jobs are better paid than more female jobs in this sample. While the patterns are consistent with what one would expect (finance is better paid than HR), they are not striking. This is because the range of salaries within each job category is quite wide. For example, “finance service professionals” jobs range from £60,000 to £1,000,000, while HR Director jobs pay from £80,000 to £250,000.

In order to better understand whether women compete for lower paid jobs than men in this setting, we ran a series of OLS regressions examining the effect of being a female on the salary of the job for which the individual is considered (see Table 4.2).

Table 4.2: OLS regression predicting natural log of the salary of the vacancy for which the candidate is considered in Step 0 of the hiring pipeline

Variable	Ln (vacancy salary)			
	Model 1	Model 2 (individual fixed effects)	Model 3 (individual fixed effects)	Model 4 (individual fixed effects)
Candidate previously placed by the search firm (1=yes)	0.0426** [0.010]	0.0320 [0.050]	0.0468 [0.050]	0.0462 [0.050]
Candidate answered an advertisement (1=yes)	-0.2014** [0.003]	-0.0598** [0.007]	-0.0636** [0.007]	-0.0637** [0.007]
Number of the search firm's prior assignments for the hiring firm		0.0015* [0.001]	0.0015* [0.001]	0.0015* [0.001]
Number of contacts in other roles		0.0013 [0.001]	0.0009 [0.001]	0.0009 [0.001]
Number of jobs considered		0.0097** [0.002]		
Number of jobs: candidate not interested			0.0419** [0.005]	0.0426** [0.006]
Number of jobs: candidate not suitable			0.0055* [0.003]	0.0068* [0.003]
Number of jobs: the outcome unknown			0.0015 [0.006]	0.0014 [0.006]
Female (1=woman)	-0.0306** [0.004]	/	/	/
Female * number of jobs not interested				-0.0011 [0.010]
Female * number of jobs not suitable				-0.0118+ [0.006]
Constant	12.0618** [0.011]	11.8659** [0.025]	11.8581** [0.025]	11.8573** [0.025]
N observations	64,505	64,505	64,505	64,505
N of candidates	47,431	47,431	47,431	47,431
R-squared	0.328	0.158	0.161	0.161

+ p<0.10, * p<0.05, ** p<0.01

All tests are two-tailed.

Robust standard errors in parentheses are clustered by individual.

All models include the following controls: assignment industry, job function, year dummies.

R-squared reported for model 1 is the R-squared for the OLS regression. R-squared reported for models 2, 3 and 4 are within R-squared.

Model 1 includes only gender, year dummies, assignment industry and job function dummies, prior placement, and the mode of entry in the consideration set (by answering an advertisement or not). In this case the female coefficient is negative, which indicates that, on average in the cross-section, women are considered for the lower paid jobs.

Models 2, 3 and 4 use the individual fixed effects specification to control for unobserved heterogeneity between individuals. In Model 2 we explore whether the candidates' prior relationship with Execo makes any difference on the pay rate of the job they are considered for. We first consider the number of jobs that the individual has been considered for in the past. We also control for their prior interactions with the search firm in roles other than candidate, as well as for the prior relationship between Execo and the hiring firm. The results show that having been considered for more jobs in the past increases the probability that the individual will be considered for higher paying jobs in the future. This could lead us to conclude that, in this setting, prior candidate relationships are positive.

However, the number of jobs considered variable does not distinguish between past jobs in which the candidate was classified as not suitable by the search firm versus those in which he/she declared not to be interested. Model 3 makes this distinction and shows that it is indeed prior candidate relationships per se that have a positive effect, as both having been not interested and not suitable have a positive effect on the salary of the vacancy a candidate is considered for. Next, we examine whether the effect of these relationships varies by gender, by introducing interaction effects between the relationship variables and the female dummy. Separate descriptive analyses indicate that, at a given point in time, women have been previously considered for 0.73 jobs, as compared to 0.77 for men ($p=0.033$). While they have been considered for fewer jobs, they are more likely to have rejected jobs (0.29 versus 0.25, $p=0.000$). Women have also been considered not suitable for fewer jobs (0.30 versus 0.37, $p=0.000$). The results in Model 4 report the interactions of prior rejections with gender. While

the main effects of the two relationship variables remain unchanged, the interaction between gender and not having been suitable in the past is negative. The coefficient is only marginally significant though ($p=0.053$)²¹. This gives some evidence that for women, being not suitable in the past has a negative effect on the probability that they will be included in the consideration sets for higher paid jobs in the future. In other words, it suggests that women (but not men) are “penalized” for not being suitable - by being considered for lower paid jobs in the future.²²

Even though the interaction between gender and prior rejections is not very strong, it still raises the question of what might be driving it. A pertinent question is whether the effect may be due to women applying for higher paid jobs at different rates than men. Fortunately, we can explore this issue because we have data on whether the candidate applied for the job or was put into the consideration set by Execo without having responded to an advertisement. In separate analyses we first analyzed whether women are more or less likely to apply for jobs. We found, both in the cross-tabs and the logistic regressions with controls, that women are less likely to apply for jobs overall. We then examined what kinds of jobs they apply for when they do apply. We found that they apply for higher paid jobs. There is thus no evidence that women are more likely to apply for lower paid jobs. When women are included in the

²¹ We also ran this model using the natural log of the relationship variables, and the coefficient on the interaction between gender and number of jobs not suitable was still negative, but marginally significant ($p=0.084$). Logarithmic transformations are often used in the studies using relationship variables to account for the assumption that the marginal effects of increases in the length or strength of the relationship decrease at the high values. In our case less than 10% of observations have the number of jobs not suitable higher than 1, so the logarithmic transformation appears unnecessary, and potentially results in the variables that less accurately represent the actual relationships.

²² The number of jobs for which people have been considered in the past in this setting is quite low on average. In separate analyses we created binary indicators for each of the relationship variables. That is, we examined the effects of having ever been considered in the past, having ever been rejected, and having ever been “not interested.” While in contrast with model 2 above just being ever considered in the past does not have a significant effect, and “ever rejected” is also not significant, the effect for the “ever rejected” is the same as the effect for the corresponding variable above. The interaction between female and having ever been rejected is not significant, suggesting that the prior rejections of women have an effect only as the number of rejections increases.

consideration set by the search firm, though, they are allocated to lower paying vacancies²³. The results thus suggest that the patterns in Table 4.2 are not driven by supply-side factors. The search firm seems to be playing at least some role in allocating women to different types of jobs in this setting (i.e. to lower paid jobs, and more so for women if they have not been suitable in the past).

Progress through the pipeline: Probability of doing an interview with Execo

Beyond the fact that there are some gender differences in the first step of the matching process (i.e. “attachment” of candidates to positions), we are interested in exploring whether men and women stand different chances of going through the next steps, and if so, whether this varies with the relationship they have with the search firm. We examine the transition from step 0 to step 1 in Figure 4.1, by analyzing the probability that the candidate will *not* make it to the interview with Execo (shaded boxes in Step 1). We look at those who drop out of the pipeline – rather than those who continue – because this allows us to better disentangle the supply-side versus demand-side reasons why candidates advance through the process. While progressing to the interview step is a combination of Execo’s and the candidate’s preferences (which we cannot disentangle in this case), the reasons for not doing an interview can be clearly identified as either the candidate’s or the search firm’s.²⁴ This allows us to gain more insight into supply versus demand factors that influence progressing through the hiring pipeline.

²³ Both results (women applying to higher paying vacancies and the search firm allocating them to lower paying vacancies) hold also in the individual fixed effects specification.

²⁴ In separate analyses we ran cross-sectional linear probability models with “making it to the search firm interview” as the dependent variable. We found that in the cross-section women are less likely to be interviewed by Execo. We also ran the individual fixed effects specifications with the Execo interview as the dependent variable. We found that women that have been unsuitable in the past (but not women that have been non-interested) are less likely to be interviewed by Execo. .

Table 4.3: Linear probability models predicting who does not go to an interview with the search firm

Variable	SUPPLY (candidate not interested)			DEMAND (candidate classified as not suitable by search firm)		
	Model 1	Model 2 (individual fixed effects)	Model 3 (individual fixed effects)	Model 4	Model 5 (individual fixed effects)	Model 6 (individual fixed effects)
Candidate previously placed by the search firm (1=yes)	-0.0937** [0.023]	0.0819 [0.088]	0.0814 [0.085]	-0.0680** [0.016]	0.3715** [0.090]	0.3717** [0.090]
Candidate answered an advertisement (1=yes)	-0.5970** [0.009]	-0.3072** [0.033]	-0.3128** [0.033]	0.1042** [0.004]	0.1081** [0.016]	0.1080** [0.016]
Number of the search firm's prior assignments for the hiring firm	-0.0012 [0.001]	-0.0052+ [0.003]	-0.0059+ [0.003]	0.0023** [0.001]	0.0015 [0.002]	0.0015 [0.002]
Number of contacts in other roles	-0.0006 [0.001]	0.0047+ [0.003]	0.0046+ [0.003]	-0.0023** [0.000]	-0.0022 [0.003]	-0.0022 [0.003]
Number of jobs considered	-0.0051+ [0.003]			-0.0054** [0.002]		
Number of jobs: candidate not interested		-0.2674** [0.015]	-0.2812** [0.017]		0.1603** [0.017]	0.1529** [0.020]
Number of jobs: candidate not suitable		0.1466** [0.013]	0.1326** [0.014]		-0.0152** [0.005]	-0.0154** [0.005]
Number of jobs: the outcome unknown		0.0514+ [0.029]	0.0516+ [0.028]		0.1574** [0.020]	0.1575** [0.020]
Female (1=woman)	0.0562** [0.009]	/	/	0.0231** [0.005]	/	/
Female * number of jobs not interested			0.0469 [0.030]			0.0314 [0.039]
Female * number of jobs not suitable			0.0832** [0.029]			0.0005 [0.013]
Constant	0.6196** [0.021]	0.5007** [0.047]	0.5042** [0.047]	0.7727** [0.013]	0.7451** [0.026]	0.7458** [0.026]
N observations	19,160	19,160	19,160	42,971	42,971	42,971
N of candidates	16,158	16,158	16,158	33,877	33,877	33,877
R-squared (within)	0.115	0.216	0.220	0.029	0.047	0.047

+ p<0.10, * p<0.05, ** p<0.01

All tests are two-tailed.

Robust standard errors in parentheses are clustered by individual.

All models include the following controls: assignment industry, job function, year dummies.

R-squared reported for models 1 and 4 are the R-squared for the OLS regression. R-squared reported for models 2, 3, 5 and 6 are within R-squared.

Table 4.3 shows a series of linear probability models (LPM), examining the factors that affect the candidate's decision not to be interviewed (Models 1-3) and Execo's decision not to interview the candidate (Models 4-6). We use a linear specification, even though we have dichotomous variables, because it is easier to interpret than the logistic regressions²⁵. Unless otherwise noted, we use the individual fixed effects specifications to control for unobserved individual heterogeneity. We re-ran all cross-sectional models with logistic regressions (available from the authors), and the results are substantively the same in all cases – except when we mention otherwise.

Basic cross-tabs indicate that women are more likely to decline an interview (20.44% versus 16.73% for men, $p=0.000$). Conversely, they are less likely to be classified as “not suitable” (52.48% versus 54.61% for men, $p=0.004$). These results, however, do not control for other factors that may affect the individual's probability not to move to the next step of the process. Model 1 in Table 4.3 examines the probability that the candidate will not be interested in an interview, controlling for prior placement, having answered an advertisement, prior contacts in the non-candidate roles, as well as including year, job function, industry and prior relationships with Execo. A number of controls are worth mentioning. Candidates who have been previously placed by Execo are less likely to decline an interview, and so are those who have applied for the job (answered an advertisement). As for our main variables of

²⁵ As described in chapter 2, the three possible outcomes when a candidate was included in the consideration set are: the candidate rejects going for the search firm interview, the candidate is dropped by the search firm, or the candidate goes for the search firm interview. Since the three outcomes are determined simultaneously, the most appropriate model to analyze their probability would be a multinomial logistic regression. The issue with the multinomial regression, particularly relevant in our case, is that it does not enable a specification with fixed effects. Since being able to run the models with fixed effects is one of the key advantages of this data, running the multinomial regressions models would not take advantage of this feature. We therefore focus on the linear probability models with fixed effects, and conduct the estimations for the rejection of the search firm interview by holding the search firm's rejection at 0 in order to compare the individual's rejection of the search firm interview to the individual's going for the interview. Similarly, we conduct the estimations for Execo's rejections by holding the candidates' rejections at 0 in order to compare the search firm's rejection of a candidate with the individual going for the interview. We employ the same approach for all steps (we use the hiring firm rejections in place of the search firm rejections for the progression from step 2 to step 3). Nevertheless, we conducted all linear probability models in cross-section, and compared the results to the results of the multinomial regressions. The results were similar, and led to the same conclusions. We are not reporting the results for the cross-section models and the multinomial models as they only partially control for unobserved individual heterogeneity, but can provide them on request.

interest, the results indicate that in the cross-section, women are more likely to refuse being interviewed, while candidates who have been considered for more jobs in the past are marginally less likely to decline the interview (the coefficient is negative, yet only marginally significant).

As we did in the prior step, we next disaggregate the “number of jobs” variable into those for which the individual was not interested versus those for which he/she was not suitable (Model 2). We find that as the individuals are considered for more jobs, once an individual has been not interested in the past he/she is less likely to be not interested in the future. Conversely, the candidates who have been not suitable in the past, when they are considered again they are more likely to be not interested. Model 3 includes interactions of these variables with gender. The negative effect of not having been interested in the past on the probability of not being interested in interviewing for the current job is not different for men versus women. In contrast, not having been suitable in the past has a positive effect on not being interested in the current job, and its interaction effect with female is positive. This means that people who have been considered not suitable in the past are more likely to reject an interview, and this is even more so for women. It seems that prior rejections are more influential for women’s decisions not to move forward in the process.

Overall, the pattern of effects of prior relationships with Execo for the candidate’s decision not to move forward in the process is as follows: those who had been not interested in the past are more likely to be interested this time around (potentially, as per results in the table 4.2, because they get allocated to the higher paying vacancies). Those who were classified as not suitable in the past are less likely to be interested this time, and this is more so for women – who, apparently, do not appreciate having been rejected in the past (and more

so than men). In separate analyses we explored whether these effects vary by levels of the salary offered for the job, and found that they do not.²⁶

We next explore Execo's decision to reject the candidate (i.e. classify him/her as "not suitable") before the interview. That is, we examine which factors lead the search firm not to even invite the individual for an interview. Model 4 indicates that in the cross-section Execo is more likely to declare women unsuitable. It is less likely to declare as unsuitable the candidates that are more "connected" to it (coefficients on the previous placement of the candidate, number of prior jobs considered and number of contacts in the other roles are all negative). Models 5 and 6 are estimated using individual fixed effects. Model 5 splits the number of jobs into those for which the candidate had not been interested and those for which he/she had not been suitable. The former has a positive effect on the probability of being rejected by Execo, while the latter has a negative effect. In other words, Execo tends to reject people who rejected interview invitations in the past, but it tends to bring for the interviews people who were not suitable before.

Finally, Model 6 explores whether these relationship effects vary by gender. Adding the interaction terms does not change the main effects of the relationship variables, and the interaction with female is not significant. That is, women are no more or less likely than men to be rejected by the search firm if they were not interested or not suitable before. We therefore find little evidence that Execo tends to treat prior rejections by women in a different way from those by men. For women, not being suitable in the past does not significantly decrease their probability of being considered for a higher paid position in the future (Table 4.2), and at the same time it does not affect the probability that they will be granted an interview with Execo in the future in a different way than it does for men (Table 4.3).

²⁶ We split the sample at the median vacancy salary of £130,000, and ran the estimation using the two subsamples.

We also analyzed the progression from step 1 to step 2 of Figure 4.1 (going to an interview with the hiring firm after having been interviewed by Execo), using the individual fixed effects specification. We found the same results as for the progression from step 0 to 1 on the supply side, and similar results for the effect of the prior relationships on the demand side. In the interest of brevity we do not report the results for that stage, but do summarize the differences on the demand side. The effect of the prior candidates' rejections is insignificant in that step, meaning that Execo is no more or less likely to send for the hiring firm interviews the candidates who have been not interested in the past. Also, as in the previous step, candidates who have not been suitable before are less likely to be dropped by Execo before the interview in the current vacancy (in this step the hiring firm interview). However, in this step the likelihood is even lower for women than for men. That is, women are less likely than men to be dropped before the hiring firm interview if both have been not suitable in the past. Women therefore seem to find it easier to progress through the Execo interview if they get to it.

Progress through the pipeline: Probability of getting an offer from the hiring firm

We next analyze what happens from step 2 to step 3 in Figure 4.1. Table 4.4 shows a set of regressions predicting the probability that the candidate will not get an offer from the hiring firm, either because he/she refuses to move ahead with the process after the interview (Model 1), or because the hiring firm rejects him/her at this step (Models 2-4). For this analysis we only look at the set of candidates who made it to the hiring firm interview (3,978 observations).

Table 4.4: Linear probability models predicting who does not get an offer from the hiring firm

Variable	SUPPLY (candidate not interested)	DEMAND (candidate classified as not suitable by the hiring firm)		
	Model 1	Model 2	Model 3 (individual fixed effects)	Model 4 (individual fixed effects)
Candidate previously placed by the search firm (1=yes)	-0.0771 [0.085]	-0.0077 [0.037]	1.0021** [0.098]	0.9995** [0.099]
Candidate answered an advertisement (1=yes)	-0.0918+ [0.052]	0.0061 [0.023]	-0.0159 [0.076]	-0.0070 [0.080]
Number of the search firm's prior assignments for the hiring firm	0.0078 [0.008]	0.0059+ [0.003]	0.0108 [0.009]	0.0123 [0.009]
Number of contacts in other roles	-0.0013 [0.002]	-0.0022* [0.001]	-0.0169* [0.007]	-0.0166* [0.006]
Number of jobs considered	0.0096 [0.012]	0.0142** [0.005]		
Number of jobs: candidate not interested			0.2103** [0.080]	0.2157+ [0.110]
Number of jobs: candidate not suitable			-0.0466+ [0.025]	-0.0453+ [0.025]
Number of jobs: the outcome unknown			-0.1536 [0.112]	-0.1562 [0.111]
Female (1=woman)	-0.0366 [0.043]	-0.0347+ [0.021]	/	/
Female * (ln) number of jobs not interested				0.0228 [0.165]
Female * (ln) number of jobs not suitable				-0.0543 [0.119]
Constant	0.3857 [0.239]	0.7806** [0.046]	0.2504 [0.182]	0.2372 [0.184]
N observations	1,046	3,070	3,070	3,070
N of candidates	1,034	2,903	2,903	2,903
R-squared (within)	0.072	0.016	0.501	0.501

+ p<0.10, * p<0.05, ** p<0.01

All tests are two-tailed.

Robust standard errors in parentheses are clustered by individual.

All models include the following controls: assignment industry, job function, year dummies.

R-squared reported for models 1 and 2 are the R-squared for the OLS regression. R-squared reported for models 3 and 4 are within R-squared.

Model 1 in table 4 estimates the probability that the candidate will decline being offered a job by the hiring firm. As the results indicate, in the cross section there is no difference by gender, while the number of prior jobs coefficient is also not significant. We unfortunately do not have enough observations to conduct the estimation of the candidates' decisions at this stage with individual fixed effects. We therefore are not able to run the equivalent of models 2 and 3 in table 4.3 for this step, and do not report the results for the cross-sectional models for that stage as well as they would not be comparable with the results for the demand side (for which we have enough observations)²⁷. We therefore now move to the analysis of the hiring firm's decision of not extending an offer to the candidate after having interviewed him/her.

Model 2 indicates that in the cross-section women are less likely to be rejected by the hiring firm (the effect is marginally significant, though)²⁸. Model 3 shows the number of jobs that the candidate has previously rejected increases the probability that the hiring firm will not give an offer to him/her. Conversely, the number of jobs for which the candidate has been classified as "non suitable" in the past decreases the probability that he/she will be rejected by the hiring firm in the current vacancy. These effects are the same for women and men. The hiring firms appear to follow the same pattern as the search firm when it comes to dropping the candidates who have previously rejected jobs/have been rejected. However, the coefficients are only marginally significant. One reason may be that the hiring firms may not

²⁷ Recall that we have constructed the comparison groups in the following way: when analyzing the candidates' decision to reject receiving an offer, we estimated the probability for those candidates for whom the outcome was not unknown for the progression from step 2 to step 3, and for whom the outcome was not the hiring firm rejection. Even though the overall number of observations for that step was 3,978, for the analysis of the candidates' rejections we used 1,046 observations (3978 - 432 unknown observations - 2500 observations where the hiring firm declared the candidate unsuitable). For the analysis of the hiring firm's rejections we used 3,070 observations (3978 - 432 unknown - 476 observations where the candidate rejected receiving an offer). For the first sample, when running the individual fixed effects estimation, Stata reported "insufficient observations" error. It did estimate the second sample though.

²⁸ No one can reject receiving an offer, or is rejected by the hiring firm, if he/she is not actually interviewed by the hiring firm in the first place. We therefore ran models that correct for selection. We used a dummy variable that takes on value 1 if the vacancy was advertised as an instrument. Results for models 1 and 2 are the same if we correct for the potential sample selection bias (and the inverse Mills ratio is not significant in either case).

necessarily be aware of the history of the rejections between the candidates and the search firm. However, this prior history may shape the way in which the search firm presents the candidates to the hiring firms. Specifically, since the search firm is more likely to drop candidates who have rejected its interviews in the past, it is perhaps less likely to thoroughly “sell” such candidates to the hiring firm if it doesn’t drop them (i.e. if it sends them for the hiring firm interview instead). Hiring firms may then act on this representation and drop candidates, without knowing where the “negative” representation by the search firm comes from. Since the underlying mechanism in this case would be indirect, it could explain why the effects for the hiring firm stage are the same as for the search firm stage, but weaker. Regardless of the underlying mechanism, this effect does not vary by gender (see Model 4)²⁹.

We do not report multivariate analysis of the very last step of the process (the candidate’s probability of accepting an offer once it is extended). We have very few observations that make it to this step, and hence cannot run the individual fixed effects analysis. In addition, in the cross-section there are no significant gender patterns – i.e. women are not more or less likely than men to reject offers.

4.5 Discussion and conclusions

In this chapter we were interested in the differential role of prior rejections in candidacies for the executive jobs for men versus women. We found that the claim that women may be disadvantaged in access to high-end jobs because of their fewer connections with executive search firms (see Brett and Stroh, 1997) only partly applies in this setting. Women candidates indeed had fewer past interactions with the search firm as candidates, but

²⁹ We considered running the individual fixed effects models with the specification that would correct for the sample selection in the similar way as we did for the cross-sectional models. However, this would result in inconsistent estimates, as described by Wooldridge in the standard econometrics textbook (Wooldridge, 2002, p. 582). We therefore estimated only standard individual fixed effects models.

more in the other roles. Nonetheless, their candidate connections matter differently for them than for men.

We find that the executive search firm's screening behaviour in the first step of the hiring process may influence whether men and women compete for different types of job. In particular, we find some evidence that women tend to be considered for lower paid positions if they have been rejected for jobs in the past. Since women are not more likely to apply for lower paid jobs in this setting, it seems that the observed pattern is at least partly due to demand-side considerations (i.e. the search firm tends to allocate women to lower paid vacancies). The advantage of the dataset is that it allows us to rule out the possibility that this effect is caused by the differences in qualifications, as we estimate our models using individual fixed effects specifications (which effectively control for unobserved individual heterogeneity).

We do find evidence that prior rejections matter differently for women later in the hiring process. The most intriguing pattern of results on the issue is as follows. First, even though women in this sample have been considered not suitable for fewer jobs than men, the effect of these prior rejections is stronger on their subsequent decisions than it is for men. In particular, women tend to be more likely to be not interested in pursuing a search if they have been rejected in the past. Second, all candidates that have been not interested in the past are more likely to be dropped by the search firm and by the hiring firms in the subsequent searches. The search firm does not seem to take lightly the rejections, and may also less thoroughly "sell" the candidates that have rejected the vacancies in the past to the hiring firms. This is consistent with the interviews with candidates. A CEO of a global software firm who was a candidate in several searches (interviewee 17) put it succinctly: "you don't want to leave any scorched earth between you and one of these captains of the network". These results suggest that prior rejections of women may create a "circle of rejections" in

terms of their access to the highest paying jobs: once they are rejected, they are more likely to subsequently reject future vacancies. After they reject these vacancies, they are more likely to be rejected by the search firm. This leads to an accumulation of rejections of women and allocation to the lower paid vacancies (the latter link is only suggestive in our data though, as the coefficient is marginally significant). These results reveal an intriguing pattern by which women's interactions with search firms may limit their access to the high paying jobs.

This study points to the need to better understand the “currencies” that flow in relationships between market actors. Prior studies of how relationships influence individuals' chances of getting jobs have focused on information-based mechanisms and/or expectations of social enrichment once the workers are placed (see Fernandez, Castilla and Moore, 2000). Although not often emphasized in the literature, we argue that past interactions may actually lead the parties to decide not to engage with each other in the future. In this setting, interactions are opportunities for each party to reject each other. Moreover, these rejections seem to matter for hiring outcomes down the road. We effectively control for the possibility that better candidates are more likely to turn down job opportunities using the individual fixed effects specifications. The presence of significant gender interaction effect, and shifts in which party rejects another through time suggest that a more complex process may be at play. Most studies that look at prior relationships are not able to distinguish between successful and unsuccessful exchanges, let alone identify which of the actors decided not to engage further. Yet our results suggest that these distinctions may offer additional insight into the mechanisms through which prior relationships affect hiring outcomes – and market processes more generally.

The question of course still remains why women's interactions with the search firm seem to be more affected by their prior rejections by the search firm and its clients. There are several potential underlying mechanisms behind this finding. First, women may have

different levels of “tolerance” for rejections, and search firms may draw different conclusions from the rejections of women. Second, it is possible that the content of interactions between the search firm and female candidates may not be exactly the same as it is for male candidates. In particular, the search process itself may work out differently depending on the gender of the candidate and the search consultant(s). Finally, women’s attitude to prior rejections could be a rational response to the fact that they may be allocated to lower paying vacancies once they have been rejected. While we found that women have actually been rejected for fewer jobs in the past (adding up rejections at any step of the process), the results provide some evidence that women are more likely to be allocated to the lower paying vacancies, and that this is not due to them applying to such vacancies. While our data do not allow us to directly probe these explanations, the results suggest that women’s response to rejections in this setting may be related to their prior experiences with the search process. If so, the usual argument that women “take themselves out” of competition for top level positions would need to be qualified with a much more complex understanding of what happens in external searches for these jobs.

With this study, we expect to make several contributions to the literature. First, we take a closer look at a key sociological phenomenon, competition for jobs in pre-hire processes, by offering insights into the social microprocesses that determine gendered access to opportunity in a relatively understudied institutional setting – brokered recruitment. Whereas some studies have analyzed the process by which headhunters fill positions, and even have hinted at the possibility that they may help homosocial reproduction (Khurana, 2002; Finlay and Coverdill, 2002), there is little we know about their contribution to women’s under-representation among high-end managers. Considering that external searches are more likely than they used to be (Cappelli and Hamori, 2005), the role of executive search firms in gender stratification may be more important than the literature has so far

acknowledged. Studies like ours, which look in detail at brokered searches, thus help open the black box of these intermediaries that are crucial in filling some of the best paid jobs in the labor market.

This study also identifies the specific points in the search process in which gender stratification is set in motion. We find that prior relationships between the search firm and candidates have different consequences for women than for men in the very first steps of the process, while gender does not matter much for their effect at the very end. A better understanding of where and how stratification may happen in a particular setting directs us to potential mechanisms of stratification that may take place in similar settings.

Finally our analysis of how relationships matter in the executive search process points to the need to examine the role of prior exchanges in more detailed ways than the literature has been able to do in the past. In particular, we highlight the need to consider the difference between past exchanges and exchanges that were attempted but never took place, which is something rarely contemplated by students of long-term market relationships. It is also important to further understand how mutual rejections in past relationships affect the probability of subsequent interactions and, in particular, whether and why this may vary depending on who is “doing the rejecting.” A better understanding of these mechanisms will allow us further insight into the role of enduring relationships between actors in structuring market competition.

CHAPTER 5: VALUE CAPTURE IN EXECUTIVE SEARCH

5.1 Introduction

The previous chapter analyzed how prior rejections of progress through the hiring pipeline affect progress in subsequent searches. Rejections by both the candidates and the search firm/the hiring firm have been considered. Specifically, the chapter focused on whether these rejections have different consequences for men vs. women. The chapter before the previous one analyzed the role of trust in executive search while considering the economic incentives of the candidates. In both of these chapters the focus was on how prior relationships with a broker affect the outcomes of a competitive process for individuals. Specifically, the outcomes of interest were the individuals' willingness to compete in the brokered recruitment process, and the progress of women and men in the same process. By examining who pursues executive vacancies, and who is pursued by the search firm chapters 3 and 4 have shed light on the central question of the thesis, i.e. how the content of the relationships between labor market participants and their labor market contacts shapes which labor market opportunities the former pursue.

In this chapter I turn to the economic implications of the differential impact of ties with different content. Prior research has explored economic consequences of relationships in the labor market (e.g. Fernandez, Castilla and Moore, 2000; Fernandez-Mateo, 2007; Bidwell and Fernandez-Mateo, 2010), but there the explanatory variables were either the presence or absence of the relationship, or relationship duration. In contrast I use explanatory relationship variables that vary in their content. In terms of the outcome of interest I focus on one dimension of economic implications, the value that the search firm captures from searches. I relate the search firm value capture to all ties introduced in the previous chapters, and assess

how value captured in each search assignment is affected by the content of its ties with candidates in the consideration set. The aim of the analysis is to uncover how value capture is affected by both ties that are characterized by different levels of trust, and by ties built through rejections. This would provide some insight into the relationship between tie content and value capture in the labor market, a relationship that is not often studied, but important for the understanding of the distribution of economic benefits in labor market.

5.2 Trusting ties, ties based on rejections, and value capture

Findings in chapter 3 indicated that among non-candidate ties (client, source and general) only the latter are characterized by the executive's development of trust in the search firm. This would suggest that the search firm value capture from general ties should be greater than value capture from the other non-candidate ties as the search firm would need to expend less effort on persuading the candidates to participate in the brokered recruitment process when it has some pre-established "goodwill" with them. However, the findings also indicate that the relationship might be more complex. Firstly, general ties indicate that at some point in the past the search firm "did not quite know what to do" with an executive, and hence established a general tie with him/her. Results in chapter 3 indicate that the past "profile ambiguity" may transfer into future searches, where candidates with general ties may not be as well matched to the vacancies. Since the quality of the match may be more apparent after the search firm interview than before, including candidates with general ties in the consideration sets for vacancies may prolong search assignments as the search firm may spend time interviewing candidates with general ties who were the ones that didn't reject the interview invitations, only to discover after the interviews that the quality of the match is not optimal, which leaves it little choice but to drop these candidates (or they may pull out themselves, as per results in chapter 3). The effect of general ties on assignment length, and

thus search firm value capture can therefore not be predicted as it is not clear when the positive effect of trust is outweighed by the opposite effect of a modest match quality.

The relationship between prior rejections and value capture is similarly ambiguous. On the one hand results in chapter 4 indicate that the rejections have a tendency to reverse. Candidates who were not interested in the vacancies in the past at some point become less likely not to be interested. Also, candidates who have been not suitable in the past at some point become less likely to be considered not suitable by the search firm (both outcomes refer to step 1, i.e. before the search firm interview). This would suggest that prior rejections should eventually lead to shorter assignments, because the search firm does not have to work with a number of candidates whom it eventually rejects (or are rejected by them), but with candidates at lower risk of rejection. On the other hand, rejections also work in the opposite way. Specifically, candidates that have been unsuitable in the past are more likely to be not interested in the future, while at the same time such candidates are less likely dropped by the search firm. This effectively means that those candidates that are less likely to be dropped by the search firm are more likely to pull out themselves. To the extent that the search firm wants to keep such candidates in the process (as indicated by their lower unsuitability), it needs to spend time and effort on persuading them to stay in the process. This would increase the length of search assignments. The dynamic for candidates who have previously been unsuitable is similarly asymmetric: such candidates are less likely to be not interested, but the search firm is more likely to drop them. The extent to which such candidates prolong or shorten searches depends on how fast the search firm decides to drop them. Since I don't have information on what tends to come sooner in searches (search firm rejections or candidate rejections), I cannot make predictions on the effect of prior rejections by candidates on the length of the assignments.

The preceding discussion suggests that the effect of different kinds of prior ties in value capture in executive search is an open empirical question. I therefore move to a multivariate analysis to statistically examine the relationship between general ties, client ties, source ties, candidate ties built through candidate rejections, and candidate ties built through search firm/hiring firm rejections, with the length of search assignments. Note that since I don't have information on the search firm's monetary costs per unit of time (which I would then multiply with the length of the search assignment), I unfortunately cannot express the amount of value capture in monetary terms. I can, though, assess whether more or less value is captured, and relate this dependent variable to the independent variables – the different kinds of ties between the search firm and the candidates for the executive jobs.

5.3 The sample and variables

I do the analysis in this chapter in two parts. For the first part, I take advantage of the fact that Execo conducted more than one search assignment for several hiring firms during the observation period. I therefore construct a panel where the unit of analysis is each “hiring firm by vacancy” (see the description of the structure of this panel in chapter 2). This allows me to estimate the models using the hiring firm fixed effects, and thereby effectively control for unobserved hiring firm heterogeneity (which might affect the length of the assignments). The goal of the first part is to establish the relationship between the different kinds of prior ties between the search firm and the candidates, and the length of the search assignment. The dataset consists of 998 vacancies. As I am studying the assignment length and am hence interested in the assignments that have concluded, I have excluded from the analysis 65 vacancies that have been marked as “ongoing” in Execo's database. I have also excluded 9 vacancies for which the assignment length was negative (which means that either the

assignment start date or the assignment end date, or both, have been incorrectly entered into Execo's database). The sample for the first part of the analysis therefore consists of 924 observations (924 vacancies/search assignments).

The dependent variable in the first part of the analysis is assignment length. I operationalize assignment length as time elapsed between the start date of the assignment (as indicated in Execo's database), and the date a candidate was placed, or if no one was placed, Execo's last contact with a candidate on the assignment. This variable represents an approximation of the costs that the search firm has on a given assignment. Execo's consultants explained that for them it is important the assignment length is as short as possible (while the quality is not compromised), as that reduces their costs. A quote by senior associate at Execo (interviewee 16) illustrates this point:

“Our motivation is to deliver the short-list as quickly as possible, we have to do that because we want us to be profitable, so we want to always be as efficient as we can on resources from our own firm, to get that short-list to the client as quickly as possible, and then get someone placed as quickly as possible using as little partner time on it as we can without compromising the service level. Because if we compromise service level the whole search goes out of kilter and we have to do it again. And if we do have to do it again, we call it re-run, which is what we do for free, and then you're really not motivated, and the whole thing becomes disjointed.”

The following variables are used in the first part of the analysis:

Independent variables

- *Average number of jobs: not interested*: the aim of the first part of the analysis in this chapter is to analyze the relationship between different kinds of prior Execo-candidate relationships, and the length of the assignment. To get at the vacancy-level variable that would capture the extent of the prior rejections by the candidates I first summed the number of rejections across the candidates in the consideration set, and then divided that number by the number of candidates. I followed the same procedure for the other kinds of prior

relationship (*average number of jobs: not suitable, average number of jobs: the outcome unknown, average number of general contacts, average number of client contacts and average number of source contacts*).

Control variables

In this chapter I used the following controls not used in the previous chapters:

- *Vacancy advertised*: this is a dummy variable that takes on value 1 if the assignment was advertised, and 0 if it was not.
- *Margin*: this variable is calculated as a ratio between the fee that the hiring firm agreed to pay to the search firm to conduct the search assignment, and the maximum salary for the vacancy.
- *Number of candidates in the consideration set*: this variable is a count of the number of candidates that were included in the consideration set for the vacancy (the total number, i.e. including both the candidates that were included by the search firm, and those that have responded to an advertisement)
- *The search firm made a prior placement for the hiring firm*: this variable takes on value 1 if Execo has between 2005 and the start date of the focal vacancy placed someone in the hiring firm
- *Execo a preferred supplier to the hiring firm*: sometimes professional service firms, including the labor market intermediaries, sign a preferred supplier agreement with their clients (a detailed description of such arrangements is provided in Fernandez-Mateo, 2007). This variable is a dummy that takes on value 1 if at some point in the past Execo and the hiring firm had a preferred supplier agreement in place.

I also used as controls the following previously described variables: *Prior search firm relationship with the hiring firm*, *Maximum salary for the position*, *Assignment industry controls*, *Job function controls* and *Year controls* .

In the second part of the analysis I explore whether the costs in terms of time and effort associated with “putting through” candidates affect the search firm’s allocation of candidates to the vacancies. Specifically, I explore whether the search firm allocates “costlier” candidates to the vacancies where its margin (fee/vacancy salary) is higher. I consider as "costlier" those candidates that the search firm wants to put through (i.e. are less likely dropped by the search firm), and that are at the same time candidates that are more likely to reject going forward with the search process. This amounts to exploring whether the search firm simply attempts to recoup the costs of a greater effort, or it attempts to allocate the candidates in the manner of a traditional broker, i.e. seeking rents in the process. That is, I analyze whether the search firm allocates to the higher-margin vacancies only those candidates that are costlier to “put through” the hiring pipeline, or also some others. The dependent variable for this part of the analysis is margin, which is described in the variables used in the first part of the analysis. The independent and the control variables used in this part of the analysis are the same as in the second part.

Table 5.1 below summarizes the descriptive statistics. For the individual level variables I present descriptives for the portion of the sample that was allocated to vacancies by the search firm (i.e. the candidates that have entered the consideration set by responding to an advertisement are excluded). I use this reduced sample as in this chapter I am interested in how the search firm adjusts the allocation in response to the different value capture implications of ties with different content.

Table 5.1: Summary of descriptive statistics⁽¹⁾

Individual level variables (N= 51,764); only individuals attached to vacancies by the search firm				Vacancy salary	160,217	154,627	0.322
Variable	Mean		Std. Dev.	Vacancy advertised (1=yes)	.276	.436	0.000
Candidate previously placed by the search firm (1=yes)	.022		.148	Margin (fee/vacancy salary)	.322	.326	0.516
Number of jobs: candidate not interested	.302		.800	Number of candidates in the consideration set	60.086	68.572	0.016
Number of jobs: candidate not suitable	.347		.812	The search firm made a prior placement for the hiring firm (1=yes)	.546	.518	0.394
Number of jobs: the outcome unknown	.073		.297	The search firm a preferred supplier to the hiring firm (1=yes)	.114	.101	0.540
Number of general contacts	.610		1.405	Average number of jobs: not interested	.259	.298	0.060
Number of client contacts	1.225		4.642	Average number of jobs: not suitable	.343	.389	0.025
Number of source contacts	.997		2.309	Average number of jobs: the outcome unknown	.062	.066	0.470
Vacancy level variables (N=924); only concluded assignments (i.e. ongoing assignments omitted)				Average number of general contacts	.553	.623	0.108
Variable	No one placed (N=463)	Someone placed (N=461)	p-value	Average number of client contacts	.906	1.272	0.015
Assignment length (months)	8.336	7.896	0.254	Average number of source contacts	.850	.951	0.112
Number of the search firm's prior assignments for the hiring firm	1.645	1.180	0.006				
Job industry (%)				Job function (%)			
Media	3.46	Infrastructure	6.28	Financial services professional	10.39	Marketing director	7.14
IT	7.47	Government	9.42	Consultant	1.08	Sales director	3.35
Pharmaceuticals	7.90	Health	0.76	Board member	4.33	Divisional sales director	2.38
Agriculture	0.11	Education	0.32	CFO	7.03	Legal and governance	5.84
Finance	33.87	NGO	1.30	Divisional finance director	3.25	Non-executive director	1.41
Professions	8.23	Other	3.79	CEO	7.14	Pharmaceuticals scientist	1.19
Leisure	6.28	Missing	0.00	Divisional managing director	13.74	Management – other	21.00
Engineering and manufacturing	12.23			Operations director	2.71	Other	14.61
Retail	15.26			Government	4.00	Missing	0.32
Energy	2.16			HR director	4.65		

(1) Statistics for job industry and job function calculated at the level of the vacancy. A given vacancy may correspond to several industries and job functions.

5.4 The analysis and results

I first analyze whether assignments that conclude successfully (i.e. with someone placed) differ in some way from the unsuccessful ones (i.e. no-one placed). The difference in the length of the unsuccessful and successful assignments is not statistically significant. The average length of the first is 8.33 months, while the second is 7.89 months. The assignments that were advertised and those with more candidates in the consideration set were more likely to result in a placement, which suggests that the mere size of the pool to choose from may increase the probability of a successful assignment. The assignments that ended in a placement also included candidates who have been rejected more in the past, and who had more contacts in the client role with Execo. These results underline that not all kinds of priorities are likely to lead to a successful completion of an assignment (i.e. a placement). The means for the individual level variables are very similar to the means described in chapters 2 and 4, so I don't describe them in detail here. The distribution of vacancies across industries and vacancies is also very similar to the distribution of the whole sample, as described in chapter 2, so I omit the description of this distribution as well.

In the first part of the analysis I examined how the different kinds of prior relationships affect the length of the assignment (and thereby the value that is captured by the search firm). I control for the value that is created in the assignment by using the search firm's margin as a control. The search firm has conducted multiple assignments for several hiring firms during the observation period. I therefore constructed a panel where the unit of the analysis is each "hiring firm by vacancy" pair. I estimated the OLS models in Table 5.2 using the hiring firm fixed effects, effectively controlling for unobserved hiring firm heterogeneity.

Table 5.2: OLS regressions predicting the length of the search assignment

Variable	Model 1 (hiring firm fixed effects)	Model 2 (hiring firm fixed effects)	Model 3 (hiring firm fixed effects)
Margin (fee/vacancy salary)	3.2731 [3.019]	2.6852 [3.089]	3.0015 [3.106]
Number of candidates in the consideration set	0.0271** [0.005]	0.0265** [0.005]	0.0271** [0.005]
Vacancy salary (ln)	1.4148+ [0.823]	1.1193 [0.917]	1.3667 [0.904]
Vacancy advertised (1=yes)	-0.5407 [0.750]	-0.6247 [0.796]	-0.3274 [0.846]
Number of the search firm's prior assignments for the hiring firm	-0.2100+ [0.118]	-0.2329+ [0.124]	-0.2342+ [0.126]
The search firm made a prior placement for the hiring firm (1=yes)	2.2142* [1.020]	1.9575+ [1.023]	2.2069* [1.008]
The search firm a preferred supplier to the hiring firm (1=yes)	0.1722 [0.411]	0.2503 [0.455]	0.0683 [0.483]
Someone placed (1=yes)	-0.0204 [0.648]	-0.0905 [0.647]	0.8336 [0.835]
Average number of jobs: not interested		-2.9188 [1.826]	-3.5443+ [2.023]
Average number of jobs: not suitable		-0.8869 [1.194]	0.4911 [1.717]
Average number of jobs: the outcome unknown		4.4809 [3.440]	5.8230+ [3.472]
Average number of client contacts		0.9199 [1.178]	0.9459 [1.153]
Average number of general contacts		-0.0777 [0.217]	-0.0623 [0.215]
Average number of source contacts		0.2672 [0.488]	0.2156 [0.482]
Someone placed * number of jobs: not interested			2.0030 [1.732]
Someone placed * number of jobs: not suitable			-4.0843* [1.844]
Constant	-13.6087 [10.519]	-10.9938 [11.119]	-14.9543 [10.899]
N observations (vacancies)	924	924	924
N of hiring firms	489	489	489
R-squared	0.178	0.186	0.194

+ p<0.10, * p<0.05, ** p<0.01

All tests are two-tailed.

Robust standard errors are clustered by the hiring firm.

All models include the following controls: assignment industry, job function, year dummies.

Number of jobs and contacts is measured as an average per candidate in the consideration set.

Reported R-squared is within R-squared.

Model 1 includes the controls. As expected the assignments with a larger number of candidates in the consideration set take longer to conclude. A somewhat more surprising finding is that the assignments take longer when the search firm has already placed someone in the hiring firm. I do not have information on why this may be so. One potential explanation, given that several interviewees have stressed the aim for repeat business in executive search might be that once the search firm places someone in the hiring firm, it gets the opportunity to work on more complex and potentially more profitable searches, which may also take longer. Model 2 includes the relationship variables. It turns out that none of the relationship variables significantly impacts the length of the assignment. In model 3 I consider whether the impact of the prior relationships may differentially affect the length of the successful versus unsuccessful assignments (i.e. the ones that end in a placement and those that don't). I therefore include the interaction of the "someone placed" variable with the prior relationships. In the interest of brevity I am presenting the results with the interactions of "someone placed" with the average number of jobs not interested, and not suitable. The results are the same if I include the full set of the interaction terms (and none of the other interaction terms are significant).

Model 3 contains an interesting result: the assignments with candidates who have been rejected before, and that result in a placement, are shorter than the other assignments where someone was placed³⁰. Apparently inclusion of the candidates with a history of prior

³⁰ As the dependent variable in table 5.2 is the length of time to an event (conclusion of an assignment), I conducted a robustness check using an event history analysis technique. Specifically, I used a Cox regression, which estimates a proportional hazard rate of an event taking place. The complication for the estimation was that the Cox regression that would use dummy variables for each panel variable (in my case hiring firm), would produce regression coefficients that are biased by 30-90% away from zero (Allison, 2009, p. 74). I followed the approach suggested by Allison (2009, p. 74-75) and estimated the Cox regression model by implementing the Stata `stcox` command with the option `strata (case id)` (in my case `strata (company number)`). Such specification estimates the hazard function for each subgroup (in my case for each hiring firm). This results in different subgroups (i.e. hiring firms) having different baseline hazard functions, while it constrains the coefficients to be the same across the subgroups (Allison, 2009, p. 74). I also employed the estimation with robust errors clustered by the hiring firm. This procedure resulted in results that are consistent with the results reported in table 5.2, yet weaker. Specifically, in the replication of model 3, the interaction term on someone placed * number of jobs: not suitable was insignificant. I then split the sample and ran the Cox estimation separately using the sample with assignments that ended in a failure (i.e. no-one was placed, N=463), and the sample with assignments that ended

search firm/hiring firm rejections leads to a “quick success” – a placement which takes less time to arrange. According to this result the prior relationships with the candidates, built through the search firm/the hiring firm rejections, enable the search firm to capture more value in the search assignments.

Putting together this result and the results in chapter 4 described in the previous section uncovers a tension associated with the prior relationships built on the unsuitability of the candidates: the search firm appears to go first for the candidates that have been unsuitable before, and when such candidates are well represented in the consideration set, the search assignment is more likely to conclude quickly and with a placement. However, such candidates are harder to talk into participation as they are more likely to reject going forward in the search process. So whilst on one hand the search firm may capture more value when there are more previously unsuitable candidates in the consideration set, this might be balanced out by the costs of persuading unwilling candidates. I unfortunately do not have data that would allow me to identify and measure these costs, so I can only speculate on their exact impact. Nevertheless, the findings point to a more general point: the same prior relationships may both enable and constrain value capture.

The search firm may therefore find the recruitment of the candidates that it wants to “put through” costly. In the second part of the analysis I analyze whether the search firm learns which candidates are costlier to recruit, and adapts its allocation strategy accordingly. In the analyses in chapter 4 I found that the previously unsuitable candidates may be costlier (from the perspective of the search firm) to recruit, as they are more unwilling to participate in the process, while the search firm is less likely to drop them, and by extension, more

in a success (i.e. someone was placed, N=461). None of the two sub-samples converged when I ran them with the strata option. I then ran them without the strata option. In the second subsample the main effect of the average number of jobs: not suitable was positive and significant. This means that the assignments where candidates have been previously not suitable, and that end with a placement, have a lower “survival rate”, i.e. take less time to conclusion. In other words, they lead to a “quicker success” (earlier placement). Still, the result with the event history approach is weaker than with the linear probability model using individual fixed effects.

willing to keep them in the pipeline. One way to examine these questions is to analyze whether prior relationships affect which candidate gets allocated to the high-, as opposed to the low-margin vacancies. A margin is the search firm's fee for the whole assignment, which is set at the beginning and remains fixed. The costs of the assignment therefore shape the value that the search firm captures from a particular assignment. I explored whether the search firm matches the "costliness" of the candidates it allocates to a vacancy with the margin of that vacancy. To address that question I ran OLS regressions predicting the margin of the vacancy a candidate is allocated to by the search firm. I included in the sample for that analysis only candidates that were placed in the consideration sets by the search firm (N=51,764), to uncover the search firm's allocation strategy. I estimated the models using the individual fixed effects specification to control for unobserved individual heterogeneity. Table 5.3 reports the results.

Table 5.3: OLS regressions predicting margin of the vacancy a candidate is allocated to by the search firm

Variable	Model 1 (individual fixed effects)	Model 2 (individual fixed effects)
Candidate previously placed by the search firm (1=yes)	-0.0051 [0.026]	-0.0011 [0.026]
Number of the search firm's prior assignments for the hiring firm	-0.0003 [0.000]	-0.0003 [0.000]
Vacancy salary (ln)	-0.1456** [0.003]	-0.1462** [0.003]
Vacancy advertised (1=yes)	0.0036* [0.002]	0.0043* [0.002]
Number of jobs: candidate not interested		0.0057** [0.002]
Number of jobs: candidate not suitable		0.0035** [0.001]
Number of jobs: the outcome unknown		0.0066** [0.002]
Number of general contacts		0.0013 [0.002]
Number of client contacts		-0.0005 [0.000]
Number of source contacts		-0.0004 [0.001]
Constant	2.0104** [0.040]	2.0198** [0.040]
N observations	51,764	51,764
N of candidates	38,927	38,927
R-squared	0.273	0.274

+ p<0.10, * p<0.05, ** p<0.01

All tests are two-tailed.

Robust standard errors in parentheses are clustered by individual.

All models include the following controls: assignment industry, job function, year dummies.

Reported R-squared is within R-squared.

Model 1 includes the controls. The interpretation of the results is not straight-forward as the two significant coefficients are for variables that are not at the individual level. One way to interpret them would be that the advertised vacancies have higher margins, while the vacancies with higher salaries have lower margins associated with them. This is consistent with the simple pair-wise correlations. The results of interest are in model 2. Candidates that have previously been unsuitable are allocated to the higher-margin vacancies. This could be interpreted as the search firm's realization that such candidates are harder to recruit, yet are at the same time more desirable to the search firm (as per the results in table 4.3 in chapter 4). If these candidates are costlier to recruit, the rational allocation strategy would be to allocate them to the higher-margin vacancies. This strategy is rational in the sense that it matches the high-cost candidates with the high-value vacancies, and in that way keeps the value captured by the search firm at a certain level.

The other significant coefficient appears somewhat more puzzling. The search firm also allocates to the high-margin vacancies candidates that have been not interested before. As per results in table 4.3 in chapter 4 such candidates are no costlier to recruit as they are keen to go for the search firm interview. Allocating the "low cost" candidates to the high-margin vacancies seems like a pure rent-seeking strategy by the search firm. Since the only paying party in the context of executive search is the hiring firm, this alludes to the possibility that the search firms might be able to capture rents from their clients (the hiring firms) even after the payment for the search assignment has been determined. This may be a particular strategy adopted by the one search firm that I am studying, so I can make no claim about the generalizability of that finding. However, it is interesting to consider how Execo's search consultants describe what the clients are paying for. The quote below is from a senior associate at Execo (interviewee 16):

“There will always be searches where you have to start from a blank piece of paper, and they take much longer and are much harder and less profitable vs. searches where we do this all the time, we know the candidate pool, we know them really well, we’ll be able to get people in very quickly. So the client will pay the same for different things in those instances, because in the first case they’re paying us to go and do a whole piece of research; in the second instance they’re paying the same for our knowledge, our intellectual property.”

The compensation structure in executive search therefore incentivizes the search firm to “buffer” for the future instances where it might incur costs that would not be fully covered by the margin for the particular assignment. Allocating candidates that have previously been not interested, and are currently low cost candidates because of their greater willingness to progress through the pipeline, may be one way in which this “buffering” takes place. Together, the results and the quote suggest that the search firm may be actually capturing some extra value from some clients, to compensate for the value they have to forego with the clients for whom they need to expend more effort to complete the search.

5.5 Discussion and conclusions

In this chapter I investigated how ties with different content affect the search firm’s ability to capture value. I first discussed the likely impact of the three non-candidate ties (general, client and source), and two kinds of candidate ties (candidate ties built on candidate rejections and those built on the search firm/hiring firm rejections of the candidates). I concluded, on the basis of results in chapters 3 and 4 that for no tie an effect can be predicted, but that it is rather an empirical question. I therefore proceeded with the multivariate analysis. I found that of all the kinds of tie, only candidate ties built on the search firm/hiring firm rejections have a significant relationship with assignment length. The relationship is negative for the assignments that ended up in placement, meaning that the more the search firm or its clients have rejected individuals in the consideration set, the faster the focal assignment

concludes with a placement. This indicates that the effect of “reversal” in the search firm rejections outweighs the effect of “grudge-holding” by the candidates. This means that the effect that the search firm at some point stops rejecting candidates that it has rejected in the past (provided it keeps on including these candidates in the consideration sets) outweighs the effect of the previously rejected candidates rejecting the search firm once they again become candidates. I also found that the search firm appears to take value capture considerations into account when it constructs the consideration sets. In that way the internal organizational process within the search firm shapes the labor market outcome, as it affects which candidates will be introduced to which hiring firms. This represents another way in which the differences in content of ties shape the functioning of the labor market.

CHAPTER 6: CONCLUSIONS

6.1 Summary of findings

In this thesis I study how the content of the relationships between labor market participants and their labor market contacts shapes which labor market opportunities they pursue. More specifically, I focus on the role of non-resource content that flows through relationships (Podolny and Baron, 1997). I situated my study in the executive labor market and studied the role that different kinds of prior ties between candidates and the search firm play in the executives' decisions about which vacancies they will pursue, and the search firm's and hiring firms' decisions about which candidate they will pursue.

I found that the differences in the content of prior ties with the search firm do indeed affect which job opportunities introduced by the search firm the executives pursue. Specifically, I found that they are more likely to pursue job opportunities when their prior interactions lead to development of executives' trust in the search firm. Trust development, in turn, is affected by the role that search consultants play when the executives interact with them. Roles that are closer in character to pure brokerage roles have been found to inhibit trust development in this context, while roles closer to boundary spanning ones have been found to stimulate it. I've explained this relationship on the basis of the difference in role autonomy between these two roles: pure brokerage roles in this setting are characterized by less role autonomy than boundary spanning roles. Consistent with prior literature that demonstrated a positive relationship between role autonomy and trust development, the former appear to be associated with less trust development than the latter.

I also found that the effect of ties that are developed through prior rejections of exchanges depends on the identity of the rejecting party. Specifically, the probability of

candidates' rejections in a given vacancy differs depending on whether it was him/her that previously rejected the search firm or the hiring firms, or it was the latter that rejected him/her. This differential effect of prior rejections also holds for the search firm's rejections of candidates in the current vacancy. Patterns identified at the search firm stage also appear at the hiring firm stage, but they are weaker (i.e. the coefficients are in the same direction, but marginally significant). I unfortunately do not have information that would allow me to precisely identify the mechanism underlying the latter result. It does appear, though, that prior rejections affect how the search firm presents the candidates to the hiring firms, as the stable individual characteristics of the candidates are controlled for (i.e. the effect is not caused by individual heterogeneity in candidates). Prior rejections by different parties, which represent another aspect of differences in tie content in the labor market, therefore shape which labor market opportunities the labor market participants pursue. The study also found that the effect of tie content may be conditional on the attributes of actors. A particular focus has been on gender, and the results indicated that prior rejections have a differential effect on progress through the hiring pipeline for women and men. For women prior rejections exhibit a dynamic that could be described as a "circle of rejections". The dynamic results in a more constrained access to the executive vacancies for women than for men.

Finally I explored the economic implications of the different kinds of ties (both trusting and rejections-based). I focused on the search firm value capture in search assignments. I found, somewhat counter-intuitively, that trusting ties in this setting are not associated with more value capture, even though they reduce the probability of rejections of candidates sought after by the search firm. Ties based on prior search firm/hiring firm rejections, however, have been found to increase the search firm value capture. These results indicate that the differences in tie content affect not only when labor market participants will

act on the information about labor market opportunities (either vacancies, or availability of particular job candidates), but that they also have identifiable economic implications.

In the remainder of this chapter I summarize the overall contributions of the thesis. At the end I outline the directions for further research.

6.2 Discussion and contributions

In this thesis the phenomenon of interest is the competition for jobs and individuals in the high-end executive labor market. This market has two key properties: first, it is triadic in the sense that the executive search firms intermediate the interactions and exchanges between executives and the hiring firms; and second, the search firms repeatedly interact with the same individuals, and in that way develop an array of different kinds of relationship with them. Many of these executives at some point become candidates, and when they do, the prior relationships may exert a powerful influence on both their and the search firm's decisions in the recruitment process.

Studies of the executive labor market have often conceptualized its triadic structure as one where the search firm plays the role of the broker (Khurana, 2002; Finlay and Coverdill, 2002). The focus of these studies was on the activities and functions of the executive search firms. They uncovered a host of roles, some of them traditionally associated with brokers (e.g. intermediation, information dissemination etc.), and some of them more specific to the context (e.g. "buffering" between the candidates and the hiring firms, provision of legitimacy to the process etc.). These studies have also indicated that the search firms "shift" the executives into different roles, such as those of a client, of a candidate etc. Due to the qualitative nature of these studies, they have stopped short of exploring in detail what kinds

of relationship develop through the repeated interactions between the search firms and the executives, and how these interactions affect the subsequent interactions and exchanges.

These studies have re-started the study of brokerage as a process, the notion that was implicit in Marsden's early definition (1982, p. 202), but gained little follow-up in the empirical literature. Studies on brokerage as a process have also emerged outside of the labor market domain (e.g. innovation, Obstfeld, 2005). This was an important development as the prevailing focus in the brokerage literature has traditionally been on the structural aspects of brokerage (Burt, 1992). Regardless of whether we conceptualize brokerage as a process or as a form of social structure, it is fundamental that we develop an understanding of the very first step in brokerage – the decision of the brokered parties to participate in brokerage. Namely, brokerage process may only be set in motion if the brokered parties agree to participate in brokerage. Similarly, even if two actors have no ties to each other (are “disconnected”), while they are both connected to a third party (the broker), the benefits of brokerage described in the literature on structural holes (Burt, 1992) can only materialize if the brokered parties participate in brokerage.

In this thesis I have started to take steps towards a better understanding of this important phenomenon. I demonstrate that the prior relationships that result in a development of trust of the brokered parties in the broker can set the brokerage process in motion. From the perspective of the traditional brokerage literature (Burt, 1992) this is a rather striking finding as that literature depicts the broker as an actor that realizes the benefits of brokerage by “playing off” the brokered parties (yet there are constraints, see Rider, 2009). This raises a reasonable doubt as to why the brokered parties would trust the broker and enter the brokerage process, in particular in the market settings. This study suggests that it is important to consider the history of interactions and exchanges between the broker and the brokered parties. The points of contact at which the brokered parties get in touch with the brokerage

organizations play a particularly important role from the perspective of future participation in brokerage. While much of the literature conceptualizes brokerage organizations as monolithic actors, in this study I disaggregated the points of contact into those where individuals at the boundary perform pure brokerage roles, and those where they perform boundary-spanning roles. In this way I was able to identify one mechanism through which the brokered parties' trust in a broker may develop. By illuminating when this happens, I also indicated that the brokerage process may take place only when the brokered parties have the tendency "to join", not just when the actors in the brokerage position have the tendency to "join together" (Obstfeld, 2005).

The second key aspect of executive search, the wealth of different kinds of interactions between executives and the search firm, means that the exchanges in that context are to some degree "embedded". The empirical literatures on embeddedness of economic exchanges (e.g. Uzzi, 1997; 1999), and on trust and familiarity in the business settings (e.g. Gulati, 1995), tend to focus on the strength and/or the scope of the relationships. They report that a greater embeddedness of the economic actors tends to lead to more exchanges in the future. I argue that repeated interactions not only provide a platform for exchange of private information and favors, but also a venue for observation of behavior of the counter-parties. I specifically highlight the relevance of roles in embedded relationships, and explicate how they can affect the content that flows through relationships. I provide empirical evidence that the differences in content have an effect on the probability of future exchanges, even when the strength and scope of the relationship do not.

While my findings are developed in the context of the labor market, they may apply to the other contexts where the roles of the actors shift. Two interesting settings where roles shift and the rejections of exchanges are rife are investment banking and private equity. These settings have already been studied by the economic sociologists (Podolny, 2001; Rider,

2009). They may constitute fertile grounds for the further development of our understanding of the relationship content and its impact on market exchanges.

The studies in this thesis also make an empirical contribution to the sociological study of labor markets. In this thesis I have emulated the approach of the detailed quantitative case studies that have sought to illuminate the mechanisms that shape the important labor market outcomes such as hiring (Fernandez, Castilla and Moore, 2000; Petersen, Saporta and Seidel, 2000), wages (Fernandez-Mateo, 2007, 2009; Bidwell and Fernandez-Mateo, 2010), employee performance (Castilla, 2005), and gender sorting (Fernandez and Mors, 2008; Fernandez-Mateo and King, 2010). However, the studies in this thesis are, to the best of my knowledge, the first ones that have been conducted in the context of the very highest end of the market (recall that the average vacancy salary is within the top 1% of the income distribution in the United Kingdom). Individuals in this segment of the population occupy positions that allow them to play different roles in the labor market at different points in time. In this way they may develop ties with different content with the same counter-parties (e.g. the same search firm, as in my case). This is in contrast to the studies above, where the members of the population generally only had ties to their labor market contacts (referrers or labor market intermediaries) as job candidates. This constrained prior studies in exploring the impact of tie content. Such constraint was not present in my data, and I was hence able to illuminate the mechanisms associated with tie content.

In addition, the population I studied in the thesis consists of a group of individuals that have many more employment options, as well as broader networks, than the individuals in the previous studies. However, I was only able to measure the relationships these individuals had with the particular search firm, and observe their decisions about the vacancies that this firm introduced to them. This means that the deck was stacked against the significant impact of those relationships that I was able to measure. The fact that I found

significant impact of the relationship variables on several outcomes of interest attests to the potency of social relations in the labor markets even at the highest level.

Finally, the studies in the thesis contribute to the study of strategic management, in particular to the sub-field concerned with the strategic management of human resources. With the emergence of knowledge economy human capital has become a key element in obtaining competitive advantage (Grant, 1996). As it takes time to earn returns on investments in human capital, it represents an instance of a valuable, rare and inimitable resource (Barney, 1991). Human capital also moderates the relationship between strategy and performance, as it facilitates (or limits) the implementation of strategies (Hitt, Bierman, Shimizu and Kochhar, 2001).

Since human capital is critical to firm strategy and performance, inter-organizational mobility of key individuals importantly affects competitive advantage of firms. Mobility of key individuals shapes competitive advantage through its impact on organizational founding and survival (Sorensen, 1999, 2004; Phillips, 2002), inter-organizational transfer of capabilities and routines (Lacetera, Cockburn and Henderson, 2004; Wezel, Cattani and Pennings, 2006; Aime, Johnson, Ridge and Hill, 2010) and inter-organizational transfer of social capital (Dokko and Rosenkopf, 2010). While extant research focused on the consequences of human capital mobility, we know little of how individuals at the very high end of the income distribution compete for jobs. In particular, while literature explored individual-level determinants of the job searches of employed individuals (Bretz, Boudreau and Judge, 1994), little is known about the impact of external factors. Specifically, there is little empirical work on the role that labor market intermediaries play in initiating job searches of employed (or highly employable) individuals (but see Hamori, 2010). This is an important gap as labor market intermediaries are increasingly shaping careers at all levels of the labor market (Cappelli, 2008), including at the level of executives (Khurana, 2002). I

contribute to this literature by providing a systematic and detailed empirical account of the role search firms play in both initiating mobility of the highly paid executives and in constraining their labor market attainment.

The findings in this thesis have important practical implications for acquisition and retention of the top-end human capital. Organizations often use the services of executive search firms. It is estimated that in the United States they are involved in the recruitment of 35% of all managerial hires at a salary higher than \$200,000 (McCool, 2008, p. 22). Many search firms advertise themselves by claiming that they have access to a number of high-ranked executives. The findings in the thesis highlight that what really matters in that setting is the prior history with the executives, rather than mere access. That is, when organizations decide which search firms to work with, they would be well advised to ask them about the specifics of the prior history they have with the executives that they intend to include in the consideration set.

The findings are also instructive with respect to talent retention. Organizations are often concerned about losing their key individuals to their competitors. These key individuals are likely to be on the search firms' books at any point in time, and there is little organizations can do to prevent their contact with the search firms wanting to "poach" them. However, the findings in the thesis suggest that organizations may go about the contacts of their top employees with the search firms in a strategic manner. Specifically, I demonstrated in chapter 3 that when an executive interacts with the search firm in the role of a client, he/she seems to develop a measure of distrust in the search firm. An interesting talent-retention strategy might therefore be the following: the organizations that are working with the search firms may consider including their most promising junior executives in the teams working with the search firms. In that way these executives would get exposed to the search process from the client side, and develop the above-described distrust in the search firms.

This could in turn make them “immune” to the approaches by other search firms. This scenario is a speculation as the data in the thesis does not allow me to assess to what extent this pattern holds for search firms in general. Nevertheless, it would be a valuable element of the organizational talent-retention strategies to consider how different kinds of contacts with the search firms affect their key employees’ perceptions of these intermediaries.

6.3 Limitations

This thesis also has limitations. The first one is that the dataset is drawn from a single organization. As described in chapter 1 and in the previous section, there is an increasing number of single-organization case studies in the labor market literature in economic sociology (see cites listed above). The main strength of such studies is that they use very detailed data from organizational settings, and can thus identify and isolate mechanisms that could not be uncovered with large-scale studies where data is inherently less fine-grained. These case studies therefore contribute by providing insights into novel mechanisms that can subsequently be tested in the multi-organization studies. An additional benefit is that the findings of the single-organization studies are by design grounded in the actual day-to-day practices of organizations, which endows them with a large dose of realism. Nevertheless, the fact that data in such studies is drawn from a single organization means that they cannot make any claims about generality of the results. This holds for my study as well. I can thus make no claims about how my findings generalize to the search industry, or executive labor market, as a whole. This remains a task for future research.

Another limitation is that I have little information on the hiring firms (apart from the industries in which they operate). I also have little information on the firms the candidates are working for when they are attached to the vacancies. This is so as the search firm’s condition

for sharing the data was that the hiring firms, as the search firm's clients, remain anonymous. A helpful aspect of executive search in that respect, however, is that the search firm divulges the information about the hiring firm sequentially, and consequently the candidates know little about the hiring firm in the early stages of the process. As I focus on the early stages of the process in the thesis (i.e. the stages when the candidates get to know little about the hiring firm) this limitation is somewhat lessened on the executive side. It is still present on the search firm side though. It remains a task for the future research to assess how the effects presented in this thesis might be affected by the source firm/hiring firm characteristics.

6.4 Future research

In this section I outline some of the questions raised by the findings in the thesis that point to the fruitful areas for further research. The studies in the thesis depict a picture of executive search as a process shaped by the differential impact of different kinds of relationship between the search firm and the candidates. In other words, the impact of prior relationships depends on the specific content of the interactions through which the relationships have been built. My dataset, although rich, does not contain data that would allow me to adjudicate whether prior relationships contribute to, or detract from, the efficient matching of candidates to jobs. To be able to do that, I would need detailed information on the profiles that the hiring firms are looking for, and comparable information for all candidates in the consideration set. This would allow me to assess the extent to which the candidates that progress through the pipeline, and are eventually placed in the hiring firms, are a good match with the initial requirements. Collecting and analyzing such data remains a task for future research.

Even with such data, the question of whether prior relationships facilitate or detract from the efficiency of matching would only be partially answered. Namely, such data would only allow the assessment of whether different kinds of prior relationship enable the search firm to fulfil the brief, and “put through” to the hiring firms the kinds of candidates that the latter have set out to hire. This would constitute a test of the match at the point of hire. However, a better match at hire may or may not lead to a continued high performance (Castilla, 2005). Therefore, to be able to fully address the question of how prior relationships affect matching and its performance implications one would need to supplement data presented in this thesis with both information from the search briefs and the individual profiles, as well as with the post-hire individual performance measures. Such data was not available to me as Execo is not systematically tracking post-move outcomes of its placements, such as performance or turnover. In that Execo was not very different from other search firms I have interviewed³¹. Future research that would complement the kind of data used in this dissertation with the data on the quality of matches and post-placement outcomes could make a significant contribution to our understanding of the role of social relations during and after the hiring process.

Data described above could be brought to bear on the questions raised by the recent literature on the portability of star employees (Groysberg, Lee and Nanda, 2008; Groysberg, 2010). This literature demonstrates that above-average performance is not portable, i.e. that above-average performers, once they move to another organization, perform worse than before the move. This finding is particularly striking as it was obtained in the context of the

³¹ When I probed why this might be the case, the interviewees did not provide a unified answer. Some of them suggested that the search consultants are content with informally checking on their placements. Others commented that it is hard to define and systematically measure “success” of the placement, and/or subsequent performance, as the reasons people are hired in the first place vary. One senior consultant for instance remarked that an often cited measure of a good placement, the “stick rate” (the amount of time the person stays in the job), is not the most appropriate metric. He explained that in some cases, e.g. corporate turnarounds, a person leaving after a year may constitute a success because it may mean the turnaround has been completed. It appears that this lack of consensus on the appropriate post-placement metrics may be a part of the explanation as to why search firms do not systematically record post-placement data.

Wall Street analysts, whose activities and clients change very little when they change employers. The authors provided evidence for several boundary conditions for this effect (e.g. it is weaker when star employees move along with their teams, it is weaker when the new employer has better capabilities than the old employer etc.). However, these boundary conditions largely consider the *post*-move factors. Since that literature didn't observe how the individuals that move are selected in the first place, it couldn't disentangle the relative impact of the post- and the *pre*-move factors. However, an appealing alternative explanation would be that mobility of stars is driven by some kind of "adverse selection" process, whereby people that shouldn't move do move, or move at the "wrong" point in time. A research design that would combine the kind of relational pre-move data used in this thesis with the detailed information on the post-move workplace integration and performance would provide important insights into how social relationships shape mobility and performance in organizations.

REFERENCES

- Aime, F., Johnson, S., Ridge, J.W., Hill, A.D. 2010. The routine may be stable but the advantage is not: Competitive implications of key employee mobility. Strategic Management Journal, 31(1): 75-87.
- Allison, P.D. 1996. Fixed-effects partial likelihood for repeated events. Sociological Methods & Research, 25(2): 207-222.
- Allison, P.D. 2009. Fixed effects regression models. Thousand Oaks, CA: Sage Publications.
- Autor, D.H. 2009. The economics of labor market intermediation: An analytic framework. In Autor, D.H. (Ed.), Studies of labor market intermediation. Chicago: University of Chicago Press.
- Barney, J.B. 1991. Firm resources and sustained competitive advantage. Journal of Management. 17(1): 99-120.
- Barron, J.M., Berger, M.C., & Black, D.A. 2006. Selective counteroffers. Journal of Labor Economics. 24(3): 385-409.
- Beckman, C.M., & Haunschild, P.R. 2002. Network learning: The effects of partners' heterogeneity of experience on corporate acquisitions. Administrative Science Quarterly. 47(1): 92-124.
- Bertrand, M., Goldin, C., & Katz, L.F. 2009. Dynamics of the gender gap for young professionals in the corporate and financial sectors. NBER Working Paper 14681. National Bureau of Economic Research, Inc.
- Bidwell, M. & Fernandez-Mateo, I. 2010. Relationship duration and returns to brokerage in the staffing sector. Organization Science. 21(6): 1141-1158.

- Bielby, W.T. & Bielby, D.D. 1999. Organizational mediation of project-based labor markets: Talent agencies and the careers of screenwriters. American Sociological Review, 64(1): 64-85.
- Blau, P.M. 1964. Exchange and power in social life. New York: John Wiley.
- Boswell, W.R., Boudreau, J.W., & Dunford, B.B. 2004. The outcomes and correlates of job search objectives: Searching to leave or searching for leverage? Journal of Applied Psychology, 89(6): 1083-1091.
- Brass, D.J. 2009. Connecting to brokers: strategies for acquiring social capital. In Bartkus, V.O. and Davis, J.H. (Eds.), Social capital: Reaching out, reaching in. Cheltenham: Edward Elgar.
- Brett, J.M. & Stroh, L.K. 1997. Jumping ship: Who benefits from an external labor market career strategy? Journal of Applied Psychology, 82(3): 331-341.
- Bretz, R.D., Boudreau, J.W., & Judge, T.A. 1994. Job search behavior of employed managers. Personnel Psychology, 47(2): 275-301.
- Bridges, W. P. & Villemez, W. J. 1986. Informal Hiring and Income in the Labor-Market. American Sociological Review, 51(4): 574-582.
- Broschak, J.P. 2004. Managers' mobility and market interface: The effect of managers' career mobility on the dissolution of market ties. Administrative Science Quarterly, 49(4): 608-640.
- Bull, C., Ornati, O., & Tedeschi, P. 1987. Search, hiring strategies, and labor-market intermediaries. Journal of Labor Economics, 5(4): S1-S17.
- Burt, R. S. 1992. Structural holes: The social structure of competition. Cambridge, MA: Harvard University Press.
- Burt, R. S. 1997. The contingent value of social capital. Administrative Science Quarterly, 42(2): 339-365.

- Burt, R. S. 1998. The gender of social capital. Rationality and Society. 10(1): 5-46.
- Burt, R. S. 2000. The network structure of social capital. Research in Organizational Behavior, 22: 345-423.
- Burt, R. S. 2004. Structural holes and good ideas. American Journal of Sociology, 110(2): 349-399.
- Cappelli, P. 2008. Talent on demand: Managing talent in an age of uncertainty. Boston, MA: Harvard Business School Press.
- Cappelli, P., & Hamori, M. 2005. The new road to the top. Harvard Business Review, 83(1): 25-32.
- Castanias, R.P. & Helfat, C.E. 2001. The managerial rents model. Journal of Management. 27: 661-678.
- Castilla, E. J. 2005. Social networks and employee performance in a call center. American Journal of Sociology, 110(5): 1243-1283.
- Castilla, E.J. 2007. Dynamic analysis in the social sciences. London, UK: Academic Press.
- Degraaf, N. D. & Flap, H. D. 1988. With a little help from my friends - social resources as an explanation of occupational-status and income in West-Germany, the Netherlands, and the United-States. Social Forces, 67(2): 452-472.
- Dokko, G., Rosenkopf, L. 2010. Social capital for hire? Mobility of technical professionals and firm influence in wireless standards committees. Organization Science. 21(3): 677-695.
- Emerson, R.M. 1962. Power-dependence relations. American Sociological Review. 27(1): 31-41.
- Fernandez, R. M., Castilla, E. J., & Moore, P. 2000. Social capital at work: Networks and employment at a phone center. American Journal of Sociology, 105(5): 1288-1356.

- Fernandez, R. M. & Gould, R. V. 1994. A dilemma of state power - brokerage and influence in the national-health policy domain. American Journal of Sociology, 99(6): 1455-1491.
- Fernandez, R.M., & Mors, M.L. 2008. Competing for jobs: Labor queues and gender sorting in the hiring process. Social Science Research, 37(4): 1061-1080.
- Fernandez, R.M., & Sosa, L. 2005. Gendering the job: Networks and recruitment at a call center. American Journal of Sociology, 111(3): 859-904.
- Fernandez-Mateo, I. 2007. Who pays the price of brokerage? Transferring constraint through price setting in the staffing sector. American Sociological Review, 72(2): 291-317.
- Fernandez-Mateo, I. 2009. Cumulative gender disadvantage in contract employment. American Journal of Sociology. 114(4): 871-923.
- Fernandez-Mateo, I., & King, Z. 2011. Anticipatory sorting and gender segregation in temporary employment. Management Science, Articles in Advance.
- Finlay, W. & Coverdill, J. 2002. Headhunters: Matchmaking in the labor market. Ithaca, NY: Cornell University Press.
- Fleming, L., & Waguespack, D.M. 2007. Brokerage, boundary spanning, and leadership in open innovation communities. Organization Science, 18(2): 165-180.
- Gould, R.V. & Fernandez, R.M. 1989. Structures of mediation: a formal approach to brokerage in transaction networks. Sociological Methodology, 19: 89-126.
- Granovetter, M.S. 1973. Strength of weak ties. American Journal of Sociology, 78(6): 1360-1380.
- Granovetter, M. S. 1974. Getting a job: A study of contacts and careers. Cambridge, MA: Harvard University Press.

- Grant, R.M. 1996. Toward a knowledge-based theory of the firm. Strategic Management Journal. 17(Special Issue): 109-122.
- Groysberg, B. 2010. Chasing stars: The myth of talent and the portability of performance. Princeton, NJ: Princeton University Press.
- Groysberg, B., Lee, L. E., & Nanda, A. 2008. Can they take it with them? The portability of star knowledge workers' performance. Management Science, 54(7): 1213-1230.
- Gulati, R. 1995. Does familiarity breed trust? The implications of repeated ties for contractual choice in alliances. Academy of Management Journal. 38(1): 85-112.
- Hamilton, B.H. & Nickerson, J.A. 2003. Correcting for endogeneity in strategic management research. Strategic Organization, 1(1): 51-78.
- Hamori, M. 2010. Who gets headhunted—and who gets ahead? The impact of search firms on executive careers. Academy of Management Perspectives, 24(4): 46-59.
- Hargadon, A. & Sutton, R. I. 1997. Technology brokering and innovation in a product development firm. Administrative Science Quarterly, 42(4): 716-749.
- Harris, D. & Helfat, C.E. 1997. Specificity of CEO human capital and compensation. Strategic Management Journal, 18: 895-920.
- Herron, M.C. 1999. Postestimation uncertainty in limited dependent variable models. Political Analysis, 8(1): 83-98.
- Hitt, M.A., Bierman, L., Shimizu, K., & Kochhar, R. 2001. Direct and moderating effects of human capital on strategy and performance in professional service firms: A resource-based perspective. Academy of Management Journal, 44(1): 13-28.
- Ibarra, H. 1992. Homophily and differential returns: Sex differences in network structure and access in an advertising firm. Administrative Science Quarterly, 37(3): 422-447.

- Ibarra, H. 1997. Paving an alternative route: Gender differences in managerial networks. Social Psychology Quarterly, 60(1): 91-102.
- Khurana, R. 2002. Searching for a corporate savior: The irrational quest for charismatic CEOs. Princeton, NJ: Princeton University Press.
- Lacetera, N., Cockburn, I.M., & Henderson, R. 2004. Do firms change capabilities by hiring new people? A study of the adoption of science-based drug discovery. In Baum, J.A.C. & McGahan, A.M. (Eds.) Business Strategy Over the Industry Life Cycle. Elsevier: Amsterdam.
- Lazear, E.P. 1986. Raids and offer matching. In Ehrenberg, R. (Ed.) Research in Labor Economics, 8, JAI Press: Greenwich, CT.
- Lin, N., Ensel, W. M., & Vaughn, J. C. 1981a. Social resources and strength of ties - structural factors in occupational-status attainment. American Sociological Review, 46(4): 393-403.
- Lin, N., Vaughn, J. C., & Ensel, W. M. 1981b. Social resources and occupational-status attainment. Social Forces, 59(4): 1163-1181.
- Long Lingo E., & O'Mahony, S. Nexus work: Brokerage on creative projects. Administrative Science Quarterly, 55(1): 47-81.
- Marsden, P. V. 1982. Brokerage behavior in restricted exchange networks. In P. V. Marsden & N. Lin (Eds.), Social structure and network analysis. Beverly Hills: Sage.
- Marsden, P. V., & Hurlbert, J. S. 1988. Social resources and mobility outcomes - a replication and extension. Social Forces, 66(4): 1038-1059.
- McCool, J.D. 2008. Deciding who leads: How executive recruiters drive, direct and disrupt the global search for leadership talent. Mountain View, CA: Davies-Black Publishing.
- McEvily, B., Perrone, V., & Zaheer, A. 2003. Trust as an organizing principle. Organization Science, 14(1): 91-103.

- Milkovich, G.T. & Newman, J.M. 2005. Compensation (8th ed.). New York: McGraw-Hil.
- Morgan, L.A. 1998. Glass-ceiling effect or cohort effect? A longitudinal study of the gender earnings gap for engineers, 1982-1989. American Sociological Review, 63(4): 479-493.
- Murray, S. O., Rankin, J. H., & Magill, D. W. 1981. Strong Ties and Job Information. Sociology of Work and Occupations, 8(1): 119-136.
- Obstfeld, D. 2005. Social networks, the tertius iungens and orientation involvement in innovation. Administrative Science Quarterly, 50(1): 100-130.
- Padgett, J.F. & Ansell, C.K. 1993. Robust action and the rise of the Medici, 1400-1434. American Journal of Sociology, 98(6): 1259-1319.
- Perrone, V., Zaheer, A., & McEvily, B. 2003. Free to be trusted? Organizational constraints on trust in boundary spanners. Organization Science, 14(4): 422-439.
- Petersen, T., & Morgan, L.A. 1995. Separate and unequal: Occupation-establishment sex segregation and the gender wage gap. American Journal of Sociology, 101(2): 329-365.
- Petersen, T., Saporta, I., & Seidel, M. D. L. 2000. Offering a job: Meritocracy and social networks. American Journal of Sociology, 106(3): 763-816.
- Phillips, D.J. 2002. A genealogical approach to organizational life or chances: The parent-progeny transfer among Silicon Valley law firms, 1946-1996. Administrative Science Quarterly, 47(3): 474-506.
- Podolny, J.M. 2001. Networks as pipes and prisms of the market. American Journal of Sociology, 107(1): 33-60.
- Podolny, J. M. & Baron, J. N. 1997. Resources and relationships: Social networks and mobility in the workplace. American Sociological Review, 62(5): 673-693.

- Pollock, T. G. 2004. The benefits and costs of underwriters' social capital in the US initial public offerings market. Strategic Organization, 2(4): 357-388.
- Pollock, T. G., Porac, J. F., & Wade, J. B. 2004. Constructing deal networks: Brokers as network "architects" in the USIPO market and other examples. Academy of Management Review, 29(1): 50-72.
- Polzer, J.T. 1995a. Role. In Nicholson, N. (Ed.), Encyclopedic dictionary of organizational behavior. Cambridge, MA: Blackwell Publishers.
- Polzer, J.T. 1995b. Role set. In Nicholson, N. (Ed.), Encyclopedic dictionary of organizational behavior. Cambridge, MA: Blackwell Publishers.
- Puranam, P., Singh, H., & Zollo, M. 2006. Organizing for innovation: Managing the coordination-autonomy dilemma in technology acquisitions. Academy of Management Journal, 49(2): 263-280.
- Reitzig, M., & Puranam, P. 2009. Value appropriation as an organizational capability: The case of IP protection through patents. Strategic Management Journal, 30(7): 765-789.
- Rider, C. I. 2009. Constraints of the control benefits of brokerage: A study of placement agents in U.S. venture capital fundraising. Administrative Science Quarterly, 54(4): 575-601.
- Shapiro, S.P. 1987. The social control of impersonal trust. American Journal of Sociology, 93(3): 623-658.
- Simmel, G. 1950. The sociology of Georg Simmel. Toronto, Ontario: Free Press.
- Sorensen, J.B. 1999. Executive migration and interorganizational competition. Social Science Research. 28(3): 289-315.
- Sorensen, J.B. 2004. Recruitment-based competition between industries: a community ecology. Industrial and Corporate Change. 13(1): 149-170.

- Sturman, M., Walsh, K., & Cheramie, R.A. 2008. The value of human capital specificity versus transferability. Journal of Management, 34(2): 290-316.
- Tushman, M. 1977. Special boundary roles in the innovation process. Administrative Science Quarterly, 22(4): 587-605.
- Uzzi, B. 1996. The sources and consequences of embeddedness for the economic performance of organizations: The network effect. American Sociological Review, 61(4): 674-698.
- Uzzi, B. 1997. Social structure and competition in interfirm networks: The paradox of embeddedness. Administrative Science Quarterly. 42(1): 35-67.
- Uzzi, B. 1999. Embeddedness in the making of financial capital: How social relations and networks benefit firms seeking financing. American Sociological Review. 64: 481-505.
- Waguespack, D.M. & Sorenson, O. 2010. The ratings game: Asymmetry in classification. Organization Science, Articles in Advance.
- Wezel, F.C., Cattani, G., & Pennings, J.M. 2006. Competitive implications of interfirm mobility. Organization Science. 17(6): 691-709.
- Williams, M. 2001. In whom we trust: Group membership as an affective context for trust development. Academy of Management Review, 26(3): 377-396.
- Wooldridge, J.M. 2002. Econometric analysis of cross section and panel data. Cambridge, MA: MIT Press.
- Xiao, Z. X. & Tsui, A. S. 2007. When brokers may not work: The cultural contingency of social capital in Chinese high-tech firms. Administrative Science Quarterly, 52(1): 1-31.
- Yakubovich, V. 2005. Weak ties, information, and influence: How workers find jobs in a local Russian labor market. American Sociological Review, 70(3): 408-421.

- Yakubovich, V. & Lup, D. 2006. Stages of the recruitment process and the referrer's performance effect. Organization Science, 17(6): 710-723.
- Yavas, A. 1994. Middlemen in bilateral search markets. Journal of Labor Economics, 12(3): 406-429.
- Zaheer, A., & Venkatraman, N. 1995. Relational governance as an interorganizational strategy: An empirical test of the role of trust in economic exchange. Strategic Management Journal, 16(5): 373-392.
- Zaheer, A., McEvily, B., & Perrone, V. 1998. Does trust matter? Exploring the effects of interorganizational and interpersonal trust on performance. Organization Science, 9(2): 141-159.
- Zucker, L.G. 1986. Production of trust: Institutional sources of economic structure, 1840-1920. In Staw, B.M., & Cummings, L.L. (Eds.), Research in organizational behavior, Vol. 8. Greenwich, CT: JAI Press.
- Zuckerman, E. W. 1999. The categorical imperative: Securities analysts and the illegitimacy discount. American Journal of Sociology, 104(5): 1398-1438.

APPENDIX

Interview questions – search consultants:

1. Project sourcing:

- What brings in assignments?
- What makes clients hire a particular consultant or a particular firm?
- How do you build relationships that bring in assignments?
- How do you build firm and individual reputation?
- What's the relationship between consultants and search firms?

2. Candidate sourcing:

- How do you approach candidates?
- What do you do to convert individuals into candidates?
- What are the criteria on the basis of which you create candidate shortlists?

3. Search process:

- Is there a particular process in search?
- When assignments are taken on, what do you think are the factors that make assignments go well, or not so well?
- Do you think assignments can go wrong because of something search consultants do?
- How do you mediate between clients and candidates?
- What does fit between client and the selected candidate consist of?
- Do you follow up and check on people you have placed?
- Do you face some problems with the HR departments?

- Does disunity in the client hiring team occur? How do you deal with it?
- Do investors or other stakeholders, such as media, affect your work?
- How are you involved in the act of resigning?
- How are international searches conducted?
- In what way are approaches of individual consultants different?
- How much of a factor is price?

4. Definitions of success and measures of performance:

- How do you define success?
- How do you track success of your placements?
- Do you think there are factors on the side of the search consultant that impact how long people stay, or how successful people are in their jobs?
- Can you put a number on how many searches that you take up actually end up as successes?
- What differentiates search consultants and search firms?

5. Search consultants' roles:

- Do you have any responsibility to candidates, or to the clients only?
- Are search consultants legitimizing agents?