KAUFFMAN DISSERTATION EXECUTIVE SUMMARY

KAUFFMAN
ENTREPRENEURSHIP SCHOLARS

KAUFFMAN DISSERTATION FELLOWSHIP

Part of the Ewing Marion Kauffman Foundation's Entrepreneurship Scholars initiative, the Kauffman Dissertation Fellowship recognizes exceptional doctoral students and their universities. The annual program awards Dissertation Fellowship grants to Ph.D., D.B.A., or other doctoral students at accredited U.S. universities to support dissertations in the area of entrepreneurship.

Since its establishment in 2003, this program has helped to launch world-class scholars into the exciting and emerging field of entrepreneurship research, thus laying a foundation for future scientific advancement. The findings generated by this effort will be translated into knowledge with immediate application for policymakers, educators, service providers, and entrepreneurs as well as high-quality academic research.

EWING MARION
KAUFFMAN
FOUNDATION

www.kauffman.org/kdf

Essays in the Economics of Technology and Innovation

Raviv Murciano-Goroff

Abstract

For innovative ventures and startups, recruiting and organizing talented workers is a critical challenge. Research on high-tech labor markets and human resource practices have been hampered by a dearth of data available for study. The first two chapters of my dissertation examine detailed data on the recruiting of highly skilled engineers at both established firms and young entrepreneurial startups as well as information gathered from a novel survey on the formation of innovative teams. Among the counterintuitive facts highlighted in my dissertation, I find differences by gender in the advertising of verified technical skills on digital resumes, which recruiters at both established and startup firms do not appear to be adjusting for in their recruiting decisions. In addition, I show that despite the increasing usage of technology for sustaining and managing collaborations across geographies, the founding and formation of innovative teams occurs predominately when the team members are collocated.

Category: Economics: Background of Entrepreneurs: Gender and Teams.

Keywords: Hiring, Human Capital and Skills, Team Formation.

Executive Summary

The production of innovative goods is pivotal to economic growth. In recent decades, digital tools and technologies have changed the nature of high-tech production. Whereas the process of creating high-tech innovations and products had been driven by investments in physical capital, digitization has increased the relative importance of human capital. Many companies say that finding and bringing together the right workers with the right skills is one of their most important challenges. Indeed, finding the right workers is absolutely critical to the success of entrepreneurial and early-stage firms. The first two chapters of my dissertation study how companies find and organize highly-skilled engineers and scientists. Attention is given to contrasting the strategies and patterns of recruiting at larger companies with startups in the earlier stages of formation. Specifically, the first chapter of the dissertation speaks to the challenge a firm faces in trying to develop a diverse workforce, a particular challenge for smaller or startup companies. The second chapter discusses how teams of scientists and inventors come together to develop collaborations, a pivotal step in the instigation of new ventures.

The final chapter of my dissertation pivots to examine state and federal policies that influence the uptake and diffusion of new technologies. One of the interesting questions that arises in the study of technology adoption and diffusion is the role of small and young ventures. This chapter reveals that while state laws on data security have influenced the investment of larger, established firms in Information Technology (IT) security, smaller startups have largely not responded. Counterintuitively, this result is not because younger firms exclusively adopt newer IT technologies from their inception.

Chapter 1. Missing Women in Tech: The Labor Market for Highly Skilled Software Engineers

The first chapter of my dissertation examines the labor market for software engineers. One of the most frequently discussed questions regarding technology companies is why their workforces are persistently gender imbalanced. Despite concerted efforts, many tech companies have been unable to increase the representation of women among their engineering staff. In 2015, tech giants such as Google, Facebook, and Twitter had a mere 17%, 15%, and 10% of their respective technical staff positions filled by female engineers.[1] With fewer resources, smaller firms and startups face particular challenges to finding, recruiting, and hiring diverse workers.

In an effort to increase diversity among the workers they hire, many employers consider gender-blind recruiting methods, such as removing first names and photos from resumes. While these methods can eliminate the effect of biases by hiring staff, if there are differences in the way in which male and female job seekers for software engineering positions portray their skills on their resume, gender-blind recruiting methods can actually have the unintended consequence of decreasing the representation of women within the hiring pipeline. Because of these potentially countervailing effects, determining what can be done to improve diversity in recruiting and hiring at tech firms—particularly at smaller startups and new ventures—requires answering two questions: 1) do gender differences in the behaviors of job seekers exist, and 2) do recruiters adjust based on such differences in ways that could increase the diversity of the job applicant pool?

---

[1] The data for these statistics comes from the equal-opportunity data websites for these firms as well as news reports, such as https://www.huffingtonpost.com/2015/03/27/women-in-tech_n_6955940.html.

I examine the initial screening and recruiting of candidates for software engineering positions using data from a large online recruiting platform. On this platform, job seekers for software engineering positions posted digital resumes with a list of skills they feel proficient in. For a subsample of those candidates, I am able to find actual previous computer code they created and uploaded online. Thus, I am able to compare the programming skills individuals claim proficiency in with some of their actual previous coding work. This comparison enables me to quantify the extent of gender differences in the advertising of programming abilities. In addition, recruiters from tech companies—both large and small— subscribe to this platform in order to find and contact potential hires. I observe which candidates on the platform recruiters expressed interest in contacting. Thus, I can examine if the self-reporting of technical skills, such as programming language knowledge, predicts different probabilities of recruiters showing interest in male and female candidates. Finally, I can differentiate if recruiters from smaller, entrepreneurial firms approach the recruitment process differently than those from larger, more established firms.

The analysis in this chapter reveals three counterintuitive results. First, recruiters on the platform appear very interested in the self-reported skills that job seekers list on their digital resumes. Recruiters could have treated these self-reported skills as larger cheap-talk. Instead, recruiters are approximately 30% more likely to contact a candidate who self-reports knowing a technical skill even when the recruiter is presented with other objective evidence of the candidate's demonstrated ability in that skill. One possible reason for recruiters' interest in self-reported skills may be because the listing of a skill indicates that the candidate is interested in jobs using that skill and will be more likely to accept a job involving that skill. Because of their more limited resources for hiring and thus need to be more targeted in whom they contact,

recruiters at smaller firms are particularly inclined towards candidates who self-report knowledge of a technical skill relative to candidates with similar demonstrated abilities but who do not also self-report their proficiency in that skill.

Given that recruiters pay so much attention to these self-reports, one might assume that all job seekers would make sure to diligently list each of their known technical skills. Surprisingly, I find that the female software engineers are less likely to self-report their demonstrated technical skills than their male counterparts even after controlling for measures of their coding experience in those languages. In particular, given similar previous coding experience in a programming language, female coders are approximately 9% less likely to self-report knowledge of that language on their digital resume.

The finding that gender differences in the propensity of male and female coders to self-report their known skills might not matter if recruiters differentially respond to self-reported skills on resumes. Indeed, if recruiters are looking for the most experienced coders then we would hypothesize that they would be more likely to show interest in female coders who self-report knowledge of a programming language than a male coder with a similar resume who also self-reports that language. Instead, I find that female candidates on average receive about 12% less attention than male candidates with otherwise similar skills on the platform, while the predicted response of recruiters to self-reported programming skills are statistically indistinguishable. This implies that on average recruiters treat the self-reporting mechanism in similar ways for male and female candidates. Because of the gender difference in self-reporting, recruiters are likely overlooking some qualified female coders. Furthermore, recruiters with less experience are more likely to fail to adjust for supply side differences in the propensity to self-

report. Therefore, startups and young ventures seeking to create diverse workforces should pay particular attention to how recruiters respond to self-reported skills on digital resumes.

Chapter 2. Why and Wherefore of Increased Scientific Collaboration

Increasingly, the creation of innovative research, products, and entrepreneurial ventures depends on teams of highly skilled individuals working together. At the same time, modern internet and communication technologies enable collaborators to more easily connect across geographies and distance. The second chapter of my dissertation, coauthored with Richard B. Freeman and Ina Ganguli, explores how and why scientific teams form and what factors contribute to their success. We focus particular attention on the question of whether or not internet based communication technologies have enabled teams of innovators to find and initiate projects across geographic distances.

For this investigation, we study the connections among innovators using data from an original survey of corresponding authors from scientific journal articles and the *Web of Science* citation database. We focus on journal articles in which at least one scientist is based in the United States and working in the fields of Particle and Field Physics, Nanoscience and Nanotechnology, or Biotechnology and Applied Microbiology. In addition to analyzing the patterns of collaborations and subsequent citations, we designed and distributed a novel, adaptive survey which we sent to these innovators. For each scientist surveyed, we asked the respondent about how they met each of their collaborators, why he or she choose to work with the other individuals, and how the team members communicated with each other throughout the development of their ideas and projects.

We uncovered four main findings. First, we find that collaborations among innovators within the United States have increased across cities domestically as well as across international borders. In particular, the nature of collaborations across cities resembles that of across countries, with geographic distance explaining far more variation in innovative team formation than geo-political considerations. Second, face-to-face meetings are important in collaborations: most collaborators first met working in the same institution and communicate often through meetings with coauthors from distant locations. Third, the main reason for most collaborations is to combine the specialized knowledge and skills of coauthors, but there are substantial differences in the mode of collaborations between small lab-based science and big science, where international collaborations are more prevalent. Fourth, for biotech, we find that citations to international papers are higher compared to papers with domestic collaborators only, but not for the other two fields. Moreover, in all three fields, papers with the same number of coauthors had lower citations if they were international collaborations.

Overall, our findings suggest that all collaborations are best viewed from a framework of collaborations across space broadly, rather than in terms of international as opposed to domestic collaborative activity. In addition, the results of this analysis suggest that in order to increase collaborations across diverse groups, funding agencies and foundations should focus their attention on developing programs that physically bring these groups together. Finally, rather than emphasizing the founders and entrepreneurs as jacks-of-all-trades, our research indicates that the increased prevalence of collaborations for innovative activity is largely driven by the need for complimentary specialized knowledge.

Chapter 3. Data Breach Disclosure Laws and Digital Security Conscientiousness

A long standing question regarding technology diffusion asks about the relative roles of young startups and established firms in the uptake of new technologies. Do established companies adopt new technologies first, testing them in contained ways and establishing new technological standards for others to follow? Alternatively, are younger startups the early adopters of new technologies with established firms learning from the results of experimentation done across entrepreneurial ventures?

The third chapter of my dissertation provides insights into this question while exploring how particular legal changes encourage firm investments in adopting new web technologies. I evaluate if data breach notification laws—legal mandates that companies who have had their data accessed by unauthorized individuals inform their customers of this breach—encourage firms to invest in information technology security. I pay particular attention to the heterogeneity in the impact of these laws and technology adoption patterns by the size and age of the firm. In doing so, I provide descriptive facts about whether large or small firms are the early adopters of new web technologies.

To evaluate the patterns of investment in IT technology as well as the impact of data breach notification laws, I collected data on the decisions of over 2,185 companies regarding which server software to install, when to update their web server software, and when to apply security patches during the two years before and after California legislators created the first data breach notification law in the United States. Comparisons are made with groups of companies that were not directly affected by the this law: companies without a physical presence in California, companies whose websites were unlikely to store personal information, and finance companies that were already subject to previously enacted federal legislation on data breaches.

This chapter begins by documenting the patterns of web server technology usage by companies. Examining the technological vintages of software used by firms, I explore if startups and entrepreneurial companies utilize the latest versions of technologies at the time of their founding. Furthermore, I ask if this drives changes in the average technological age of server software used across firms. Somewhat surprisingly, I find that startups do not necessarily start out by using the latest web server technology. Indeed, possibly because of risk aversion due to the more frequent changes in the newer editions and because of the often higher investment required to use cutting edge versions the average technology age at founding is not statistically different than the technology age among more established firms.

Regarding the impact of legal liability changes, I find that the data breach notification law in California caused firms based there to use web server software that was 4.88-12.06% newer. As newer server software tends to address the known security problems within previous versions, this finding implies that the notification law encouraged firms to use server software that would be less vulnerable to hacking. The effect of this law was most pronounced for companies using open source rather than proprietary web server software. The reason for this difference is most likely because open source technologies require firms to actively download new versions of software, whereas proprietary server technologies had functions to automatically update. As larger and more established firms are more likely to use proprietary software, this could help explain why the distribution of technology ages are similar for established and startup ventures.

My research also revealed that larger and more established firms responded far more to the legal change in California than smaller and younger ventures. Whereas firms with tens of thousands of workers responded to the increased liability due to the legislative change by

installing patches and adopting the latest versions of server software, entrepreneurial and young firms did not change their habits in statistically discernable ways. As documented, this is not because the younger firms were already using the latest software.

One reason for the differential response to the legal changes could be because of differences in the level of threat to their public facing websites. Companies with higher trafficked websites were more likely to respond to the legal change. As higher trafficked websites were also more likely to be targeted for hacking, this might have motivated these firms to be more strict in their security conscientiousness and tech adoption. Indeed, after controlling for the relative amount of web traffic, the heterogeneity in the effect of the law becomes attenuated although still statically significant.

The California data breach law became the template for all other states to subsequently enact data breach notification laws. Somewhat surprisingly, however, these subsequent laws had little if no meaningful impact on the web server security and updating behavior of firms. One possible reason for this is because despite the laws being enacted individually by states, each state's law had far reaching implications. For example, if a company based in Massachusetts had a data breach in which data about customers in California fell into unauthorized hands, the firm could still be held liable in a California court if they failed to disclose the breach to the California residents. Because of this, the initial California law may have acted as a defacto national legal policy.

Subsequent research will be required to understand more about how startups and established firms differentially respond to legal liability changes and their roles in the diffusion of new technologies. I plan to continue this work in the coming months as a I transition into a postdoc and assistant professorship.