

**Essays on Organizing Human Capital, Automation and Innovation**

By

Harshvardhan J. Ketkar

A dissertation submitted in partial fulfillment  
of the requirements for the degree of  
Doctor of Philosophy  
(Business Administration)  
in the University of Michigan  
2020

Doctoral Committee:

Associate Professor Seth Carnahan, Co-Chair, Washington University in St Louis  
Associate Professor Felipe Csaszar, Co-chair  
Professor Gautam Ahuja, Cornell University  
Professor Jerry Davis  
Professor Kentaro Toyama

Harshvardhan J. Ketkar

[hketkar@umich.edu](mailto:hketkar@umich.edu)

ORCID iD: 0000-0002-7192-9680

© Harshvardhan J. Ketkar 2020

## **DEDICATION**

*To Dad and Aai,  
for making me who I am.*

## ACKNOWLEDGMENTS

It takes a village to raise a child. And quite frankly, a dissertation is nothing less than a child. It gives you extreme joys and extreme despair; it makes you ill-tempered as well as exuberant; it wanders all around, yet you must not go mad at it and cajole it back to the right path; and most importantly, it makes you, the parent, most insufferable to everyone else by making you incapable of speaking about anything else (“*Did you say underwater basket weaving? My research deals EXACTLY with that!*”).

It certainly does take an entire village.

In my case, I had an entire tribe on my side. Firstly, I was incredibly lucky to have a wonderful dissertation committee who not only served as an inspiration for me but were also always available whenever I sought out their help. Seth Carnahan, my adviser, was my friend, philosopher and guide in this journey. Despite the distance (he had moved to Washington University in St Louis), he was always there whenever I needed advice. When it comes to dissertation advisers, they don’t make ‘em like him very often. It is seldom that people manage to have a good heart and a sharp mind at the same time; Seth is without a slightest doubt, one of those. It is this rare combination of openness to ideas and thoroughness in execution, coupled with a big heart to boot, that always made me feel jealous of myself at the end of our Skype meetings. Thank you, Seth. Having you as my adviser was the best decision I made during my PhD.

Felipe Csaszar, my co-adviser, was a source of stability for me at Ross; he sat down with me, delved deeply into the finer aspects of my dissertation and job-market paper, and taught me to take rigorous approaches towards solving problems. Gautam Ahuja has always been an inspiration for me, not just as a scholar, but also as a person, and for the sincere interest he has in the welfare of PhD students. Jerry Davis has been much more than just a committee member for me. I have often said to him that he is the one who provided the gateway drug for my research topic through his seminar in my first term here. To say that he has had a huge impact in shaping my ideas would be an understatement; if it were not for him, I would be working on something else. But more importantly, he was always there to provide me support in difficult times. I have also been extremely lucky to have Kentaro Toyama on my committee; his feedback on my job market paper has been invaluable, and like a surgeon, he used to get down to the core issue within minutes.

But even apart from my wonderful dissertation committee, I have been especially fortunate to have been blessed with mentors who took an active interest in my work and shaped my ideas. In the short time that we overlapped at Ross, I received invaluable help from Chris Rider, be it for developing my job market paper and presentation, or through his advice on being a scholar and conducting research, and beyond. His words have considerably impacted the way I view academic research in general. Thank you, Chris, for your support. Phanish Puranam has been a guardian angel for me. I drew extensively upon his work, and his ideas and feedback were instrumental in helping me develop my ideas. In the various conversations I have had with him over the years, I have always felt tremendously lucky to have his insights, and to have him as a role model. Thanks, Phanish, for your support and guidance. A special mention also for my co-

author Maciej Workiewicz; he has been a great peer-level mentor for me, and I learnt a lot from him.

It has been a true privilege to be a part of the Strategy group at Ross. Through numerous seminars and hallway discussions with the faculty in the department, I got a chance to observe what it means to be a first-rate researcher and scholar. I developed a bias toward impactful research and tackling them with rigor, and for that alone, I am indebted to the department. I want to sincerely thank Jordan Siegel and Jim Westphal for being so supportive and helpful through my years at Ross. I have always looked up to them as scholars, and I have never felt wanting for support from their side. They were always forthcoming and offered help, even before I reached out to them. Likewise, it was also great to have Justin Frake, Cheng Gao, Derek Harmon, Michael Jensen, Brian Wu and Maggie Zhou on my side.

I have also been incredibly fortunate to have wonderful friends in the Ross Strategy PhD program who are brilliant and kind, and always available for coffee and conversations on research and beyond. They patiently sat through my mock job talk presentations (thrice in three weeks) and provided excellent feedback, every single time. A big shout out to Cha Li, Aseem Sinha, Christine Choi, Diana Jue Rajasingh, Mana Heshmati, Reuben Hurst, Yun-ha Cho, David Chung and Jusang Lee. You guys are awesome; thank you so much for being there! Apart from them, I also want to thank other friends and office mates at Ross: Cassandra Chambers, Teddy Dewitt, Frank Li, George Malikov, Dylan Nelson, Prashant Rajaram, Jordan Rhodes, Guy Shani, Steve Shaw, Laura Sunday, Casidhe Troyer and Dana Turjeman. Finally, I would also like to thank Brian Jones and Ashley Stauffer for their support through the Doctoral Studies Office, as well as Susan Booth for administrative support.

I am also thankful to my friends in Ann Arbor without whom it would not have been possible to get through those years. In particular, I want to thank Juilee Rege and Gautam Sule, Apurva Gokhale and Harsh Karve, Pooja Panwalkar and Atul Khire, Manaswini Sarangi and Koustubh Vaze, Aditi Kulkarni, Upasana Kulkarni, and others in India and elsewhere. Those memorable dinner parties and game nights, stimulating conversations over coffee, and wonderful trips went a long way in keeping my sanity intact.

Lastly, but most importantly, I am grateful to my family for supporting me throughout this process. I was merely having fun doing what I loved (research), but it is they who made sacrifices so that I could pursue my dreams. Ashwini, thank you for being so patient and understanding, and standing firmly beside me; you stood with me even when things were not going the way I wanted them to, and I used to get cranky. My dissertation owes as much to you as it does to me. I am also thankful to my parents-in-law, Sharayu and Sateesh Bhide (*Aai* and *Baba*), for being a source of strength for Ashwini and me.

Finally, I am indebted to my dad, Jayant Ketkar, and *Aai*, Manjiri Ketkar, for everything. Thank you for instilling in me a love for learning; teaching me to be curious, to ask questions, and also providing me the means to attempt answering them; and wholeheartedly supporting my decision to pursue a career in research. You provided me a strong moral foundation that I was able to draw upon in difficult times; having you was the unfair advantage I had over others. I dedicate my dissertation to them.

## TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGMENTS	iii
LIST OF TABLES	ix
LIST OF FIGURES	xii
LIST OF ILLUSTRATIONS	xiv
LIST OF APPENDICES	xv
ABSTRACT	xvi
Chapter I: Introduction	1
Theory	2
Dissertation Chapters	4
Conclusion	7
Chapter II: <i>Automation and its Discontents</i> : The Impact of Automating Knowledge Integration on Innovation	8
Introduction	8
Theory	12
Empirical Setting	22
Data	28
Results	32
Discussion and Conclusion	41
Chapter III: Automation and the Evaluation of Men and Women’s Work Product: Evidence from Software Contributions on Github	46



Introduction	46
Theory	49
Data and Empirical Design	53
Results	59
Discussion and Conclusion	75
Chapter IV: <i>Power to the People</i> : Decentralized Project Selection and Employee Self-Allocation in Organizations	78
Introduction	78
Theory	82
Model	91
Results	96
Discussion	115
Conclusion	120
Chapter V: Discussion	122
Future Directions	124
Conclusion	125
APPENDIX	126
Appendix A: Robustness Checks for Chapter II	126
Appendix B: Additional Analyses for Chapter III	128
Appendix C: Additional Robustness Analyses for Chapter IV	133
BIBLIOGRAPHY	138

## LIST OF TABLES

Table II.1: Descriptive Statistics (post-matching)	31
Table II.2. Effect of automation on contribution scope (DV = HHI)	34
Table II.3. Age at adoption as moderating variable	37
Table II.4. Effect on contribution scope (DV = HHI) for Natural Experiment	40
Table III.1 (a): Descriptive Statistics for Control Projects (prior to Travis adoption)	55
Table III.1 (b): Descriptive Statistics for Control Projects (after Travis adoption by matched Treatment Project)	56
Table III.1 (c): Descriptive Statistics for Treatment Projects (prior to Travis adoption)	57
Table III.1 (d): Descriptive Statistics for treatment Projects (after Travis adoption)	58
Table III.2: Descriptive statistics for computed dependent variable for all projects	58
Table III.3: Acceptance rate for female contributors before and after Travis-CI (Period wise dataset)	61
Table III.4: Rejection rate for female contributors before and after Travis-CI (Period-wise dataset)	62
Table III.5: Acceptance rate for female contributors before and after Travis-CI(consolidated observations)	63

Table III.6: Rejection rate for female contributors before and after Travis-CI (collapsed observations)	64
Table III.7: Acceptance rate for male contributors before and after Travis-CI (period-wise observations)	65
Table III.8: Rejection rate for male contributors before and after Travis-CI (period-wise observations)	66
Table III.9: Acceptance rate for male contributors before and after Travis-CI (collapsed observations)	67
Table III.10: Rejection rate for male contributors before and after Travis-CI (collapsed observations)	68
Table III.11: Percentage of male contributors before and after Travis-CI (period-wise observations)	69
Table III.12: Percentage of female contributors before and after Travis-CI (collapsed observations)	70
Table III.13: Percentage of male contributors before and after Travis-CI (period-wise observations)	71
Table III.14: Percentage of male contributors before and after Travis-CI (collapsed observations)	72
Table III.15: Percentage of neutral contributors before and after Travis-CI (period-wise observations)	73

Table III.16: Percentage of neutral contributors before and after Travis-CI (collapsed observations)	74
Table A1. Herfindahl Index calculated for a four week period	126
Table B1: Share of All Pull Requests Made by Female Programmers	117
Table B2: Share of All Pull Requests Made by Male Programmers	118
Table B3: Share of All Pull Requests Made by Neutral Programmers	118

## LIST OF FIGURES

Figure II.1: Leads and Lags for Main Results	35
Figure II.2: Leads and Lags on Natural Experiment (DV Match precision = 30%)	38
Figure IV.1: Performance Comparison Between Open and Centralized Allocations	98
Figure IV.2: Projects Launched and Active Workers for Each Form	99
Figure IV.3: Selection Errors	103
Figure IV.4: Allocation Errors	104
Figure IV.5: Allocation Commission Errors (ACEs) per project	106
Figure IV.6: Performance Impact of Human Resource Misallocation	107
Figure IV.7: Policy Levers	110
Figure A1. Month level HHI over 320 weeks	127
Figure B1: Acceptance rate for female programmers over 320 weeks	131
Figure B2: Rejection rate for female programmers over 320 weeks	131
Figure B3: Acceptance rate for male programmers over 320 weeks	132
Figure B4: Rejection Rate for Male Programmers over 320 weeks	132
Figure C1: Organizational performance with reduced level of attention ( $r_i = 8$ )	134

Figure C2: Performance when each worker can see only 50% (8 out of 16) of projects or 25% (4 out of 16), while the manager sees all of the projects (100%) 136

Figure C3: Comparison of performance for different levels of  $\sigma$  137

## LIST OF ILLUSTRATIONS

Illustration I.1: Overview of the theory	3
Illustration I.2: Overview of the chapters in the dissertation	6
Illustration II.1: Github Contribution Workflow	24
Illustration II.2: Github Contribution Workflow with Travis-CI	25
Illustration II.3: Example of Github project conversations before and after automation	27
Illustration IV.1: Calculating the project's profit	96
Illustration IV.2: Example of Employee Allocation	100

## **LIST OF APPENDICES**

Appendix A: Robustness Checks for Chapter II	126
Appendix B: Additional Analyses for Chapter III	128
Appendix C: Additional Robustness Analyses for Chapter IV	133



## ABSTRACT

In this dissertation, I explore how different modes of organizing and automation shape processes and outcomes of knowledge-based work within organizations. With advances in artificial intelligence and automation, firms are increasingly automating knowledge-based work which were previously carried out by humans alone. The use of such cognitive tools presents implications for adopting organizations in terms of how these tools shape knowledge production, or how they affect bias and inclusion against organizational members belonging to certain groups. Correspondingly, firms are also experimenting with new forms of organizing human capital across its various projects, since deploying it judiciously becomes a source of competitive advantage. Knowledge intensive tasks that were previously carried out by experts or managers are now either being delegated to lower-level employees, or automated using algorithms, or carried out by using a combination of the two. This dissertation focuses on tradeoffs associated with new forms of organizing and automation in the context of knowledge-based tasks such as human resource allocation, integration of new knowledge, and socialization of newcomers and outsiders. I explore these questions through the three chapters of this dissertation. In chapter II, I look at how automating the integration of new knowledge affects what kind of knowledge contributions get integrated into an organization's knowledge base. I use the context of software projects hosted on GitHub, some of which automated the process of evaluating code contributions. I find that projects that adopt automation tend to integrate narrower, component-

level contributions, rather than broader, systemic contributions, perhaps because the algorithm crowds out unstructured coordination among contributors and maintainers, which is necessary for systemic contributions. In chapter III, using the same context, we look at how automation shapes inclusion and discrimination against female code contributors on GitHub. We find that after automating code review, projects tend to attract code contributions from female programmers at a greater rate than before (and compared with non-adopters). They are also more likely to eventually accept code contributions from female programmers. Finally, in chapter IV, we provide a computational model of how organization structure (open allocation or hierarchical allocation) affects allocation of human resources to available. We find that the relative balance between the organization's human resources and the number of opportunities it faces is critical in determining the advantages of open allocation, which performs better when human resources are scarce relative to opportunities. Conversely, hierarchical allocation (with a manager) performs better when opportunities are scarce. Methodologically, I employ a combination of agent-based models and large datasets consisting of fine-grained, contribution level data to present my findings across the three chapters. Overall, this dissertation aims to better understand the impact using of different ways of organizing human capital and algorithmic automation for knowledge work within organizations.

## **Chapter I: Introduction**

Traditionally, organizations have been viewed as entities that facilitate the storage, manipulation and production of knowledge (Kogut and Zander, 1992; 1996; Grant, 1996). Tasks related to these activities require close coordination among its members and are more efficiently carried out within the boundaries of the organization, where managerial fiat or shared contexts alleviate conflict and misunderstanding (Conner and Prahalad, 1996).

Recent advances in technology, however, have made it possible to carry out knowledge-intensive tasks in new, non-traditional ways. Increasingly, firms are choosing to decentralize or automate core organizational functions such as allocating human resources to projects or carrying out knowledge-intensive evaluation-based tasks, which were previously carried out by groups of human employees. The use of algorithms and AI has helped reduce the costs of making knowledge-intensive decisions and predictions (Carr, 2015; Ford, 2015; Agrawal, Gans and Goldfarb, 2018). This has led to significant changes in 1) the way these knowledge intensive functions are carried out (Agrawal, Gans and Goldfarb, 2018) 2) the outcomes of these knowledge intensive processes (Furman and Teodoridis, 2020), and 3) the participation of individual members within organizations. Correspondingly, firms are widely experimenting with highly decentralized, new forms of organizing knowledge-based tasks, wherein core organizational functions previously carried out by experts or managers are now being delegated to lower level employees (Puranam, Alexy and Reitzig, 2014; Lee and Edmondson, 2017).

These developments give rise to some important questions. Firstly, if employees can carry out managerial tasks, why do organizations still need managers? What are the tradeoffs of choosing hierarchical control or decentralized organizing? Secondly, if robots and algorithms can automate complex knowledge-based tasks, why do organizations still need human employees? Are there tradeoffs associated with these choices, and if yes, what are they? Given that these considerations lie at the heart of organizations' existence (and their strategic outcomes), it is necessary to better understand and theorize about the consequences of these choices of organizing knowledge work. This dissertation sets out to do exactly that.

### **Theory**

The theory outlined in this dissertation helps understand how different ways of organizing knowledge work can shape various organizational processes and outcomes. The dissertation takes the view that knowledge work can either be carried out through decentralized or centralized forms of organizing, or through algorithmic automation. Firstly, I provide a theoretical perspective rooted in organization design theory that views automation as a form of structured coordination. I posit that the use of algorithms for knowledge-intensive tasks that were previously carried out by teams of humans transforms patterns of coordination among these workers from unstructured, organic and emergent to structured, predetermined and predictable. This changes the way these organizations draw upon knowledge stored within themselves, and consequently shapes the characteristics of new knowledge produced by them. Furthermore, automating the evaluation of knowledge contributions (either partly or wholly), changes *whose* contributions are more likely to be accepted, and therefore *who* is more likely to contribute within those organizations. Thus, the outlined theory provides insights into what implications automation and algorithms could present for knowledge-intensive work within organizations.

Secondly, organizations could carry out knowledge work through decentralized or centralized organizational forms. The outlined theory states that the efficacy of either form is determined by the relative balance of the available human resources and available opportunities to which these resources must be allocated. While prior theory views the efficacy of hierarchy or decentralization as stemming from managerial span of control (Woodward, 1965; Simon, 1997; Rajan and Wulf, 2006), this dissertation provides an additional perspective on the link between organization structure and performance.

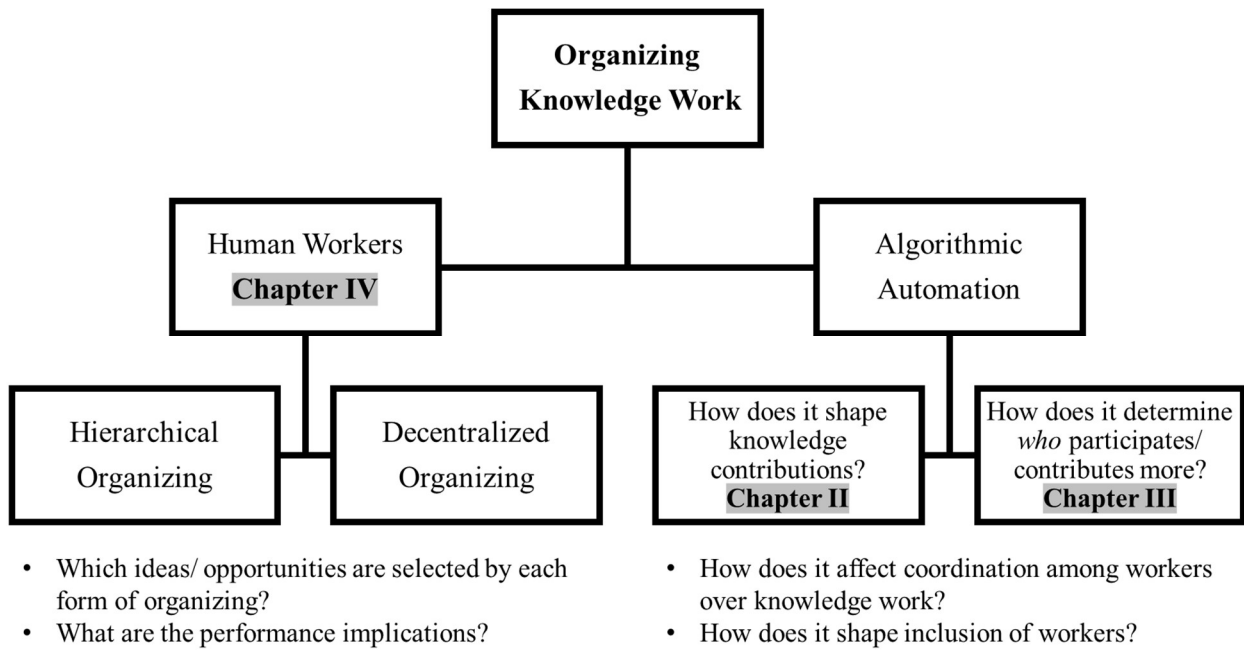


Illustration I.1: Overview of the theory

Thus, the chapters presented here explore various tradeoffs associated with choosing either decentralization, or automation, or a combination of these forms for carrying out knowledge-based tasks. Through this dissertation, I aim to make both theoretical and empirical contributions. From the standpoint of theory, this dissertation contributes to a vast literature

related to, more broadly, 1) the knowledge-based view of the firm (Kogut and Zander, 1992, 1996; Grant, 1996; Conner and Prahalad, 1996) , absorptive capacity (Cohen and Levinthal, 1990; Zahra and George, 2002; Volberda, Foss and Lyles, 2010). Furthermore, I contribute to an emerging body of literature that explores new forms of organizing (Puranam, Alexy and Reitzig, 2014), and the impact of AI and automation (Cowgill, Seamans and Ziv, 2017; Autor, 2015; Agrawal, Gans and Goldfarb, 2018; Von Krogh, 2018; Choudhury, Starr and Agarwal, 2019).

Lastly, I also make an empirical contribution. While there have been a few recent studies on the effects of automation (e.g. Beane, 2018; Oliver, Calvard and Potocnik, 2017) on organizational outcomes, nearly all of them employ qualitative methodologies<sup>1</sup>. Most of the current empirical work is at the aggregate level (e.g. Acemoglu and Restrepo, 2018), however there is no dataset that tracks adoption of AI and automation at the organizational or individual level (Raj and Seamans, 2019). My empirical setting, that of automation within software projects hosted on GitHub, allows me to access micro-level data related to code contributions made to those projects, as well as exploit a natural experiment. Thus, in my research, I provide one of the first large-scale causally identified results using a large dataset with fine grained data at the transaction level.

## **Dissertation Chapters**

The three chapters of my dissertation explore conceptually separate but theoretically interrelated themes within literature related to organization design, knowledge-based work and human capital. Chapter II is titled, *Automation and its Discontents: The Impact of Cognitive Task Automation on Organizational Innovation*. The idea for this chapter originates from the

---

<sup>1</sup> Notable exceptions are is that of Kleinberg, et al (2017) and Furman and Teodoridis (2019). However, they do not consider the organization-level effects of automation, which this dissertation addresses.

observation that recent advances in AI and Machine Learning are making possible the automation of white-collar jobs (Susskind and Susskind, 2017; Autor, 2015; Raj and Seamans, 2019) which are characterized by cognitively intensive, knowledge-based tasks. Given that organizations serve to store, maintain and create new knowledge, how would automating knowledge-based tasks shape processes and outcomes for such tasks? This is an unaddressed question which presents important strategic level implications, since the capability to produce knowledge is a source of competitive advantage for organizations. In this chapter, I explain the impact of automating cognitive tasks on the integration of new knowledge within organizations. I develop a theory that explains why the automation of knowledge integration within organizations reduces the ability of its members to integrate broad, systemic knowledge. I test the outlined theory using the empirical setting of GitHub projects, some of which adopt tools that partially automate the review, testing and integration of new code contributions, a process that was previously carried out manually by project owners and maintainers. To provide a causal explanation for my main findings, I exploit a natural experiment in my setting in which a troll bot sent activation requests quasi-randomly to software projects listed on a public registry. Results obtained strongly support the outlined hypotheses.

Chapter III, co-authored with Seth Carnahan, is titled *Automation and the Evaluation of Men and Women's Work Product: Evidence from Software Contributions on Github*. This chapter builds upon recent work that looks at how the use of algorithmic automation affects bias and discrimination against women and minorities (Cowgill and Tucker, 2019). Like chapter II, we use the context of GitHub projects, and we investigate how the use of automation shapes inclusion and discrimination against female code contributors on GitHub. We find that after automating code review, projects tend to attract code contributions from female programmers at

a greater rate than before (and compared with non-adopters). They are also more likely to eventually accept code contributions from female programmers. By looking at how automating evaluation of work contributions shapes the decisions of women programmers to participate and contribute, this study aims to make theoretical contributions to literature, and provide practical policy implications for managers.

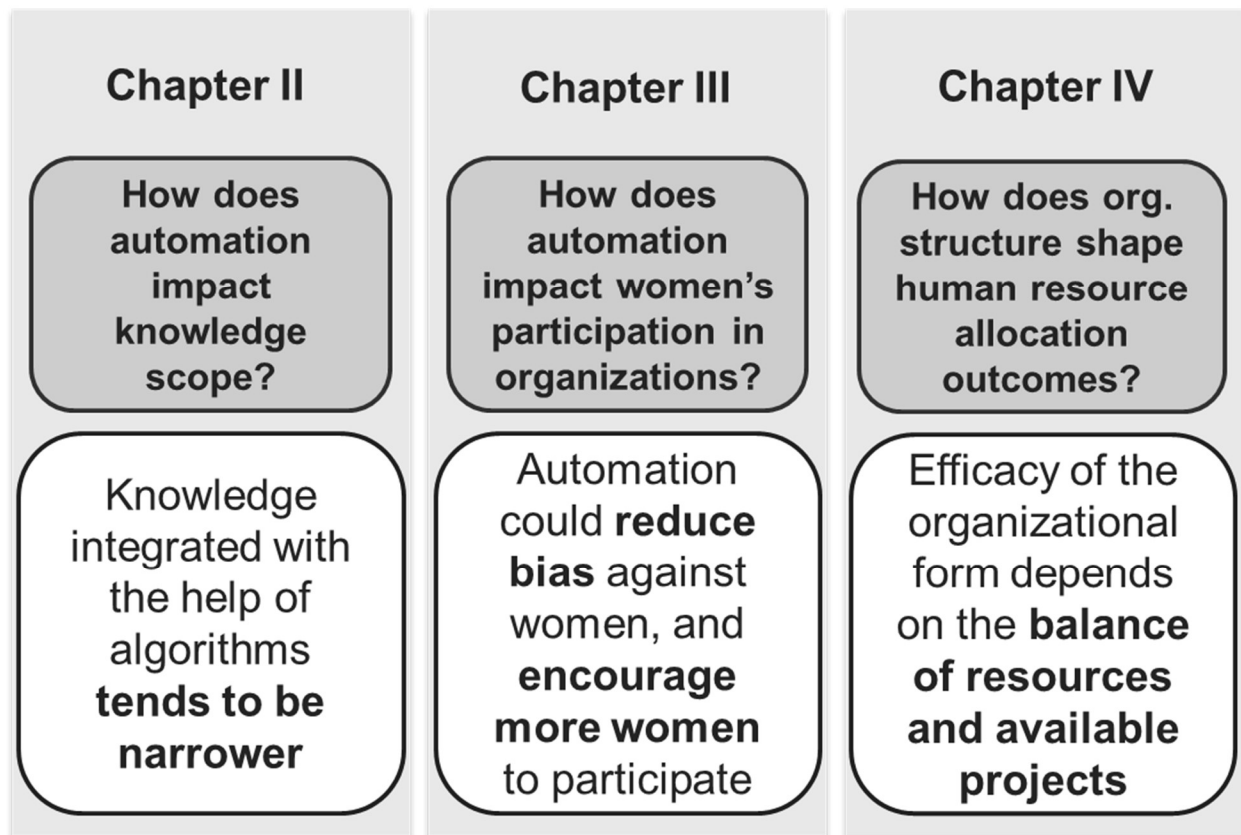


Illustration I.2: Overview of the chapters in the dissertation

Chapter IV, co-authored with Maciej Workiewicz, is titled *Power to the People: Decentralized Project Selection and Employee Self-allocation in Organizations*. This chapter explores the tradeoffs associated with choosing centralized or decentralized control over allocating human resources to the most lucrative available opportunities. This tradeoff emerges over a dilemma associated with allocating human resources in knowledge-based firms: on the



one hand, employees in such firms are engaged in knowledge-based tasks, and they value autonomy highly; on the other hand, it is necessary to ensure that they do not engage in wasteful hobby projects or crowd the extremely attractive projects and reduce organizational payoffs. In this chapter, using an agent-based model, we examine whether decentralized (open) allocation can help firms achieve better allocation of human resources than that offered by a traditional hierarchical solution. Our results suggest that the relative balance between the organization's human resources and the number of opportunities it faces plays a critical role in determining the advantages of open allocation, which performs better when human resources are scarce relative to opportunities. We also examine a set of common managerial interventions and find that depending on the balance between available resources and opportunities, these policies may produce opposite results to those intended.

## **Conclusion**

Through my dissertation (and my research in general), I hope to uncover the effects of new forms of organizing and automation on knowledge-based work and associated organizational outcomes. I believe that my dissertation provides insights into the underlying processes related to human resource allocation, integration of new knowledge, and assimilation of outsiders into the organization under new ways of organizing knowledge work such as decentralization and automation. Taken together, the chapters in this dissertation also offers practical implications for managers when choosing between hierarchy, decentralized forms or algorithms to organize knowledge work. From both theory and phenomenon point-of-view, a lot of exploration remains to be done in these areas, and I believe that my dissertation provides a starting point to pursue exciting opportunities for further research.

## **Chapter II: *Automation and its Discontents: The Impact of Automating Knowledge Integration on Innovation***

### **Introduction**

A vast body of prior literature views organizations and firms as repositories of knowledge and routines (Grant, 1996; Nelson and Winter 1982; Kogut and Zander, 1996). Among other things, organizations combine existing knowledge spread across its members, integrate it and draw upon it to produce new knowledge and engage in problem-solving (Kogut and Zander, 1992). This is even truer for creative, knowledge-based organizations, which engage in collaborative, complex problem solving by drawing upon their existing knowledge sets (Nickerson and Zenger, 2004). A core task for managers is to integrate the work outputs of employees into a coherent whole; for example, managers in consulting teams integrate individual contributions into a well-structured report. The ability to draw widely upon existing organizational knowledge and integrate new knowledge into it is an important source of competitive advantage (Cohen and Levinthal, 1990; Zhou and Li, 2012; Nagle and Teodoridis, 2017; Tortoriello, 2015; Kim and Anand, 2018; Cockburn and Henderson, 1994). Through sustained and repeated collaborative knowledge production, organizations develop coordinative capabilities (Teece, Pisano, and Shuen, 1997; Eisenhardt and Martin, 2000; Helfat and Martin, 2015) which are instrumental in integrating new knowledge.

Specifically, close coordination within members is critical for integrating new knowledge (Okhuysen and Bechky, 2009; Garicano and Wu, 2012; Tortoriello, 2012; Ben-Menahem, *et al.*

2016; Gardner, Gino, and Staats, 2012). These processes related to coordination and integration of knowledge are usually carried out by members through roles, defined routines, or shared communication channels (Okhuysen and Bechky, 2009). Organizations either engage in structured coordination (through well-defined routines, tasks, roles, and colocation of interdependent agents) or unstructured coordination (through informal communication and mutual adjustment). This unstructured coordination to integrate organizational knowledge can be achieved through emergent processes characterized by mutual adjustment or “heedful interrelating” among members (Weick and Roberts, 1993). Alternatively, organizations can also deploy generalists who can facilitate the integration of different knowledge bases (Nagle and Teodoridis, 2017) or project managers who are instrumental to this function by enabling mutual adjustment and coordination across departments (Stan and Puranam, 2016).

With the advances in artificial intelligence (AI) and machine learning, however, managers are increasingly using the assistance of automated bots, algorithms, and AI-based agents to carry out these coordinative and integrative functions (Agrawal, Gans, and Goldfarb, 2018). These interfaces provide structure and increase predictability and consistency in executing these integrative functions (Puranam, Raveendran, & Knudsen, 2012; Claggett and Karahanna, 2019). For example, legal firms are increasingly using AI-based automation to review documents, conduct legal research, or even strategize, activities that collaborating teams of lawyers previously executed manually (Remus and Levy, 2017). Similarly, according to a recent article in *Nature*, journals have begun adopting AI-based peer review processes, which assists reviewers and editors in assessing contributions made with respect to an existing body of literature (Heaven, 2018). Such tasks are unstructured, undefined, and require teams to draw upon

knowledge spread across the organization. Indeed, the use of automation might improve predictability for jointly conducted knowledge integration.

The present study looks closely at the tradeoffs associated with choosing highly structured coordination interfaces such as bots or algorithms or unstructured coordination processes and routines involving humans when it comes to integrating knowledge. The use of automating bots reduce uncertainty by making the execution of tasks more predictable. As such, bots should improve overall coordination among collaborating knowledge workers. At the same time, however, it could stymie workers' ability to engage in unstructured coordination through mutual adjustment. This involves continuous exchange of information "during the process of action" (Thompson, 1967:56), which is needed to integrate new knowledge that is broad or systemic. Consequently, the decision to use bots or humans for such tasks presents major implications related to collaborative knowledge production for organizations (Von Krogh, 2018). This leads to a pair of research questions: given these opposing tendencies, how will automating knowledge integration impact the knowledge produced and integrated by the organization's members? How will it affect routines associated with unstructured coordination?

In this chapter, I develop a theory that explains why automation might make organizations integrate narrower, component-level knowledge rather than broader, systemic knowledge. I argue that automation enables structured coordination, but diminishes the ability to carry out unstructured coordination, which is necessary to integrate broader, systemic knowledge. Narrower contributions become more likely to be integrated because: (1) automation crowds out opportunities for an organization's members to engage in jointly evaluating and integrating new knowledge, and (2) members might make narrower knowledge contributions, because they are more likely to be accepted by the bot's highly structured integration rules. Furthermore, this

effect is moderated by *when* the focal organization adopts automation to integrate new knowledge. I hypothesize that this effect is reduced for late adopters, because organizations develop routines for unstructured coordination over time.

I test the outlined theory using the empirical setting of GitHub, a social networking platform that programmers use to collaborate on software projects. Some of these projects have automated the review, testing, and integration of new code contributions using customized scripts or bots, a process that project owners and maintainers previously performed. This review process is similar to the journal peer review process in academia. More specifically, I look at the adoption and use of Travis-CI, a bot software projects used to automate reviewing and integrating code contributions, as well as the impact of this automation on the scope of the contributions. To address concerns regarding the endogeneity of the decision to adopt automation, I exploited a natural experiment on a separate sample of JavaScript GitHub projects, in which a “troll” bot randomly chose projects in alphabetical order from a public registry and sent them requests to activate automation software on them. I found that projects that adopt automation tend to integrate narrower and component-based code contributions, as opposed to broader, systemic. Furthermore, I found that this effect is diminished for projects that are older at the time of adoption.

I aim to make theoretical and empirical contributions, as well as offer managerial implications. First, this study highlights a tradeoff associated with using highly structured coordination interfaces such as bots or algorithms for complex tasks involving the integration of new knowledge into the existing knowledge base. Although such interfaces might enhance coordination efficiency by improving the predictability of complex, integrative tasks, it could instead reduce employees’ ability to engage in unstructured coordination, which could aid in

integrating broader, systemic knowledge. Second, few extant studies provide empirical evidence of the impact of automating cognitive tasks (e.g., Kleinberg *et al.*, 2017). The prime reason for this deficiency is that fine-grained secondary data for econometric analysis at the group or team level are either unavailable or difficult to access (Raj and Seamans, 2018; 2019). Thus, this study aims to provide some of the first empirical evidence regarding the effects of cognitive task automation within organizations.

Finally, this study presents implications related to an organizational design choice that managers increasingly face: which tasks can humans carry out alone and which ones should humans be assisted by bots? Alternatively, where should the boundaries of automation be drawn when it comes to organizing knowledge-based tasks within firms? What are the tradeoffs associated with these choices? By viewing automation as an interface that shapes coordination among collaborating organizational members, this chapter provides a theoretical framework that can help guide future research that examines the impact of automation on organizational design, coordination, and knowledge production. This theoretical contribution helps advance the recent and burgeoning body of scholarly work about the role of AI within organizations and the strategic performance implications it holds for them (Von Krogh, 2018; Beane, 2018; Shestakosfsky, 2017; Oliver, Calvard, and Potocnik, 2017; Agrawal, Gans, and Goldfarb, 2018).

## **Theory**

A vast body of prior literature has viewed organizations and firms as repositories of knowledge and routines (Grant, 1996; Nelson and Winter, 1982; Kogut and Zander, 1996), and serving to integrate knowledge spread among its members (Kogut and Zander, 1992). Organizational knowledge is not just the sum of the knowledge contained within individuals; it is also knowledge contained in the interactions within its constituent members and its routines (Nelson

and Winter, 1982; Weick and Roberts, 1993; Nahapiet and Ghoshal, 1998; Srikanth and Puranam, 2011). The ability to draw on existing knowledge to produce new knowledge is realized through close interactions among its members (Katila and Ahuja, 2002; Paruchuri and Awate, 2017; Fonti and Maoret, 2015; Ben-Menahem *et al.*, 2016). This capability to draw on existing knowledge sets is especially important for creative, knowledge-based organizations that engage in collaborative, complex problem solving (Nickerson and Zenger, 2004).

Over time, an organization's members develop knowledge-sharing routines that facilitate knowledge transfer among themselves (Cohen and Bacdayan, 1994; Szulanski, 1996; Kogut and Zander, 1996). With sustained, repeated, and joint practice of routines related to knowledge production, organizations develop capabilities (Teece, Pisano, and Shuen, 1997; Eisenhardt and Martin, 2000; Helfat and Martin, 2015), as well as social and intellectual capital (Weick and Roberts, 1993; Nahapiet and Ghoshal, 1998), all of which contribute to an organization's competitive advantage. Thus, organizations maintain knowledge bases, develop routines to draw upon that knowledge, and coordinate efforts among its members to extend it and produce new knowledge.

### **Integrating Knowledge within Organizations**

A well-established body of prior literature views innovation and new knowledge production as an outcome of recombining existing knowledge (Schumpeter, 1942; Hargadon and Sutton, 1997; Fleming, 2001; Rosenkopf and Nerkar, 2001). A firm's innovativeness is determined in part by its ability to draw on its existing knowledge base and integrate new knowledge (Henderson and Cockburn, 1994; Galunic and Rodan, 1998; Carlile, 2004; Gardner, Gino, and Staats, 2012; Carnabuci and Operti, 2013). Thus, a firm's existing knowledge base, combined with its ability

to draw on it extensively, is instrumental in determining its absorptive capacity (Cohen and Levinthal, 1990), and is strategically important, especially for knowledge-intensive firms.

In this paper, I study the phenomenon of integrating new knowledge contributions into an organization's existing knowledge base (Zhou and Li, 2012; Cockburn and Henderson, 1994; Okhuysen and Eisenhardt, 2002). The phenomenon of new knowledge integration is general in the sense that it is widely observed across knowledge-based organizations ranging from startups to academic collaborations (Nagle and Teodoridis, 2017). A key source of competitive advantage is the ability to integrate new knowledge flexibly across knowledge boundaries within the organization (Henderson and Cockburn, 1994). It can be viewed as composed of two interrelated organizational processes: (1) evaluating new knowledge and (2) assimilating the new knowledge into the existing knowledge base. (Zhou and Li, 2012; Tortoriello, 2015).

Examples of integrating new knowledge abound, and the underlying processes of evaluation and integration are found consistently across different types of organizations carrying out knowledge-based tasks. For example, teams of attorneys collaborate when writing legal briefs, and each team draws on individual expertise and experience to provide intellectual contributions. Collectively, they deliberate over which of the contributions may be integrated (with or without modifications) into a coherent whole. Another example is that of journal editors and reviewers who jointly perform the cognitively intensive task of reviewing and modifying submitted articles in collaboration with authors. The challenge here is not only to evaluate the merits of the article itself, but also determine how new contributions could fit into debates within the extant literature. The editors and the reviewers must bring together skill, experience, and judgement to determine whether a particular knowledge contribution is worthy in terms of quality or relevance to be integrated into the existing body of knowledge. A final example is that



of software development teams, who collaborate and integrate new code features into the existing code base after reviewing it and ensuring it does not adversely affect existing functionality. In all of these examples, collaborating members draw on existing, shared knowledge; collectively evaluate, review, and modify new knowledge contributions; and finally integrate them into a coherent, collective output.

### **The Role of Coordination in Integrating Knowledge**

A sizeable body of prior literature acknowledges that drawing on knowledge spread across the organizations is nontrivial (Szulanski, 1996) and is a source of competitive advantage for firms that are proficient at it (Henderson & Cockburn, 1994; Zhao and Anand, 2012; Kim and Anand, 2018). Difficulties arise from the fact that an organization's members are boundedly rational, are often interdependent, and must coordinate their efforts in order to integrate their work output (March and Simon, 1976). Furthermore, they must also maintain consistency with existing knowledge and predictability to ensure they avoid "glitches" (Hoopes and Postrel, 1999).

Organizations carry out knowledge integration by facilitating coordination among individual members with respect to work output in a variety of ways (Puranam, 2018; Okhuysen and Bechky, 2009). More specifically, to coordinate the integration of individual work outputs, they use a combination of roles and accountability, routines that improve the predictability of task execution, as well as common understanding to coordinate individual efforts (Okhuysen and Bechky, 2009). For example, co-authors on an academic paper might use an online folder with version tracking, to ensure consistency across collaborators or a clear division of labor by assigning separate tasks. Furthermore, by collaborating repeatedly over time, co-authors also develop shared expectations regarding the quality and quantity of contributions other co-authors

make. This enables them to plan and execute their own contributions with respect to others' contributions.

Organizations coordinate individual work efforts by reducing epistemic interdependence with respect to the organization's other members (Puranam, Raveendran, and Knudsen, 2012), thus increasing the predictability of the successful execution of tasks. According to Puranam, Raveendran, and Knudsen (2012), epistemic interdependence arises if an agent's task execution depends on predicting what the other agent will do. To coordinate with each other, the agents must have predictive knowledge about each other, which can be formed through communication, mutual adjustment, and joint decision-making. Organizations can reduce epistemic interdependence by employing: (1) formal, structured coordination and/or (2) informal, unstructured coordination (Puranam, 2018). Structured coordination is established within organizations through prespecified routinized processes and plans (Simon, 1957; March and Simon, 1958; Okhuysen and Bechky, 2009; Claggett and Karahanna, 2018). Examples of structured coordination include plans, standard operating procedures, and rules (Nelson and Winter, 1982; Feldman and Rafaeli, 2002). Alternatively, it can also be established using formal, mechanistic organization structures such as hierarchies, in which task allocation is carried out through managerial control (Puranam, Alexy & Reitzig, 2014). Structured coordination is common in organizations that seek to focus on activities characterized by predictability and exploitation rather than uncertainty and exploration (Nickerson and Zenger, 2004).

In contrast, unstructured coordination is driven by processes that are emergent, unprogrammed, and characterized by improvisation (Bechky and Okhuysen, 2011). Examples of unstructured coordination include mutual adjustment and feedback among team members, open communication, common language, and informal meetings. Unstructured coordination is

emergent and organic and is realized through extensive interactions and communication regarding shared tasks (Puranam and Jacobides, 2006). More generally, through repeated practice over time, organizations develop shared understandings of how their individual work outputs fit into a coherent whole, as well as who is likely to perform certain tasks well (Austin, 2003, Feldman and Rafaeli, 2002). Armed with this common, shared understanding, members can dynamically adjust their behavior with respect to each other (Claggett and Karahanna, 2018), a process termed “heedful interrelating,” which leads to forming a collective mind (Weick and Roberts, 1993). Alternatively, individual managers can facilitate unstructured coordination (Stan and Puranam, 2016), as can generalists who possess knowledge that transcends different organizational knowledge bases, thus enabling them to play an integrative role (Nagle and Teodoridis, 2017). These individuals can foster flexibility across members, which improves mutual adjustment and coordination across the organization (Okhuysen and Eisenhardt, 2002).

Choosing between structured and unstructured coordination is problematic for organizations. On the one hand, structured coordination routines makes executing knowledge-based tasks more predictable and aids integration of efforts. On the other hand, it reduces opportunities for unstructured coordination characterized by mutual adjustment, which enables flexibility. Unstructured coordination, however, is costly in terms of time and effort: it requires an organization’s members to engage in rich and open communication, and entails collaborating through trial and error and mutual adjustment.

### **Coordination with Humans and Automated Agents**

Coordination in the past has been carried out using centralized IT systems that reduce the cost of information transfer (Brynjolfsson and Hitt, 2000). Recent advances in artificial intelligence (AI), however, have enabled automating algorithms that can support humans, such

as middle managers, in coordinating roles (Autor, 2015), but can also perform a wide range of tasks that previously required managerial expertise and skill. For example, firms are automating a large portion of core organizational functions such as hiring, audit, resource allocation, and so on, which were previously carried out by teams of humans and required close coordination and exchange of knowledge (Ford, 2015; Brynjolfsson and McAfee, 2014).

Bots, algorithms, and AI-based agents are instances of interface objects that provide highly structured coordination for complex, knowledge-intensive tasks such as problem-solving and decision-making (Agrawal, McHale, Oettl, 2018; Agrawal, Gans, Goldfarb, 2018). Thus, automation is an example of a structured coordination interface. For example, physicians are increasingly using clinical decision support systems (CDSS) in conjunction with electronic medical records (EMR) for patient diagnosis. This partially automates decision-making and knowledge sharing, thus reducing the possibility of errors and increasing efficiency (Aron *et al.*, 2011). CDSS and EMR make coordination more structured by providing standardized, predictable decision-making and reducing uncertainty (Claggett and Karahanna, 2019). Similarly, journals are adopting AI-based peer review to analyze manuscripts, extracting main conceptual contributions, and extract connections between different disciplines (Heaven, 2018). Thus, bots, algorithms, and AI-based agents enforce predetermined decision rules and make coordination among members more structured (as opposed to being characterized by improvisation and mutual adjustment). Indeed, it makes executing those decision rules more predictable and consistent.

Because bots and algorithms execute predetermined rules, they crowd out close communication and unstructured coordination among organizational members, leading to diminished collective learning and performance of knowledge-intensive tasks. Beane (2018)

illustrated this phenomenon in the context of hospitals. The adoption of robotic surgery methods reduced coordination between senior surgeons and junior residents and thus stymied the transfer of knowledge between them. Similarly, Oliver, Calvard, and Potocnik (2017) illustrated this phenomenon with the case of increased automation in airplane cockpits. The added automation diminished pilots' ability to collectively coordinate and respond to a crisis effectively and led to a plane crash. Both examples demonstrate situations in which introducing automation led organizational members to coordinate their efforts in a structured manner and crowded out close coordination.

A similar effect might manifest itself when it comes to automating the integration of knowledge. I argue that adopting automation to integrate new knowledge might lead an organization's members to integrate conceptually narrower outputs. This narrowing, or what Heath and Staudenmayer (2000) refer to as "component-focus," could occur through two mechanisms: (1) by diminishing close coordination among members and (2) by nudging contributors to submit narrower contributions. Even without formal interventions (Okhuysen and Eisenhardt, 2002), when it comes to integrating knowledge, organizations often exhibit a natural tendency to take a component-level view rather than a systemic one (Heath and Staudenmayer, 2000), primarily because of bounded rationality (March and Simon, 1958). This narrower component focus can be mitigated through heedful interrelating (Weick and Roberts, 1993) and by developing commonly held social capital through repeated collaboration (Nahapiet and Ghoshal, 1998). These mitigating actions enable organizational members to collectively develop a system-level understanding of the existing knowledge base. Because of reduced need—and consequently fewer opportunities—to coordinate in an unstructured way with other members, automation could exacerbate the tendency of knowledge producers to focus on components

rather than the system level. Automation reduces the cognitive load of coordination for individuals by increasing predictability and reducing epistemic interdependence. The reduction in opportunities for interpersonal interaction and improvisation, however, might make organizations less likely to integrate systemic knowledge. To summarize, the use of automating bots would crowd out unstructured coordination among organizational members (engaged in integrating new knowledge), exacerbate a component-focus, and thus lead to narrower knowledge contributions getting integrated.

Furthermore, algorithms cannot engage in mutual adjustment with human contributors through trial and error, which is necessary when new knowledge needs to be integrated into an existing knowledge base. For example, a new knowledge contribution might prompt integrators to think deeply about the interconnections, purpose, and future possibilities associated with the existing knowledge base (Henderson and Cockburn, 1994; Zhou and Li, 2012) and, consequently, how the new contribution could be accommodated more effectively. Unlike generalists or managers (Nagle and Teodoridis, 2017; Stan and Puranam, 2016), bots are unable to take such a broad view of the existing knowledge base or engage in mutual adjustment; therefore, they are more likely to fail at integrating such systemic contributions. For example, a bot that aids the academic review process might rate a submission poorly that draws widely on several different literatures, because the review rules written for the bot did not anticipate submissions of this type. As a result, knowledge contributors are more likely to restrict themselves to making less systemic, narrow contributions. This stems from a component-level focus (Heath and Staudenmayer, 2000) that emerges in response to the bot's inability to process broader, systemic knowledge contributions, as well as reduced unstructured coordination among members.

Thus, it follows that if diminished coordination (due to automation) leads to a lesser likelihood of broadly drawing on the existing knowledge base, then the new knowledge integrated likewise would be more likely to be narrower and concentrated and, therefore, less likely to be systemic (Baldwin and Clark, 2000). Therefore,

*Hypothesis 1: Organizations that automate knowledge integration tasks will integrate narrower knowledge than organizations that do not automate such tasks.*

Furthermore, I argue that the outlined effect will be higher for younger adopting organizations because the shared social and intellectual capital, which are required for coordination competence, will fail to develop such competence the first place. In contrast, late adopters have well-developed social and intellectual capital in place, which aids them in collaborating over integration (Weick and Roberts, 1993; Nahapiet and Ghoshal, 1998). Organizations develop persistent routines through repeated performance (Nelson and Winter, 1982). Furthermore, discarding existing routines and capabilities is difficult; organizations tend to depend on them (Leonard-Barton, 1992). This effect is seen even when new technology is introduced, leading to hindrances in successful adoption (Zuboff, 1988; Edmondson, Bohmer, and Pisano, 2001). In general, older organizations are more rigid and are loath to change (De Figueireido, Rawley, and Rider, 2015; Bakker and Josefy, 2018). Although the effects of automating knowledge integration could be present, they would be less pronounced for older adopters, because such organizations would depend less on automation; indeed, they would already have prior, well-established routines for knowledge integration that persist and are still used, even with an automating algorithm in place. Therefore,

*Hypothesis 2: Organizational age will negatively moderate the narrowing effect of automation on the scope of new knowledge integrated.*

## **Empirical Setting**

I test the outlined theory and hypothesis using data from software projects hosted on GitHub, an online social repository used by software programmers to collaborate on writing software. Furthermore, several of Github's features make it an ideal setting to test the theoretical arguments presented here. Firstly, software projects on Github can be considered to be organizations, in which members collaborate over evaluating, reviewing and integrating contributions toward a common knowledge artifact, i.e. software. Secondly, this setting also provides a discrete unit of a knowledge contribution made to the software project, namely the 'commit', and a wide range of detailed micro-level data that can be used to measure the size, scope and complexity for each commit. A commit and its associated information is recorded every time the contribution is 'pushed' into the code base. Thirdly, the setting offers a clearly identifiable instance of the adoption of a tool, Travis-CI, that automates cognitive tasks related to evaluating, testing and integrating contributions. Thus, I can construct extremely fine-grained panel data using which I can track changes induced by the adoption of automating tools in the code contributed to projects. In this way, the empirical setting offers a 'fruit-fly' instance of automation and knowledge contributions made.

## **GitHub Projects and Contributions**

As of 2017, there are about 25 million software projects, both public and private, hosted on GitHub<sup>2</sup>. The primary functionality provided by GitHub is code version control, and a social

---

<sup>2</sup> These statistics were retrieved from "The State of the Octoverse 2017", a report that gave an overview of the activity on GitHub in 2017. (<https://octoverse.github.com/>). Retrieved on May 31, 2018.



media platform that allows contributors to maintain an online presence which they can use to follow project updates for popular projects, or to keep track of contributions made by other individuals. It also provides a shared repository to store code and associated documentation for project management, allowing individuals separated in time and space to contribute to the same project, and owners/ maintainers to review and accept contributions asynchronously. While most of these projects consist of code, a significant chunk of these also consist of non-code projects to organize information such as wikis and open encyclopedias. A GitHub project can be viewed as a type of organization: it consists of a group of individuals collaborating to integrate new knowledge contributions (in the form of code contributions) to an existing knowledge base (in the form of the software being created). A large body of prior literature has looked specifically at organizational processes within such self-organizing projects (e.g. Puranam, 2018; Belenzon and Schankerman, 2015; Dahlander and O'Mahony, 2011; Foss, Frederiksen & Rullani, 2015, etc.).

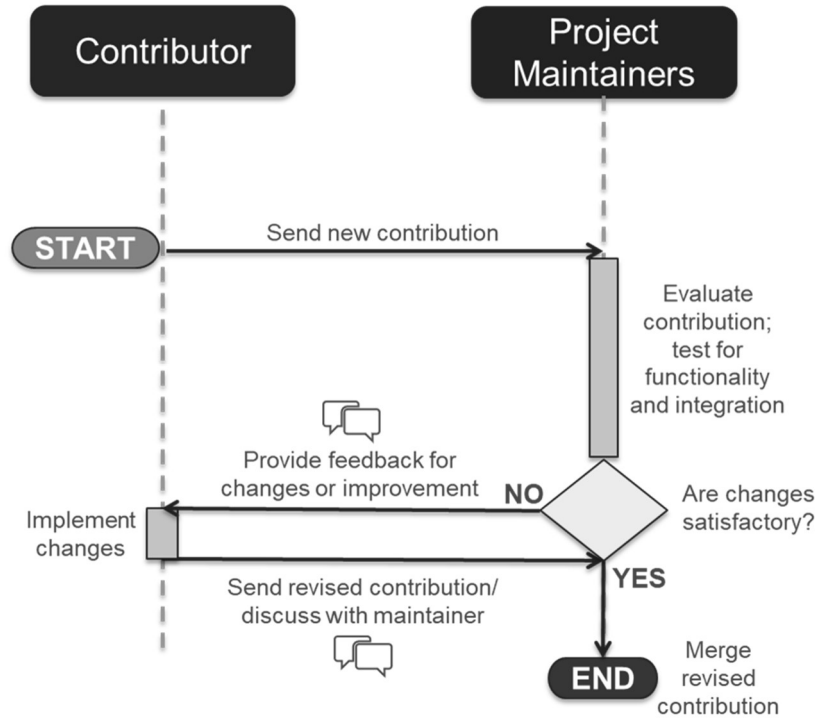


Illustration II.1: Github Contribution Workflow

Illustration II.1 outlines the typical GitHub workflow, simplified for understanding. When a user is interested in making a new contribution, or implementing improvements to an existing feature, she first downloads a copy of that software project on her local machine. She proceeds to implement a code contribution, and tests it on her local copy. Once she is sure that it works properly as intended, she sends a request to the project’s owner/ maintainer to integrate her contribution into the project. The maintainer creates a local copy of the project, integrates the incoming contribution, and tests it to ensure that it does not break the existing functionality, and additionally, functions as it is claimed to. If there are discrepancies, the maintainer discusses it with other project maintainers and the contributor and provides feedback on changes as required. After the original contributor implements some modifications to the original contribution as per

feedback provided, and maintainers confirm that the modified code now meets the requirements shared by everyone involved, it is finally integrated into the final, ready-to-download ‘master’ version. In this way, the workflow is not dissimilar to the peer review process for research articles submitted to a journal. This review process is a cognitively intensive one, and it requires the owner/ maintainer to maintain a detailed, expert’s view of the project in mind, one which can only be developed by actively working on integrating contributions and contributing code to the project in a sustained manner.

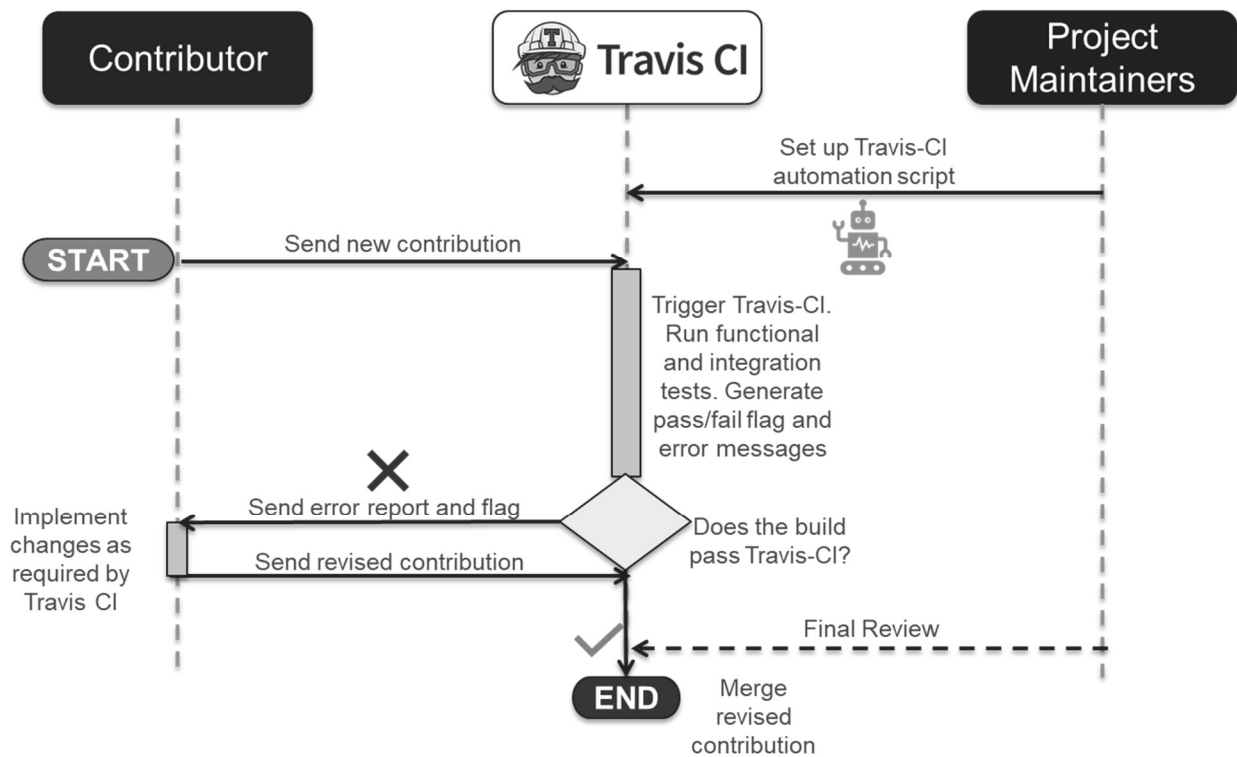


Illustration II.2 Github Contribution Workflow with Travis-CI

## **Travis-CI and Automation**

Several software projects on GitHub have adopted Travis-CI, a software tool which automates the testing of functionality and integration of the incoming contribution. Illustration II.2 shows a simplified version how the contribution workflow is modified by the adoption of Travis-CI.

Maintainers and owners usually review, test, and integrate incoming contributions manually; if an error shows up then the maintainer provides suggestions to the contributor on how to correct the error or improve the contribution. There is an extensive feedback process wherein maintainers and contributors engage in rich conversations and exchange thoughts on how the new contribution should be integrated. This is very akin to a journal's review process, in which authors, reviewers and editors engage in a conversation to integrate the intellectual contribution into extant debates.

In the case of Travis-CI, this process is automated to a great extent. As soon as the contribution is made, the testing and integration is done automatically; if the process passes without errors, a green flag is shown, and if it fails then a red flag (with a list of errors) is shown to the contributor and maintainer. The maintainer then directs the contributor to make corrections as indicated by Travis-CI and resubmit the contribution. Though the maintainer might review the contribution before integrating it, human involvement in the review is relatively peripheral as compared to the earlier review process.

### Bugs related to loading local resource from the bundle #453

Merged blakewatters merged 2 commits into restkit:master from unknown repository on Dec 1, 2011

Conversation 5 Commits 2 Checks 0 Files changed 2

f1337 commented on Dec 1, 2011 Contributor + @ ...

Hello! I made two minor changes to support loading JSON files packaged in the app's resource bundle. One bug was in RKURLm where escaped spaces (%20) in URLs were stripped by [NSURL path] and not restored, causing request failures. Obviously this applies to more than just bundle resources, but since bundle resources use "Application%20support" in their paths, I stumbled across it last night. I added a stringByAddingPercentEscapesUsingEncoding at the end of the completePath building statement, and it restores the escapes.

The other bug was in RKResponse, where local requests for bundle resources returned a NSURLResponse instead of a NSHTTPURLResponse, and therefore doesn't respond to statusCode or allHeaderFields. I added a respondsToSelector: check on the statusCode and allHeaderFields getters.

f1337 added some commits on Dec 1, 2011

- Fixed bug where escaped spaces (%20) in URLs were stripped by [NSURL path] 6391e59
- Fixed bug where local requests for bundle resources returned a NSURLResponse 0a8c3ac

blakewatters added a commit that referenced this pull request on Dec 1, 2011

Merge pull request #453 from fantasticnf/master 0e5db75

blakewatters merged commit 0e5db75 into RestKit:master on Dec 1, 2011

blakewatters commented on Dec 1, 2011 Member + @ ...

Thanks for the fixes!

f1337 commented on Dec 2, 2011 Contributor + @ ...

Thank you for RestKit! It's good stuff! :)

blakewatters commented on Dec 13, 2011 Member + @ ...

I have just reverted the the RKURLm change as it was introducing some additional issues. I should not have merged this without more solid unit tests -- RKURL is a soft underbelly. Will look at getting some tests and a fix in place to cover the %20 replacement specifically.

blakewatters commented on Dec 13, 2011 Member + @ ...

Why was your URL going through RKURL? Did you not get back a full NSURL from NSBundle, i.e. - (NSURL \*)URLForResource:(NSString \*)name withExtension:(NSString \*)ext NS\_AVAILABLE(10\_6, 4\_0);

How were you firing the request?

f1337 commented on Dec 14, 2011 Contributor + @ ...

It was a local URL for a file in the Resource path. I'll look up the approach I used and get back to you as soon as possible. Thank you! :)

(A) Before automation

### Avoid Flushing In-Memory Managed Object Cache while Accessing #2370

Merged Adlai-Holler merged 4 commits into development from dontflushmemorycache on Dec 12, 2015

Conversation 5 Commits 4 Checks 0 Files changed 6

Adlai-Holler commented on Dec 12, 2015 Contributor + @ ...

- Elevate UIApplicationDidReceiveMemoryWarningNotification observation up to RKEntityCache from RKEntityByAttributeCache.
- Add beginAccessing and endAccessing methods to RKEntityCache which will defer any calls to flush until all accesses are completed.
- Snuck in an improvement of using dispatch\_group\_notify to observe when dispatch groups end rather than holding up a thread with dispatch\_group\_wait.

This works around #1613.

Adlai-Holler added some commits on Dec 12, 2015

- Prevent flushing an RKEntityByAttributeCache while accessing it, to w... a104914
- Use dispatch\_group\_notify to avoid consuming a thread in RKEntityCach... 1cd0f771

segiddins commented on Dec 12, 2015 Member + @ ...

Can you look into the test failures?

Adlai-Holler added some commits on Dec 12, 2015

- Import UIKit if needed in RKEntityCache.m 428ad25
- Move cache flush test from RKEntityByAttributeCache to RKEntityCache 66e3cae

Adlai-Holler commented on Dec 12, 2015 Contributor + @ ...

@segiddins Sorry I was tired so I ran the OSX test locally and pushed. Tests are fixed.

(B) After automation

Illustration II.3: Example adapted from Hilton *et al*, (2016). These examples in panels A and B can be accessed at <https://github.com/RestKit/RestKit/pull/453> and <https://github.com/RestKit/RestKit/pull/2370> respectively.

Illustration II.3 illustrates this phenomenon through an example. Panel A shows the conversation between the contributor and maintainers over a contribution made by the former; it is characterized by the presence of rich feedback. Panel B contains an instance of the

communication between the two after the adoption of Travis CI for the same project. Clearly, the discussion is richer in the contribution made before automation, with greater incidence of joint problem-solving and feedback, as compared with that made after automation.

## **Data**

I define the treatment set as consisting of all Java-based GitHub projects which adopted and used Travis-CI at some point in their existence. Restricting my sample to Java projects allows me to automatically control for any possible language specific effects; furthermore, Java is a mature and stable programming platform for a wide variety of applications and has been in use for a reasonably long period of time. Additionally, Java is also the most represented language platform on Travis-CI. I identify projects using Java as their programming language by using data from GHTorrent, an online archive that records all events for nearly all projects hosted on GitHub (Gousios, 2013). After filtering out all forked projects (which are merely copies of the original), actively developed projects, and excluding projects which are marked as deleted and no longer available on Github, I arrive upon a list of 1.78 million projects. From these projects, I identify the ones which have adopted Travis-CI by checking whether they have a Travis-CI configuration file. By tracking when the file got created, I can arrive upon the date when the focal project adopted automation. I define this set of projects as the treatment set. Conversely, I define the control set as consisting of projects that never adopted Travis-CI at any point in their respective lifespans. Note that each treatment project adopts Travis-CI at different points of time. For all the projects, I download data that describes a number of project-level parameters such as age and date when the project was last updated.

To ensure balance between treatment and control sets, I match them on certain covariates in two stages. In stage 1, I retain treatment-control pairs of projects both of which were created

within six months of each other, and the last contribution on the control project is after the treatment project has adopted Travis-CI. This ensures that both treatment and control projects have similar ages at the time of adopting automation by the former, and also that activity on either of the two is not driven by idiosyncrasies due to certain features introduced on to the GitHub platform at a certain time. In stage 2, I match treatment and control projects on the number of commits before and after Travis-CI is implemented on the treatment project. More specifically, if there is a pair of projects TX and CTL (corresponding to treatment and control projects), I compare the number of commits on TX and CTL *before* TX adopts Travis-CI and carry out the same comparison for post-Travis-CI commits for those projects. I retain those projects for whom the number of commits is within  $\pm 20\%$  of each other. At this stage, I obtained 352,641 treatment-control matched pairs, consisting of 7594 treatment projects, and 89241 control projects. From this set of 89241 projects, I carried out stratified random sampling to arrive upon a set of 6686 control projects and 7594 treatment projects (7594 pairs in all) by matching on the number of commits for treatment projects at the time of their adoption of Travis-CI.

I carried out further matching of treatment and control projects on the dependent variable for the periods before the treatment project adopted Travis-CI. After obtaining the initial set of 7594 matched treatment-control pairs, I downloaded all commits for both treatment and control projects. For each commit, I retrieved the date of commit, number of files updated in the commit, the number of additions and deletions per file in terms of lines of code, and details pertaining to the creator of the commit. I computed an Herfindahl-Hirsch index (HHI) - based measure (explained in greater detail in the following section) for each commit in each project. Then I aggregated these commits at the project-month level (by taking month-level means of the HHI

measure). I carried out a further round of matching on the dependent variable in which I retained all treatment-control pairs whose average HHI for 10 periods before adoption by treatment project was within  $\pm 20\%$  of each other. I dropped projects which showed negative ages at the time of adopting Travis-CI, owing to discrepancies in the project creation date retrieved from GitHub. This left 881 treatment-control pairs consisting of 847 treatment projects and 918 control projects, which together formed 46,590 project-month level observations used to carry out analyses.

### **Main Dependent Variable- Herfindahl-Hirsch Index**

In order to construct a measure that reflects how systemic or component-oriented a given code contribution is, I compute a Herfindahl-index based measure that measures how distributed are the changes made over files for the contribution made in each commit. The intuition of the measure is that if commit consists of code changes distributed across several files, then it is more systemic in nature and requires broader knowledge on the part of the contributors and maintainers of the existing code base. For instance, if the contribution is small in terms of lines of code, and restricted to one file, then it is less systemic, and requires a relatively marginal knowledge of the underlying code base (Tsay, Dabbish, and Herbsleb, 2014a;2014b).

The measure, denoted by  $H$  is represented as:

$$H = \sum_{i=1}^N s_i^2$$

Here,  $s$  is the proportion of changes (with respect to the total changes in that commit) made to the  $i$ th file, and  $N$  is the number of files updated in that commit. Since the unit of analysis is commit- week, I use averages of the measure per project per four-week period.



## Independent and Control Variables

The main independent variable is a dummy variable that takes the value of 0 before Travis-CI adoption, and 1 after. In addition, I use project-level and time fixed effects to account for specific idiosyncratic effects. Project level variables such as age and activity in terms of number of commits are already accounted for earlier during matching of treatment and control projects, and hence not included in the econometric specification. However, project age at adoption is used as a moderating variable.

### Treatment Projects (N=847)

	Sum	Mean	Std. Dev.	Min.	25%	50%	75%	Max.
No. of commits pre-Treatment	294877	334.7	1215.5	3	37	101	260	28775
No. of commits post-Treatment	195231	217	599	3	16	53	155	6563
Age at Travis-CI adoption (in days)		456.3	456.5	1	128.5	294	639	2306
Normalized HHI before adoption		0.584	0.210	0	0.46	0.58	0.71	1

(a) Treatment Projects

### Control Projects (N=918)

	Sum	Mean	Std. Dev.	Min.	25%	50%	75%	Max.
No. of commits pre-Treatment	326799	331.5	1173.7	3	39	103	271.7	29100
No. of commits post-Treatment	213282	215.2	589	3	18	58	162.7	7270
Age at Travis-CI adoption (in days)		478.1	473.7	1	129	311	686.5	2306
Normalized HHI before adoption		0.65	0.23	0	0.51	0.67	0.82	1

(b) Control Projects

Table II.1: Descriptive Statistics (post-matching)

## **Descriptive Statistics**

Tables II.1a and II.1b outline the summary statistics for treatment and control projects respectively and are recorded at the time of adoption. These treatment-control pairs are obtained after matching on age, and pre and post-treatment activity in terms of commits. Thus, we can observe balance on these two variables. We also match on values of the dependent variable for 10 periods prior to treatment. In addition, I also include statistics for project size in terms of number of files and number of distinct contributors. Taken together, it can be observed that the treatment and control projects are well matched at the time of adoption.

## **Results**

In this section, I provide a description of the econometric specification used and describe the econometric model employed. Subsequently, I lay out the results in three parts. Firstly, I present results for the main effect related to hypothesis 1 and describe the dynamics of the treatment effect. Next, I provide an extension of the main result in which I explain the moderating effect of project age at adoption on the main effect. Finally, I provide results from a natural experiment in which a spam bot on GitHub named travis4all randomly chose JavaScript projects and sent requests to activate Travis-CI on them. I also describe the dynamics of the treatment effect as done previously for the main results.

## **Econometric Specification**

I use a difference-in-differences specification to evaluate the treatment effect of adopting automation on GitHub projects. The estimating equation relates the effect of adoption of Travis-CI by project  $i$  on the main dependent variable  $H$  in time  $t$ .

$$E[H_{it}|X_i] = \beta_0 + \beta_1 \text{TREATED*POST} + \delta_t + \gamma_i$$

... (1)

Here,  $H$  is the Herfindahl Index measure (previously described) computed for project  $i$  in period  $t$ ,  $TREATED$  is a dummy variable that is set to 0 for observations corresponding to control projects and 1 for treatment projects,  $POST$  is a dummy variable which is set to 0 for observations before adoption of Travis-CI and 1 otherwise,  $\delta_t$  is a set of quarter-year indicator variables, and  $\gamma_i$  are project level fixed effects. The project fixed effects control for idiosyncratic project level factors that stay constant during the entire lifespan of the project, such as specific functionality of the project. Time fixed effects allow controlling for broader macro level factors that vary with time, such as evolution in the functionality offered by the Github platform or the programming language, or seasonal variation in contributions to projects.

Number of observations = 46590

	<b>Model 1</b>	<b>Model 2</b>	<b>Model 3</b>
TREATED=1	-0.069*** (0.01)	0 (.)	0 (.)
POST=1	0.013 (0.01)	0.022*** (0.01)	0.032*** (0.01)
TREATED=1*POST=1	0.050*** (0.01)	0.044*** (0.01)	0.044*** (0.01)
Period Fixed Effects Included	No	No	Yes
Project Fixed Effects Included	No	Yes	Yes
constant	0.654*** (0.01)	0.618*** (0)	0.700*** (0)
R-squared	0.019	0.273	0.275

\* p<0.05, \*\* p<0.01,\*\*\* p<0.001

Table II.2. Effect of automation on contribution scope (DV = HHI)

### Main effect

Table II.2 outlines the main results with HHI as the dependent variable. Recall that hypothesis 1 stated that the automation of integrating new knowledge will lead to the generation of more incremental ideas. The value of the dependent variable ranges from 0 to 1, with lower values indicating more spread out, complex contributions, and higher values denoting more incremental, concentrated contributions. If the adoption of Travis-CI leads to projects having a higher HHI, then it means that adopting automation is associated with the integration of more incremental

knowledge that draws only upon a smaller part of existing knowledge contained in the project. Results outlined in table II.2 indicate that this seems to be the case. Model 1 consists of a basic specification without time or project fixed effects; model 2 consists of project fixed effects, and model 3 consists of both project and time fixed effects. The coefficient of interest is the interaction between the TREATED and POST variables, and in all three specifications, the coefficient obtained is significant and positive, thus supporting hypothesis 1. The results indicate that on average, adopting Travis-CI leads to an increase of 7.5% in the HHI after treatment relative to the mean HHI for treated projects. Standard errors are clustered at the project level.




Figure II.1: Leads and Lags for Main Results

A core assumption of the differences-in-differences framework is that the treatment and control projects are similar to each other over a period of time prior to treatment (Meyer, 1995; Bertrand, Duflo and Mullainathan, 2004). To investigate further as to whether the treatment effect lasts over time or stays constant, I plot trends for the dependent variable in both treatment and control projects. The results for this analysis are presented in figure II.1. The coefficients for the dependent variable are computed relative to time 0, i.e. the period when Travis-CI is adopted. The length of each time period is 28 days or 4 weeks. The effect of adopting Travis-CI is evident in the discontinuous jump for the treatment projects right after the period of adoption, while no such effect is present for the control projects<sup>3</sup>. Furthermore, this effect persists and goes on increasing with time for treatment projects, which further bolsters the theory that automation is leading to an erosion of project specific knowledge that is structural and transcends various modules.

---

<sup>3</sup> The high value of the HHI coefficient in period 1 is because of the concentration of contributions to one part of the project right after Travis adoption, i.e. in the Travis configuration file. This is expected, as the adoption of Travis leads to changes in project-level contribution processes, and often lead to initial teething problems.

Number of observations = 44127

	<b>Model 1</b>	<b>Model 2</b>	<b>Model 3</b>
TREATED_POST=1	0.028** (-0.01)	0.097*** (-0.01)	0.096*** (-0.01)
TREATED_POST=1*AGE	-0.003* (0)	-0.005*** (0)	-0.005*** (0)
constant	0.632*** (-0.01)	0.621*** (0)	0.694*** (-0.00)
Period Fixed Effects Included	No	No	Yes
Project Fixed Effects Included	No	Yes	Yes
R-squared	0.001	0.271	0.273

\* p<0.05, \*\* p<0.01, \*\*\* p<0.001

Age in days at adoption rescaled to 1/100 of the original value.

Table II.3. Age at adoption as moderating variable

### **Moderating Effect of Age at Adoption**

To test the moderating effect of project age on how adoption affects HHI, I use a slightly different specification.

$$E[H_{it}|X_i] = \beta_0 + \beta_1 \text{TREATED\_POST*AGE} + \delta_t + \gamma_i \quad \dots (2)$$

Here, AGE stands for the age of the project in days at the time of adoption. Additionally, a new dummy variable, TREATED\_POST is generated, which is set to 1 for observations corresponding to treatment projects in periods after adoption, and 0 otherwise. The rest of the specification is the same as before. The results are presented in table II.3. The interaction effect

for TREATED\_POST and AGE<sup>4</sup> has a small negative value but is strongly significant. This means that the older a project is at the time of adoption, the lesser is the effect of automation on the structural complexity of contributions made in that project. Recall that hypothesis 2 stated that the impact of automation on older adopters of automation will be lesser than that on younger adopters. Thus, the results support this hypothesis.




Figure II.2: Leads and Lags on Natural Experiment (DV Match precision = 30%)

<sup>4</sup> The computed age of some projects was negative, presumably because of an error in the recording of date of creation on GitHub; therefore, I dropped observations for those projects. Since Travis-CI can only be deployed for projects that use the GitHub platform, age can never be negative. It was clear that this was because of an error in recording the creation date. This led to a reduction of 20 projects from the sample. This discrepancy did not affect earlier analyses as age is not used as a moderating variable there.



## The travis4all Natural Experiment

Although the matching research design employed in the main analysis controls for several important covariates, there could be a possibility that projects that require more incremental, less complex contributions would be more likely to automate the integration of new contributions. This could give rise to concerns about possible endogeneity stemming from self-selection. To address these concerns, I present results from a natural experiment arising in this setting. A set of JavaScript projects were randomly nudged<sup>5</sup> into adopting Travis-CI in 2012 by a ‘troll’ bot named travis4all which directly added the configuration file to those projects<sup>6</sup>. The projects were chosen by the bot in alphabetical order of their names. A total of 1230 projects received requests from this bot to activate Travis-CI for them. Following complaints about spamming, the bot was shut down by GitHub after this point. However, most of the projects accepted these requests and continued using Travis-CI<sup>7</sup>. After applying filtering criteria as before, and matching on covariates such as age, pre-treatment number of commits and HHI<sup>8</sup> as previously described I arrived upon 41 treatment and 454 control projects, with 463 treatment-control pairs in all. Finally, I obtained 13031 project-month level observations.

---

<sup>5</sup> In further work I plan to use a two-stage least squares design to add further nuance to my empirical findings.

<sup>6</sup> McMillan, Robert. 2012. Github says 'no thanks' to bots — even if they're nice. *Wired*. <https://www.wired.com/2012/12/github-bots/> (Retrieved on 31<sup>st</sup> January 2019)

<sup>7</sup> I confirm this using archival data from GHTorrent.

<sup>8</sup> The matching was done on 1) pre-Treatment commits, and 2) average HHI for 10 periods before treatment at 30% precision level.

Number of Observations = 13031

	<b>Model 1</b>	<b>Model 2</b>	<b>Model 3</b>
TREATED=1	-0.062** (0.02)	0 (.)	0 (.)
POST=1	0.020 (0.01)	0.027* (0.01)	0.035 (0.02)
TREATED=1*POST=1	0.085*** (0.02)	0.073*** (0.03)	0.079** (0.03)
Period Fixed Effects Included	No	No	Yes
Project Fixed Effects Included	No	Yes	Yes
constant	0.740*** (0.01)	0.705*** (0.01)	-0.067*** (0.01)
R-squared	0.030	0.271	0.299

\* p<0.05, \*\* p<0.01, \*\*\* p<0.001

Table II.4. Effect on contribution scope (DV = HHI) for Natural Experiment- Projects matched on number of commits before treatment, and average value of dependent variable HHI for 10 periods before treatment at precision=30%.

Table II.4 outlines the results. The main results are preserved: we can see that the coefficient on the interaction between TREATED and POST is strongly significant. Figure II.2 plots the dynamics of the treatment effect; although the effect appears less prominently as compared with the main results, the trends are similar, including the noticeable jump in the coefficient right after the adoption of Travis-CI. These results suggest a causal interpretation of the relation between the adoption of automation and the decrease in the systemic nature of the contributions made.

Taken together, the following can be inferred from the results presented here. Firstly, automating even a part of jointly carried out cognitive tasks such as the integration of new knowledge is associated less systemic knowledge contributions being integrated. Secondly, this effect of automation is moderated by the age of the adopting organization at the time of adoption: in this context, adopting projects that are older will experience a lesser impact of automation on the systemic focus of the new contributions being integrated. Finally, results from the natural experiment setting provide support for a causal explanation for the main results.

## **Discussion and Conclusion**

This chapter looks closely at the tradeoffs associated with choosing highly structured coordination interfaces such as bots or algorithms, or unstructured coordination processes and routines involving humans when it comes to integrating knowledge. In this chapter, I develop a theory that explains why automation might make organizations integrate narrower, component-level knowledge as opposed to broader, systemic knowledge. I argue that automation enables coordination by reducing unpredictability of tasks executed, but it diminishes the ability to carry out unstructured coordination, which is necessary to integrate broader, systemic knowledge. I argue that this happens primarily because: 1) automation crowds out opportunities for organizational members to engage in joint evaluation and integration of new knowledge, and 2) members might make narrower knowledge contributions since they are more likely to be accepted by the integrating bot. Furthermore, this effect is moderated by *when* the focal organization adopts automation for integrating new knowledge. I find that this effect is reduced for late adopters since organizations develop routines for unstructured coordination over time.

I analyzed code contributions made to software projects hosted on Github and used them as an example of knowledge contributions. I compared the contributions to software projects that

automated the integration and testing, with those projects in which this process was carried out by human maintainers. I found that automation adopting projects integrated narrower, component level contributions, as opposed to non-adopters, who integrated broader, systemic contributions. This happens because automation crowds out close, unstructured coordination, and nudging members towards adopting an increased component focus at the cost of systemic focus. Furthermore, I find that this effect is lesser for projects that are older at the time of adoption since routines for unstructured coordination are already well-established.

Using this context has several advantages. Firstly, the context provides a simple and parsimonious instance of automating a knowledge-intensive task which helps isolate mechanisms of interest. Secondly, it provides rich and fine-grained data at the level of individual contribution, which allows me to closely track the impact of automation within projects over time. This not only provides me with a large number of observations, but also allows me to match on a set of covariates, thus allowing for identification. Lastly, the context also provides a unique natural experiment, which addresses concerns of endogeneity associated with the decision to automate certain activities, and suggests a causal interpretation of the outlined results.

The results present a number of broader implications from the point of view of both theory and practice. Firstly, conventional wisdom and prior suggests that automation is efficiency-enhancing, and should improve firm performance (Aron, et, al, 2011); however, this study posits that when it comes to knowledge-intensive tasks (such as integrating new knowledge), it might lead to undesired performance outcomes. Secondly, the use of automated bots for integrating knowledge should lead to improved coordination among organizational members (and drawing upon existing knowledge more broadly) as it reduces epistemic interdependence by increasing the predictability of tasks carried out (Puranam, Raveendran &

Knudsen, 2012). However, we see that it *reduces* unstructured coordination among members through mutual adjustment and exacerbates component focus. Finally, the results offer a set of interesting implications for managers: should they relegate such knowledge-intensive tasks to humans, or should they automate it? Or alternatively, where should boundaries of automation lie when it comes to organizing knowledge-based tasks within firms?

From a theoretical standpoint, the findings outlined in this study can be linked to literature on the role of coordination within organizations with respect to knowledge intensive work. In particular, this study looks at the impact of reducing epistemic interdependence on the component or systemic focus of activities being carried out. While structured coordination can be realized through differentiation and integration of tasks (Lawrence and Lorsch, 1967), or through modular organization design (Baldwin and Clark, 2000), recent advances in automation and AI enable it through increased predictability of task execution. More generally, this chapter views bots and algorithms as instances of structured coordination that reduce epistemic interdependence and looks closely at the tradeoffs associated with choosing automation over unstructured coordination by humans.

Furthermore, this study also speaks to extant literature on knowledge integration. While prior literature on has looked at the role of organization structure (Puranam, 2018), and at the role played by ‘integrators’ such as knowledge generalists (Nagle and Teodoridis, 2017), and middle managers (Stan and Puranam, 2016), this chapter adds another dimension: the use of algorithms that enable structured coordination among knowledge workers.

Finally, this chapter also responds to calls for research on how automation and AI are shaping work within organizations, and consequently, how they organize and operate (Cowgill, Seamans and Ziv, 2017; Raj and Seamans, 2019). While current debates in this literature revolve

primarily around how automation and AI is changing the way these tasks are being carried out (Ford, 2015; Carr, 2015; Agrawal, Gans and Goldfarb, 2018; Choudhury, Starr and Agarwal, 2019), or how it will affect societal outcomes (Susskind and Susskind, 2015; Bostrom, 2014; Kleinberg, *et al*, 2017), or economic outcomes (Acemoglu and Restrepo, 2018), there is little discussion about how automation will shape organizational processes and performance outcomes related to knowledge-based work, and consequently, strategic advantage. This chapter situates bots, algorithms and more broadly, artificial intelligence in existing research on coordination and knowledge integration by viewing it as an interface that increases the predictability of tasks executed.

In addition to extending theory, I also make empirical contributions through this chapter. There are few extant studies that provide empirical evidence of the impact of automation of cognitive tasks (e.g. Kleinberg *et al*, 2017), the prime reasons for this deficiency being that fine-grained secondary data for econometric analysis at the group or team level are either unavailable or difficult to access (Raj and Seamans 2019). Thus, this study also aims to provide some of the first empirical evidence of the effects of automation on outcomes of knowledge-based work.

At the same time, this study is not without limitations, and I am currently in the process of addressing those as I develop it further. Firstly, there could be some limits to generalizability: though GitHub projects could be classified as organizations consisting of collaborating members, they are quite different from traditional organizations: there is limited face-to-face interaction among individuals. Nevertheless, the underlying process of integrating new knowledge into an existing knowledge base is general: it can be found widely in more conventional contexts, such as legal firms, hospitals, or even academic peer review. Furthermore, these projects are driven by individuals, or large firms, or by open consortia (such as Apache Foundation) and are widely

deployed and used in various applications. Secondly, while I present results related to features of the contributions that are integrated, I do not have a measure that directly measures performance. This shortcoming can be addressed by collecting data on project stars (akin to Facebook likes), and forks (created when someone makes a copy of the project to build further upon it<sup>9</sup>). Finally, to better flesh out mechanisms related to coordination, it is necessary to analyze forum discussions regarding individual contributions made to projects, which will allow me to get a more nuanced view of how automation shapes conversations and coordination among project members.

Nevertheless, this study suggests a number of interesting further avenues for research on the impact of automation, AI, and robots on knowledge work in organizations. Future work could also focus on how automation could impact organization structure and co-ordination within teams, and how the two interact to shape outcomes of knowledge-intensive work such as innovation and creativity. Secondly, the adoption of automation also presents important implications for the role of human capital within organizations. In this way, this study aims to stimulate further inquiry in this area.

---

<sup>9</sup> Wu, Wang and Evans (2019) use forks as a measure that reflects intellectual impact of a project, akin to a citation in academic publishing.

## **Chapter III: Automation and the Evaluation of Men and Women’s Work Product: Evidence from Software Contributions on Github**

*(Co-authored with Seth Carnahan)*

### **Introduction**

Evaluations are an integral part of life in organizations and other cooperative groups. Supervisors and managers regularly evaluate the quality of job applications, work product, and more. A large literature emphasizes that these evaluations are often more of an “art than a science” (Botelho and Abraham, 2017: 699) which require significant judgement on the part of decision makers who have incomplete information about quality of the person or output being evaluated. Because ascriptive characteristics like gender are often readily available and widely associated with status characteristics and performance expectations (Berger, 1977), evaluators often use gender as a mental shortcut, providing male contributors and their outputs with higher evaluations than female contributors and their outputs (Correll and Ridgeway, 2003). Evaluators are even more likely to provide higher ratings to men in male-dominated tasks and occupations (Eagly, 2013; Tak, Correll, and Soule, 2019), such as beer-making or software programming (Correll and Mackenzie, 2016).

Given the possibility that decision makers might exhibit a pro-male bias, researchers and managers have sought to build policies and mechanisms to reduce it. These include using multiple evaluators (Campion, Palmer, and Campion, 1997), using objective criteria to inform evaluations (Maas and Torres-Gonzalez, 2011; Levashina *et al.*, 2014), and using gender-blind evaluation processes (Goldin and Rouse, 2000), among other approaches.



While this work has helped researchers and managers understand how to reduce pro-male bias among decision makers, much less attention has been paid to whether women are more willing to supply their labor to organizations or groups which have these kinds of mechanisms in place<sup>10</sup>. Given that organizations and other groups often experience a performance premium from gender diversity, group leaders increasingly wish to convince more women to join them. It is important to understand whether—and under what conditions—policies that might reduce pro-male bias in evaluations accomplish this goal.

This chapter puts forth two hypotheses. The core hypothesis of this chapter is that women will be more likely to contribute to groups that use automated, objective screening criteria to aid evaluations. Our argument is that female contributors will expect less pro-male bias from evaluations that incorporate automated, objective screening criteria. The expectation of fairer treatment during the evaluation process may encourage women to supply labor to the group. Secondly, we predict that automating some aspects of the evaluation process will increase the rate at which groups accept contributions from women.

We test these hypotheses using data from GitHub, an online social programming platform that hosts the largest open source software community in the world (Gousios, 2013). On GitHub, individuals called “project maintainers” can post software code projects in order to solicit contributions from a worldwide population of software programmers. Programmers submit changes to the project’s code base, and those changes must be approved by the project maintainer. This approval process relies on a mixture of subjective and objective criteria, where

---

<sup>10</sup> Some research in economics suggests that firms which pay their employees using piece rates tend to have more female workers (e.g. Goldin, 1986). Jirjahn and Stephan (2004) suggest that women might be attracted to piece rate payment schemes in part because such schemes are more objective and less likely to be influenced by gender biased evaluations, although Goldin (1986) describes other reasons why women might prefer piece rates (e.g. women might have shorter expected tenure and piece rates might enable more flexible work hours).

evaluators rely on their prior experience as a programmer and technical competence, as well as specific factors such as contributor attributes such as gender and prior social ties (Terrell, *et al*, 2017). In particular, we study the introduction of a robot platform called Travis-CI, which first arrived on Github in 2012. Travis changed the evaluation process by allowing evaluators to write and run scripts that automate the testing of functionality and integration of the incoming contribution. Previously the maintainer carried out those processes manually and provided feedback to contributors on what changes would be required. By contrast, Travis-CI executes them and provides objective indicators on whether the tests passed or failed.

We predict that software projects that automate evaluation tasks using Travis-CI will attract more contributions from female programmers. Furthermore, adopting project will also accept code contributions from female programmers at a higher rate after automation.

This chapter contributes to the literature by describing why organizations which screen individual contributions using automated, objective criteria might experience an increase in the gender diversity of their labor pool. This idea has implications well beyond our setting of open source software. For example, organizations might be able to attract more women job applicants if organizations make clear that performance tends to be evaluated using objective criteria. Competitions, such as startup pitch competitions, might be able to attract more women participants if they utilize and advertise objective screening criteria. Given that many organizations and other groups wish to increase gender diversity and unlock its performance enhancements (Deszo and Gaddis Ross, 2010), this contribution has important practical implications. Future versions of this study will also articulate the conditions under which using automated, scripted criteria for evaluation might increase the proportion of women in a labor pool.

## Theory

### Gender and Evaluators' Perceptions about Output Quality

A large body of literature suggests that evaluation processes are strategically important, since they are associated with choosing the best from available options (Csaszar and Eggers, 2013). While evaluations can apply to people and/or the outputs they produce (Azoulay, Stuart, and Wang, 2014), this chapter focuses on evaluations of outputs. That said, our theoretical perspective never loses sight of the reality that it is often difficult for evaluators to assess the quality of an output without reference to the identity of the output's producer (Tak, Correll and Soule, 2019). Whether it is executives choosing a strategy to pursue (Csaszar, 2013), consumers choosing a product to consume (Kovacs and Sharkey, 2014), or investors choosing an idea to fund (Lee and Huang, 2017), evaluations of productive outputs are commonplace and important. Evaluators face difficult challenges when making evaluations, because the quality of the options under consideration is rarely certain. As a consequence, evaluators sometimes rely on accessible, but potentially less relevant, indicators of expected quality (Merton, 1968), such as the race (Bertrand and Mullainathan, 2004) or gender (Brooks *et al.*, 2014) of the producer.

A variety of theoretical perspectives suggest that evaluators use ascriptive characteristics such as race, gender, age, or nationality to inform their evaluations of products, because this information is often readily available and may inform an evaluator's performance expectations. These theories differ in their precise underlying mechanisms, but each suggests that members of groups with higher status ascriptive characteristics, such as men, benefit when evaluators have limited information and rely more heavily on ascriptive characteristics.

Status characteristics theory (e.g. Ridgeway and Correll, 2004) suggests that widely held cultural beliefs imbue some ascriptive characteristics, such as male gender, with greater competence and social value. Evaluators rely on these beliefs about competence and value and

transfer them, not only to the individuals they are evaluating, but also to those individuals' outputs (Tak, Correll, and Soule, 2019). Gender role theory suggests that evaluators view some occupations or positions in society as requiring stereotypically male or female traits (Eagly, 2013). Stereotypically male traits include decisiveness, analytical rigor, and aggressiveness; roles that are often seen as requiring these traits are executive, computer programmer, or police officer. Stereotypically feminine traits include warmth, concern for others, and kindness; roles that are often seen as requiring these traits are nurse, teacher, or administrative assistant. Gender role theory suggests that evaluators will discount individuals (and their outputs) when their gender does not match the expectation for a given role (Eagly and Karau, 2002). The combination of status characteristics theory and gender role theory suggests that evaluators will be particularly skeptical of women who venture into stereotypically male roles, such as Science, Technology, Engineering, or Medicine (STEM) fields like computer programming (Correll and Mackenzie, 2016). Not only do women not fit the social expectations for the role, but their lower status means that evaluators are unlikely to give them or their products the benefit of the doubt, as they might for a man who enters a stereotypically feminine domain (Tak, Correll & Soule, 2019).

Fortunately, research suggests that providing evaluators with additional information can, in some circumstances, reduce evaluators' reliance on gender or other ascriptive characteristics. For example, Botelho and Abraham (2017), using data from an investing platform, find that bias against investment recommendations by women on the platform is lower when evaluators have lower search costs and when evaluations take place later in the evaluation process as the evaluator gains access to additional pertinent information about the recommendation. Using an experiment, Tak, *et al* (2019) find that awards matter much more for evaluators' perception of

the quality of beer brewed by women (with beer brewing being a stereotypically male role) as compared to beer brewed by men. Bohren, Imas and Rosenberg (2019) find that evaluations in online Q&A forums get biased *in favor of* female contributors if they have strong signals of quality, in the form of high reputation scores.

### Digital Algorithms as a Source of Information about Output Quality

A relatively new but important source of information about the potential quality of workers' outputs are algorithms which compare the characteristics of workers' outputs against a pre-defined set of criteria (Kellogg, Valentine and Christin, 2020). Algorithms that can evaluate worker output are seeing increased adoption across several fields with a wide variety of applications (Autor, 2015). For example, Amazon uses algorithms to track the output and productivity of its warehouse staff (Lecher, 2019), and these algorithms inform remuneration and termination decisions.

These algorithms might reduce gender bias in evaluations of output quality, particularly in settings where women occupy roles which are stereotypically masculine, such as computer programming. For instance, Cowgill (2017) finds that algorithms are better than humans at evaluating job candidates on hard-to-measure attributes such as leadership and cultural fit, where bias might be at play. Kleinberg (2017) show that when an algorithm was used for bail decisions, it outperformed human judges through lower crime rates or reduction in jailing rates, while reducing racial disparities. Algorithms provide an additional, structured source of information about output quality that might reduce the evaluator's reliance on the gender of the producer as a signal of quality<sup>11</sup>.

---

<sup>11</sup> An important consideration relates to whether a given algorithm has pro-male bias baked into itself. For instance, the algorithm might be written by a man, and in such a way that the algorithm evaluates the outputs of men more

Prior work suggests that gender bias in evaluation declines when evaluators have more information about output quality. For example, Botelho and Abraham (2017) find that evaluators of investment recommendations display less pro-male bias when evaluators have more information about the recommendation, and Tak, Correll and Soule (2019) find that evaluators of beer experience less bias against a female-brewed beer when it has won an award, so too might evaluators experience less pro-male bias against outputs that have been pre-evaluated by a digital algorithm.

Prior work also suggests that gender gaps in evaluation tend to be lower when workers are evaluated using criteria with more structure. For example, job interview scholars find that women and under-represented minorities tend to receive better employability ratings when job interviews are structured such that all candidates are asked the same questions (Campion, Palmer, and Campion, 1997; Levashina, *et al*, 2014). When all candidates are asked the same questions, there is less scope for interviewers' tastes and biases to influence the content of the interview. Unlike humans, algorithms cannot "relate" better to someone because they have shared cultural capital with the person they are evaluating (Rivera, 2012). Thus, when algorithms (which uniformly apply on the same set of impersonal criteria to all outputs) provide evaluators with information on output quality, there may be lesser scope for evaluators' tastes and biases to influence their ultimate evaluation. In this way, with algorithms evaluating, females are more likely to have their work evaluated positively; correspondingly, males will experience an increased incidence of negative evaluations.

---

favorably than the outputs of women, on account of the designer's latent assumptions baked into the algorithm's design (Anthony, 2018). For example, a male call center manager might write a call script that fits much more naturally with a male communication style, rather than a female communication style. Alternatively, the algorithm might also be based on data that contain a pro-male bias (Cowgill and Tucker, 2020). For example, an algorithm that uses natural language processing to assign a quality score to a piece of writing might be trained using pieces of writing that were produced or rated by men.

*Hypothesis 1a: Groups that automate evaluation tasks are more likely to accept (less likely to reject) contributions from females compared to groups that rely on human evaluation.*

*Hypothesis 1b: Groups that automate evaluation tasks are more likely to reject (less likely to accept) contributions from males compared to groups that rely on human evaluation.*

To the extent that the above hypotheses are accurate and groups which use algorithms to aid evaluation are fairer to women, we would expect women to respond positively and supply more of their labor to such groups.

*Hypothesis 2: Groups that use algorithms to aid evaluation tasks will attract more female contributors.*

## **Data and Empirical Design**

We test our hypotheses by using data from software projects hosted on GitHub, an online, social programming platform. In particular, we study the use of a robot, Travis-CI, for automating evaluation of code contributions. A detailed description of the code contribution process and evaluation workflow, as well as Travis-CI can be found in Chapter 2. The context allows us to set up a difference-in-differences design to track outcomes within both adopting and non-adopting projects over time, as well as compare them with each other. Projects that adopt Travis-CI are designated as belonging to the ‘treatment’ group, and those which do not, are part of the ‘control’ group.

Our sample consisted of a broader set of matched projects used in Chapter II<sup>12</sup> (for a detailed account of how the matched pairs were obtained, please refer to Chapter II). For each project we downloaded all of its ‘pull requests’ using the GitHub API. A pull request is initiated by a programmer, and is submitted to a project, when they wish to make a code contribution to that project. The pull request is then reviewed by the project’s maintainer, and depending on whether it is accepted or rejected, it is either merged into the project’s code base, or closed respectively. We also downloaded a range of variables associated with the pull request, such as the date when the pull request was created, date when it was closed, date when it was merged, name of the programmer who submitted the pull request, etc. We use the open-source API genderize.io to identify the gender for each programmer in our data by using their first names. Programmers whose names were missing from their profiles, or those whose names could not be identified as male or female were coded as ‘neutral.’

After filtering out projects with just one code contribution, our final dataset consisted of 18436 treatment-control project pairs, consisting of 2828 treated and 6728 control projects. Each pair had four observations, corresponding to values of the variables before and after adopting Travis, for both treated and control projects. To carry out further analysis that tracks the variables over time, we also construct another dataset wherein observations are recorded for each 4-week period. Doing this allows us to track how the adoption of Travis changes outcomes over time, and compare them with those for non-adopters.

### Dependent Variables

We use a variety of dependent variables to study the gender composition of programmers submitting code contributors, as well as the rate at which their contributions are accepted or

---

<sup>12</sup> Note that in chapter II we took a stratified random sample of a broader matched sample for convenience. Here we used the broader sample for greater statistical power.



rejected. The variable *percent\_female\_contributors* is calculated as the ratio of number of female contributors in a period to the number of all contributors in the same period. *Female\_pr\_share* represents the proportion of all pull requests made in a period that are made by female programmers. *Female\_pracceptance\_rate* and *female\_prrejection\_rate* represent the proportion of pull requests made by female programmers that are accepted or rejected respectively. The corresponding variable for programmers with male and neutral public identities are computed similarly.

	<b>count</b>	<b>mean</b>	<b>std</b>	<b>min</b>	<b>25%</b>	<b>50%</b>	<b>75%</b>	<b>max</b>
<b>Number of PRs</b>	10976	8.98	20.82	1	2	3	8	440
<b>Number of PRs by Women</b>	10976	0.35	2.37	0	0	0	0	62
<b>Number of PRs by Men</b>	10976	6.17	16.17	0	1	2	5	369
<b>Number of PRs by Gender Neutral Contributors</b>	10976	2.47	8.27	0	0	0	2	235
<b>Number of Unique Contributors</b>	10976	3.01	4.64	1	1	2	3	94
<b>Number of Accepted PRs</b>	10976	7.16	17.73	0	1	3	6	401
<b>Number of Rejected PRs</b>	10976	1.69	6.74	0	0	0	1	205
<b>Number of Accepted PRs (Men)</b>	10976	4.92	13.75	0	0	2	4	356
<b>Number of Rejected PRs (Men)</b>	10976	1.17	5.52	0	0	0	1	192
<b>Number of Accepted PRs (Women)</b>	10976	0.28	2.09	0	0	0	0	62
<b>Number of Rejected PRs (Women)</b>	10976	0.06	0.56	0	0	0	0	20
<b>Number of Accepted PRs (Gender Neutral)</b>	10976	1.96	7.25	0	0	0	1	223
<b>Number of Rejected PRs (Gender Neutral)</b>	10976	0.46	2.16	0	0	0	0	72

Table III.1 (a): Descriptive Statistics for Control Projects (prior to Travis adoption)

	<b>count</b>	<b>mean</b>	<b>std</b>	<b>min</b>	<b>25%</b>	<b>50%</b>	<b>75%</b>	<b>max</b>
<b>Number of PRs</b>	16796	12.55	29.25	1	2	5	12	1021
<b>Number of PRs by Women</b>	16796	0.54	2.95	0	0	0	0	83
<b>Number of PRs by Men</b>	16796	8.25	22.72	0	1	3	8	737
<b>Number of PRs by Gender Neutral Contributors</b>	16796	3.76	11.07	0	0	1	3	312
<b>Number of Unique Contributors</b>	16796	4.50	7.15	1	1	2	5	247
<b>Number of Accepted PRs</b>	16796	9.54	23.89	0	2	3	9	719
<b>Number of Rejected PRs</b>	16796	2.04	9.78	0	0	0	2	935
<b>Number of Accepted PRs (Men)</b>	16796	6.36	18.94	0	1	2	6	651
<b>Number of Rejected PRs (Men)</b>	16796	1.30	7.33	0	0	0	1	693
<b>Number of Accepted PRs (Women)</b>	16796	0.42	2.65	0	0	0	0	76
<b>Number of Rejected PRs (Women)</b>	16796	0.08	0.55	0	0	0	0	17
<b>Number of Accepted PRs (Gender Neutral)</b>	16796	2.76	9.38	0	0	0	2	306
<b>Number of Rejected PRs (Gender Neutral)</b>	16796	0.65	3.06	0	0	0	0	233

Table III.1 (b): Descriptive Statistics for Control Projects (after Travis adoption by matched Treatment Project)

	<b>count</b>	<b>mean</b>	<b>std</b>	<b>min</b>	<b>25%</b>	<b>50%</b>	<b>75%</b>	<b>max</b>
<b>Number of PRs</b>	13097	8.73	19.19	1	2	4	9	761
<b>Number of PRs by Women</b>	13097	0.24	1.43	0	0	0	0	32
<b>Number of PRs by Men</b>	13097	6.63	15.43	0	1	3	7	709
<b>Number of PRs by Gender Neutral Contributors</b>	13097	1.86	7.48	0	0	0	1	254
<b>Number of Unique Contributors</b>	13097	2.92	4.30	1	1	2	3	80
<b>Number of Accepted PRs</b>	13097	7.16	16.37	0	1	3	7	729
<b>Number of Rejected PRs</b>	13097	1.48	4.82	0	0	0	1	199
<b>Number of Accepted PRs (Men)</b>	13097	5.50	13.10	0	1	2	6	679
<b>Number of Rejected PRs (Men)</b>	13097	1.06	4.13	0	0	0	1	196
<b>Number of Accepted PRs (Women)</b>	13097	0.19	1.21	0	0	0	0	30
<b>Number of Rejected PRs (Women)</b>	13097	0.05	0.37	0	0	0	0	11
<b>Number of Accepted PRs (Gender Neutral)</b>	13097	1.47	6.52	0	0	0	1	242
<b>Number of Rejected PRs (Gender Neutral)</b>	13097	0.37	1.58	0	0	0	0	29

Table III.1 (c): Descriptive Statistics for Treatment Projects (prior to Travis adoption)

	count	mean	std	min	25%	50%	75%	max
<b>Number of PRs</b>	16132	14.69	35.69	1	2	5	13	856
<b>Number of PRs by Women</b>	16132	0.52	5.23	0	0	0	0	228
<b>Number of PRs by Men</b>	16132	10.79	26.75	0	1	4	10	703
<b>Number of PRs by Gender Neutral Contributors</b>	16132	3.38	14.46	0	0	1	3	702
<b>Number of Unique Contributors</b>	16132	5.16	9.20	1	1	2	5	227
<b>Number of Accepted PRs</b>	16132	10.76	27.99	0	2	4	10	680
<b>Number of Rejected PRs</b>	16132	2.80	10.33	0	0	0	2	476
<b>Number of Accepted PRs (Men)</b>	16132	8.25	21.74	0	1	3	8	644
<b>Number of Rejected PRs (Men)</b>	16132	1.81	6.51	0	0	0	1	424
<b>Number of Accepted PRs (Women)</b>	16132	0.34	3.09	0	0	0	0	215
<b>Number of Rejected PRs (Women)</b>	16132	0.15	3.59	0	0	0	0	214
<b>Number of Accepted PRs (Gender Neutral)</b>	16132	2.16	11.42	0	0	0	2	562
<b>Number of Rejected PRs (Gender Neutral)</b>	16132	0.84	4.15	0	0	0	0	129

Table III.1 (d): Descriptive Statistics for treatment Projects (after Travis adoption)

	<b>Before Treatment</b>		<b>After Treatment</b>	
	<b>Control</b>	<b>Treated</b>	<b>Control</b>	<b>Treated</b>
<b>PR Acceptance Rate for Women</b>	81%	79%	78%	66%
<b>PR Rejection Rate for Women</b>	17%	20%	15%	28%
<b>Women's Share of all PRs submitted</b>	4%	3%	4%	4%
<b>PR Acceptance Rate for Men</b>	80%	83%	77%	76%
<b>PR Rejection Rate for Men</b>	19%	16%	16%	17%
<b>Men's Share of all PRs submitted</b>	69%	76%	66%	73%

Table III.2: Descriptive statistics for computed dependent variable for all projects

### Independent and Control Variables

The main independent variable tracks whether a given observation falls in the pre- or post-treatment period within the matched pair, and is denoted by POST. TREATED represents whether the project belongs to the treated or control group. Since the treated and control projects are already matched for activity levels and age, we do not include those control variables. We also use project-level and time fixed effects (at the year and quarter level) to account for idiosyncratic effects.

Table 1a and 1b outline summary statistics for treatment and control projects. Taken together, it can be seen that treated and control projects have fairly similar distributional properties.

### **Results**

In this section, we provide a description of the econometric specification used to test the outlined theory. Subsequently we provide results related to how Travis affects the propensity to contribute for female, male and neutral programmers, and its effect on acceptance or rejection rates for those programmers.

### Econometric Specification

We use a difference-in-differences specification to evaluate the effect of adopting Travis on participation and acceptance/ rejection rates for projects.

$$E[Y_{it} | X_i] = \beta_0 + \beta_1 \text{TREATED} * \text{POST} + \delta_t + \gamma_i \quad \dots (3)$$

Here, Y represents various dependent variables as defined in the previous section, TREATED is a dummy variable that is set to 0 for observations corresponding to control projects and 1 for

treatment projects,  $POST$  is set to 0 for observations before adoption of Travis-CI and 1 otherwise,  $\delta_t$  is a set of quarter-year indicator variables, and  $\gamma_i$  are project level fixed effects. The project fixed effects control for idiosyncratic project level factors that stay constant during the entire lifespan of the project, such as specific functionality or unobserved properties of the project. Time fixed effects allow to adjust for time-varying factors, such as changes in the GitHub platform or the programming language over time, or seasonal variation in contribution rates to projects (for instance, projects might attract more contributions during certain periods of the year such as summer, when undergraduate students might have more free time on their hands).

#### Effect on Acceptance and Rejection Rates

Recall that hypothesis 1 states that projects are more likely (less likely) to accept (reject) contributions from female programmers after they adopt automation for evaluating contributions. Table III.3 outlines results from regressions that use *female\_pracceptance\_rate* as the dependent variable. As in previous results, model 1 is the most basic specification, and model 4 is the most stringent specification. While the coefficient on the interaction is in the hypothesized direction only for models 3 and 4, none of the results are statistically significant. Similarly, when *female\_prrejection\_rate* is used as the dependent variable (table III.4), the coefficients on the interaction are in the hypothesized direction.

**Female PR Acceptance Rate (Period Wise)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
TREATED=1	-0.029* (0.01)	-0.029 (0.04)	0.000 (.)	0.000 (.)
POST=1	-0.024* (0.01)	-0.024 (0.02)	-0.016 (0.01)	0.013 (0.01)
TREATED=1* POST=1	-0.030 (0.02)	-0.030 (0.05)	0.008 (0.07)	0.022 (0.06)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.795*** (0.01)	0.795*** (0.02)	0.767*** (0.02)	0.643*** (0.09)
R-squared	0.006	0.006	0.706	0.719

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.3: Acceptance rate for female contributors before and after Travis-CI (Period wise dataset)

**Female PR Rejection Rate (Period wise)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
TREATED=1	0.041** (0.01)	0.041 (0.04)	0.000 (.)	0.000 (.)
POST=1	-0.020 (0.01)	-0.020 (0.02)	-0.010 (0.01)	-0.019 (0.01)
TREATED=1* POST=1	-0.005 (0.02)	-0.005 (0.04)	-0.017 (0.07)	-0.015 (0.06)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.182*** (0.01)	0.182*** (0.02)	0.197*** (0.02)	0.243** (0.08)
R-squared	0.004	0.004	0.675	0.687

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.4: Rejection rate for female contributors before and after Travis-CI (Period-wise dataset)

Since performing analyses using collapsed pre-post treatment observations helps deal with serial correlation in difference-in-difference estimates (Bertrand, Duflo and Mullainathan, 2004), we conducted additional analyses (outlined in Table III.5 and III.6) for the same dependent variables as a robustness check, in which we collapse period-wise observations into single pre-Travis and post-Travis observations. Thus, for each project in a treated-control pair, we have one pre-Travis value and one post-Travis value for each of the treated and control projects. Thus



Here, we introduce time fixed effects for the year and quarter in which Travis-CI was adopted. We find that the coefficient of interaction between *POST* and *TREATED* are in the hypothesized direction for all models, and are even significant for some specifications. Our results are conceptually consistent with our hypotheses.

**Female PR Acceptance Rate (collapsed observations)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
POST=1	-0.077*** (0.02)	-0.077*** (0.02)	-0.031 (0.02)	-0.034 (0.02)
TREATED=1	-0.061*** (0.02)	-0.061 (0.04)	0.000 (.)	0.000 (.)
TREATED=1* POST=1	0.018 (0.02)	0.018 (0.06)	0.003 (0.12)	0.006 (0.12)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.784*** (0.01)	0.784*** (0.02)	0.728*** (0.04)	0.759*** (0.05)
R-squared	0.009	0.009	0.923	0.923

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.5: Acceptance rate for female contributors before and after Travis-CI (consolidated observations)

**Female PR Rejection Rate (collapsed observations)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
POST=1	-0.004 (0.01)	-0.004 (0.02)	-0.004 (0.02)	-0.008 (0.02)
TREATED=1	0.083*** (0.02)	0.083 (0.04)	0.000 (.)	0.000 (.)
TREATED=1* POST=1	-0.067*** (0.02)	-0.067 (0.05)	-0.071 (0.08)	-0.067 (0.08)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.182*** (0.01)	0.182*** (0.02)	0.224*** (0.03)	0.217*** (0.04)
R-squared	0.007	0.007	0.929	0.929

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.6: Rejection rate for female contributors before and after Travis-CI (consolidated observations)

Correspondingly, we carry out similar analyses for male contributions, for both period-wise and consolidated data, which are outlined in Tables III.7, III.8, III.9 and III.10. These results are consistent with our hypotheses and prior results, and even significant for the most stringent specification when it comes to the PR rejection rate.

**Male PR Acceptance Rate (period-wise observations)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
TREATED=1	0.021*** (0.00)	0.021 (0.01)	0.000 (.)	0.000 (.)
POST=1	-0.037*** (0.00)	-0.037*** (0.01)	-0.033*** (0.01)	0.015** (0.01)
TREATED=1* POST=1	-0.027*** (0.00)	-0.027* (0.01)	-0.029* (0.01)	-0.020 (0.01)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.783*** (0.00)	0.783*** (0.01)	0.793*** (0.00)	0.914*** (0.09)
R-squared	0.004	0.004	0.394	0.399

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.7: Acceptance rate for male contributors before and after Travis-CI (period-wise observations)

**Male PR Rejection Rate (period-wise observations)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Standard Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
TREATED=1	-0.013*** (0.00)	-0.013 (0.01)	0.000 (.)	0.000 (.)
POST=1	-0.024*** (0.00)	-0.024*** (0.01)	-0.023*** (0.00)	-0.007 (0.00)
TREATED=1* POST=0	0.022*** (0.00)	0.022* (0.01)	0.002 (0.01)	0.005 (0.01)
TREATED=1* POST=1	0.193*** (0.00)	0.193*** (0.01)	0.192*** (0.00)	0.157* (0.07)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	-0.013*** (0.00)	-0.013 (0.01)	0.000 (.)	0.000 (.)
R-squared	0.000	0.000	0.349	0.352

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.8: Rejection rate for male contributors before and after Travis-CI (period-wise observations)

In the case of collapsed observations, while the results remain consistent for acceptance rate, the results lose significance for the rejection rate for male contributors.

**Male PR Acceptance Rate (collapsed observations)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
POST=1	-0.058*** (0.00)	-0.058*** (0.01)	-0.046*** (0.01)	-0.048*** (0.01)
TREATED=1	0.030*** (0.00)	0.030** (0.01)	0.000 (.)	0.000 (.)
TREATED=1* POST=1	-0.027*** (0.01)	-0.027* (0.01)	-0.030 (0.02)	-0.027 (0.02)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.791*** (0.00)	0.791*** (0.01)	0.801*** (0.00)	0.897*** (0.05)
R-squared	0.014	0.014	0.786	0.787

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.9: Acceptance rate for male contributors before and after Travis-CI (collapsed observations)

**Male PR Rejection Rate (collapsed observations)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
POST=1	-0.026*** (0.00)	-0.026*** (0.01)	-0.031*** (0.01)	-0.033*** (0.01)
TREATED=1	-0.016*** (0.00)	-0.016 (0.01)	0.000 (.)	0.000 (.)
TREATED=1* POST=1	0.008 (0.00)	0.008 (0.01)	0.003 (0.01)	0.005 (0.01)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.183*** (0.00)	0.183*** (0.01)	0.178*** (0.00)	0.168 (0.09)
R-squared	0.002	0.002	0.764	0.764

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.10: Rejection rate for male contributors before and after Travis-CI (collapsed observations)

**Percentage of Female Contributors (period-wise observations)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
TREATED=1	-0.010*** (0.00)	-0.010* (0.00)	0.000 (.)	0.000 (.)
POST=1	0.006*** (0.00)	0.006* (0.00)	0.004 (0.00)	-0.001 (0.00)
TREATED=1* POST=1	-0.005*** (0.00)	-0.005 (0.00)	0.000 (0.00)	-0.001 (0.00)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.036*** (0.00)	0.036*** (0.00)	0.031*** (0.00)	0.004 (0.01)
R-squared	0.002	0.002	0.403	0.404

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.11: Percentage of male contributors before and after Travis-CI (period-wise observations)

**Percentage of Female Contributors (collapsed observations)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
POST=1	0.006*** (0.00)	0.006* (0.00)	0.003 (0.00)	0.003 (0.00)
TREATED=1	-0.004* (0.00)	-0.004 (0.01)	0.000 (.)	0.000 (.)
TREATED=1* POST=1	-0.004 (0.00)	-0.004 (0.01)	0.001 (0.01)	0.001 (0.01)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.036*** (0.00)	0.036*** (0.00)	0.034*** (0.00)	0.032*** (0.00)
R-squared	0.001	0.001	0.761	0.761

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.12: Percentage of female contributors before and after Travis-CI(collapsed observations)

Effect on Participation by Programmers

According to hypothesis 2, groups that adopt automation should attract a greater proportion of female contributors after adoption. Table III. 11 outlines the main results with *percent\_female\_contributors* as the dependent variable. Model 1 is the most basic econometric specification where we regress independent variables on the dependent variable. In model 2, we cluster standard errors at the project level, in model 3 we include project fixed effects in addition, and finally in model 4 we also include time fixed effects. While the coefficient of the interaction



of *POST* and *TREATED* is in the hypothesized direction for models 3 and 4, it is not significant. Similar patterns are observed when collapsed observations are used (table III.12)

**Percentage of Male Contributors (period-wise observations)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
TREATED=1	0.075*** (0.00)	0.075*** (0.01)	0.000 (.)	0.000 (.)
POST=1	-0.030*** (0.00)	-0.030*** (0.01)	-0.027*** (0.01)	0.003 (0.01)
TREATED=1* POST=1	0.011** (0.00)	0.011 (0.01)	-0.005 (0.01)	0.001 (0.01)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.693*** (0.00)	0.693*** (0.01)	0.736*** (0.00)	0.848*** (0.04)
R-squared	0.010	0.010	0.451	0.453

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.13: Percentage of male contributors before and after Travis-CI (period-wise observations)

<b>Percentage of Male Contributors (collapsed observations)</b>				
	<b>Model 1: Basic</b>	<b>Model 2: Clustered Standard Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project/ Time FE, Clustered errors</b>
POST=1	-0.033*** (0.00)	-0.033*** (0.01)	-0.037*** (0.01)	-0.039*** (0.01)
TREATED=1	0.088*** (0.00)	0.088*** (0.01)	0.000 (.)	0.000 (.)
TREATED=1* POST=1	-0.023*** (0.01)	-0.023 (0.01)	-0.009 (0.01)	-0.007 (0.01)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.675*** (0.00)	0.675*** (0.01)	0.718*** (0.00)	0.777*** (0.04)
R-squared	0.016	0.016	0.798	0.798

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.14: Percentage of male contributors before and after Travis-CI (collapsed observations)

However, when we carry out the same analyses with *percent\_male\_contributors* as the dependent variable (tables III.13 and III.14), we find that the coefficient on the interaction between the two independent variables is negative and strongly significant in the basic specification. Nevertheless, the significance vanishes as we include fixed effects and cluster standard errors. Correspondingly we also find that the coefficient is positive and significant when *percent\_neutral\_contributors* is used as a dependent variable (tables III.15 and III.16), with the significance vanishing similarly as we include fixed effects and clustered standard errors.

**Percentage of Neutral Contributors (period-wise observations)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
TREATED=1	-0.065*** (0.00)	-0.065*** (0.01)	0.000 (.)	0.000 (.)
POST=1	0.024*** (0.00)	0.024*** (0.01)	0.023*** (0.01)	-0.001 (0.00)
TREATED=1* POST=1	-0.006 (0.00)	-0.006 (0.01)	0.005 (0.01)	-0.000 (0.01)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.270*** (0.00)	0.270*** (0.01)	0.233*** (0.00)	0.147** (0.05)
R-squared	0.008	0.008	0.449	0.451

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.15: Percentage of neutral contributors before and after Travis-CI (period-wise observations)

**Percentage of Neutral Contributors (collapsed observations)**

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
POST=1	0.026*** (0.00)	0.026*** (0.01)	0.034*** (0.01)	0.036*** (0.01)
TREATED=1	-0.084*** (0.00)	-0.084*** (0.01)	0.000 (.)	0.000 (.)
TREATED=1* POST=1	0.027*** (0.01)	0.027* (0.01)	0.008 (0.01)	0.007 (0.01)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.290*** (0.00)	0.290*** (0.01)	0.248*** (0.00)	0.191*** (0.04)
R-squared	0.014	0.014	0.801	0.801

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table III.16: Percentage of neutral contributors before and after Travis-CI (collapsed observations)

Taken together, these results suggest that post-automation, projects might be attracting a marginally larger proportion of female contributors, reducing the proportion of male contributors, and also increasing the proportion of neutral contributors. Similar results are seen when *female\_pr\_share*, *male\_pr\_share*, and *neutral\_pr\_share* (which represent the share of pull requests made by female, male or neutral programmers respectively), and these are outlined in appendix B. Furthermore, as analyses with acceptance and rejection rates suggests, automation might be making adopting projects more attractive for female programmers, since their

contributions are less likely to be rejected right away, and even more likely to be accepted, albeit marginally so. These results could be an artifact of the fact that there are very few female programmers contributing to projects in our sample. Going ahead, we plan to collect more data and expand our sample in order to obtain stronger results.

## **Discussion and Conclusion**

This chapter looks at how the use of algorithms that (partially or completely) automate evaluation of work contributions may shape the extent to which females contribute to groups, and how their contributions are received by the group. We argue that the use of such algorithms may induce more female individuals to contribute to adopting groups. The reason for this is that the use of algorithms reduces cognitive loads for evaluators, and thus making it less likely for them to rely on biased mental shortcuts for evaluating work contributions. Furthermore, this effect is moderated by the gender of the evaluator; the bias-reducing effect of the algorithm is higher if the adopting evaluator is male.

We tested our theory by using the context of software projects hosted on Github (some of which automated the code evaluation process), and code contributions made by programmers with male, female and gender-neutral public identities. We found that automation adopting projects tended to attract more contributions from female programmers because they tended to accept their contributions. Through a difference-in-differences empirical design, we found that there was no difference between matched adopting and non-adopting projects before automation; however, adopting projects showed higher acceptance rates for female contribution compared to non-adopting projects after adopting automation. Correspondingly, contribution acceptance rates for male programmers reduced after adoption, while they stayed the same for non-adopting

projects. Furthermore, this bias-reducing effect is greater for projects whose maintainers are male.

The context of GitHub is especially well-suited to test our hypotheses. Firstly, the problem of lesser female participation, within STEM fields in general, and in software and tech firms in particular is widely documented (Terrell, *et al*, 2017; Murciano-Goroff, 2018). Secondly, the context provides a discrete instance of automation of cognitively intensive evaluation tasks, which require the use of subjective as well as objective criteria; furthermore the outcomes of these evaluation tasks can be tracked over time (before and after automation), and also for entities that automate or do not automate these tasks. Finally, the rich, fine-grained data allow us to control for a range of attributes for the unit under consideration, i.e. software projects.

The results outlined present several implications for both theory and practice. A recent body of literature that studies the use of AI-based tools and its effect on bias and fairness has presented mixed results: while some studies suggest that they could reduce bias and improve performance outcomes (Kleinberg *et al*, 2017), some other studies suggest that they could reinforce existing biases (Lambrecht and Tucker, 2019). However, bias can manifest in various ways within AI-based tools and algorithms (Choudhury, Starr and Agarwal, 2020), and most studies focus on bias stemming from data used by those algorithms (Cowgill and Tucker, 2019). This study contributes to this emerging literature by focusing on automation through algorithms that reduce the cognitive load of evaluators by automating evaluation tasks, and its effects on participation by females in those groups. Secondly, this study also aims to contribute to a burgeoning literature on evaluation bias against females within organizations (Botelho and Abraham, 2017; Carnahan and Greenwood, 2018; Uribe and Carnahan, 2018; Rivera and Tilcsik, 2019; Witteman, *et al*, 2019), and how it may affect their participation within them (Cardador, 2017).

Prior work in this stream looks at different policy measures that might mitigate bias within organizations, such as the Rooney rule (Rider *et al*, 2019), blind reviews (Goldin and Rouse, 2000), and designing appropriate evaluation tools (Rivera and Tilcsik, 2019). In this chapter, we present how automating evaluation tasks might affect evaluation bias within organizations.

At the same time, this study also has limitations which we plan to address as we develop it further. Firstly, the dataset currently includes only Java projects. Although Java is a stable and mature programming language, and is one of the most widely used programming languages in software development, including projects using other programming languages would allow us to expand our sample and increase the generalizability of our results. Secondly, GitHub projects are used extensively for online collaboration. Although virtual work and computer-mediated work is becoming increasingly widespread (Raghuram, *et al*, 2019), it still remains to be seen as to how automation would change bias and discrimination with physical co-presence.

Nevertheless, this study presents some of the first evidence of how automation could impact the decision of females (or members belonging to a marginalized community) to participate in and contribute to a collaborative group. Future work could also focus on how automation could change communication patterns within the group, and thus further affect bias. For instance, the use of algorithms could reduce the propensity for collective sensemaking through sharing of subjective opinions, and thus enable females to participate to a greater extent. Furthermore, it could also be worthwhile to explore the characteristics of groups that adopt automation. In this way, this study aims to stimulate research in the area of automation and gender bias within organizations.

## **Chapter IV: *Power to the People*: Decentralized Project Selection and Employee Self-Allocation in Organizations**

*(Co-authored with Maciej Workiewicz)*

### **Introduction**

Managing human capital is at the top of managerial agendas and one of the key drivers of firms' strategic advantage. It underpins organizational knowledge and allows organizations to reconfigure their activities to pursue new strategic initiatives (Ployhart and Moliterno, 2011; Campbell, Coff, and Kryscynski, 2012; Helfat and Peteraf, 2015). As firms increasingly engage in knowledge-intensive activities and become dependent on human capital rather than physical assets, the potential of employee misallocation to harm a firm's performance also increases. It is thus no surprise that companies are experimenting with different approaches to allocating their human capital as automation and AI further elevate the value of human skills for firms (Agrawal, Gans, Goldfarb, 2018). These approaches vary, ranging between centralization of the allocation of human capital and giving more autonomy to employees (Coff, 1997; Dobrajska, Billinger, and Karim, 2015; Levinthal, 2017). Here, organizations face an inherent dilemma: on the one hand, it is necessary to exercise a certain degree of control over employees so that they do not engage in "hobby" projects that could be wasteful for the firm's overall strategy. On the other hand, employees are hard to direct and usually have strong preferences for autonomy coupled with a distaste for formal organizational processes (Hamel and Prahalad, 1996; von Nordenflycht, 2010; Lee and Edmondson, 2017) and at the same time they often possess superior understanding of the firm's task environment (Aghion and Tirole, 1997; Hamel, 2011).



Traditionally, solutions to the centralization vs. autonomy dilemma have fallen within a broader class of hierarchical structures, like the functional form, the M-form, or the matrix form (Chandler, 1962; Davis and Lawrence, 1977; Miles and Snow, 2003; Levinthal and Workiewicz, 2018); more recently, however, firms have increasingly been adopting more decentralized and open solutions (Chesbrough, 2003; Lakhani and von Hippel, 2003). One approach that has captured the interests of scholars and practitioners is *open allocation*; i.e., allowing employees to initiate projects and join those they perceive as most promising, without interference from upper management (Puranam *et al.*, 2014; Burton *et al.*, 2017; Lee and Edmondson, 2017). This unorthodox solution to human capital allocation is an extreme example of decentralization. This approach goes beyond the bottom-up resource allocation process model wherein lower-level employees propose initiatives but still must obtain approval from at least one of their superiors to proceed (Bower, 1970; Burgelman, 1983). While open allocation was used in the past predominately in high-performing academic or research-oriented organizations, like the RAND Corporation or Princeton's Institute for Advanced Study (Augier, March and Marshall, 2015), traditional firms have also been using open allocation to complement their hierarchical approach. Some firms allow their employees to identify and join promising projects freely in some specific areas or part of the time, like Google's 20 percent time rule, 3M's 15 percent time rule (Levinthal, 2017), or Netflix with their high degree of employee autonomy (Gulati, 2018), while others like Valve or W. L. Gore and Associates use open allocation exclusively (Hamel, 2011; Puranam and Håkonsson, 2015; Puranam, 2018). Although open allocation is most commonly found in knowledge-intensive firms such as software development, it has also been introduced in other, more conventional settings like the manufacturing, food, and retail industries (Hamel, 2011). These examples suggest that this novel approach to human capital allocation may not only

apply to a small set of radical organizations, but could also be applied more broadly as an alternative or a complement to existing hierarchical forms, particularly as companies automate repetitive and physical activities, focusing on higher-order cognitive tasks (Agrawal, Gans, Goldfarb, 2018).

To realize the benefits and avoid the pitfalls of a new management innovation like open allocation, it is necessary to acquire a deeper understanding of the phenomenon. Academic research on open allocation processes, however, has been scarce, and key mechanisms and boundary conditions remain poorly understood (Lee and Edmondson, 2017). Furthermore, despite initial enthusiasm for this decentralized form of human capital allocation, anecdotal accounts point to limits to its efficacy and difficulties with scaling up (Foss, 2003; Puranam and Håkonsson, 2015; Augier, March and Marshall, 2015; Burton *et al.*, 2017). As some adopting firms gained more employees, they chose either to switch to a hierarchical structure, as GitHub did, or downscale its employees, as Valve did (Burton *et al.*, 2017).

While giving full autonomy to employees can improve organizational performance through a variety of mechanisms, like increased motivation, creativity, or organizational commitment (Lee and Edmondson, 2017; Puranam, 2018), we focus strictly on the problem of task allocation (Puranam, Alexy, and Reitzig, 2014). Specifically, we examine whether the task of identification of project opportunities and the allocation of employees to these projects should be carried out by a central planner or delegated to employees themselves and, if the latter, then under what circumstances. In answering these questions, we explore the mechanisms and boundary conditions of open allocation and compare its efficacy to that of centralized allocation. By doing so we aim to deepen our understanding of how can open allocation benefit organizational performance and how can organizations unlock its potential.

To address our research question, we build on the literature on project screening in hierarchies and polyarchies (Sah and Stiglitz, 1986; 1988; Christensen and Knudsen, 2010; Csaszar, 2013) and develop an agent-based computational model that captures the key elements of the process of project selection and resource allocation within organizations. We use this stylized model to identify the mechanisms and boundary conditions that give open allocation an advantage over a hierarchical allocation, or vice versa. One popular claim is that companies like Google or Valve give autonomy to their employees because they have excess resources (slack) and thus can afford to sacrifice some efficiency to keep their employees happy (Burton *et al.*, 2017). Thus only human-resource rich companies can afford this approach. Our results suggest just the opposite. We argue that in the presence of an imperfect evaluation of new opportunities, companies with slack in human resources would benefit by adopting centralized, rather than decentralized, approaches to organizing resource allocation. Organizations with insufficient human resources relative to opportunities may benefit from allowing their employees to self-allocate to projects, while those with slack human resources may find that open allocation leads to the “winner’s curse” (Thaler, 1988), wherein too many self-directed agents pursue a limited number of attractive opportunities, which in turn leads to overcrowding and adversely affects overall organizational performance. We argue that, in such situations, it is more beneficial to employ centralized project selection and human capital allocation, which, while subject to its own inefficiencies, reduces misallocation by limiting the overcrowding of projects.

In our subsequent computational experiment we explore the efficacy of key solutions proposed in the literature to mitigate some of the common pathologies associated with this form: a) stipulating the minimum number of employees required to launch a project, b) allowing employees to quit projects freely, c) setting up a minimum profit threshold, and d) requiring

approval from a superior to launch an employee-initiated project. Our analysis suggests that the efficacy of these managerial interventions also depends on whether an organization experiences either insufficient or slack human capital vis-à-vis available opportunities. In sum, our findings help explain why, as organizations scale up, the pressure to assume more hierarchical modes of employee allocation increases.

## **Theory**

Management scholars have long recognized that resource allocation processes and resources themselves play an important role in how organizations adapt to their environments (Bower and Gilbert, 2005; Levinthal and Wu, 2010; Folta, Helfat, and Karim, 2016). One of the central strategic paradigms, the resource-based view, identifies resources as the key factor determining the competitive advantage of companies (Barney, 1991). Similarly, the literatures on organization design (Galbraith, 1977), resource allocation processes (Bower, 1970; Burgelman, 1983), and dynamic capabilities (Eisenhardt and Martin, 2000) complement the resource-based view by focusing on processes of allocation and reallocation of organizational resources to respond to opportunities.

Scholars have defined organizational resources quite broadly to include any tangible or intangible assets controlled by the firm, like brands, knowledge, information, technology, machinery, plants, capabilities, and human capital (Barney, 1991). Human capital in particular has attracted significant scholarly attention as a part of the microfoundations of dynamic capabilities (Helfat and Peteraf, 2015) and as a source of sustainable competitive advantage due to the tacit nature and causal ambiguity of employees' knowledge and social complexity (Barney, 1991; Ethiraj and Garg, 2012). These characteristics of human capital protect firms against imitation by competitors but also make it difficult for firms to adjust the level of their

human capital in the short term (Coff, 1997; Campbell, Coff, and Kryscynski, 2012). On the one hand, many organizations, particularly in the early stages of growth, have insufficient levels of human capital vis-à-vis the opportunities available to them and struggle to hire and train new employees (Penrose, 1959; Birley and Westhead, 1990). Penrose (1959) in particular argues that the rate of growth of organizations has limits and thus many companies won't pursue all profitable opportunities. On the other hand, some more mature firms may maintain slack in their human capital as a buffer against environmental turbulence (Thompson, 1967; Lecuona and Reitzig, 2014; Bentley and Kehoe, *forthcoming*), even when this slack has a negative effect on their performance (Vanacker, Collewaert, and Zahra, 2017). Due to this "stickiness" of human capital, researchers have turned their attention to how firms may optimize the deployment of the human capital they already possess.

The increasing importance of this topic is reflected in researchers' renewed interest in human capital and its allocation (e.g., Lecuona and Reitzig, 2014; Mawdsley and Somaya, 2016; Sevchenko and Ethiraj, 2018). Allocating human capital can be divided into two tasks; a) evaluation of available opportunities, and b) allocating human capital to the most promising opportunities. With respect to the first task, research on organizational design has focused on the role of structure on the efficacy of information aggregation, with scholars studying how, by reconfiguring the structure of an organization, managers can channel and aggregate information to improve detection of good opportunities (Sah and Stiglitz, 1986; 1988; Knudsen and Levinthal, 2007; Christensen and Knudsen, 2010; Csaszar, 2013). With respect to the second task, allocation of human capital has been at the center of organizational design and the broader management research (Puranam, Alexy, and Reitzig, 2014; Levinthal, 2017).

The allocation of human capital can present significant challenges for organizations (Coff, 1997; Chatain and Meyer-Doyle, 2017). At a more general level, challenges arise because of the non-scale-free nature of human capital: while some resources, like patents or brands, can be applied to new projects without preventing their use in those already under way, most organizational resources, human capital included, possess an opportunity cost, as their use in one area precludes their use in another (Levinthal and Wu, 2010). This means that it is not enough to evaluate which opportunities are profitable, but also to be able to rank them relative to each other. The challenge is further exacerbated by several additional characteristics of human capital. First, employees, particularly in human-capital-intensive firms, may differ from their superiors in their evaluation and choice of available alternatives (Aghion and Tirole, 2007). Second, unlike machinery or capital, employees possess agency and often significant bargaining power (Coff, 1997). Third, employees generally have a strong preference for autonomy and tend to distrust centralized hierarchies (Herzberg, 1966; Hackman and Oldham, 1976). An organization designer thus must seek a balance between centralizing the coordination of employees to satisfy the global goals of the organization and allowing the employees to act according to their own local interests (Aghion and Tirole, 1997; Levinthal and Workiewicz, 2018).

Scholars and practitioners alike have picked up on this challenge and suggested many approaches to the efficient allocation of human capital. One trend observed among companies is the increased use of non-hierarchical forms of organization, where employees themselves are given significant latitude in initiating new projects and self-allocating to them (Daft and Lewin, 1993; Puranam, 2018). This trend has been attributed to increased uncertainty, new technologies, growth in knowledge-intensive work, and changing societal preferences (Lee and Edmondson, 2017). The autonomy of employees in such organizations extends beyond that found within

earlier frameworks like the “bottom-up” resource allocation process (Bower, 1970; Burgelman, 1983), in which employees suggested initiatives but had to obtain approval from their superiors to gain access to necessary resources. In contrast, companies practicing open allocation allow employees themselves to initiate and self-allocate to projects without approval from management—a much more radical approach.

Improving our understanding of open allocation is increasingly important as knowledge-based firms are playing an increasingly larger role in the economy (von Nordenflycht, 2010). However, prior work on such radically decentralized organizations has not addressed many key research questions and we still know very little on how this organizational form impacts resource allocation (Csaszar, 2013; Lee and Edmondson, 2017).

#### Open allocation

Several recent studies have offered qualitative accounts of open allocation of human capital in organizations. Puranam and Håkonsson (2015) provided a detailed account of Valve Software, a major developer of a computer gaming platform. By design, Valve Software eschews formal authority in managing its workforce: there are no project managers, and many important decisions—such as hiring and distributing rewards, which in hierarchies are completed by top or middle management—are delegated to lower-level employees. More important from our point of view is the observation that any Valve employee has the authority to launch a new project, provided her decision is supported by at least two other colleagues (this is called “the rule of three”). In addition, any employee can join another project if she finds it more attractive than her current assignment.

Some scholars, however, have expressed doubts that open allocation can scale to a larger number of employees or be implemented in more traditional sectors such as manufacturing (e.g.,

Burton *et al.*, 2017). Some manufacturing firms, such as W. L. Gore or Morningstar, use open allocation but impose certain restrictions on the number of employees allowed to work in a single team (Shipper and Manz, 1992; Hamel, 2011). Similarly, Google learned the value of middle managers when it tried to remove them from its hierarchy and gave greater autonomy to its engineers (Sutton and Rao, 2013). The company subsequently reversed its decision and reinstated middle managers, realizing that these managers provided an important interface between its executives and engineers.

A Danish manufacturer of hearing aids, Oticon, also experimented with open allocation. While its employees were allowed to initiate projects and allocate themselves freely, management retained project evaluation and monitoring rights. The company subsequently experienced significant growth in the 1990s but, due to the increasing interference of upper management into the allocation of resources, Oticon eventually reverted to a more conventional matrix form (Foss, 2003).

Another software company, GitHub, touted the value of open allocation for years but in 2014 it surprised everyone and decided to abandon it in favor of a strict hierarchy following a rapid growth in the number of employees (Burton *et al.*, 2017).

The literature in organizational economics has identified several factors pushing organizations to either greater centralization or towards autonomy. In their review of this literature, Aghion, Bloom, and Van Reenen (2014) associate delegation of decision rights with a) similarity in preferences between superior and subordinates; b) greater level of trust; c) greater industry heterogeneity; d) higher skills of workers; e) higher communication costs; f) greater use of information technologies; and g) higher product market competition. More important from our



perspective, however, is the finding that a firm's larger size has been found to lead to greater decentralization of decision rights, just the opposite to what the examples above suggest.

This leads to the following question: with all else being equal and in the absence of limits to the span of control and communication, can the growth in the number of employees force an organization to centralize the decision rights with respect to human capital allocation? Our study points to one such mechanism: the effect that the balance between a firm's human capital and opportunities it faces has on the joint task of identifying profitable projects and allocating employees to the most promising ones.

### Evaluating opportunities and allocating human capital in organizations

To be able to theorize clearly about the process of human capital allocation in organizations we need to first define the tasks that make up this activity. Generally, it can be conceptualized as consisting of three stages: 1) searching for alternatives, 2) evaluating them, and 3) implementing the chosen alternative (Csaszar and Eggers, 2013). We will first discuss the literature on search and evaluation, and then focus on the allocation of human resources in organizations.

#### Search and evaluation

The literature has extensively studied the role of organizational structure in shaping organizational search (e.g., Siggelkow and Levinthal, 2003; Ethiraj and Levinthal, 2004; Siggelkow and Rivkin, 2005). These studies have modelled firms as searching for the best combination of interdependent variables to explore how organizational structure affects this search process and the quality of the selected combination.

Prior researchers have also extensively studied how organizations evaluate available alternatives—a subject more central to the present chapter. Marschak and Radner's (1972) economic theory of teams was one of the first attempts to develop a formal model of

organizational decision-making with a focus on information processing. Building on their work, Sah and Stiglitz (1986, 1988) advanced a formal model of polyarchies and hierarchies composed of individuals who evaluate incoming projects using their own private noisy estimates. They examined how organizational structure determines the number of omission errors (failing to select a value-positive project) and commission errors (selecting a value-negative project) a given organizational form produces, finding that polyarchies (flat, decentralized organizations) produced more errors of commission, whereas vertical hierarchies committed more errors of omission. The intuition for this result is relatively simple: in a polyarchy, an organization in which a positive verdict from any decision-maker is sufficient to pursue a given project, too many projects will be launched, including some with negative values. Conversely, in a hierarchy, where approval at each level of the organization is needed to launch a project, the organization makes fewer errors of commission but misses more positive opportunities. More errors of omission occur in a hierarchy because a single mistake at any point in the chain of command will eliminate the proposal.

Christensen and Knudsen (2010) and Csaszar (2013) provided a more general version of the Sah and Stiglitz (1986, 1988) model by considering additional configurations of decision-makers. Using different approaches to calculate errors of omission and commission, both studies demonstrated that, for a given number of agents, different configurations of decision-makers in an organization can work to control the number of errors of commission and omission that the organization commits. Christensen and Knudsen (2010) offered further mathematical proof that the reliability of project screening can be adjusted by varying the organizational structure and the number of agents, thereby allowing the creation of an arbitrarily reliable organization out of unreliable parts (decision-makers).

Bridging the literature on hierarchies and polyarchies with the literature on organizational search and structure, Knudsen and Levinthal (2007) compared the performances of hierarchies and polyarchies, where a local search in a complex task environment was directed by the noisy estimates of individual agents. They suggested that hierarchies became stuck at local maxima faster than polyarchies because the latter generate positive errors that let them get unstuck from these local peaks.

### Allocation of human capital

In the above-cited approaches, employees in decentralized structures are assumed to have the authority to evaluate projects, yet they are viewed strictly as passive information processors who can be arranged in different static configurations to produce desired organizational outcomes related to evaluation. Their job is reduced to evaluating incoming opportunities and making a binary accept/reject decision about whether to pursue an opportunity. What has been relatively less studied is the subsequent task of resource allocation, which could be carried out either by the employees themselves or by the manager. In other words, the implicit assumption in this literature is that an organization always has sufficient resources to pursue all selected projects, and only the screening efficacy is important. Firms, however, do not stop at simply identifying the most promising projects among those available: they must also judiciously assign limited resources to maximize their performance (Noda and Bower, 1996; Bardolet, Fox, and Lovallo, 2011; Klingebiel and Rammer, 2014). Resource misallocation, which occurs when too many or too few resources are allocated to projects, is an important consideration after all available projects have been identified (Levinthal and Wu, 2010).

Prior work has shown that firms do not operate in equilibrium when it comes to human capital—they either maintain a certain amount of slack resources (Vanacker *et al.*, 2017) or seek

to hire new employees to meet increasing demands, since hiring and firing employees is costly, especially for human-capital-intensive, knowledge-based firms (Penrose, 1959; Lecuona and Reitzig, 2014).

While there has recently been a resurgence in research on human capital allocation, to our knowledge there are no formal models exploring the mechanisms and boundary conditions of open allocation in organizations. Although a large body of literature uses simulations to explore the effects of organization design on organizational behavior and adaptation (see Baumann, Schmidt, and Stieglitz, 2018 for a recent review), or how adding employees in certain structural positions can impact screening efficacy in particular (e.g., Knudsen and Levinthal, 2007; Csaszar, 2013), there is little formal research on the role of organizational structure in shaping the process of resource allocation in organizations (Sengul, Almeida-Costa, and Gimeno, 2018). Similarly, Csaszar (2013) argued that incorporating resource constraints into the screening model would be an important extension of that literature.

The few existing formal models of resource allocation in management (Burton and Obel, 1984; Coen and Maritan, 2010; Hutchison-Krupat and Kavadias, 2015), organizational economics (e.g. Athey and Roberts, 2001; Rajan and Zingales, 2001; Hart and Moore, 2005), or finance (e.g. Stein, 2002) focus on financial capital or material resources in general and examine only hierarchical solutions to the problem. Recognizing this gap, Gertner and Scharfstein (2013) specifically called for more research to address issues of human resource allocation such as “assignment of workers to various jobs within firms” (Gertner and Scharfstein, 2013:674).

Our aim is to help fill this lacuna in the research on open allocation by exploring the efficacy of human capital allocation in self-managed organizations relative to those with hierarchical structures with the help of a formal computational model. In this chapter, we

consider the processes of project evaluation and selection and human resource allocation jointly. Our goal is to study how open allocation and a hierarchy each perform in this dual task and identify the boundary conditions for their efficacy. In doing so, we highlight the importance of the relative balance between available opportunities and the human capital needed to realize gains.

## **Model**

We set up an agent-based model to examine the efficacy of human capital allocation in organizations. We consider two canonical types of employee allocation: 1) open allocation, in which the employees individually evaluate and self-allocate to projects; and 2) centralized allocation, in which a single manager (superior) evaluates opportunities and allocates employees (subordinates) to selected projects.<sup>13</sup> For both types of allocation, all agents (including the manager in a centralized organization) possess inaccurate estimates of the projects' true values, which sets up the problem of identifying value-creating opportunities and assigning sufficient resources to maximize gains.

In representing open and centralized allocation of human capital, we adopt the polyarchy and hierarchy definition introduced by Sah and Stiglitz (1986), who described the two concepts as follows:

[Polyarchy is] a system in which there are several (and possibly competing) decision makers who can undertake projects (or ideas) independently of one another. In contrast, decision-making authority is more concentrated in a hierarchy in the sense that only a few individuals (or only one individual) can undertake projects while others provide support in decision making. (Sah and Stiglitz, 1986: 716)

---

<sup>13</sup> Although the literature on project screening and the role of organizational structure has explored the role of different hierarchical arrangements in project evaluation (for a detailed analysis, see Christensen and Knudsen, 2010, and Csaszar, 2013), we start with two canonical forms. This simplified approach allows for greater tractability, because we consider a smaller number of interactions while preserving the key mechanism of interest.

In terms of project selection, our conceptualization of open allocation is equivalent to a polyarchy. Similarly, centralized allocation is represented by a hierarchy, where a single manager selects projects and allocates employees to them. In the next sections we introduce the key elements of the model; a) an organization, b) its task environment, c) the allocation process, and d) the approach to measuring performance of our stylized organizations.

### Organization and its task environment

We begin by specifying the task environment. In each round  $t$ , both organizations—open and centralized—face a fixed number of projects,  $r$ . For each of the  $r$  projects, we randomly draw a number that determines the true revenue potential of that project, denoted as  $\beta_r$ . This parameter is independently distributed according to  $U(-10,10)$ , with the expected payoff potential of a single project thus set to 0.

Each organization has  $m$  employees. The goal of an organization is to allocate the available  $m$  employees across the projects to maximize the organization's overall profit. Each of the  $r$  projects in turn has a *carrying capacity*, which is the number of workers needed to profitably execute a project. When the number of workers allocated to a project exceeds (is smaller than) its *carrying capacity*, then the project is overstaffed (understaffed). By summing all *carrying capacities* available in a given round, we obtain the carrying capacity of the task environment in a given round, which we denote by  $L_t$ .

The key variable of interest is *resource load*, a ratio between the overall number of workers  $m$  and  $L_t$  given the number of projects and their value distribution.<sup>14</sup> This variable determines whether on average an organization is able to meet the opportunities offered by the environment.

---

<sup>14</sup> Our parameter is similar to the *net energy load* from the Garbage Can Model (Cohen, March, and Olsen, 1972). However, unlike in the Garbage Can Model, we are interested in the efficacy of project selection and human capital allocation rather than in completion rate of projects.

When *resource load* is less than 1, it means that an organization cannot pursue all the profitable opportunities and in the long run will be seeking to increase the number of employees, i.e., grow (Penrose, 1959). Conversely, when this variable is more than 1, an organization has a slack in human capital. It may maintain it for strategic reasons (for more detailed discussion see Lecuona and Reitzig, 2014), or in the long run it will seek to reduce the number of employees. With *resource load* equal to 1, an organization has just the right amount of human capital to allocate to all profitable opportunities.

### Allocation process

Open allocation proceeds as follows. First, we generate noisy estimates of projects' revenue potential for each of the  $m$  workers. Specifically, for each of the  $r$  projects, each worker observes the true revenue potential  $\beta_r$  with some noise  $\sigma$ , which we draw from a normal distribution  $N(0,2)$ . We denote the worker's noisy estimate as  $\beta_{r,m}$ . Next, we randomly select one worker who then identifies the best allocation, i.e., the project with the highest marginal profit  $P_{r,m}$ . The marginal profit for each project is a function of the estimated revenue potential of the project ( $\beta_{r,m}$ ), the number of workers already attached to that project ( $n_r$ ), and the marginal costs of adding one more worker to the project ( $C$ ). The payoff potential of a given project is calculated as:

$$P_{r,m} = \beta_{r,m} \cdot n_r - C_r \quad (4)$$

where:

$$\beta_{r,m} = \beta_r + \sigma \quad (5)$$

$$C_r = \frac{n_r(n_r-1)}{2} \quad (6)^{15}$$

---

<sup>15</sup> Increasing employee cost captures coordination difficulty in teams. Thus, the first worker generates no cost, the second worker generates a marginal cost of 1, the third a marginal cost of 2, and so on. This conceptualization

Thus, with each additional worker, the marginal profit of a project diminishes.<sup>16</sup>

Diminishing returns are a necessary condition for a resource allocation problem to arise. The presence of diminishing returns has been confirmed in a wide variety of settings, like R&D projects (Scherer, 1967), allocation of talent in manufacturing (Murphy, Shleifer, and Vishny, 1991), human capital in general (Chatain and Meyer-Doyle, 2017), and software projects (Brooks, 1995; Boehm, Abts, and Chulani, 2000). It also underpins neoclassical theories of production in general (Maksimovic and Phillips, 2002). Diminishing returns from subsequent allocations of human capital to a given project and the limited nature of human capital change the nature of the allocation problem. Instead of focusing only on appropriately identifying value-positive and value-negative projects, as is customary in models of project screening (Sah and Stiglitz, 1986, 1988; Knudsen and Levinthal, 2007; Christensen and Knudsen, 2010), the organization designer also needs to consider opportunity costs of resources.

Having calculated the payoff potential of each project, the worker selects the project with the highest positive marginal value and joins that project. We then proceed by randomly selecting the next worker from those remaining. The second worker performs the same calculation, and so on, until all workers have had an opportunity to self-allocate to a project. A project is considered launched when at least one worker is assigned to it. Workers who cannot find a project with a net positive payoff remain idle for that round. Thus, following the earlier discussion, we implicitly assume that in the short term the organization does not hire or fire employees; consequently, wages of workers are not taken into account.

---

results in a situation in which assigning only one worker, and thus launching a randomly selected project, will generate an expected profit of 0.

<sup>16</sup> We can implement diminishing returns in several ways. For example, we could keep the cost constant and lower the revenue per each additional worker. Our results largely hold for other, alternative implementations of decreasing marginal return from additional workers.



In centralized allocation, we place a single manager in charge of screening and allocation decisions. The manager also possesses noisy estimates of the projects' revenue potential. Following the literature on screening, we hold the distribution of the noise term constant between the workers in open allocation and the manager in centralized allocation, as we are interested in the effects of structure and not screening ability (Knudsen and Levinthal, 2007; Csaszar, 2013). Given her own estimates of project values, the manager then proceeds to allocate the workers using the same formula the workers use in open allocation. The manager proceeds until either no more workers remain, or she cannot find any positive marginal profit allocations within the existing projects.

#### Organizational performance

At the end of each round, when all the workers have had a chance to self-allocate or have been allocated to a project by their manager, we calculate the actual profit per project by summing the profits or losses that each active worker generates. Specifically, to calculate organizational performance, we use the true revenue potential  $\beta_r$ , instead of the noisy estimate of the project's revenue potential  $\beta_{r,m}$ . The profit calculation is illustrated in illustration IV.1.




Illustration IV.1: Calculating the project's profit<sup>17</sup>

## Results

We start our simulation by setting the number of projects available per round to  $r = 16$  for both open and hierarchical allocation and vary the number of workers between 20 and 80, specifically  $m = \{20,30,40,60,80\}$ . By doing so, we can observe how well the two organizational forms allocate employees to opportunities in environments in which the resources available to each organization may or may not be sufficient to pursue all profitable projects. We selected this range for  $m$  for a specific reason. In an environment where the number of workers  $m = 40$ , both of our organizations, on average, have exactly the number of workers needed to pursue all positive-value projects.<sup>18</sup> Subsequently, when  $m = 60$  and  $m = 80$ , the organization's *resource*

<sup>17</sup> Note: The figure illustrates project's profit calculation. Each additional worker adds a fixed amount of revenue equal to the revenue potential associated with that given project (here, equal to 4.5). The cost of labor, which increases with the number of workers attached to the project, is deducted from the total revenue to arrive at total profit. Increasing the number of workers beyond the point where marginal cost equals marginal revenue (here five workers) leads to negative marginal profit.

<sup>18</sup> In an environment with  $r = 16$ , the expected number of value-positive projects is 8. Because the value of the project is distributed uniformly, in each of these 8 projects, the average expected revenue potential,  $\beta_{r,m} = 5$ . Thus, according to our discussion regarding the model, each of the projects can accommodate, on average, 5 workers, which results in a total carrying capacity of 40 workers.

*load* is equal to 150 percent and 200 percent, respectively. In other words, the organization has a slack in human capital. Similarly, for  $m = 20$  and  $m = 10$ , the *resource load* is 75 percent and 50 percent, respectively. The organization has on average insufficient human capital to pursue all opportunities.

We simulate both organizations over 20 time periods and run 10,000 iterations of the simulation, presenting average results to eliminate artifacts of random sampling. For ease of comparison, we normalize the results by comparing the performance of the open and centralized allocation to that of the *optimal allocation*, which represents the best possible performance under given circumstances. We present the results achieved by the open and centralized allocation as a percentage of that optimal value.

#### Organizational performance

We begin our analysis by examining the performance of both organizational forms. Figure IV.1 shows the relationship between *resource load* and the normalized firm performance. While the performance of centralized allocation remains largely unchanged with respect to the increase in the *resource load* ratio, the performance of open allocation decreases as *resource load* increases beyond 100 percent. In other words, the performance of open allocation sharply deteriorates as the organization's size increases.



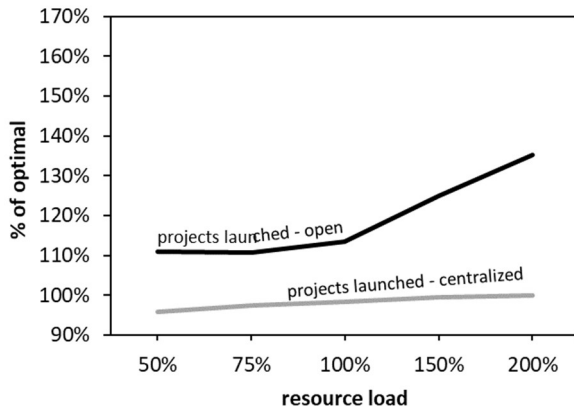


Figure IV.1: Performance Comparison Between Open and Centralized Allocations

To understand the mechanism behind this result, we need to first examine a) the ability of each form to select correct projects and b) the efficacy of resource allocations to those projects. We begin by examining the number of projects launched and the number of active workers for both allocation modes, as shown in Figure 8. We find that for all values of *resource load*, open allocation initiates more projects than centralized allocation, which in turn launches roughly the same number of projects as optimal allocation. This is consistent with prior findings that associate polyarchy (which is analogous to open allocation) with accepting more projects (Sah and Stiglitz, 1986, 1988; Csaszar, 2013; Christensen and Knudsen, 2010). The difference in projects launched increases with *resource load*. For a *resource load* of only 50 percent, centralized allocation launches 96 percent of the optimal number of projects, and open allocation launches 111 percent. When examining the number of allocated workers, we notice that open allocation systematically uses too many workers, and overstaffing increases with *resource load*.



a) Percentage of projects launched (w.r.t. optimal allocation)

(b) Percentage of workers active (w.r.t. optimal allocation)

Figure IV.2: Projects Launched and Active Workers for Each Form

However, these facts alone do not explain our main result presented in Figure IV.2. For low levels of resource load, open allocation outperforms centralized allocation even though it makes more errors in both project selection and resource allocation. To understand this result, we must explore not only the quantity of project selection and resource allocation errors, but also the efficacy of allocation.<sup>19</sup>

#### Commission and Omission Errors in Project Selection and Resource Allocation

While the concepts of commission and omission errors are key in the screening literature, our addition of resource constraint fundamentally changes the selection problem and requires us to consider other types of errors a decision maker faces when available resources must be taken into account. Here it may be helpful to consider a stylized example of these challenges. Figure 4 presents such a hypothetical scenario. For clarity of exposition, we use only six projects and six employees.

<sup>19</sup> By “accurate allocation decisions,” we mean allocations that result in maximum overall profit for the organization given its human capital.

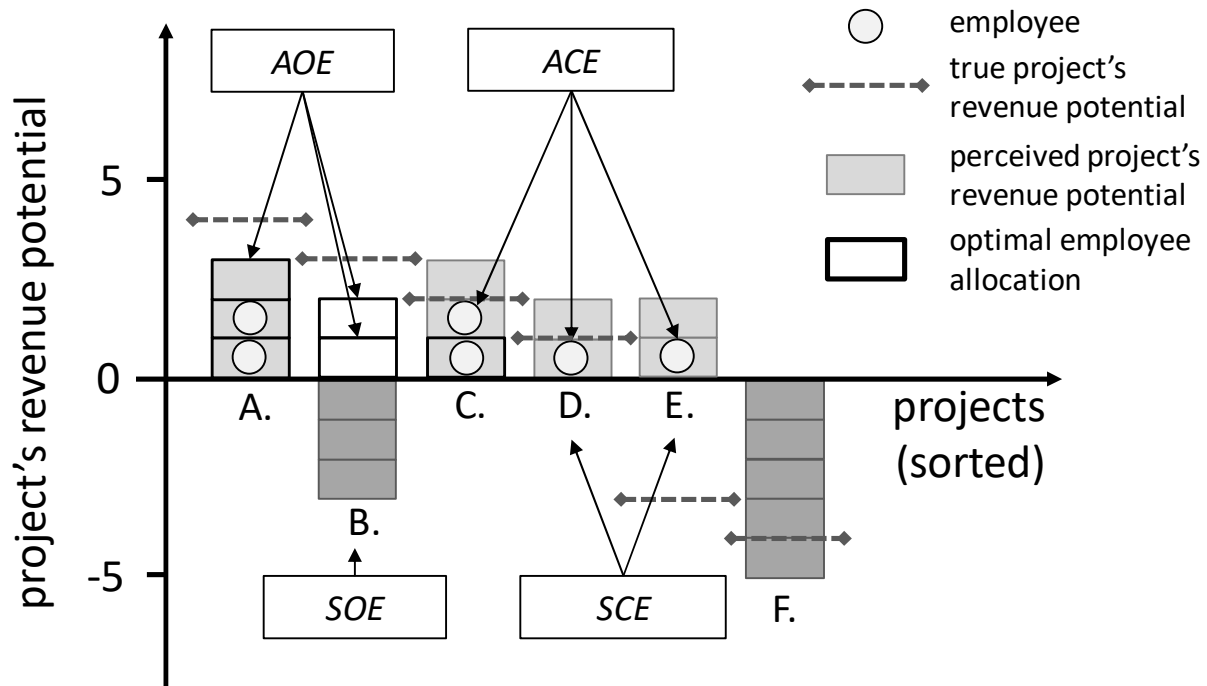


Illustration IV.2: Example of Employee Allocation

To help interpret the figure, we indicate the *true revenue potential* of each project with a dashed line and use thick black frames to mark where workers should be assigned according to optimal allocation. Because the manager and the workers have noisy estimates of projects' *true revenue potential*, they can make two types of errors, omission and commission, in the projects they launch (selection) or in the number of workers they assign to each project (allocation). We denote selection omission errors as *SOE* and selection commission errors as *SCE*. In the case of *SOE*, the organization may fail to launch a project that would have been launched with optimal allocation (Illustration IV.2, project B). We denote the projects selected by the optimal allocation as optimal projects.<sup>20</sup> The organization may also launch a project that should not have been launched, thereby committing an *SCE* (Figure 4, projects D and E). However, here we encounter

<sup>20</sup> By definition, all optimal projects are value-positive

the first difference between our model and the models considered in the extant literature on polyarchies and hierarchies. Instead of only one type of *SCE*, there are two: 1) launching a value-negative project (Illustration IV.2, project E), and 2) launching a value-positive project that should not have been launched due to resource constraint (Figure 9, project D). The second type of *SCE* arises because a project selector not only must consider whether a project is value-positive or value-negative but also must prioritize the projects to maximize the use of limited resources. When resources are scarce only the best value-positive projects should be selected.

An organization can also make allocation errors. We distinguish between allocation omission errors, which we denote as *AOE*, and allocation commission errors, which we denote as *ACE*. An organization commits an *AOE* if it assigns a less-than-optimal number of workers to a project (Illustration IV.2, projects A and B). Note that while all *SOEs* are *AOEs* by definition, the converse is not true. An organization may launch an appropriate project but allocate too few resources (Illustration IV.2, project A).

An *ACE* is committed when an organization assigns more workers to the project than optimal allocation would dictate. Just as there are different types of selection errors, there are three types of allocation errors when resources are limited: 1) allocating excessive resources to an optimal project (Figure 4, project C), 2) allocating resources to a value-positive but not optimal project (Figure 9, project D), and 3) allocating resources to a value-negative project (Figure 9, project E). Thus, all *SCEs* are *ACEs*, but the converse is not true.

Thus, in the presence of limited resources, an organization can make omission and commission errors not only with respect to project selection, but also with respect to resource allocation. In the presence of limited resources, there are several types of commission errors (for both selection and allocation problems) and two types of omission errors with respect to allocation. The implicit assumption in the classical literature is that an organization possesses

enough resources to address all accepted projects; thus, it focused only on the efficacy of the project evaluation (Sah and Stiglitz, 1986; Knudsen and Levinthal, 2007; Christensen and Knudsen, 2010; Csaszar, 2013; Csaszar and Eggers, 2013) or the speed of decision (Radner, 1993; Garicano, 2000). In our model, however, the amount of human capital available to allocate to opportunities is an important factor. It may be worthwhile to forgo some positive-value projects if the available employees would be better deployed on other opportunities. Thus, allocation decisions are just as important as project screening in determining organizational performance. Our model thus suggests that an organization designer should consider the balance between human capital and opportunities when choosing an appropriate organizational structure.

#### Mechanisms Underlying Selection and Allocation Errors

To understand the source of differences in performance between open and hierarchical allocation, we examine the two organizational forms' propensity to commit selection and allocation errors, and we calculate how each type of error contributes to the overall performance. Figure IV.3 shows how project selection errors vary with changes in the *resource load* for both allocation modes. Unsurprisingly, panel IV.3. a) shows that open allocation rarely fails to launch an optimal project, while centralized allocation misses on average almost one optimal project when *resource load* is low. For both allocation modes, the number of SOEs drops when the number of employees available grows. In terms of SCEs, panel IV.3.b) shows that open allocation launches more non-optimal projects (both value-positive and value-negative) than centralized allocation, and this tendency increases with *resource load*. In other words, when it comes to project selection, open allocation is more risk taking, while centralized allocation is more conservative. These findings are in line with the propositions found in the literature (Sah and Stiglitz, 1986; Csaszar, 2013). In fact, when resource constraint is not an issue (high



*resource load*), the allocation problem is reduced to a project screening task; i.e., simply identifying and launching value-positive projects. Note that the number of value-positive but not optimal projects approaches zero as *resource load* increases. As mentioned above, when resources cease to be a binding constraint, all value-positive projects should be undertaken; i.e., they become optimal.




Figure IV.3: Selection Errors

This does not necessarily imply, however, that when *resource load* is low, the relative advantage of open allocation over centralized allocation stems exclusively from the lower number of SOEs. Examining how each allocation mode assigns workers to projects yields further insights into the relative performance of the two forms and the mechanisms at play.

Figure 11 compares the allocation commission and omission errors (AOEs and ACEs) between open and centralized allocation. For the AOEs presented in Figure IV.4. a), the picture is similar to the one we observed for SOEs; the ACEs, however, offer a surprising and interesting insight. Contrary to the SCEs and the intuition offered by the literature on screening, when resources pose a constraint (low *resource load*), centralized allocation commits *more* commission errors (Figure IV.4, panel b). Specifically, centralized allocation overstaffs optimal

projects by two workers per round on average (3.28 vs. 1.29) for *resource load* = 50 percent and continues to overstaff optimal projects until *resource load* passes 100 percent. Centralized allocation also allocates more workers to projects that are value positive but not optimal; however, the difference is more modest (0.82 vs. 0.72 workers for *resource load* = 50 percent). Thus, we observe that when resource constraints are not an issue (high *resource load*), both allocation modes commit selection and allocation errors in line with what we would expect from the literature on hierarchies and polyarchies. However, when an organization must also consider its limited amount of human capital, the centralized form not only commits more SOEs but also commits more ACEs, which is a novel insight not yet documented in the literature to the best of our knowledge.




Figure IV.4: Allocation Errors

The difference is particularly stark when we consider the number of allocation errors per project. Figure 12 shows the relative number of ACEs per project. The first admittedly obvious observation is that as *resource load* increases, the number of ACEs per project committed by open and centralized allocation increases. This is to be expected since the more resources are available, the more can be wrongly allocated.

More importantly, the other trend helps to illustrate the key mechanism at play. Centralized allocation results in more ACEs when *resource load* is both below and above 100 percent. In contrast, open allocation starts with relatively few ACEs per project when *resource load* is low, but ACEs per project increase sharply as *resource load* increases. In other words, with open allocation, if the human capital significantly exceeds the number of projects available, the excess employees will overstaff the few projects available, thereby reducing the projects' overall profit. In open allocation, after the carrying capacity (i.e., the maximum number of workers a given project can profitably accept) of all projects is met, the remaining workers may still join if their estimates for a given project are sufficiently positively biased. This mechanism leads open allocation to overstaff projects when *resource load* is high. The result is similar to the winner's curse (Thaler, 1988), where the last worker to join a project is likely to have an extremely positively skewed estimate of the project's true value. Hierarchy, with its centralized allocation procedure, does not experience open allocation's problem of compounding errors. The single manager still makes mistakes (performance never reaches 100 percent for either of the two forms), but errors are made only once per project. Regardless of the number of arriving projects, the manager is equally likely to err in her estimates.

The situation is different, however, when *resource load* is low. The effect of overcrowding is reduced for open allocation because workers spread themselves across many projects, which reduces the compounding of errors. When evaluating and selecting projects, each worker draws a preferred ordering of the available projects based on her private estimates. Because all workers act independently, using their own erroneous estimates, these orderings are unlikely to coincide exactly. With an abundance of understaffed positive-value projects, the carrying capacity of projects will rarely (if ever) be reached. Furthermore, because the private

estimates of workers are correlated with the projects' true potentials, workers are more likely to make profitable decisions. This effect is absent in centralized allocation because only the manager's erroneous estimates are used to formulate a single preference ordering, which is used to assign workers to projects. With only one person allocating workers, the manager's omission errors will not be ameliorated, and commission errors will be fully realized. In open allocation, an omission error committed by one worker is limited by the amount of human capital that the worker represents. The next worker is not bound by the beliefs of preceding workers and may select a project that the previous one deemed unattractive.

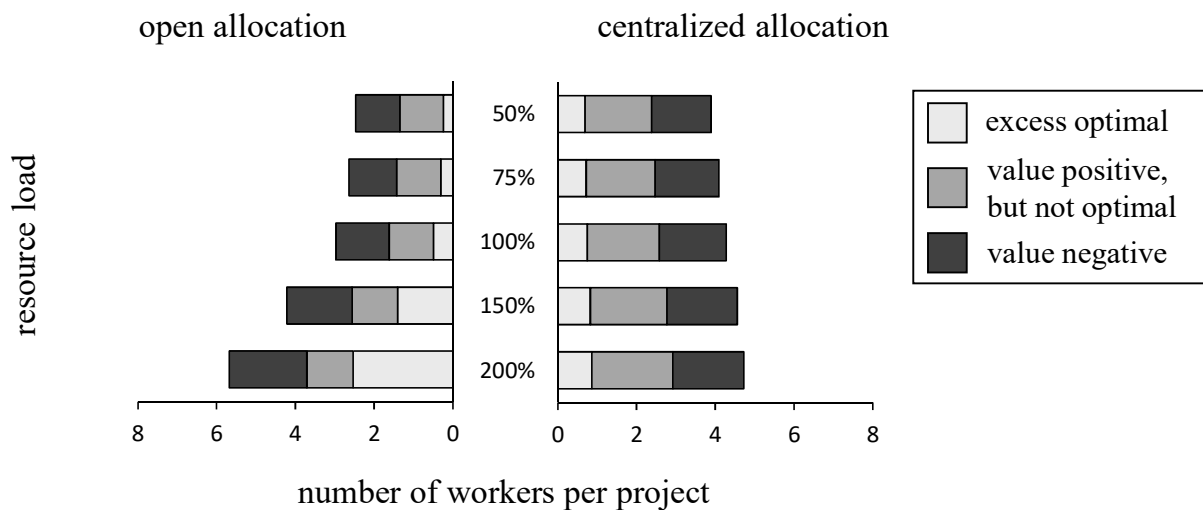


Figure IV.5: Allocation Commission Errors (ACEs) per project

### Impact of Errors on Organizational Performance

Collectively, our results suggest that the relative balance between available human capital and opportunities—and thus the degree to which an organization's resources pose a constraint—presents important implications for organization design. By considering how each form performs relative to project selection and resource allocation in regimes characterized by different levels

of *resource load*, we identified the mechanisms responsible for the performance discrepancy displayed in Figure IV.1.

To complete our analysis, we calculated the average performance change from the optimal level due to each of the resource allocation errors (AOEs and ACEs). Figure IV.6, panels a) and b) presents the effect of the different types of resource misallocation on performance for open allocation and centralized allocation, respectively. The relatively steady performance of centralized allocation, as outlined earlier in Figure IV.1, masks two trends.

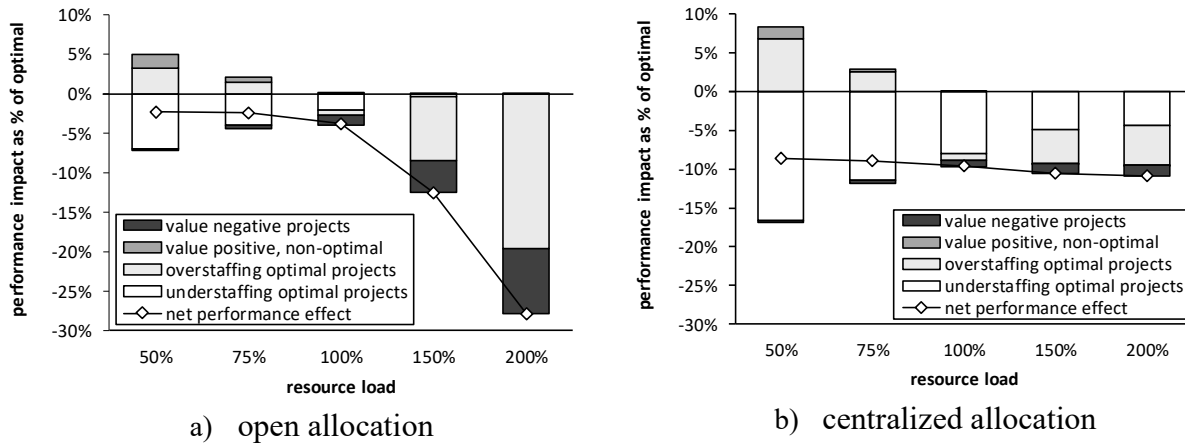


Figure IV.6: Performance Impact of Human Resource Misallocation

First, for low *resource load* (50 and 75 percent), the performance discount is higher for centralized allocation because it not only tends to understaff certain optimal projects, but also tends to overstaff other optimal projects. This is another illustration of the idea that the allocation errors made by a single manager are fully realized. Consequently, this forces centralized allocation to commit AOE for optimal projects, which brings down its overall performance. Overstaffing optimal projects might increase performance, but it is a suboptimal use of resources that could be better deployed on other projects.

Second, for open allocation with high *resource load*, ACEs are the main source of the performance discount. Specifically, workers' errors compound, leading to overcrowding of projects and to a significant drop in performance. Simply put, there are too many "fools" (i.e., noisy evaluators) chasing too few opportunities. When there are fewer opportunities relative to the number of workers, in the absence of managerial oversight, overly optimistic workers are more likely to crowd optimal projects or assign themselves to net-negative projects, bringing down overall performance.

### Policy levers

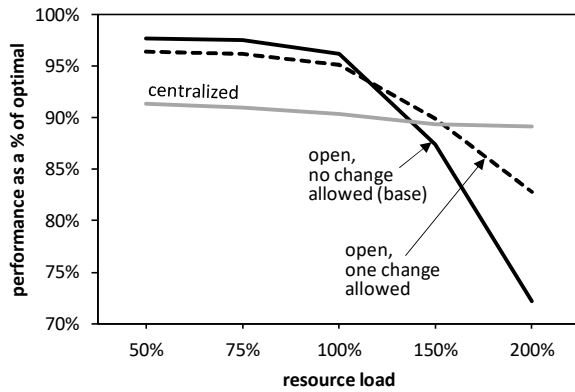
After the main analysis, we examine the key decision rules that firms employ to address some of the shortcomings of open allocation. These policy levers, often presented in extant literature, are meant to help firms reduce misallocation due to evaluation errors. They include 1) allowing employees to leave a project they selected previously and move to another project if they evaluate the latter as more attractive than the former, 2) setting a minimum number of workers a project must attract before the organization will sanction it, 3) introducing a minimum profit threshold that proposed projects must meet, and 4) introducing a manager who must approve projects proposed by employees. We examine these policy levers in turn and use our model to examine their efficacy and potential limitations. We find that these seemingly reasonable and prudent interventions may in fact, under certain conditions, beget opposite effects to those intended.

### Allowing Employees to Quit Projects Freely

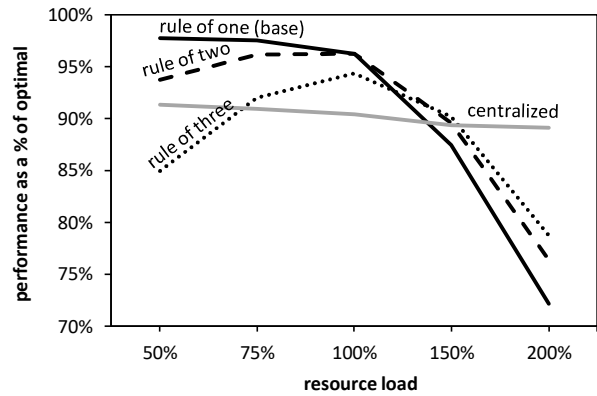
In the base version of our model, a worker evaluates projects by considering her personal estimate of the project's potential and the number of workers currently staffed to it. Each worker is drawn randomly from the pool of unassigned workers and decides which project, if any, to

join. This arrangement, however, may lead to an unfavorable situation for a worker who joins the project too early. If a worker gets an early chance to choose an initially attractive project, other workers may join thereafter, making the project unattractive to the initial worker. Companies like Valve allow employees to freely change projects without any repercussions if they find one that is more appealing (Puranam and Håkonsson, 2015). In an extended version of our model, we introduce a new variable, *maxswaps*, which controls the number of times any worker may reconsider her decision. By setting *maxswaps* to one, we allow each worker to change her mind one time per round and either join another project or remain idle. The workers are selected in random order and given the opportunity to revise their prior decision following the procedure outlined in the Model section. In Figure IV.7, panel a), we present the results and compare them with those of the original model.

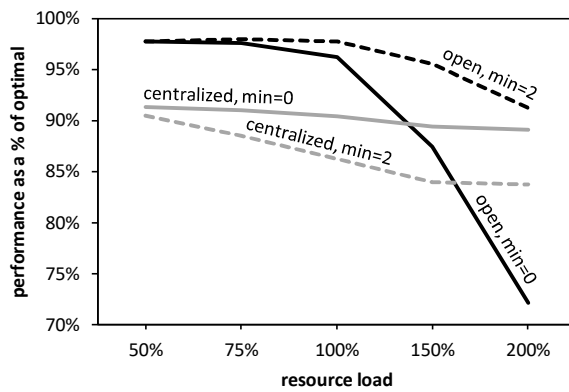
The effect of setting *maxswaps* to allow one change is most evident for high *resource load* and when overcrowding of projects leads to a performance drop for open allocation. In a situation where too many workers chase too few opportunities, open allocation benefits from allowing workers to reevaluate their choice. Without the ability to change their mind, workers who joined a given project early may find themselves in a situation where they no longer find the project attractive (because it is overstaffed). Such mistakes would affect overall firm performance adversely in two ways. First, launching an overcrowded project directly reduces organizational performance. Second, it imposes an opportunity cost: the worker could have joined a more attractive project and increased that project's outcomes instead. By allowing employees to reevaluate their decisions, we minimize such inefficiencies.



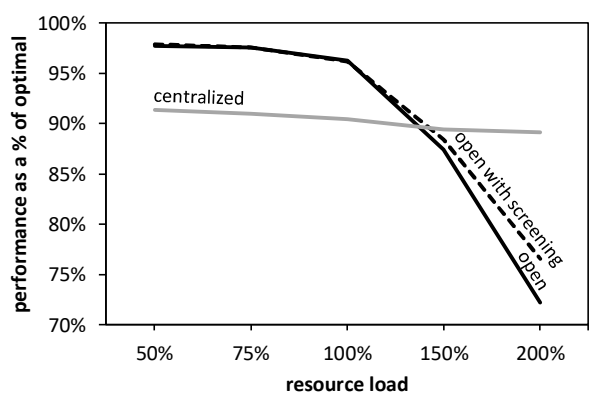
a) Allowing employees to quit projects freely



b) Setting a minimum number of workers to launch



c) Setting a minimum threshold for profits



d) Requiring an approval from a manager to launch a project

Figure IV.7: Policy Levers

However, allowing employees to change their mind has a negative effect when *resource load* is low. This negative effect does not disappear with an increase in the *maxswaps* variable and it reveals yet another mechanism. In short, allowing employees to change their minds after everyone has had a chance to allocate to a project may lead to a situation where pessimists will unnecessarily leave good projects. Consider a simple example. Let's assume that there are two employees and only one project to allocate to. The project's *true revenue potential* ( $\beta_r$ ) is 1.5, which means that it can profitably accommodate two workers. The first worker, who is a



pessimist, evaluates the project at 0.5, while the second worker, who is an optimist, evaluates it at 2.5. Let's assume further that the workers joined the project with the pessimist deciding first and the optimist second. Both would decide to join and the organization would realize optimal profit. However, if we allow workers to reconsider, the pessimist will leave deciding that the project is overcrowded (remember that with her private evaluation of 0.5 she thinks the project can only accommodate one person). This results in a loss for the organization.

This effect occurs with greatest intensity when the number and value of opportunities available just about matches the number of available human capital. When *resource load* is high, this negative effect is overtaken by the positive effect of lessening overcrowding of projects; but with low resource load this does not happen, and an organization is left with only a negative effect of pessimists leaving projects they should stay on. Thus, contrary to common wisdom, it may be better for an organization to prohibit employees from leaving projects when human capital is constrained.

#### Setting a Minimum Number of Workers for Launching

Firms such as GitHub and Valve impose a “minimum viability” condition to start a new project. This ensures that employees do not pursue personal interest projects at the cost of the firm's resources. For example, Valve imposes a “rule of three,” which mandates that any new project attract at least three employees to be sanctioned (Puranam and Håkonsson, 2015). Similarly, GitHub imposed a “rule of two” (Burton *et al.*, 2017). Because workers determine a project's attractiveness based on their erroneous estimate of its value, some workers with an extremely high positive bias may eventually join a value-negative project because of their positively skewed evaluations. Setting a minimum number of workers necessary to launch a project makes

it more difficult for biased employees to start one, as they must find other employees who share their extreme (and possibly misplaced) optimism about that project.

To simulate this policy intervention in our model, we implement a variable *rule of*, which specifies the minimum number of workers required to launch a project. When deciding whether to join a given project, each worker estimates the project's value to confirm that it can support the minimum number of workers required. After all workers have made their choice, projects that do not attract the required number of individuals (the minimum) are dropped. Note that the base model has the *rule of* set to one. We again find that the efficacy of this rule depends on *resource load*; thus, its use should be considered with this balance in mind.

Figure 14, panel b) shows that a high *rule of* has a positive effect on firm performance for high *resource load*. When the *rule of* value increases, workers launch fewer projects. The rule has a positive impact primarily because it prevents the launching of value-negative projects, but it does not prevent the overstaffing of optimal projects. When *resource load* is low, a higher *rule of* results in the rejection of too many good projects due to the insufficient number of workers. Along with some value-negative projects, some attractive projects also get culled, which reduces overall organizational performance. Consider an extreme case in which there are many good projects available but few resources (low *resource load*). To reach optimum organizational performance, many of the projects should be staffed by only a single worker. In this case, the effect of *rule of* turns negative because it prevents resources from being spread effectively over many projects. Thus, a higher *rule of* helps reduce errors of commission when resource constraints are not an issue, but it can turn negative by increasing errors of omission when resources must be spread thinly.

### Minimum Threshold for Profits

Setting a minimum profit threshold or payback time, a common practice in firms (Goold and Campbell, 1987), is like the *rule of* condition in our model. Each worker can start or join a project only if its profit potential is greater than or equal to a minimum value, here a variable called *threshold*. Figure 14, panel c) outlines the results for *threshold* = 0 and *threshold* = 2. While the effect of the performance threshold significantly increases performance for open allocation, it reduces the performance of centralized allocation, where a single manager makes allocation decisions subject to the same minimum value condition. The effects become more pronounced as *resource load* increases. The boost in performance for open allocation occurs because it takes a more erroneous estimate to decide to join the project. The condition limits overcrowding of workers and selection of bad projects. Because there is only one decision maker in centralized allocation, overcrowding is not an issue. However, by setting *threshold* > 0, we reduce the pool of projects the decision maker can consider. With many resources to spare, it makes sense to launch even the smallest value-positive projects. Therefore, adding a profit threshold condition for centralized allocation tends to reduce overall performance irrespective of resource load. As per our earlier discussion, when resources are abundant, all value-positive projects are also optimal.

### Requiring Approval from a Superior to Launch a Project

In the final manipulation, we allow employees under open allocation to launch new projects only after they receive their superior's approval. Foss (2003) documents the implementation of a "spaghetti organization" within Oticon, in which employees were given the freedom to launch their own and join existing projects, but senior managers retained their right to veto projects. We operationalize this policy by adding a manager with erroneous estimates of the projects who

blocks the allocation of employees to a project she considers value-negative. Figure 14, panel d) presents the effects of introducing this manipulation. While performance of the new open allocation with managerial intervention is similar to that of “pure” open allocation when the *resource load* is low, adding managerial intervention does offer a boost in performance as the *resource load* increases. As expected, the manager ensures that wildly optimistic employees who evaluate bad projects as good ones do not launch and assign themselves to such projects, bringing down the overall organizational performance. This is the case particularly when the number of available projects is small relative to human capital, and thus the risk of overstaffing is the greatest. When *resource load* is low, however, there is no effect, as employees can easily find a new opportunity when their current favorite is vetoed.

#### Robustness checks

We ran several robustness checks. To confirm the validity of our measure of *resource load*, we ran our model for different values of number of workers and different arrival rates of projects per round. This did not qualitatively impact the results. We also examined different values of  $\sigma$  (estimation error), different types of distributions of project values, and found that our results were consistent with the base model. In addition, as expected, when we set the estimation error to zero, open allocation and centralized allocation achieved the same level of performance, which equaled that of the optimal allocation for all values of  $r$ . In other words, with no estimation errors, the choice of an organizational form does not matter from the perspective of project selection and evaluation. We also examined higher values of *maxswaps*, *rule of*, and differences in screening ability, and the results held.

## **Discussion**

In this chapter, we focus on a specific role of organizational structure: the allocation of decision rights related to project evaluation and human capital allocation. We consider two opposing modes of addressing this dual challenge: open allocation and centralized allocation. The key insight from our analysis is that the balance between resources and opportunities is an important factor to consider when choosing between open and centralized resource allocation regimes. Whereas open allocation performs better in environments rich in opportunities relative to available resources, centralized allocation performs better when the environment is less munificent; that is, when opportunities are scarce relative to available resources.

We also find that when the resources available for allocation are limited, looking only at the efficacy of project screening yields an incomplete picture of how the two organizational forms function. Under resource constraints, one must account not only for project selection errors but also for resource allocation errors. Furthermore, selection-related commission errors arise not only from choosing value-negative projects but also from choosing value-positive but not optimal projects. Commission errors also arise from assigning excessive resources to optimal projects and assigning resources to value-positive but not optimal projects, as well as value-negative projects. Each of these errors creates different levels of inefficiencies, depending on the extent to which the human capital available fall short given the size and number of available opportunities.

Our study offers several practical implications. First, we highlight the dilemma that firms (particularly knowledge-intensive firms) face because they depend on human capital: the struggle to strike a balance between tightly controlling and directing employees and allowing them to direct their own efforts. Second, we conclude that organizational design choices are

dictated not only by size or span of control, but also by the relative availability of human capital and opportunities. Under certain conditions, centralized allocation could be a better-performing, more suitable choice than open allocation. Lastly, we argue that under certain conditions organizations can mitigate the weaknesses of open allocation by implementing certain elements of centralized allocation, like mandating a minimum number of workers per project, imposing a profit threshold, or requiring managerial approval. Employing such policy levers could mitigate the need to change the organizational structure completely, an exercise fraught with risk of failure (Foss, 2003).

Our analysis suggests that as a firm grows or its available opportunities diminish, it may no longer be feasible to maintain open allocation. In fact, a company's very success as a decentralized organization may eventually lead it to replace open allocation with some form of hierarchy. Anecdotal evidence supports this observation. Small technology startups with few employees often use open allocation to match human capital to projects. Many of these firms, however, face pressures to adopt a more hierarchical structure as the number of employees grows. These pressures are particularly salient in the case of GitHub's transition from a decentralized "boss-less" structure to a hierarchy (Burton *et al.*, 2017). Although it is customary to attribute the pressure to adopt hierarchy to the increasing span of control and communication difficulties (Chandler, 1962; Galbraith, 1977), our model suggests a new mechanism which can contribute to this outcome: the compounding of independent evaluators' errors.

Our findings can also be applied to changes in organizational forms throughout an industry's trajectory. The literature on industry evolution generally assumes that an industry experiences distinct stages (Agarwal and Tripsas, 2008). After the initial burst of opportunities, when new applications and market segments are being discovered, the market's growth slows,

and readily available opportunities for growth diminish. Our model thus suggests that the rationale for choosing open allocation over centralized allocation changes with the industry's growth phase, even if we keep the size of the firm constant. Early in the industry lifecycle, the decentralized form is more efficacious, as the high number of available opportunities limits the risk of commission errors. Later in the industry lifecycle, however, as the number of opportunities shrinks, centralized allocation becomes a better alternative, because it helps avoid the winner's curse produced by self-allocating employees.

At the same time, our results do not necessarily suggest that large organizations should avoid open allocation completely. While maintaining a hierarchical structure, large organizations can 1) use open allocation selectively, either in separate units in which the available resources and opportunities are favorable, or 2) allow employees to dedicate some fraction of their time to pursue projects of their choosing, or 3) regulate screening and allocation within open allocation through the use of policy levers that alleviate some of the pathologies of open allocation. However, while prior literature suggests that organizations should use polyarchy for exploratory activities and keep hierarchy to pursue exploitation (e.g., Knudsen and Levinthal, 2007; Csaszar, 2013), we argue that this may not be universally true. Organization designers should consider the amount of human capital dedicated to exploratory efforts relative to that required. Overstaffed skunk works that embrace employee autonomy and open allocation may in fact be less beneficial for the firm than centralized units when pursuing exploration.

From a theoretical perspective, we use a formal model to examine the phenomenon of open allocation of human capital, precisely define the underlying mechanisms, and explore the boundary conditions. The present study extends the literature on project screening in hierarchies and polyarchies (e.g., Sah and Stiglitz, 1986, 1988; Christensen and Knudsen, 2010; Csaszar,

2013) by considering the subsequent process of resource allocation (Noda and Bower, 1996; Bower and Gilbert, 2005; Bardolet, Fox, and Lovallo, 2011). Evaluation of opportunities and allocation of resources go hand-in-hand, as both processes have information processing at their core (Christensen and Knudsen, 2010; Knudsen and Levinthal, 2007; Csaszar and Eggers, 2013). This study addresses this previously identified gap in the literature (Csaszar, 2013). More specifically, we focus on how the choice of allocation mode mediates the evaluation of available opportunities and the subsequent allocation of human capital, a question that is especially important to knowledge-intensive, human-capital-rich firms (Coff, 1997; Lee and Edmondson, 2017). We highlight the importance of balance between the firm's resources and opportunities by demonstrating how organizational structure (centralized and open allocation) affects the joint process of project evaluation and resource allocation.

Having established the mechanisms behind the advantages of open allocation over centralized allocation under certain conditions, we explain the rationale behind four policy levers that can be employed under open allocation: 1) allowing employees to change the project on which they are working, 2) imposing a condition on the minimum number of employees required to initiate a project, 3) instituting a threshold for selecting projects, and 4) introducing a mid-level manager who approves or rejects projects selected by workers. Our analysis reveals that the efficacy of these policy levers also depends on the ratio between resources and opportunities.

Because our model design choices are driven by simplicity, we make certain assumptions in the model setup and abstract away certain real-world features of open and centralized allocation. While this is an inevitable tradeoff between parsimony and external validity that each modeler has to make (Page, 2018; Knudsen, Levinthal, and Puranam, 2019), it is important to highlight these assumptions and discuss potential extensions of our model.



First, our model assumes that the incentives of individual workers and the organization are aligned. While there could be a difference between what *really* is a good project and what a worker *perceives* to be a good project, we assume that this difference is because of evaluation error rather than the worker's opportunism or maliciousness. This is a common assumption in the economic theory of teams (Marschak and Radner, 1972; Van Zandt, 1999), and it stems from the study's focus on information processing rather than incentive alignment. While this assumption is consistent with the literature on project screening (Sah and Stiglitz, 1986), we can foresee scenarios in which workers' preferences and the firm's goals are misaligned. Workers may seek to join only the projects they find attractive for personal reasons, which could be detrimental to the firm's overall performance.

Second, following the literature on screening, we abstract away the effect of increased autonomy on employee motivation and creativity. Studies of employee motivation have found that having control over one's job increases effort and dedication, which in turn may lead one to generate more ideas (Herzberg, 1966; Hackman and Oldham, 1976; Lee and Edmondson, 2017). Such an effect could reduce the effects of overcrowding in boss-less organizations by generating more opportunities and could be an interesting extension of the current model.

Third, we assume that each worker can evaluate all existing projects. While this may be a plausible assumption for a relatively small number of projects, we also relax this assumption in the additional analyses shown in appendix C, where we model cases in which each worker can evaluate only a subset of all available projects. While the results are qualitatively similar, open allocation receives an additional boost by allowing workers to collectively evaluate and pursue more projects than centralized allocation.

Fourth, following the literature, we assume homogenous screening ability. However, one could argue that managers are promoted precisely because of their superior skills. A manager could, therefore, either have more precise estimates of projects' true values or be able to evaluate more projects than an average worker. We discuss and examine relaxing the homogeneity assumption in appendix C. In short, the results show that the central manager should be significantly better than an average employee to outweigh the advantages of open allocation.

Fifth, all projects in our model are similar in type and differ only in terms of payoff; we also assume that all workers have homogenous skills but may have different evaluation abilities. This assumption is common in prior models from the screening literature and stems from the primary focus on the role of structure (Knudsen and Levinthal, 2007). One could, however, consider a case where workers have idiosyncratic skills. The question of matching a particular worker to an appropriate project would be an important and interesting extension to our study.

Finally, we assume that each project lasts only one period and that a project will be completed, regardless of the number of employees working on it. It could be valuable to remove this simplification to examine the effects that open allocation policies may have on project completion rates and time.

## **Conclusion**

The present chapter picks up the suggestion that we can use existing theories to study novel forms of organizations (Puranam *et al.*, 2014). We extend the literatures on human capital allocation and project screening by highlighting the importance of the balance between the number of opportunities an environment presents and the amount of human capital available. Examining the results of an agent-based model, we demonstrate that open allocation performs better when available human capital is constrained. Centralized allocation, on the other hand,

allows organizations to avoid overcrowding; that is, when too many employees pursue a small number of good opportunities.

While the economic theory of teams connects the choice of organizational form to the characteristics of the firm's task environment, our model suggests that the rationale for using open versus centralized allocation will change with the relative size of the firm's resources available to pursue existing opportunities. Even if the number of opportunities the environment generates remains constant and the increasing span of control and increasing communication challenges do not yet demand change, an organization designer may decide to migrate from open allocation to centralized allocation as the organization grows.

With open allocation becoming popular among startups and established firms alike, further theorizing on the various features of such firms is needed to identify their benefits and shortcomings. By illustrating the importance of the ratio of human capital to opportunities, we hope to contribute to the literature on human capital allocation and on novel organizational designs.

## **Chapter V: Discussion**

Through this dissertation, I aim to uncover the strategic impact of new forms of organizing human capital, and automation for knowledge-based work. From the perspectives of both theory and phenomenon alike, significant exploration remains to be accomplished in these areas. I believe that my dissertation provides an exciting starting point for further research, as well as important implications for practitioners. From the point of view of theory, I aim to contribute to the exciting new field of how the use of automation shapes organizational outcomes, especially for knowledge-intensive tasks (Autor 2015). At the same time, as artificial intelligence based tools and algorithms are increasingly being used to automate knowledge work (Agrawal, Gans and Goldfarb, 2018), identifying the best opportunities from those available, and allocating human capital to them achieves increased importance. As such, firms are increasingly experimenting with novel organizing forms to achieve better performance under varying environments (Puranam, Alexy and Reitzig, 2014).

Broadly, this dissertation takes the view that knowledge work can either be carried out through decentralized or centralized organizational forms, or by employing the use of algorithms. Firstly, I posit that algorithmic automation fosters structured coordination over knowledge work, leading to conceptually narrower knowledge being produced and integrated by adopting organizations. Secondly, if novel organizational configurations (such as decentralized, bossless forms) are used, their efficacy over traditional hierarchical forms is determined by the balance of available human resources and opportunities. Taken together, this dissertation also offers

practical implications for managers when choosing between hierarchy, decentralized forms or algorithms to organize knowledge work.

Chapters II and III explore the different impacts that automation could have on various aspects of knowledge production: not only might it change the kind of knowledge produced, it could also change *who* contributes, and *how frequently* they might do so. These are important considerations that significantly shape the direction and outcomes of innovation within organizations. In chapter II, I argue that automating the integration of new knowledge contributions could lead to the adopting organization to integrate narrower knowledge contributions. This occurs because the use of automation hampers interpersonal coordination among organizational members. Thus, I argue that automation fosters structured co-ordination and stymies unstructured co-ordination, which in turn hinders the development of shared knowledge. Therefore, organizations tend to integrate new knowledge that is conceptually narrow. This chapter addresses a nascent literature on how automation shapes knowledge production (Furman and Teodoridis, 2020). In chapter III, we turn our attention to how algorithms might shape women's participation within collaborative communities. We find that when algorithms are used for evaluating knowledge contributions, more women contribute to such communities, and they submit more contributions on average. This is because women's contributions are more likely to be accepted, and less likely to be rejected as compared to the case where humans carry out the evaluation. This chapter addresses another new stream of literature on the implications of algorithms on inclusion and fairness (Cowgill and Tucker, 2020). While previous work tends to look at whether algorithms are substitutes or complements to humans, this dissertation takes a slightly different view, and focuses on how automation shapes processes and outcomes relating to knowledge production.

In Chapter IV we develop an agent-based model to analyze the efficacy of open allocation, a novel form of organizing human capital within organizations. We demonstrate that open allocation performs better than centralized allocation when available human capital is constrained. Centralized allocation, on the other hand, allows organizations to avoid overcrowding; that is, when too many employees pursue a small number of good opportunities.

### **Future Directions**

This dissertation is not without limitations. For instance, the sample used in chapters II and III consists of only Java projects, which I try to address by using a separate sample consisting of Javascript projects for affirming the robustness of the results. A better test would be to use projects drawn from various languages. Secondly, it would be worthwhile to provide specific evidence of the change in coordination patterns among members due to automation, and this can be done by analyzing conversations on discussion forums used for exchanging feedback on code contributions. Finally, an interesting extension of this context would study how individual project members use the extra time liberated by the automation of tasks. More broadly, it remains to be seen as to how this theory could extend to other, more traditional organizations. For instance, increased automation could differently affect more traditional knowledge intensive settings, such as research centers or hospitals, where face-to-face interactions between individuals are common. Additionally, extending the agent-based model in chapter IV to incorporate matching of individual skills to specific projects, as well as giving the agents an ability to work on multiple projects would provide some interesting insights into the phenomenon of open allocation.

## **Conclusion**

In conclusion, this dissertation extends our understanding of how different forms of organizing knowledge work (either by automating it, or arranging human resources in different structural forms) can shape the processes and outcomes associated with it. The advent of artificial intelligence and automation of knowledge work has led firms to experiment with different, more effective ways of organizing employees. This dissertation thus serves as a starting point for further enquiry into this exciting and emerging area of inquiry.

## APPENDIX

### Appendix A: Robustness Checks for Chapter II

The following robustness checks are used to ascertain if the results obtained still hold when all contributions in a project for a given period of four weeks are considered together. In the dataset, the treated and control projects are matched on age, and number of contributions. I calculated a ‘consolidated’ Herfindahl for each project where I used all updates made to a project across all files in a given four-week period.

DV = Month Level Herfindahl Index

	<b>Model 1: Std Errors clustered at project level</b>	<b>Model 2: Project FE only, clustered errors</b>	<b>Model 3: Project and Period FE, clustered errors</b>
TREATED = 1	-0.009 (0.00)	0.000 (.)	0.000 (.)
POST = 1	0.071*** (0.00)	0.076*** (0.00)	0.030*** (0.00)
TREATED = 1 # POST=1	0.032*** (0.01)	0.053*** (0.00)	0.052*** (0.00)
Project Fixed Effects	No	No	Yes
Time Fixed Effects	No	No	No
constant	0.321*** (0.00)	0.310*** (0.00)	0.946*** (0.01)
R-squared	0.017	0.265	0.270

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table A1. Herfindahl Index calculated for a four week period



Results are robust for this newly calculated dependent variable (table A1). The coefficient on the interaction of treated\_flag and prepost indicator is strongly significant, for all econometric specifications, even when fixed effects are introduced. On plotting the coefficients (figure A1), we can see that although pre-treatment trends for the treatment group diverge slightly from those for the control group, there is a clear jump for the former after treatment. This is expected because these checks are carried out on the sample of projects which self-select into the automation regime.

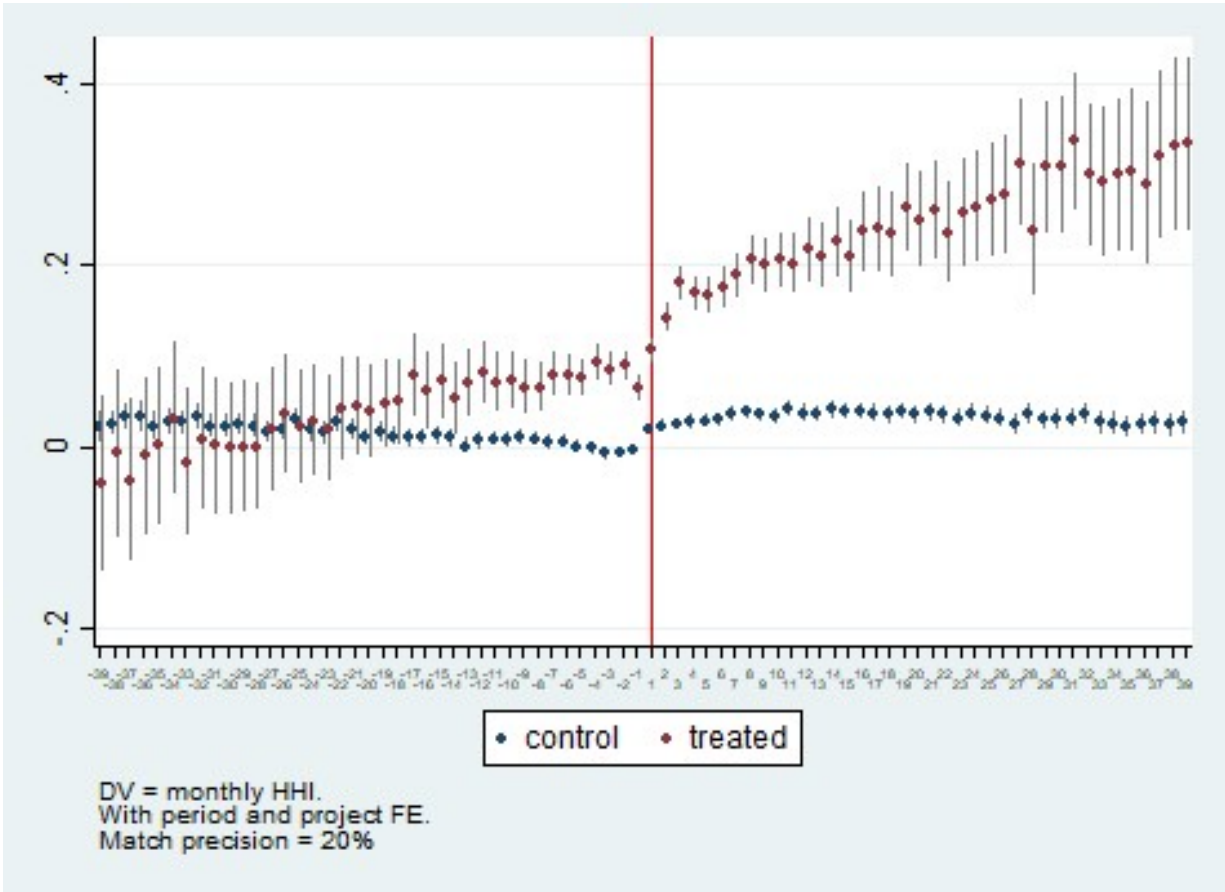


Table A1. Month level HHI over 320 weeks; Travis adoption occurs at t=0, corresponding to red vertical line.

## Appendix B: Additional Analyses for Chapter III

Female PR Share (collapsed observations)

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
POST=1	0.008 <sup>***</sup> (0.00)	0.008 <sup>**</sup> (0.00)	0.003 (0.00)	0.003 (0.00)
TREATED=1	-0.007 <sup>***</sup> (0.00)	-0.007 (0.01)	0.000 (.)	0.000 (.)
TREATED=1* POST=1	-0.005 (0.00)	-0.005 (0.01)	0.002 (0.01)	0.002 (0.01)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.036 <sup>***</sup> (0.00)	0.036 <sup>***</sup> (0.00)	0.033 <sup>***</sup> (0.00)	0.031 <sup>***</sup> (0.00)
R-squared	0.002	0.002	0.777	0.777

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table B1: Share of All Pull Requests Made by Female Programmers

Male PR Share (collapsed observations)

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
POST=1	-0.028*** (0.00)	-0.028*** (0.01)	-0.029*** (0.01)	-0.031*** (0.01)
TREATED=1	0.092*** (0.00)	0.092*** (0.01)	0.000 (.)	0.000 (.)
TREATED=1* POST=1	-0.022*** (0.01)	-0.022 (0.01)	-0.007 (0.02)	-0.005 (0.02)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.678*** (0.00)	0.678*** (0.01)	0.721*** (0.00)	0.764*** (0.03)
R-squared	0.015	0.015	0.803	0.803

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table B2: Share of All Pull Requests Made by Male Programmers

Neutral PR Share (collapsed observations)

	<b>Model 1: Basic</b>	<b>Model 2: Clustered Errors</b>	<b>Model 3: Project FE, Clustered errors</b>	<b>Model 4: Project FE, Time FE, Clustered errors</b>
POST=1	0.021 <sup>***</sup> (0.00)	0.021 <sup>**</sup> (0.01)	0.026 <sup>***</sup> (0.01)	0.028 <sup>***</sup> (0.01)
TREATED=1	-0.085 <sup>***</sup> (0.00)	-0.085 <sup>***</sup> (0.01)	0.000 (.)	0.000 (.)
TREATED=1* POST=1	0.027 <sup>***</sup> (0.01)	0.027 <sup>*</sup> (0.01)	0.005 (0.01)	0.003 (0.01)
Project Fixed Effects	No	No	Yes	Yes
Time Fixed Effects	No	No	No	Yes
constant	0.287 <sup>***</sup> (0.00)	0.287 <sup>***</sup> (0.01)	0.246 <sup>***</sup> (0.00)	0.205 <sup>***</sup> (0.04)
R-squared	0.013	0.013	0.806	0.806

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

Table B3: Share of All Pull Requests Made by Neutral Programmers

The following figures are created from period-wise data.

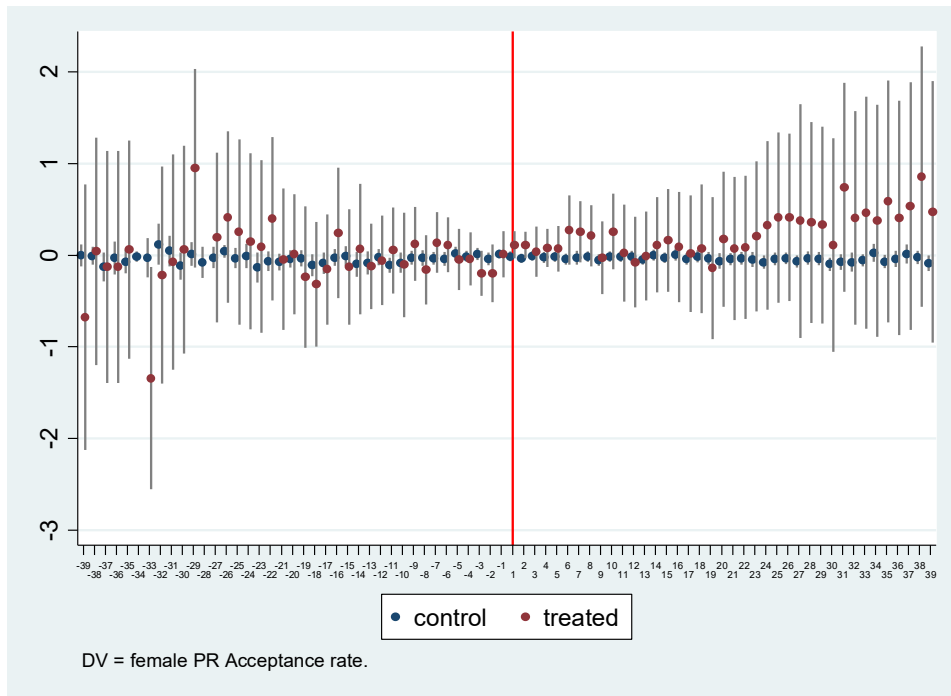


Figure B1: Acceptance rate for female programmers over 320 weeks. Travis adoption occurs at  $t=0$ , corresponding to the red vertical line.

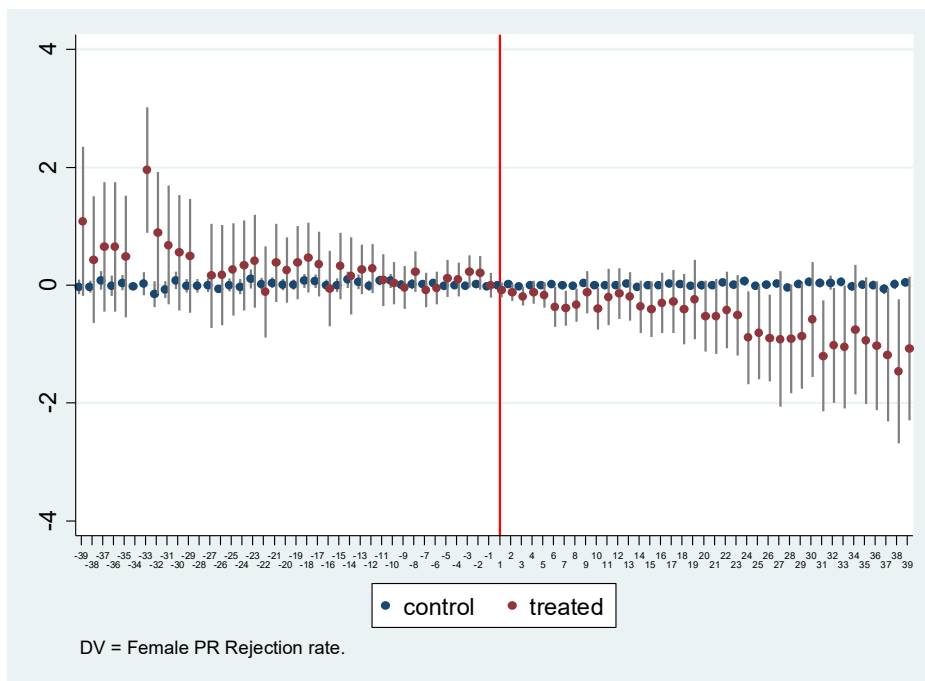


Figure B2: Rejection rate for female programmers over 320 weeks. Travis adoption occurs at  $t=0$ , corresponding to the red vertical line.

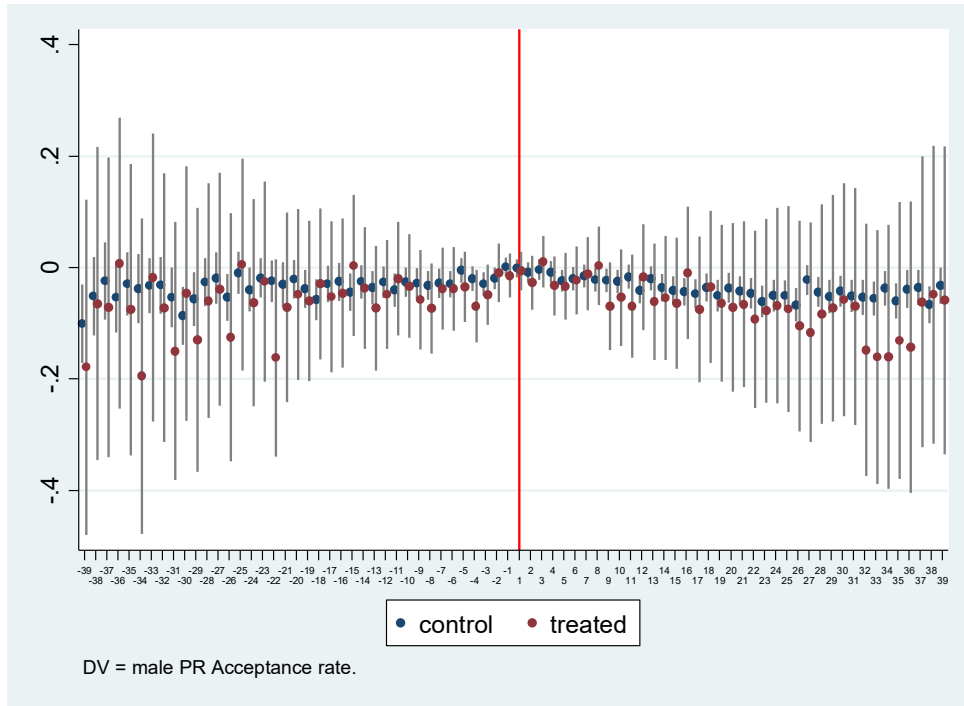


Figure B3: Acceptance rate for male programmers over 320 weeks. Travis adoption occurs at  $t=0$ , corresponding to the red vertical line.




Figure B4: Rejection Rate for Male Programmers over 320 weeks. Travis adoption occurs at  $t=0$ , corresponding to the red vertical line.

## Appendix C: Additional Robustness Analyses for Chapter IV

This appendix contains the following robustness checks referenced in the chapter: 1) the limits to an individual's attention, and 2) superiority of the manager in evaluating projects.

### 1. Alternative ways of defining bounded rationality (attention) of agents

In this set of analyses we modify the way we model bounded rationality for the agents in the model by defining it in terms of the proportion of projects that each is able to consider for evaluation.

In the model presented in the chapter, we assume that the evaluating manager and workers are similar to each other in terms of their evaluation function, and both are able to consider all available projects to evaluate and choose from. In its effort to study the role of structure, the literature on screening (Sah and Stiglitz, 1986, 1988; Knudsen and Levinthal, 2007) abstracts from such concepts as motivation or attention. Here we relax this assumption and examine the effects of making agents in our model bounded with respect of the maximum number of projects that they can review per round.

Let us consider an organization operating in an environment, where there are  $r$  projects available each round to pursue. Each agent  $i$  can learn about  $r_i$  opportunities per round, where  $r_i$  reflects the degree of cognitive boundedness such that  $r_i \leq r$ . We keep  $r_i$  the same for the manager in the centralized allocation and the employees in the open allocation. In other words, both are equally boundedly rational. Running the model with  $r_i < r$  reveals another benefit of open allocation: it allows the organization to “see” more opportunities, compared with centralized allocation. Figure A1 shows the results of the model where  $r_i = 8$ . We observe that while centralized allocation immediately suffers from reduced attention, open allocation

manages to roughly maintain its level of performance compared with  $r_i = r$  (the base case scenario in the main chapter).

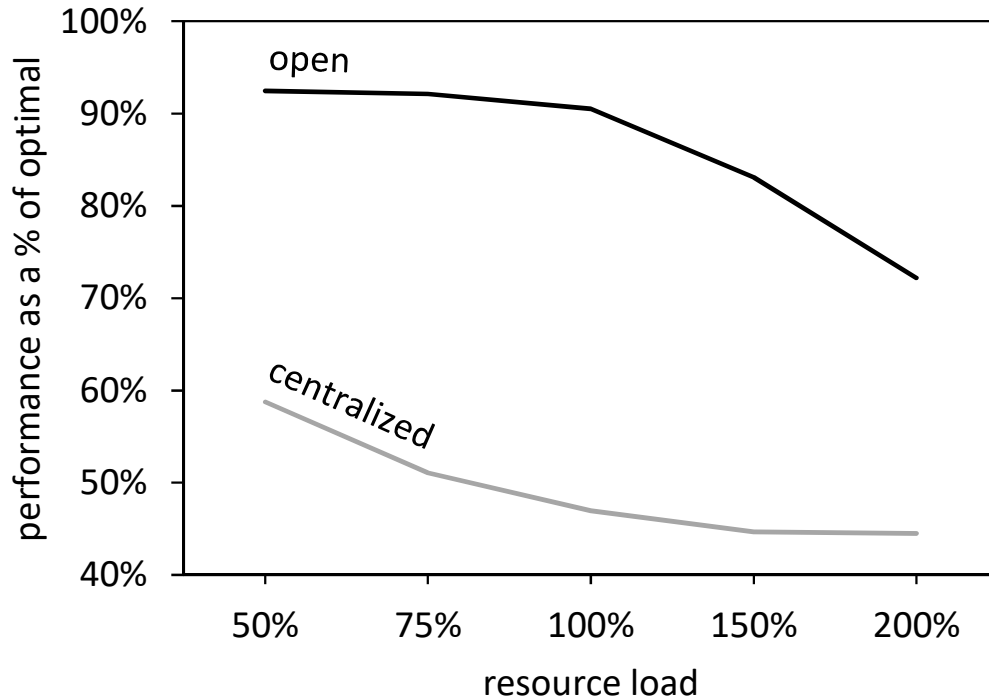


Figure C1: Organizational performance with reduced level of attention ( $r_i = 8$ )

## 2. Superior skills of the manager

When it comes to starting assumptions, the model presented in the study is anchored firmly in the screening literature (Sah and Stiglitz, 1986, 1988; Csaszar, 2013) and the economic theory of teams (Marschak and Radner, 1972) more broadly. This helps us better situate our results to existing literature and provides a firmer foundation for our model and its assumptions. In particular, we draw upon some fundamental behavioral assumptions relating to knowledge, attention and motivation of agents that are used for building models within this literature. With



respect to knowledge, as mentioned earlier, all agents (managers and subordinates) are assumed to possess the same average ability, since the goal of that literature is to analyze the effects of organizational structure on the efficacy of screening, as it is in our case (Knudsen and Levinthal, 2007).

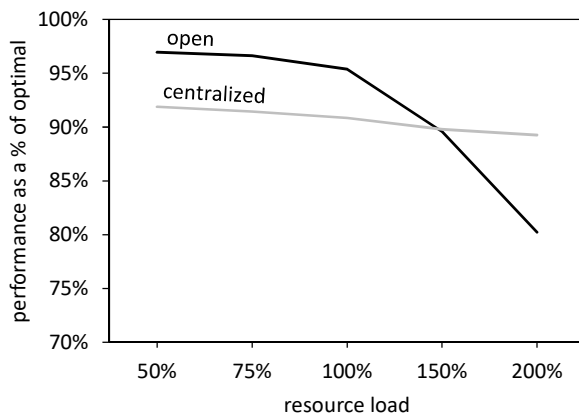
This assumption is also plausible for another reason. It is not immediately clear whether a superior would indeed be better or worse positioned to evaluate projects' potential. On the one hand, we can envision that a superior might have more experience evaluating the projects or be promoted to the position precisely for such good performance. On the other hand, a superior may not be as familiar with the details of the technology or be as close to customers as her subordinates. Assuming homogeneous skill leaves this question open to further study and allows us to focus on the effects of organizational structure. Thus, the main model in the chapter is close in its approach to that of March and March (1977), where an executive, while occupying a position of power and influence, is indistinguishable from her subordinates in terms of skill.

However, a literature stemming from work in transaction cost economics (Williamson, 1975), organizational economics (Garicano, 2000; Rotemberg and Saloner, 2000) and resource allocation process (Bower, Doz and Gilbert, 2005) assumes that managers might have been put in their role precisely because of their high competence or greater experience. By virtue of their superior skill and their vantage point within the firm, they possess a broader view of all available opportunities and are therefore able to consider more projects and are more precise in their estimates. We examine these two possibilities in turn.

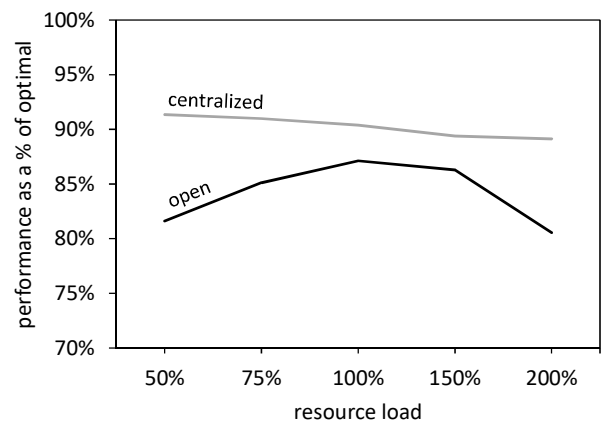
### Screening Breadth

We conduct experiments where we keep the screening competency equal between the manager and workers, but allow the manager to see and evaluate all available projects; at the same time,

the workers were able to see only a half (8 out of 16 ) or a quarter (4 out of 16) of all the projects. Results show that when employees can only see half the projects compared to the manager, open allocation can still outperform a single manager as collectively workers see nearly all of the projects. Here results deviate very little from those presented in the chapter. Only when workers see four times fewer projects than the manager, the performance of open allocation drops below that of centralized allocation. Figure A2 illustrates the results for both discussed cases. This further illustrates the additional power of open allocation and suggests that a manager would have to be significantly more skilled to outweigh the advantage of open allocation when it comes to the breadth of screening.



a) employees see only 50% of projects



b) employees see only 25% of projects

Figure C2: Performance when each worker can see only 50% (8 out of 16) of projects or 25% (4 out of 16), while the manager sees all of the projects (100%)

### Screening Accuracy

Another possibility is that a manager is more accurate in evaluating the projects. In another variant of the base model, we relax our previous assumption that the manager and the workers

have the same screening accuracy. More specifically, we vary the standard error of true value estimation ( $\sigma$ ) for workers and the manager. The results presented in figure A3 demonstrate that at least in the context of our model, the manager should be significantly more accurate than workers in order to perform at a comparable level as open allocation. For example, for low *resource load*, the manager should be twice as precise as her employees (see open allocation at  $\sigma = 2$  and centralized allocation at  $\sigma = 1$ ).

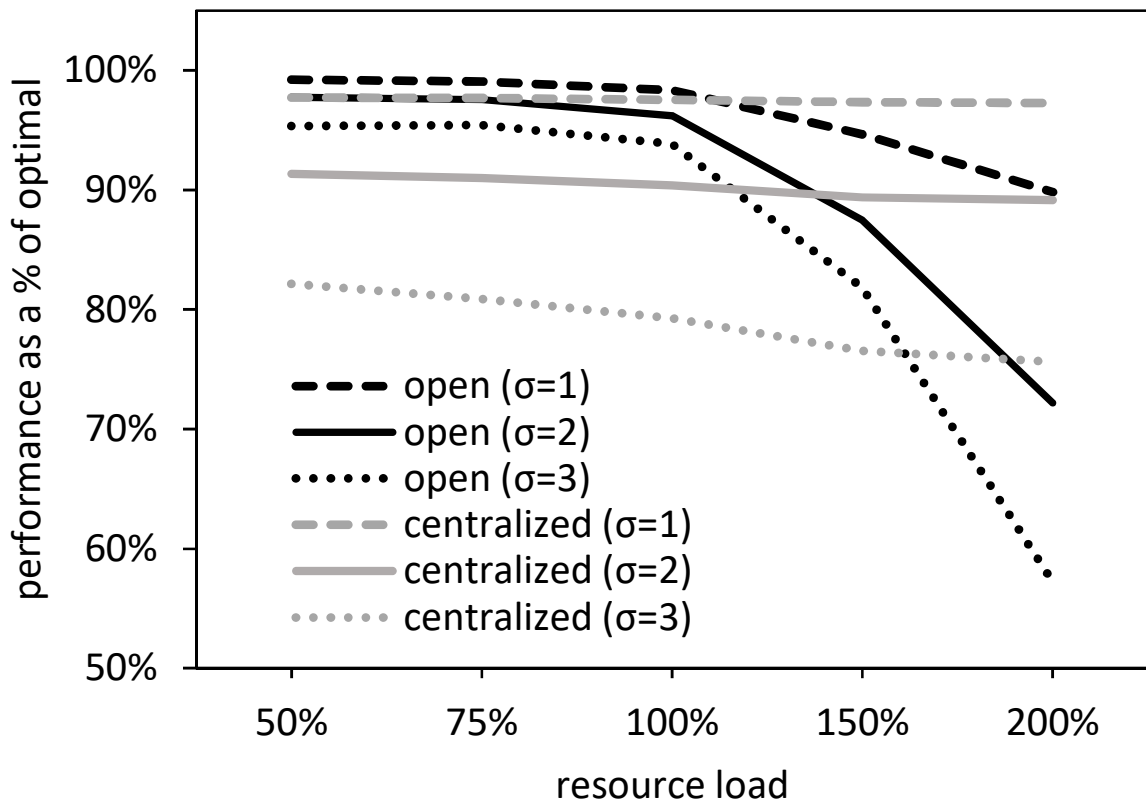


Figure C3: Comparison of performance for different levels of  $\sigma$

## **BIBLIOGRAPHY**

- Abraham M. 2020. Gender-role incongruity and audience-based gender bias: an examination of networking among entrepreneurs. *Administrative Science Quarterly* **65**(1): 151–180
- Acemoglu D, Restrepo P. 2017. Low-skill and high-skill automation. *NBER Working Paper no. 24119*.
- Acemoglu D, Restrepo P. 2018. Artificial intelligence, automation and work. *NBER Working Paper No. 24196*.
- Agarwal R, Tripsas M. 2008. *Technology and industry evolution. The Handbook of Technology and Innovation Management* **1**: 1–55. John Wiley & Sons, Inc., Hoboken, NJ, USA.
- Aghion P, Bloom N, Van Reenen J. 2014. Incomplete contracts and the internal organization of firms. *The Journal of Law, Economics, & Organization* **30**(suppl\_1): i37–i63.
- Aghion P, Tirole J. 1997. Formal and real authority in organizations. *Journal of Political Economy* **105**(1): 1–29.
- Agrawal A, Gans J, Goldfarb A. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Press: Boston, MA.
- Agrawal AK, Gans JS, Goldfarb A. 2018b. *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Review Press, Boston, MA, USA.
- Agrawal AK, McHale J, Oettl A. 2018. Finding needles in haystacks: artificial intelligence and recombinant growth. NBER Working Paper No. 24541.
- Aldrich H, Herker D. 1977. Boundary spanning roles and organization structure. *Academy of Management Review* **2**(2): 217–230.

- Anthony, C., 2018. To question or accept? How status differences influence responses to new epistemic technologies in knowledge work. *Academy of Management Review* **43**(4), pp.661-679
- Aron R, Dutta S, Janakiraman R, Pathak PA. 2011. The impact of automation of systems on medical errors: evidence from field research. *Information Systems Research* **22**(3):429–46.
- Athey S, Roberts J. 2001. Organizational design: Decision rights and incentive contracts. *American Economic Review* **91**(2): 200-205.
- Augier M, March JG, Marshall AW. 2015. The flaring of intellectual outliers: An organizational interpretation of the generation of novelty in the RAND Corporation. *Organization Science* **26**(4): 1140–1161.
- Austin JR. 2003. Transactive memory in organizational groups: The effects of content, consensus, specialization, and accuracy on group performance. *Journal of Applied Psychology* **88**(5): 866–878.
- Autor DH. 2015. Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives* **29**(3): 3–30
- Autor DH. 2015. Why are there still so many jobs? The history and future of workplace automation. *Journal of Economic Perspectives* **29**(3): 3–30.
- Azoulay P, Stuart T, Wang Y. 2014. Matthew: effect or fable? *Management Science* **60**(1): 92–109
- Bakker RM, Josefy, M. 2018. More than just a number? The conceptualization and measurement of firm age in an era of temporary organizations. *Academy of Management Annals* **12**(2): 510–536.

- Baldwin CY and Clark KB. 2000. Design rules: The power of modularity (Vol. 1). MIT Press, Cambridge, MA, USA.
- Barach MA, Leung MD, Lu S. 2017. Foot in the door or head and shoulders above the rest? How algorithmic recommendations affect employer search and hiring behavior. *Working paper*.
- Bardolet D, Fox CR, Lovallo D. 2011. Corporate capital allocation: A behavioral perspective. *Strategic Management Journal* **32**(13): 1465–1483.
- Barney J. 1991. Firm resources and sustained competitive advantage. *Journal of Management* **17**(1): 99–120.
- Baumann O, Schmidt J, Stieglitz N. 2019. Effective search in rugged performance landscapes: a review and outlook. *Journal of Management* **45**(1): 285–318.
- Beane M. 2018. Shadow learning: building robotic surgical skill when approved means fail. *Administrative Science Quarterly* (OnlineFirst).
- Bechky BA, Okhuysen GA. 2011. Expecting the unexpected? How SWAT officers and film crews handle surprises. *Academy of Management Journal*. **54**(2): 239-261.
- Belenzon S, Schankerman M. 2015. Motivation and sorting of human capital in open innovation. *Strategic Management Journal* **36**(6): 795–820.
- Ben-Menahem SM, von Krogh G, Erden Z, Schneider A. 2016. Coordinating knowledge creation in multidisciplinary teams: evidence from early-stage drug discovery. *Academy of Management Journal* **59**(4): 1308–1338.

- Bentley FS, Kehoe RR. *forthcoming*. Give them some slack—they're trying to change! The benefits of excess cash, excess employees, and increased human capital in the strategic change context. *Academy of Management Journal*.
- Berger J. 1977. *Status characteristics and social interaction: An expectation-states approach*. Greenwood Publishing Group, Santa Barbara, CA, USA
- Bertrand M, Duflo E, Mullainathan S. 2004. How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics* **119**(1): 249–275
- Bertrand M, Duflo E, Mullainathan S. 2004. How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics*. **119**(1): 249–275.
- Bertrand M, Mullainathan S. 2004. Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review* **94**(4): 991–1013
- Birley S, Westhead P. 1990. Growth and performance contrasts between 'types' of small firms. *Strategic Management Journal* **11**(7): 535–557.
- Boehm B, Abts C, Chulani S. 2000. Software development cost estimation approaches—A survey. *Annals of software engineering* **10**(1–4):177–205.
- Bohren JA, Imas A, Rosenberg M. 2019. The dynamics of discrimination: theory and evidence. *American Economic Review* **109**(10): 3395–3436
- Bostrom N. 2014. *Superintelligence: Paths, Dangers, Strategies*. Oxford University Press, Oxford, UK



- Botelho TL, Abraham M. 2017. Pursuing quality: how search costs and uncertainty magnify gender-based double standards in a multistage evaluation process. *Administrative Science Quarterly* **62**(4): 698–730
- Bower JL, Gilbert CG. eds. 2005. *From resource allocation to strategy*. Oxford University Press: Oxford, UK.
- Bower JL. 1970. *Managing the resource allocation process: A study of corporate planning and investment*. Harvard University Press, Boston, MA, USA.
- Brooks AW, Huang L, Kearney SW, Murray FE. 2014. Investors prefer entrepreneurial ventures pitched by attractive men. *Proceedings of the National Academy of Sciences* **111**(12): 4427–4431
- Brooks FP. 1995. *The Mythical Man-Month: Essays on Software Engineering, 20th Anniversary Edition*. Addison-Wesley, Reading, MA, USA.
- Brynjolfsson E, Hitt LM. 2000. Beyond computation: information technology, organizational transformation and business performance. *Journal of Economic Perspectives* **14**(4): 23–48.
- Brynjolfsson E, McAfee A. 2014. *The second machine age: Work, progress, and prosperity in a time of brilliant technologies*. WW Norton & Company, New York, USA.
- Burgelman RA. 1983. A process model of internal corporate venturing in the diversified major firm. *Administrative Science Quarterly* **28**(2):223–244.
- Burns TE, Stalker GM. 1961. *The Management of Innovation*. Tavistock: London, UK.

Burton RM, Håkonsson DD, Nickerson J, Puranam P, Workiewicz M, Zenger T. 2017. GitHub: exploring the space between boss-less and hierarchical forms of organizing. *Journal of Organization Design* **6**(10): 1–19.

Burton RM, Obel B. 1984. *Designing efficient organizations: Modelling and experimentation* (Vol. 7). North Holland, New York, NY, USA.

Campbell BA, Coff R, Kryscynski D. 2012. Rethinking sustained competitive advantage from human capital. *Academy of Management Review* **37**(3): 376–395.

Campion MA, Palmer DK and Campion JE. 1997. A review of structure in the selection interview. *Personnel Psychology* **50**(3): 655–702

Cardador MT. 2017. Promoted up but also out? The unintended consequences of increasing women's representation in managerial roles in engineering. *Organization Science* **28**(4): 597–617

Carlile PR. 2004. Transferring, translating, and transforming: an integrative framework for managing knowledge across boundaries. *Organization Science* **15**(5): 555–568.

Carnabuci G, Operti E. 2013. Where do firms' recombinant capabilities come from? Intraorganizational networks, knowledge, and firms' ability to innovate through technological recombination. *Strategic Management Journal* **34**(13): 1591–1613.

Carnahan S, Greenwood BN. 2018. Managers' political beliefs and gender inequality among subordinates: does his ideology matter more than hers? *Administrative Science Quarterly* **63**(2): 287–322

- Carnahan S, Greenwood BN. 2018. Managers' political beliefs and gender inequality among subordinates: does his ideology matter more than hers? *Administrative Science Quarterly* **63**(2): 287–322
- Carr N. 2015. *The glass cage: how our computers are changing us*. WW Norton & Company, New York, USA.
- Chandler A. 1962. *Strategy and structure: chapters in the history of the American industrial enterprise*. MIT Press, Cambridge, MA, USA.
- Chatain O, Meyer-Doyle P. 2017. Alleviating managerial dilemmas in human-capital-intensive firms through incentives: Evidence from M&A legal advisors. *Strategic Management Journal* **38**(2): 232–254.
- Chesbrough H. 2003. Open innovation: how companies actually do it. *Harvard Business Review* **81**(7): 12–14.
- Choudhury P, Starr E, Agarwal R. 2020. Machine learning and human capital complementarities: experimental evidence on bias mitigation. Forthcoming in *Strategic Management Journal*
- Choudhury P, Starr E, Agarwal, R. 2019. Machine learning and human capital complementarities: experimental evidence on bias mitigation. Available at SSRN: <https://ssrn-com.proxy.lib.umich.edu/abstract=3185022>
- Christensen M, Knudsen T. 2010. Design of decision-making organizations. *Management Science* **56**(1): 71–89.

- Claggett JL, Karahanna E. 2018. Unpacking the structure of coordination mechanisms and the role of relational coordination in an era of digitally mediated work processes. *Academy of Management Review* **43**(4): 704–722.
- Cockburn IM, Henderson R, Stern S. 2018. The impact of artificial intelligence on innovation, in: *Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, Chicago, IL, USA (forthcoming).
- Coen CA, Maritan CA. 2010. Investing in capabilities: the dynamics of resource allocation. *Organization Science* **22**(1): 99–117.
- Coff RW. 1997. Human assets and management dilemmas: coping with hazards on the road to resource-based theory. *Academy of Management Review* **22**(2): 374–402.
- Cohen MD, Bacdayan P. 1994. Organizational routines are stored as procedural memory: evidence from a laboratory study. *Organization Science* **5**(4): 554–568.
- Cohen WM, Levinthal DA. 1990. Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly* **35**(1):128–152.
- Conner KR, Prahalad CK. 1996. A resource-based theory of the firm: Knowledge versus opportunism. *Organization Science* **7**(5): 477–501.
- Correll S, Mackenzie L. 2016. To succeed in tech, women need more visibility. Harvard Business Review. Retrieved from <https://hbr.org/2016/09/to-succeed-in-tech-women-need-more-visibility> on March 30, 2020

- Correll SJ, Ridgeway CL, Delamater J. 2003. *Handbook of social psychology*: 29-51. Kluwer Academic Publishers, New York, NY, USA
- Cowgill B, Seamans R, Ziv O. 2017. The future of work. *State of the Field (Kauffman Foundation)* <https://www.kauffman.org/microsites/state-of-the-field/topics/future-of-work> Retrieved on February 1, 2019.
- Cowgill B, Tucker C. 2020. Algorithmic Fairness and Economics. *Working Paper*
- Cowgill B. 2017. Automating judgement and decision making: theory and evidence from resume screening. *Working Paper*
- Cowgill B. 2018. Bias and Productivity in Humans and Algorithms: Theory and Evidence from Resume Screening. *Working Paper*.
- Csaszar FA, Eggers JP. 2013. Organizational decision making: an information aggregation view. *Management Science* **59**(10): 2257–2277.
- Csaszar FA. 2013. An efficient frontier in organization design: organizational structure as a determinant of exploration and exploitation. *Organization Science* **24**(4): 1083–1101.
- Daft RL, Lewin AY. 1993. Where are the theories for the “new” organizational forms? An editorial essay. *Organization Science* **4**(4): i–vi.
- Dahlander L, O’Mahony S. 2011. Progressing to the center: Coordinating project work. *Organization Science* **22**(4): 961–979.
- Davis SM, Lawrence PR. 1977. *Matrix*. Addison-Wesley, Reading, MA, USA.

- de Figueiredo R, Rawley E, Rider CI. 2015. Why are firms rigid? A general framework and empirical tests. *Organization Science* **26**(5): 1502–1519.
- Dezsö CL, Ross DG. 2012. Does female representation in top management improve firm performance? A panel data investigation. *Strategic Management Journal* **33**(9):1072–1089
- Ding WW, Murray F, Stuart TE. 2006. Gender differences in patenting in the academic life sciences. *Science* **313**(5787): 665–667
- Dobrajska M, Billinger S and Karim S. 2015. Delegation within hierarchies: How information processing and knowledge characteristics influence the allocation of formal and real decision authority. *Organization Science* **26**(3): 687–704.
- Eagly AH, Karau SJ. 2002. Role congruity theory of prejudice toward female leaders. *Psychological Review* **109**(3): 573–598
- Eagly AH. 2013. *Sex differences in social behavior: A social-role interpretation*. Psychology Press, Oxford, UK
- Edmondson AC, Bohmer RM, Pisano GP. 2001. Disrupted routines: team learning and new technology implementation in hospitals. *Administrative Science Quarterly* **46**(4): 685–716.
- Eisenhardt KM, Martin JA. 2000. Dynamic capabilities: what are they? *Strategic Management Journal* **21**(10/11): 1105–1121.
- Eisenhardt KM, Martin JA. 2000. Dynamic capabilities: what are they? *Strategic Management Journal*. **21**(10): 1105–1121.

- Ethiraj SK, Garg P. 2011. The division of gains from complementarities in human-capital-intensive activity. *Organization Science* **23**(3): 725–742.
- Ethiraj SK, Levinthal DA. 2004. Bounded rationality and the search for organizational architecture: an evolutionary perspective on the design of organizations and their evolvability. *Administrative Science Quarterly* **49**(3): 404–437.
- Faraj S, Sproull L. 2000. Coordinating expertise in software development teams. *Management Science* **46**(12): 1554–1568.
- Feldman MS, Rafaeli A. 2002. Organizational routines as sources of connections and understandings. *Journal of Management Studies* **39**(3): 309–331.
- Fleming L. 2001. Recombinant uncertainty in technological search. *Management Science* **47**(1): 117–132.
- Folta T, Helfat C, Karim S. 2016. Resource Redeployment and Corporate Strategy. *Advances in Strategic Management Volume 35*. Emerald Group Publishing, Bingley, UK.
- Ford, M., 2015. *Rise of the robots: technology and the threat of a jobless future*. Basic Books, New York, USA.
- Foss NJ, Frederiksen L, Rullani F. 2016. Problem-formulation and problem-solving in self-organized communities: How modes of communication shape project behaviors in the free open source software community. *Strategic Management Journal* **37**(13): 2589-2610
- Foss NJ. 2003. Selective intervention and internal hybrids: interpreting and learning from the rise and decline of the Oticon spaghetti organization. *Organization Science* **14**(3): 331–349.

Furman JL, Teodoridis F. 2018. Automation, Research Technology, and Researchers' Trajectories: Evidence from Computer Science and Electrical Engineering. Available at SSRN:

<https://doi.org/10.2139/ssrn.3285286>

Galbraith JR. 1977. *Organization design*. Addison-Wesley, Reading, MA, USA.

Gardner HK, Gino F, Staats BR. 2012. Dynamically integrating knowledge in teams: transforming resources into performance. *Academy of Management Journal* **55**(4): 998–1022.

Garicano L, Wu Y. 2012. Knowledge, communication, and organizational capabilities. *Organization Science* **23**(5):1382–1397

Garicano L. 2000. Hierarchies and the organization of knowledge in production. *Journal of Political Economy* **108**(5): 874–904.

Gertner R, Scharfstein D. 2013. Internal capital markets, in: *The Handbook of Organizational Economics*, Gibbons R, Roberts J. (eds.), Princeton University Press, Princeton, NJ, USA:677–701.

Goldin C, Rouse C. 2000. Orchestrating impartiality: the impact of “blind” auditions on female musicians. *American Economic Review* **90**(4): 715–741

Goldin C. 1986. The earnings gap between male and female workers: An historical perspective. *NBER Working Paper No. 1888*

Goold M, Campbell A. 1987. *Strategies and Styles: The Role of the Centre in Managing Diversified Corporations*. Basil Blackwell, Oxford, UK



- Gousios G. 2013. The ghtorent dataset and tool suite. *Proceedings of the 10th Working Conference on Mining Software Repositories, MSR '13*. IEEE Press, Piscataway, NJ, USA, pp. 233–236.
- Grant RM. 1996. Toward a knowledge-based theory of the firm. *Special Issue: Knowledge and the Firm: Strategic Management Journal* **17**: 109–122.
- Gulati R. 2018. Structure that's not stifling. *Harvard Business Review* **96**(3): 68–79.
- Hackman JR, Oldham GR. 1976. Motivation through the design of work: test of a theory. *Organizational Behavior and Human Performance* **16**(2): 250–279.
- Hamel G, Prahalad CK. 1996. Competing in the new economy: Managing out of bounds. *Strategic Management Journal* **17**(3): 237–242.
- Hamel G. 2011. First, let's fire all the managers. *Harvard Business Review*. **89**(12): 48-60.
- Hansen MT. 1999. The search-transfer problem: the role of weak ties in sharing knowledge across organization subunits. *Administrative Science Quarterly* **44**(1): 82–111.
- Hargadon A, Sutton RI. 1997. Technology brokering and innovation in a product development firm. *Administrative Science Quarterly* **42**(4): 716–749
- Hart O, Moore J. 2005. On the design of hierarchies: coordination versus specialization. *Journal of Political Economy* **113**(4): 675–702.
- Heaven D. 2018. AI peer reviewers unleashed to ease publishing grind. *Nature*
- Helfat CE, Martin J. 2015. Dynamic Managerial Capabilities. *The Oxford Handbook of Creativity, Innovation, and Entrepreneurship*. Oxford University Press, New York, NY, USA.

- Helfat CE, Peteraf MA. 2015. Managerial cognitive capabilities and the microfoundations of dynamic capabilities. *Strategic Management Journal* **36**(6): 831–850.
- Henderson R, Cockburn I. 1994. Measuring competence? Exploring firm effects in pharmaceutical research. *Strategic Management Journal* **15**(S1): 63–84.
- Henderson RM, Clark KB. 1990. Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly* **35**(1): 9–30.
- Hengel E. 2020. Publishing while female: Are women held to higher standards? Evidence from peer review. *Working Paper*
- Herzberg FI. 1966. *Work and the Nature of Man*. World Publishing: Cleveland, OH, USA.
- Hill RA, Dunbar RI. 2003. Social network size in humans. *Human Nature*, **14**(1): 53–72.
- Hilton M, Tunnell T, Huang K, Marinov D, Dig D. 2016. Usage, costs, and benefits of continuous integration in open-source projects, in: *2016 31st IEEE/ACM International Conference on Automated Software Engineering (ASE)*: 426–437 Presented at the 2016 31st IEEE/ACM International Conference on Automated Software Engineering (ASE).
- Hoopes DG, Postrel S. 1999. Shared knowledge, “glitches,” and product development performance. *Strategic Management Journal* **20**(9): 837–865.
- Hutchison-Krupat J, Kavadias S. 2015. Strategic resource allocation: top-down, bottom-up, and the value of strategic buckets. *Management Science* **61**(2): 391–412.
- Jirjahn U, Stephan G. 2004. Gender, piece rates and wages: evidence from matched employer–employee data. *Cambridge Journal of Economics* **28**(5): 683–704

- Katila R, Ahuja, G. 2002. Something old, something new: a longitudinal study of search behavior and new product introduction. *Academy of Management Journal* **45**(6): 1183–1194.
- Kellogg KC, Valentine MA, Christin A. 2019. Algorithms at work: the new contested terrain of control. *Academy of Management Annals* **14**(1): 366–410
- Kim S, Anand J. 2018. Knowledge complexity and the performance of inter-unit knowledge replication structures. *Strategic Management Journal* **39**(7): 1959–1989.
- Kleinberg J, Lakkaraju H, Leskovec J, Ludwig J, Mullainathan S. 2017. Human decisions and machine predictions. *Quarterly Journal of Economics* **133** (1): 237–293.
- Klingebiel R, Rammer C. 2014. Resource allocation strategy for innovation portfolio management. *Strategic Management Journal* **35**(2): 246–268.
- Knudsen T, Levinthal D, Puranam P. 2019. A model is a model. *Strategy Science* **4**(1): 1–3.
- Knudsen T, Levinthal DA. 2007. Two faces of search: alternative generation and alternative evaluation. *Organization Science* **18**(1): 39–54.
- Kogut B, Zander U. 1992. Knowledge of the firm, combinative capabilities, and the replication of technology. *Organization Science* **3**(3): 383–397.
- Kogut B, Zander U. 1996. What firms do? Coordination, identity, and learning. *Organization Science* **7**(5): 502–518.
- Kovács, B, Sharkey AJ. 2014. The paradox of publicity: How awards can negatively affect the evaluation of quality. *Administrative Science Quarterly* **59**(1): 1–33

- Lakhani KR and von Hippel E. 2003. How open source software works: “free” user-to-user assistance. *Research Policy* **32**(6): 923–943.
- Lambrecht A, Tucker C. 2019. Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of stem career ads. *Management Science* **65**(7): 2966–2981
- Lawrence PR, Lorsch JW. 1967. *Organization and Environment: Managing Differentiation and Integration*, Revised edition. Harvard Business Review Press: Boston, MA, USA.
- Lecher C. 2019. How amazon automatically tracks and fires warehouse workers for ‘productivity’. *The Verge*. Retrieved from <https://www.theverge.com/2019/4/25/18516004/amazon-warehouse-fulfillment-centers-productivity-firing-terminations> on March 30, 2020
- Lecuona JR, Reitzig M. 2014. Knowledge worth having in ‘excess’: The value of tacit and firm-specific human resource slack. *Strategic Management Journal* **35**(7): 954–973.
- Lee M, Huang L. 2018. Gender bias, social impact framing, and evaluation of entrepreneurial ventures. *Organization Science* **29**(1): 1–16
- Lee MY, Edmondson AC. 2017. Self-managing organizations: Exploring the limits of less-hierarchical organizing. *Research in Organizational Behavior* **37**: 35–58.
- Leonard-Barton D. 1992. Core capabilities and core rigidities: a paradox in managing new product development. *Strategic Management Journal* **13**: 111–125.

- Levashina J, Hartwell CJ, Morgeson FP, Campion MA. 2014. The structured employment interview: narrative and quantitative review of the research literature. *Personnel Psychology* 67(1):241–293
- Levinthal DA, Workiewicz M. 2018. When two bosses are better than one: nearly decomposable systems and organizational adaptation. *Organization Science* 29(2): 207–224.
- Levinthal DA, Wu B. 2010. Opportunity costs and non-scale free capabilities: profit maximization, corporate scope, and profit margins. *Strategic Management Journal* 31(7): 780–801.
- Levinthal DA. 2017. Resource allocation and firm boundaries. *Journal of Management* 43(8): 2580–2587.
- Maas VS, Torres-González R. 2011. Subjective performance evaluation and gender discrimination. *Journal of Business Ethics* 101(4): 667–681
- Maksimovic V, Phillips G. 2002. Do conglomerate firms allocate resources inefficiently across industries? Theory and evidence. *The Journal of Finance* 57(2): 721–767.
- March J, Simon HA. 1958. *Organizations*. John Wiley & Sons, New York, NY, USA.
- Marschak J, Radner R. 1972. *Economic Theory of Teams*. Yale University Press: New Haven, CT, USA.
- Mawdsley JK, Somaya D. 2016. Employee mobility and organizational outcomes: An integrative conceptual framework and research agenda. *Journal of Management* 42(1): 85–113.

- McAfee A, Brynjolfsson E. 2017. *Machine, platform, crowd: harnessing our digital future*. W. W. Norton & Company, New York, USA.
- Merton RK. 1968. The Matthew effect in science: The reward and communication systems of science are considered. *Science* **159**(3810): 56–63
- Meyer BD. 1995. Natural and quasi-experiments in economics. *Journal of Business & Economic Statistics*. **13**(2): 151–161.
- Miles RE, Snow CC. 2003. *Organizational strategy, structure and process*. Stanford University Press, Stanford, CA, USA
- Murciano-Goroff R. 2018. Missing women in tech: the labor market for highly skilled software engineers. *Working Paper*
- Murphy KM, Shleifer A, Vishny RW. 1991. The allocation of talent: Implications for growth. *The Quarterly Journal of Economics*. **106**(2): 503–530.
- Nagle, F., Teodoridis, F., 2017. Jack of All Trades and Master of Knowledge: The Role of Generalists in Novel Knowledge Integration. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3017363>
- Nahapiet J, Ghoshal S. 1998. Social capital, intellectual capital, and the organizational advantage. *The Academy of Management Review* **23**(2), 242–266.
- Nelson RR, Winter SG. 1982. *An evolutionary theory of economic change*. Belknap Press, Cambridge, MA, USA.

- Nickerson JA, Zenger TR. 2002. Being efficiently fickle: a dynamic theory of organizational choice. *Organization Science* **13**(5): 547–566.
- Nickerson JA, Zenger TR. 2004. A knowledge-based theory of the firm—the problem-solving perspective. *Organization Science* **15**(6): 617–632.
- Noda T, Bower JL. 1996. Strategy making as iterated processes of resource allocation. *Strategic Management Journal* **17**(1): 159–192.
- Okhuysen GA, Bechky BA. 2009. Coordination in organizations: an integrative perspective. *Academy of Management Annals* **3**(1): 463–502.
- Okhuysen, G.A., Eisenhardt, K.M., 2002. Integrating knowledge in groups: how formal interventions enable flexibility. *Organization Science* **13**(4): 370–386.
- Oliver N, Calvard T, Potočnik K. 2017. Cognition, technology, and organizational limits: lessons from the air france 447 disaster. *Organization Science* (Articles in Advance).
- Page SE. 2018. *The Model Thinker: What You Need to Know to Make Data Work for You*. Basic Books: New York, NY.
- Paruchuri S, Awate S. 2017. Organizational knowledge networks and local search: The role of intra-organizational inventor networks. *Strategic Management Journal* **38**(3): 657–675.
- Penrose ET. 1959. *The Theory of the Growth of the Firm*. John Wiley and Sons, New York, NY, USA.

- Ployhart RE, Moliterno TP. 2011. Emergence of the human capital resource: a multilevel model. *Academy of Management Review* **36**(1): 127–150.
- Puranam P, Alexy O, Reitzig M. 2014. What’s “new” about new forms of organizing? *Academy of Management Review* **39**(2): 162–180.
- Puranam P, Alexy O, Reitzig M. 2014. What’s “New” About New Forms of Organizing? *Academy of Management Review* **39**(2): 162–180.
- Puranam P, Håkonsson DD. 2015. Valve’s way. *Journal of Organization Design* **4**(2): 2–4.
- Puranam P, Jacobides M. 2006. The dynamics of coordination regimes: implications for organizational design. *Working Paper*.
- Puranam P, Raveendran M, Knudsen T. 2012. Organization design: the epistemic interdependence perspective. *Academy of Management Review* **37**(3): 419–440.
- Puranam P. 2018. *The Microstructure of Organizations*. Oxford University Press, New York, NY, USA.
- Puranam P. 2018. *The microstructure of organizations*. Oxford University Press, Oxford, UK.
- Radner R. 1993. The organization of decentralized information processing. *Econometrica* **61**(5): 1109–1146.
- Raghuram S, Hill NS, Gibbs JL, Maruping LM. 2019. Virtual work: bridging research clusters. *Academy of Management Annals* **13**(1): 308–341



- Raj M, Seamans R. 2018. AI, labor, productivity, and the need for firm-level data. In *Economics of Artificial Intelligence*. University of Chicago Press, Chicago, IL, USA.
- Raj M., Seamans R. 2019. Primer on artificial intelligence and robotics. *Journal of Organization Design* 8(1)
- Rajan, RG, Wulf J. 2006. Flattening firm: Evidence from panel data on the changing nature of corporate hierarchies. *Review of Economics & Statistics* **88**(4): 759- 773.
- Rajan RG, Zingales L. 2001. The firm as a dedicated hierarchy: a theory of the origins and growth of firms. *The Quarterly Journal of Economics* **116**(3): 805–851.
- Remus D, Levy F. 2017. Can robots be lawyers: computers, lawyers, and the practice of law. *Georgetown Journal of Legal Ethics* **30**: 501–558.
- Rider CI, Wade JB. 2019. Racial disparity in leadership: Bias for whites in promotions of National Football League coaches. *Working Paper*
- Ridgeway CL, Correll SJ. 2004. Motherhood as a Status Characteristic. *Journal of Social Issues* **60**(4): 683–700
- Rivera LA. 2012. Hiring as cultural matching: the case of elite professional service firms. *American Sociological Review* **77**(6): 999–1022
- Rivera, L.A., Tilcsik, A., 2019. Scaling down inequality: rating scales, gender bias, and the architecture of evaluation. *American Sociological Review* **84**(2): 248–274.
- Rosenkopf L, Nerkar A. 2001. Beyond local search: boundary-spanning, exploration, and impact in the optical disk industry. *Strategic Management Journal*. **22**(4), 287–306.

- Sah RK, Stiglitz JE. 1986. The architecture of economic systems: hierarchies and polyarchies. *The American Economic Review* **76**(4): 716–727.
- Sah RK, Stiglitz JE. 1988. Committees, hierarchies and polyarchies. *The Economic Journal* **98**(391): 451–470.
- Schumpeter JA. 1942. *Capitalism, socialism, and democracy*. Harper & Brothers, New York, USA.
- Sengul M, Costa AA, Gimeno J. 2018. The allocation of capital within firms. *Academy of Management Annals* **13**(1): 43–83.
- Sevcenko V, Ethiraj S. 2018. How do firms appropriate value from employees with transferable skills? A study of the appropriation puzzle in actively managed mutual funds. *Organization Science* **29**(5): 775–795.
- Shestakofsky B. 2017. Working algorithms: software automation and the future of work. *Work and Occupations* **44**(4): 376–423.
- Shipper F, Manz CC. 1992. Employee self-management without formally designated teams: An alternative road to empowerment. *Organizational Dynamics* **20**(3): 48–61.
- Siggelkow N, Levinthal DA. 2003. Temporarily divide to conquer: centralized, decentralized, and reintegrated organizational approaches to exploration and adaptation. *Organization Science* **14**(6): 650–669.
- Siggelkow N, Rivkin JW. 2005. Speed and search: designing organizations for turbulence and complexity. *Organization Science* **16**(2): 101–122.

Simon HA. 1962. The architecture of complexity. *Proceedings of the American Philosophical Society* **106** (6): 467–482.

Simon HA. 1997. *Administrative Behavior*. The Free Press, New York, NY, USA.

Sparrow B, Liu J, Wegner DM. 2011. Google effects on memory: cognitive consequences of having information at our fingertips. *Science* **333**(6043): 776–778.

Srikanth K, Puranam P. 2011. Integrating distributed work: comparing task design, communication, and tacit coordination mechanisms. *Strategic Management Journal*. **32**(8): 849–875.

Stan M, Puranam P. 2017. Organizational adaptation to interdependence shifts: The role of integrator structures. *Strategic Management Journal* **38**(5):1041–1061.

Stein JC. 2002. Information production and capital allocation: decentralized versus hierarchical firms. *The Journal of Finance* **57**(5): 1891–1921.

Susskind RE, Susskind D. 2015. *The future of the professions: How technology will transform the work of human experts*. Oxford University Press, USA.

Sutton RI., Rao H. 2014. *Scaling Up Excellence: Getting To More Without Settling For Less*. Crown Publishing Group, New York, NY, USA.

Szulanski G. 1996. Exploring internal stickiness: impediments to the transfer of best practice within the firm. *Strategic Management Journal* **17**(Winter Special Issue), 27–43.

Tak E, Correll SJ, Soule SA. 2019. Gender inequality in product markets: when and how status beliefs transfer to products. *Social Forces* **98**(2): 548–577.

- Teece DJ, Pisano G, Shuen A. 1997. Dynamic capabilities and strategic management. *Strategic Management Journal* **18**(7): 509–533.
- Terrell J, Kofink A, Middleton J, Rainear C, Murphy-Hill E, Parnin C, Stallings J. 2017. Gender differences and bias in open source: pull request acceptance of women versus men. *PeerJ Computer Science* **3**:111.
- Thaler RH. 1988. Anomalies: The winner's curse. *The Journal of Economic Perspectives* **2**(1): 191–202.
- Thompson JD. 1967. *Organizations in Action*. McGraw-Hill: New York, NY, USA.
- Thompson JD. 1967. *Organizations in action: Social science bases of administrative theory*. McGraw Hill, New York, NY, USA.
- Tortoriello M. 2015. The social underpinnings of absorptive capacity: The moderating effects of structural holes on innovation generation based on external knowledge. *Strategic Management Journal* **36**(4): 586–597.
- Tsay J, Dabbish L, Herbsleb J. 2014a. Let's talk about it: evaluating contributions through discussion in GitHub. *Proceedings of the 22Nd ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2014* pp. 144–154. ACM, New York, NY, USA.
- Tsay J, Dabbish L, Herbsleb J. 2014b. Influence of social and technical factors for evaluating contribution in GitHub. ACM Press: 356–366.
- Uribe J, Carnahan S. 2018. Men check out, women pitch in: how familiarity influences effort in mixed-gender teams. *Working Paper*.

- Van Zandt T. 1999. Real-time decentralized information processing as a model of organizations with boundedly rational agents. *The Review of Economic Studies* **66**(3): 633–658.
- Vanacker T, Collewaert V, Zahra SA. 2017. Slack resources, firm performance, and the institutional context: evidence from privately held European firms. *Strategic Management Journal* **38**(6): 1305–1326.
- von Krogh G. 2018. Artificial intelligence in organizations: new opportunities for phenomenon-based theorizing. *Academy of Management Discoveries* **4**(4): 404–409.
- von Nordenflycht A. 2010. What is a professional service firm? Toward a theory and taxonomy of knowledge-intensive firms. *Academy of Management Review* **35**(1): 155–174.
- Weick KE, Roberts KH. 1993. Collective mind in organizations: heedful interrelating on flight decks. *Administrative Science Quarterly* **38**(3): 357–381.
- Williamson OE. 1975. *Markets and Hierarchies: Analysis and Antitrust Implications: A Study in the Economics of Internal Organization*. The Free Press: New York, NY, USA.
- Witteman HO, Hendricks M, Straus S, Tannenbaum C. 2019. Are gender gaps due to evaluations of the applicant or the science? A natural experiment at a national funding agency. *The Lancet* **393**(10171): 531–540.
- Woodward J. 1965. Analysis of organizations. In *Industrial Organization: Theory and Practice*. Oxford University Press, London, UK; 17–34.
- Wu L, Wang D, Evans JA. 2019. Large teams develop and small teams disrupt science and technology. *Nature*, **566**(7744): 378.

Zhao ZJ, Anand J. 2013. Beyond boundary spanners: The 'collective bridge' as an efficient interunit structure for transferring collective knowledge. *Strategic Management Journal* 34(13): 1513–1530.

Zhou KZ, Li CB. 2012. How knowledge affects radical innovation: Knowledge base, market knowledge acquisition, and internal knowledge sharing. *Strategic Management Journal* 33(9): 1090–1102.

Zuboff S. 1988. *In the age of the smart machine: The future of work and power*. Basic books, New York, NY, USA