Essays on Uncertainty in Labor and Financial Markets

Inauguraldissertation

zur Erlangung des Grades eines Doktors
der Wirtschaftswissenschaften

durch

die Rechts- und Staatswissenschaftliche Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Valentin Nicolas Kecht

aus Wasserburg am Inn

2026

Dekan:	Prof. Dr. Martin Böse
Erstreferent:	Prof. Dr. Farzad Saidi
Zweitreferent:	Prof. Dr. Thomas Dohmen
Tag der mündlichen Prüfung:	22. Juli 2026

Acknowledgements

If you happen to know me, you know that I like fun. Not exclusively in the sense of self-indulgence, but also in the sense of doing things that are meaningful to me and, hopefully, to others. By that measure, my PhD has been a lot of fun, and it would not have been the same without the people who accompanied me along the way.

First and foremost, I am grateful to my main advisor, Farzad Saidi. I have benefited immensely from your intellectual curiosity, your extremely broad knowledge of economics research, and your ability to strategize within the profession. You have never ceased to impress me, not only with your discipline and perseverance but also with your devotion to life outside of academia, which more than once helped me recalibrate my own priorities. I very much enjoy working with you and discussing, through whatever medium and at whatever hour, anything from research to personal matters to pop culture. You are truly a role model to me, and over the past years you have become a very good friend from whom I have learned a lot.

I would also like to thank my second advisor, Thomas Dohmen. Our weekly meetings gave me countless opportunities to draw on your expertise across an extremely wide range of fields and to witness your genuine excitement for your own and everyone else's research, as well as your unwavering commitment to supporting your students in whatever situation they may find themselves. I have deeply appreciated your feedback and am grateful for all the time you have invested in me.

A special thanks goes to Alessandro Lizzeri, who generously agreed to host me during my research visit at Princeton University and to serve as my third supervisor. Your brilliance and way of thinking about problems, driven by a desire to understand rather than by ego, have profoundly shaped my own perspective, not least during our walks across campus.

This dissertation would not have been possible without fruitful collaborations with my other co-authors. I want to thank Georg Schneider and Giacomo Marcolin, not only for countless exchanges and discussions, but also for your friendship.

I also want to thank Joachim Winter, who sparked my excitement for empirical research more than a decade ago, and Nicola Gennaioli, who helped me navigate PhD applications during the uncertain times of Covid.

I am grateful to the Bonn Graduate School of Economics for an excellent research environment, and to ECONtribute, the CRC TR 224, and the German Academic Exchange Service for funding and opportunities for exchange.

Clearly, a PhD is a team effort, and I was lucky to share mine with outstanding people. Importantly, Simon Cordes, Daniel Evans, and Georg Schneider, for their constant feedback, high-quality memes, and the occasional Taxiteller. I was equally fortunate to have excellent colleagues at the Institute of Finance and Statistics: Christina Brinkmann, Dominik Damast, Alexandros Gilch, Maxi Guennewig, Tobias König, and Anna Mogilevskaja. For many stimulating interactions and discussions, I also thank the earlier graduates Martin Kornejew, Timothy Meyer, Lorenzo Ranaldi, and Philipp Schirmer. During the final stretch of my PhD, I am indebted to Pietro Buri and many others for making me feel at home in the US, including the Procter Fellows and the great people I met in the Economics Department at Princeton.

This PhD journey was made all the more fascinating and enjoyable by people outside of academia. I want to thank my long-time friends Laura, Heidi, and Henri, for your encouragement and many shared moments of *dolce vita* in and around Cologne throughout those years.

I also owe a great deal to my parents, Karin and Jochen, for making me curious about the world. Thank you for your unconditional support and warmth, and for always believing in me and in my ability to stand on my own feet and make the right decisions for myself.

Most of all, I would like to thank Gloria. Thank you for your support over almost fifteen years and half a dozen long-distance living arrangements, and for trusting me to follow my passion just as you follow yours. Thank you for your love, your patience, your enthusiasm, and your impeccable sense of humor. You are the best of wives and best of women.

Valentin Kecht

Bonn, June 2026

Contents

List of Figures

List of Tables

Introduction

Human capital is central to economic production and often characterized as firms' most important asset (Becker 1964). Unlike physical capital, however, it is embodied in people, and its accumulation and deployment reflect a multitude of individual decisions. Workers choose which skills to acquire, which jobs to accept, and which risks to bear, and the value of these choices hinges on market conditions beyond their control. These conditions have shifted markedly over the past decades, most notably through technological disruption and a rise in uncertainty and economic complexity (Bloom 2014; Baker, Bloom, and Davis 2016).

This dissertation advances our understanding of the interaction between these aggregate forces and three specific labor markets, namely those for chief executives, bankers, and workers in the gig economy. The three settings could hardly be more heterogeneous in that they span the entire skill and income distribution, from executives at the very top of the corporate hierarchy to workers coping with job loss. The common feature across these environments is uncertainty, though it operates through distinct channels. At the top, it shapes firms' skill demand and CEO selection. In the financial sector, labor market conditions determine employees' implicit insurance against adverse outcomes and shape the financial risks they take. For workers facing job loss, uncertainty takes the form of income risk, against which platform work offers partial insurance.

Studying these interactions has proved fruitful for three reasons. First, it reveals how labor markets adjust to shifts in aggregate economic conditions. Demographic change and technological progress are expected to render high-skilled human capital increasingly scarce (Autor 2022), and higher-order skills such as judgment and decision-making are projected to grow in importance (Deming and Silliman 2025), implying that the questions tackled here are set to become more relevant. Second, institutional arrangements moderate these adjustments, and understanding how they do so helps delineate the scope for policy. Financial regulation, for instance, typically focuses on banks' balance sheets rather than bankers' labor markets, and the generosity of unemployment benefits affects the insurance value of new forms of work. Third, causality also runs in the opposite direction. In elite labor markets in particular, the incentives faced by employees shape their choices, which in turn impact aggregate outcomes ranging from talent allocation to credit supply. In this sense, the dissertation departs from the standard view of labor as a mere production input and emphasizes the role of human capital in decision-making and resource allocation.

A recurring theme across the three essays is the option value of labor market opportunities under uncertainty. Broad skills become more important when business conditions change frequently, outside options cushion bankers against employment risk, and gig work provides a fallback for the recently unemployed. How such options emerge and are used lies at the core of each chapter.

Chapter 1 (joint with Alessandro Lizzeri and Farzad Saidi) explores how rising aggregate uncertainty and economic complexity have reshaped the market for chief executives. We argue that, in stable environments, firms tend to select leaders based on their ability in a well-defined task. Yet, when conditions change frequently, firms place greater weight on skills that are portable across products and market conditions. Because such generalist skills are tied to accumulated experience rather than raw ability, and experience takes time to build, a growing premium on generalist human capital should be reflected in the age of the leaders firms choose.

Consistent with this hypothesis, the chapter documents that the age at which CEOs are appointed in the United States has risen by 7–8 years since 2000, with similar developments in Europe. This trend has been accompanied by fundamental changes in how executives build their careers. The age increase can be fully accounted for by additional time spent in external positions across a growing number of firms and industries, while experience within the appointing firm has remained unchanged. Demographic shifts, entrenchment, and rising market concentration cannot explain the magnitude of these changes. We therefore formalize our hypothesis in a matching model in which candidates differ in ability, which peaks in mid-career, and experience, which grows with age. The model predicts that the rising value of experience shifts top positions toward older candidates in general equilibrium, most strongly at smaller firms, where the age increase is most pronounced.

The chapter then scrutinizes both sides of this market for leadership. To isolate firms' demand for generalist skills, we leverage spatial variation in the availability of generalists in the form of strategy consultants, allowing us to study the differential response in CEO appointments to industry-level shocks across space. Using differences in travel time to offices of leading consulting firms, we find that firms exposed to higher uncertainty appoint older, more generalist CEOs where consultants are less available, with small firms reacting most. On the supply side, we show that executives adjust their careers to these changes in skill demand. Once a former coworker is appointed CEO at another firm, an event that reveals which career profiles pay off, executives become more mobile themselves and accept positions at lower seniority levels and pay in exchange for diversifying their skill portfolio. Consistent with the option value of broad skills, they bear short-run costs for human capital whose returns materialize only later.

While these patterns are potentially a rational response to a more uncertain and complex business environment, they may come at aggregate costs. Since older leadership is associated with less risk-taking and slower firm growth, and since small firms contribute disproportionately to job creation and innovation, the changing selection of corporate leaders could dampen business dynamism and the economy's capacity to adapt to long-run challenges.

Chapter 2 (joint with Georg Schneider) shifts focus from the matching of workers and firms to behavior on the job. We hypothesize that better labor market outside options raise workers' willingness to take risk by providing a fallback in case of failure. We study this insurance effect in the financial industry, where job-to-job mobility is high and where bankers conduct large, risky transactions that firms can only imperfectly monitor. In a stylized model, the outside option acts as insurance: by truncating the downside of a banker's payoff distribution, its option value lowers effective risk aversion and leads to riskier lending decisions.

Testing this prediction requires variation in outside options that is orthogonal to a bank's own trajectory. Natural proxies such as employee outflows fail this requirement, as departures themselves respond to a bank's fundamentals. To overcome this identification challenge, we construct a network-based measure of outside options from employment history data, a measure that captures hiring shocks at other institutions and is unrelated to a bank's past outflows and financials. Drawing on transaction-level data on syndicated loans, we find that banks expand lending volumes in response to improved outside options by arranging more loans. Financial risk rises in the tail of the loss distribution, as does the bank's contribution to systemic risk in distress. At the loan level, better outside options shift lending toward riskier borrowers without compensating increases in spreads, which is difficult to reconcile with yield-seeking and suggests a genuine increase in risk-taking.

Two sets of individual-level results tie these patterns to the insurance channel. Matching loans to the arranging bankers, we find the effects are concentrated among younger bankers, who are most exposed to dismissal risk. We then extend the analysis to financial advisors, for which we document that misconduct becomes more frequent when outside options improve, while the expected punishment for this behavior shrinks. Outside options reduce the probability of job separation following misconduct and, conditional on separation, raise the likelihood of reemployment within the sector.

Since outside options co-move with the business cycle, the insurance effect is most pronounced during credit expansions, and vice versa. The chapter thereby identifies a worker-level channel through which labor market conditions translate into fluctuations in credit supply, suggesting a role for labor markets in macroprudential oversight. Common measures may be limited in their effectiveness, as a bank that unilaterally defers pay or imposes clawbacks may lose talent at the time when rivals are hiring. Containing the distortion may thus require coordination, for instance through sector-wide standards.

Chapter 3 (joint with Giacomo Marcolin) moves from the upper end of the income distribution to workers coping with job loss. Since the financial crisis, digital labor platforms have emerged as a new form of employment with low barriers to entry. Such platform work may help workers absorb income shocks after a job loss, with consequences for job search and the take-up of unemployment benefits. Existing evidence, however, stems almost exclusively from ride-sharing platforms in the United States, and it is unclear whether it extends to economies with more generous welfare systems. This chapter examines whether the gig economy insures workers against the income risk following job separation in Germany.

Because platforms enter local markets selectively rather than at random, naive comparisons across districts conflate platform access with preexisting differences. We address this concern by exploiting the staggered entry of gig delivery platforms across German districts, which we combine with administrative employment data.

We show that recently separated workers are more likely to take up gig work once a platform operates in their district. Take-up is concentrated among young, male, and high-skilled workers, and we find some evidence that gig work crowds out other forms of employment in the medium run. Platform availability does not, however, reduce reliance on unemployment benefits or spur business formation. Thus, rather than substituting for public unemployment insurance, gig work serves as a supplementary source of income during job search. Two further patterns support this interpretation of gig work as a transitory arrangement rather than stable employment. The employment effects persist and grow over time, and platform entry raises the number of workers who are employed but keep searching for a job. Platform entry also increases participation in active labor market policies, which suggests a signaling role of gig work. The contrast with US evidence shows that the option value of platform work as a fallback hinges on the surrounding institutional environment.

References

Autor, David. 2022. "The Labor Market Impacts of Technological Change: From Unbridled Enthusiasm to Qualified Optimism to Vast Uncertainty." NBER Working Paper 30074. National Bureau of Economic Research. [1]

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. "Measuring Economic Policy Uncertainty." *Quarterly Journal of Economics* 131 (4): 1593–636. [1]

Becker, Gary S. 1964. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education.* New York: Columbia University Press. [1]

Bloom, Nicholas. 2014. "Fluctuations in Uncertainty." *Journal of Economic Perspectives* 28 (2): 153–76. [1]

Deming, David, and Mikko Silliman. 2025. "Skills and Human Capital in the Labor Market." In *Handbook of Labor Economics*, edited by Christian Dustmann and Thomas Lemieux, 6: 115–57. Elsevier. [1]

Chapter 1

Aging at the Very Top

Joint with Alessandro Lizzeri and Farzad Saidi

1.1 Introduction

We uncover a striking evolution in the age of CEOs over the last decades and study its underlying forces. Between 2000 and 2023, the average age of CEOs in the United States has increased by ten years (Figure 1.1), five times the increase for the college-educated labor force. Similar patterns have emerged in Europe. Longer tenures or entrenchment cannot fully explain this trend, as age at appointment has risen almost as sharply from 47–48 to 55 years over the same time span. Using comprehensive data on CEOs alongside their employment histories, we consider a variety of potential explanations and conclude that the patterns we observe are indicative of firms' increased demand for generalist human capital. We also discuss how executives alter their career paths in response to these changes in skill demand.

We document a number of important changes in the career profiles of managers and executives over the sample period. First, individuals who eventually become CEOs accumulate roughly ten years more external experience than in 2000, whereas internal experience has remained almost unchanged during this period. This pattern holds for both external and internal appointments, and is particularly pronounced in smaller firms. Second, we consider the composition of this additional experience. At appointment, today's CEOs have held a larger number of positions across more firms and industries than their counterparts in 2000. Third, turnover has increased with shorter spells in each position. Fourth, the increase in the number of positions outweighs the faster turnover, so CEOs' employment histories have lengthened, raising their age at appointment.

We evaluate several potential explanations for this aging of corporate leadership, including demographic trends, agency conflicts, and rising concentration, but conclude that none of these can account for the magnitude of the observed age increase. We then turn to changes in market forces that affect the matching of firms

Figure 1.1. Average CEO Age Over Time

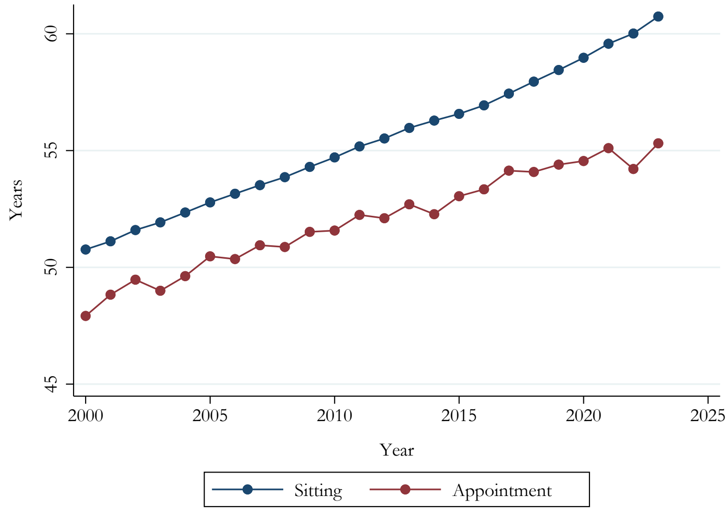

Notes: The plot shows the average age of CEOs over time, separately for sitting CEOs and CEOs at appointment. The sample includes 50,510 CEOs in the United States, for which we obtain information from BoardEx.

and CEOs. To formalize this mechanism, we develop a many-to-one matching model with the following ingredients. Workers have multidimensional skills: they vary in both age-adjusted ability (which peaks at mid-career) and experience (which increases with age). Firms differ in productivity (related to firm size) and have multiple positions with hierarchically ranked productivity levels. We provide conditions on the shape of the productivity functions under which the equilibrium assignment is consistent with the key empirical patterns. Our main result characterizes how an increase in the value of experience shifts top positions toward older workers, especially at lower-productivity, or smaller, firms. We further show that, under mild assumptions, CEOs in smaller firms are younger on average.

Aging at the very top of the corporate ladder can have substantial effects on a variety of firm outcomes, e.g., because of older CEOs' distinct management style (Bertrand and Schoar 2003; Schoar and Zuo 2016) or preferences (Jenter and Lewellen 2015). A key takeaway from this literature is that CEO age is negatively associated with business dynamism and firm risk.[1] Thus, our findings reveal not only how firms adjust to changing economic conditions, but also how these decisions in turn shape broader trends. Because our dataset spans a much wider universe than the large, listed firms typically studied, we can document systematic differences in CEO hiring patterns across the firm size distribution. Smaller firms offer fewer opportunities to accumulate generalist experience internally and rely more heavily on

1. See Appendix Section 1.D for a detailed discussion of the nexus between CEO age and firm-level outcomes.

external hiring, whereas larger firms can cultivate such skills through internal assignments. Since smaller firms tend to be younger and account for a disproportionate share of job creation and innovation (Haltiwanger, Jarmin, and Miranda 2013), and since today's large firms began as small firms, understanding how executive skill demand shapes leadership selection at smaller firms provides crucial insights for aggregate economic performance.[2]

We use granular data from BoardEx and LinkedIn to track firms' hiring decisions and executives' career trajectories in the United States. To shed light on the rising value of experience, we focus on two dimensions along which firms' operating environment has evolved. First, economic uncertainty has increased (Baker, Bloom, and Davis 2016). Second, the CEO role has become more complex, as evidenced by firms simultaneously expanding across geographies and business lines while navigating a growing set of regulatory constraints (Dessein and Santos 2006; Pan 2017). These forces can lead firms to seek leaders with generalist skills, which are more closely tied to accumulated experience than to raw ability. As executives require longer career paths to build such capabilities, firms appoint older CEOs.

In support of these hypotheses, we first establish that in the face of rising industry-level uncertainty, firms demand different profiles in their prospective CEOs. A challenge in relating industry-level uncertainty to the demand for CEO characteristics is that uncertainty itself may be correlated with the pool of available CEO candidates. To address this, we use spatial variation in the availability of generalists in the form of strategy consultants. This strategy allows us to study the differential response in CEO appointments to industry-level shocks across space, while controlling for time-varying geographic and sectoral heterogeneity. Exploiting plausibly exogenous variation in travel time to the offices of the top three strategy consulting firms—stemming from office openings and air route expansions—we find that greater industry-level uncertainty raises the demand for generalist skills, leading firms to appoint older CEOs wherever other generalists are less available. The results are stronger for small firms where generalist human capital is more difficult to accumulate internally, consistent with the larger gradient in external experience of small firms' CEOs over time that we document. Therefore, the value added of generalist consultants as substitutes for experience is arguably higher for small firms.

To assess how overlapping trends in the complexity of organizational structures may have contributed to the rising value of generalist human capital, we extend our identifying framework using the following proxies: business diversification (measured by the number of distinct VP-level roles within firms), spatial diversification within the US, and trade-induced variation in economic complexity. Across all these measures, we find qualitatively similar results.

2. These considerations are particularly concerning in light of decreasing firm entry rates, which have been traced back to increased entry costs (Kozeniauskas 2025) as well as demographic changes (Hopenhayn, Neira, and Singhania 2022).

Next, we investigate how executives have adapted their career strategies in light of firms' growing demand for generalist skills. One interpretation of prospective CEOs' increased turnover across positions, firms, and industries is that they respond strategically to the increasing premium for broad managerial capabilities. To explore this hypothesis, we start from the idea that executives learn about career-enhancing opportunities through professional networks. If the acquisition of generalist skills is revealed to be an important career ingredient, they may become more willing to accept lower-level positions and reduced pay in the short run in exchange for building a generalist skill set that enhances their long-term career prospects.

To test this, we study changes in employees' career paths after they learn about CEO appointments of former peers. Individuals who worked contemporaneously at the same firm and seniority level as well as in the same metro area with the newly appointed CEO (of another firm) respond by increasing their own job mobility more strongly than colleagues who shared the same firm and seniority level but were located in different metro areas. On average, these treated individuals accept reduced compensation in the short run in exchange for the presumed value of additional experience. This is reflected in the higher prevalence of downward switches across firms—where employees move to positions at lower seniority levels—but also in more frequent transitions across industries. These responses to information shocks are stronger for coworkers with longer shared tenure, for those that exhibit more mobility on the way to the top, and for workers in small firms.

Taken together, our findings support the idea that both the demand and the supply of generalist skills have shaped the age profile of CEO appointments since the early 2000s. We leverage this insight to decompose the overall age trend and show that around 37% of the variation in CEO age can be attributed to sectoral changes, reflecting changes in skill demand. A further 28% can be explained by geographical factors, capturing supply-side forces. The remaining share can be traced back to the interaction of these two dimensions.

Our analysis makes the following contributions. First, we provide novel systematic documentation of CEO aging trends and their relationship to changing career patterns, and link them to the increasing value of generalist human capital. Second, we provide a matching model that rationalizes and interprets these patterns. Third, we develop identification strategies that separately estimate demand- and supply-side factors contributing to these changes. Fourth, we demonstrate how professional networks transmit information about optimal career strategies in the market for CEOs, leading to widespread behavioral changes among executives.

Related Literature: Our paper speaks to several strands of the literature. A number of studies have explored the changing demand for CEO skills over time. Murphy and Zabojnik (2004, 2007) and Frydman (2019) link the rise in executive compensation and the increasing prevalence of external CEO hiring during the end of the 20^{th} century to a shift toward generalist skills.[3] Custódio, Ferreira, and Matos (2013) document a pay premium for generalist CEOs, while Adams, Akyol, and Verwijmeren (2018) highlight the importance of diversified skill sets among board members.

Other research has more directly analyzed shifts in skill requirements using data from job postings and interviews. Kaplan and Sorensen (2021) show that CEOs, compared to other executives, have higher levels of general ability, execution orientation, strategic focus, and analytical skills. Following the financial crisis, Decressin, Kaplan, and Sorensen (2025) document a shift away from interpersonal skills toward a stronger emphasis on execution orientation, which they find to be positively correlated with firm performance. Similarly, Deming (2021) emphasizes the increasing importance of decision-making skills, and Deming (2017) and Fuller et al. (2021) show that social skills have become more valued in the labor market. Kaplan and Sorensen (2021) argue that too much weight is placed on social skills in the selection of CEOs. Moreover, Deming and Silliman (2025) report that higher-level skills, such as the ability to allocate individual tasks, have gained increasing importance over time.[4]

The evolution of CEO skill demand has frequently been examined in conjunction with changes in firm characteristics, using two-sided matching models. This literature has sought to explain rising executive compensation through assortative matching of managerial talent and firm size (Gabaix and Landier 2008; Terviö 2008). This relationship may be driven by the multiplicative impact of managerial skills in large firms (Edmans, Gabaix, and Landier 2009), or by the fact that larger firms pay efficiency wages to discourage shirking, as the signal of managerial quality is more diffuse (Gayle, Golan, and Miller 2015). Using these models, closely related to our findings, Gayle, Golan, and Miller (2015) emphasize the importance of general

3. A large literature discusses the role of CEO hiring patterns, as the share of internal compared to external appointments is informative about frictions in this labor market. In the S&P 500, the majority of CEOs are promoted internally (Cziraki and Jenter 2021). Relatedly, Hamori and Kakarika (2009), Koch, Forgues, and Monties (2017) and Lyman, Barry, and Zhao (2025) underscore the importance of internal work experience for CEOs of large listed organizations. In contrast, Gompers, Kaplan, and Mukharlyamov (2023) examine private equity buyouts and find that most CEOs are hired externally. Other studies also point toward a competitive market for CEOs. Cappelli (2010) and Cappelli and Hamori (2004) date the shift from internal development programs toward external hiring to around 1980. Bai and Mkrtchyan (2023) show that outside hiring generates productivity gains, and Gao, Luo, and Tang (2015) document that companies increase executive pay after managers leave for other firms.

4. A related literature has focused on variation across individuals by exploring how managerial styles influence firm outcomes (Bertrand and Schoar 2003; Schoar and Zuo 2016; Bennedsen, Pérez-González, and Wolfenzon 2020; Nguyen 2025). These styles are malleable, shaped by early-career experiences (Schoar and Zuo 2017), and vary with firms' ownership and governance structures, e.g. in family firms (Bertrand and Schoar 2006; Bertrand et al. 2008; Mullins and Schoar 2016). Fee, Hadlock, and Pierce (2013) complement this work by demonstrating that boards deliberately select candidates who implement strategies aligned with their own objectives.

human capital in the CEO market, while Pan (2017) highlights the nexus between firm scope (i.e., diversification) and CEO experience in relevant industries. Dow and Raposo (2005) model CEO selection based on firms' strategic needs, showing how demand for change management capabilities shapes the matching of CEOs to firms. We demonstrate that such adaptive capabilities are tied to accumulated generalist experience, explaining why firms requiring strategic flexibility appoint older rather than younger CEOs.

Our work further contributes to the literature that examines how technological and institutional changes influence human capital accumulation and how human capital is traded in labor markets. This debate has been shaped by the distinction between firm-specific and general human capital (as introduced in Becker 1964). Generalist skills are often measured by work experience in multiple firms and industries, capturing skills that are transferable across different organizations or roles (Custódio, Ferreira, and Matos 2013). Conversely, as noted by Lazear (2009), few skills are inherently firm-specific, but many combinations of skills or skill requirements are. Following this reasoning, the extent to which human capital is firm-specific or general mainly depends on market thickness, as documented in Sauvagnat and Schivardi (2024) for the market for executives. Beyond the firm level, the evidence points toward human capital that is specific to industries (Neal 1995), occupations (Sullivan 2010), and individual tasks (Gibbons and Waldman 2004). Similarly, Dessein and Santos (2006) investigate firms' adaptability—defined as the level of discretion granted to workers—in an environment with uncertainty and task bundling. They show that more adaptive organizational structures, in response to increased uncertainty, tend to feature more bundled jobs, thereby increasing the demand for generalists.

Our supply-side channel also resonates with a broader literature that highlights the importance of on-the-job human capital accumulation in driving wage growth (see, e.g., Adda and Dustmann 2023; Jedwab et al. 2023). In this context, firm characteristics have been shown to play a crucial role in shaping human capital development. These characteristics include coworkers (Gong, Sun, and Wei 2018; Nix 2020; Jarosch, Oberfield, and Rossi-Hansberg 2021), export shares (Macis and Schivardi 2016), firm size (Arellano-Bover 2024), and firm connectivity (Del Prato 2023). Arellano-Bover (2024) documents large differences in firms' abilities to offer on-the-job learning opportunities, emphasizing the importance of general human capital. In line with this, Groes, Kircher, and Manovskii (2015) find that occupational mobility follows a U-shaped pattern with respect to wages, suggesting that higher earners are more likely to transition to other jobs.

Lastly, we add to the literature on age trends in high-skill professions and their consequences. Older CEOs' management styles reflect their different incentives compared to their younger counterparts, including distinct wealth concerns and career horizons (Dechow and Sloan 1991; Antia, Pantzalis, and Park 2010), future career concerns (Li, Low, and Makhija 2017), and different compensation structures involving pensions and deferred benefits (Sundaram and Yermack 2007). Empirically, CEO age has been shown to be negatively related to firm growth and innovation, for instance, in terms of investment and sales (Acemoglu, Akcigit, and Celik 2022), firm value and operating performance (Cline and Yore 2016), R&D expenditure (Serfling 2014; Li, Low, and Makhija 2017), artificial intelligence (AI) usage (Yotzov et al. 2026), and radical innovation (Acemoglu, Akcigit, and Celik 2022). Yet, given that managers can shape firms' risk exposure (Schoar, Yeung, and Zuo 2024), older CEOs tend to increase firms' diversification across business segments (Serfling 2014), while younger CEOs take on more risk by pursuing acquisitions (Yim 2013), aggressive takeovers (Levi, Li, and Zhang 2010), and strategic restructuring (Li, Low, and Makhija 2017).[5] Even though the role of CEO age in shaping firm-level outcomes has received widespread attention, changes in age structures over time have been relatively understudied. One reason may be the lack of comprehensive panel data covering firms beyond the largest ones, for which CEO age has remained relatively similar over the last century.[6]

Similar trends appear beyond corporate leadership: the average age of scientists has been increasing over the past decades, particularly among grant winners.[7] Daniels (2015) attributes this trend to longer pre-grant career paths, while Fons-Rosen, Gaule, and Hrendash (2023) point to slower hiring and delayed retirements. Jones (2010) and Jones and Weinberg (2011) show that both the age at which Nobel Prize-winning discoveries are made and the age at which the prize is awarded have increased over the course of the 20$^{\text{th}}$ century. In many fields, such as the judiciary, individual performance tends to decline with age (Ash and MacLeod 2024). One possible explanation is that aging is associated with changes in cognitive abilities (Salthouse 2012; Singh-Manoux et al. 2012; Waelchli and Zeller 2013), physiological capacity, and various forms of human capital (MacLeod 2016). Additionally, older individuals may experience motivational shifts, placing greater emphasis on maintenance and loss prevention over growth (Ebner, Freund, and Baltes 2006), and may tolerate lower effort.

Aging professionals can also have effects on others. Ash and MacLeod (2024) identify negative productivity spillovers to younger judges, likely due to shifts in

5. We provide an extensive summary of the literature in Appendix Section 1.D.1.

6. Taussig and Joslyn (1932) report that the median age of American business leaders in the early 1930s ranges between 50 and 54. Warner et al. (1967) states that the typical management hire in the early 1950s was around 50 years old. Frydman (2019) tracks the 50 largest US corporations over time, for which CEO age has remained almost unchanged throughout the second half of the 20th century. For CEOs in the Fortune 100, Cappelli, Bonet, and Hamori (2024) find a slight decline in average CEO age during 1980 and 2000, followed by a rebound by 2020. For European countries, analyses based on small samples of very large listed companies yield mixed conclusions. While Freye (2010) and Sluyterman and Westerhuis (2022) find no clear age trend in CEO age at appointment throughout the second half of the 20$^{\text{th}}$ century for Germany and the Netherlands, respectively, Adams et al. (2023) document a decline in the UK of about five years.

7. https://grants.nih.gov

workload distribution within teams. A higher concentration of older workers has been shown to hinder younger workers' career outcomes in local labor markets (Mohnen 2025) and slow down younger colleagues' promotion prospects in internal labor markets (Bianchi et al. 2023). These dynamics have important distributional consequences along demographic dimensions, including age (Bianchi and Paradisi 2024) and gender (Arellano-Bover et al. 2024). The evidence in d'Astous, Geelen, and Hajda (2025) indicates that older workers may be productive complements to younger workers, especially in industries with high physical capital intensity.

1.2 Data

The main empirical analysis relies on CEO data from BoardEx and employment biographies from Revelio Labs, which we combine with company-level covariates and information on the availability of strategy consultants. This section contains details on the data.

BoardEx: We use CEO data from BoardEx, a product offered by *Delinian*. This dataset provides comprehensive details on executives' employment histories and demographic characteristics. While BoardEx includes individuals from both North America and continental Europe, our analysis focuses primarily on the United States. We define CEOs as individuals with the job title "CEO" categorized as "Executive Director," following BoardEx's classification. This yields a sample of 50,510 CEOs in the United States (17,962 in Europe) with available age data, of whom roughly 22% (25%) are appointed at Compustat (listed) firms.

For the descriptive analysis, we transform the data into a panel, capturing employment status as of January 1st each year. In cases of overlapping employment periods, we prioritize more senior positions and those with earlier start dates.[8] As in Cziraki and Jenter (2021), we exclude CEO appointments of less than 12 months. We focus on a CEO's first appointment per firm, omitting CEO reappointments and subsequent board appointments at the same firm. Internally appointed CEOs are defined as those employed by the firm in some other capacity during the last observation prior to their appointment; all others are considered externally appointed.[9] Among internally appointed CEOs, we further distinguish those who initially joined the firm at a managerial level and those who entered at lower levels.

8. The seniority ranking applied is as follows (in descending order): *CEO*, *Other Executive Director*, *Senior Manager*, *Supervisory Director*, and *All Other Employees*. Throughout this paper, we define seniority as an individual's hierarchical position within a firm, distinguishing it from the concept of relative tenure compared to coworkers, as used for instance in Buhai et al. (2014).

9. Therefore, external insiders, i.e., returning executives who account for less than 4% of all appointments, are classified as externally appointed.

Summary statistics for the descriptive analysis for the U.S. are shown in Table 1.C.1, Panel A.[10] The average CEO age at appointment is 52. Around 20% of CEOs are appointed internally, split evenly between those first entering the firm at the managerial level and those joining below it. Given the prevalence of private firms in our sample, the median firm has 141 employees. Importantly, the dataset extends well beyond the S&P 500 or S&P 1500 firms covered in most prior work, and the resulting firm characteristics differ accordingly. For example, in the S&P 1500 the average CEO age at appointment is 54, around 50% of CEOs are appointed internally, and the median firm employs nearly 4,000 workers.[11]

Employment Biographies: We also utilize employment biography data provided by Revelio Labs, which compiles this information by scraping LinkedIn profiles. The dataset encompasses employment details for nearly 500 million workers across approximately 1.8 billion employment spells. The linked employer-employee data include the start and end dates of employment spells, alongside a seniority ranking based on job titles ranging from 1 (lowest) to 7 (highest).[12] Revelio Labs also imputes salaries for each employment spell. The data cover the period from 2019 to 2023 but are backfilled, implying that we can observe employment spells from earlier periods to the extent that they are reported by the respective individual.

Compared to BoardEx, this dataset offers the advantage of including information on rank-and-file employees. Another key strength of the data is that they assign individual employment spells to 145 metropolitan statistical areas (MSAs). We use this information to determine the headquarters of the firms included in BoardEx. For this purpose, we define the location of a firm's headquarters as the modal MSA of its executives in the two highest seniority levels, provided they are at least at seniority level 4, which corresponds to the managerial level. Among other variables, we compute firm-level employment as the number of individuals employed at a firm at the beginning of each year. To ensure that time trends are not driven by increased reporting in later periods, we exploit the overlap of our dataset with Compustat and scale the measure by the year-specific ratio of average employment in Compustat to our employment counts.

10. For the remaining analysis on the U.S., we start with a slightly larger sample of executives. Focusing on age at appointment, we calculate average CEO age at the firm-year level between 2000 and 2022, thereby accounting for Co-CEOs, individuals appointed CEO at multiple companies, and those who held the same position multiple times at the same company. We then merge this variable with the company-level variables described below. Summary statistics for this sample are reported in Table 1.C.1, Panel B.

11. Using an alternative data preparation approach in Appendix Section 1.E, the share of internally appointed CEOs increases to 32% for the overall sample, 68% for the S&P 1500, and 73% for the S&P 500. The latter figure is very close to Cziraki and Jenter (2021), who find that 72% of S&P 500 CEOs are promoted internally. In Appendix Section 1.E, we also document that our main descriptive results are insensitive to the data processing approach used.

12. For examples, see www.data-dictionary.reveliolabs.com.

To merge the LinkedIn data with BoardEx at the firm level, we proceed as follows. First, we clean company names on both sides by applying standard data normalization techniques. Second, we consecutively merge on ISIN and CIK, manually verifying the merges since these identifiers can change over time. Third, for the unmerged companies, we perform an exact merge on company names. Fourth, we apply a fuzzy merge algorithm based on *term frequency–inverse document frequency* on the remaining unmerged company names. For the employment histories, we retain only firms with more than 10 employment spells in North America. The procedure is highly efficient, as it vectorizes the data into a sparse matrix and fits a k-nearest neighbors algorithm. We then randomly select 1,000 matches, which we manually classify as good or bad. In this way, we determine a precision threshold of 0.85, at which 95% of the matches can be considered accurate.[13] Overall, we are able to merge firm-level details for 65% of the CEOs in the BoardEx sample.

Company-Level Covariates: We complement the individual-level data with firm-level information. We obtain annual balance-sheet information from Compustat, which we merge with the US BoardEx data using CRSP identifiers. We focus on the balance-sheet items revenue (*revt*) and employment (*emp*). For both variables, we drop negative values and winsorize below the 1$^{\text{st}}$ and above the 99$^{\text{th}}$ percentile.

We then construct symmetric growth rates γ_{jt} accommodating entry and exit as $\gamma_{jt} = 100 \times \frac{x_{jt}-x_{jt-1}}{0.5(x_{jt}+x_{jt-1})}$, where x_{jt} is the variable of interest for firm j at time t. Given the non-negativity restriction, this implies values between -200 and 200. Using these growth rates, we construct volatility measures defined as the cross-sectional standard deviation of γ_{jt} within NAICS-4 industries.

MSA-Level Data: Our identification strategy in Section 1.5.1 exploits plausibly exogenous variation in the supply of (former) strategy consultants, focusing on the presence of MBB (*McKinsey & Company*, *Boston Consulting Group*, and *Bain & Company*) offices across the United States between 2000 and 2023. The data on office expansions are obtained through manual searches of the *archive.org* platform and cross-checked for the years 2001 to 2013 with information from Weinstein (2025).

The identification also leverages information on flight times across MSAs. We use the *28DS-T-100 Domestic Segment Data* provided by the *Bureau of Transportation Statistics* (BTS), which include monthly information on flight connections by carrier in the US from 1990 onwards. We restrict the data to regular flight connections, defined as scheduled direct flights exceeding 50 miles in distance and, per year and direction, more than 1,000 passengers and at least one departure per week. We then aggregate all flights to the MSA-year level and calculate the travel time as the average ramp-to-ramp time across all flights between two MSAs in that period.[14]

13. For additional detail on the fuzzy merge algorithm, we refer to Van Dijcke et al. (2023).

14. To assign airports to MSAs, we obtain airport coordinates from the *Master Coordinate Table* provided by the BTS and map these coordinates into core-based statistical areas (CBSAs) using the

1.3 Descriptive Evidence and Potential Explanations

1.3.1 Aging at the Very Top

In Figure 1.1, we document a consistent increase in CEO age since the year 2000. For the overall sample of CEOs contained in BoardEx, average age rose by more than 10 years to 61 in 2023. Average age at appointment also increased noticeably from less than 48 to 55 years, suggesting that the aging trend is unlikely to arise solely from longer tenures, later retirement, or CEO entrenchment. The age increase is pervasive and not limited to specific segments of the distribution of newly appointed CEOs. Figure 1.B.1, Panel (a), plots the quartiles of the age distribution over time, revealing an upward shift throughout the sample period. Similarly, in Panel (b), we observe the same trend for all CEOs, irrespective of their order of appointment.

Next, we examine heterogeneity in the CEO age trend across the firm size distribution, with results shown in Figure 1.2. The binscatter plot in Panel (a) demonstrates that average CEO age at appointment and firm size exhibit a strongly positive relationship. In Panel (b), we separately plot age at appointment for firms in the largest size quartile and for smaller firms, showing that the latter group is primarily responsible for the overall age increase.

To place these findings from our sample period starting in 2000 into historical perspective, we draw on data on 1,000 executives from the *Great American Business Leaders of the 20th Century* database. As we show in Figure 1.B.2, the average CEO age in this sample barely changed between 1940 and 1990. However, this is due to a combination of an increase in age at appointment alongside a decline in tenure. Moreover, in contrast to the bulk of the literature (e.g., Frydman 2019), we observe an increase in average CEO age during the 1990s, possibly because our data contain substantially more annual observations and, thus, a larger number of smaller firms.

1.3.2 Potential Explanations

One explanation for the CEO age trend is that it reflects an efficient market response to changes in the importance and availability of experience among CEOs. Given that many features of firms and the economy more broadly have changed since 2000, we first consider a range of potential alternative explanations. This section reviews several competing narratives and argues that none of these convincingly account for the documented age patterns.

CEO Aging as a Demographic Inevitability: One possible explanation centers on demographic shifts, which have accompanied the age increase among executives.

2013 version of the *TIGER/Line Shapefiles*. We then manually aggregate the CBSAs into the MSAs as defined by Revelio Labs.

Figure 1.2. CEO Age at Appointment and Firm Size (Number of Employees)

(a) Binscatter Average CEO Age vs. Firm Size

(b) Average CEO Age over Time by Firm Size

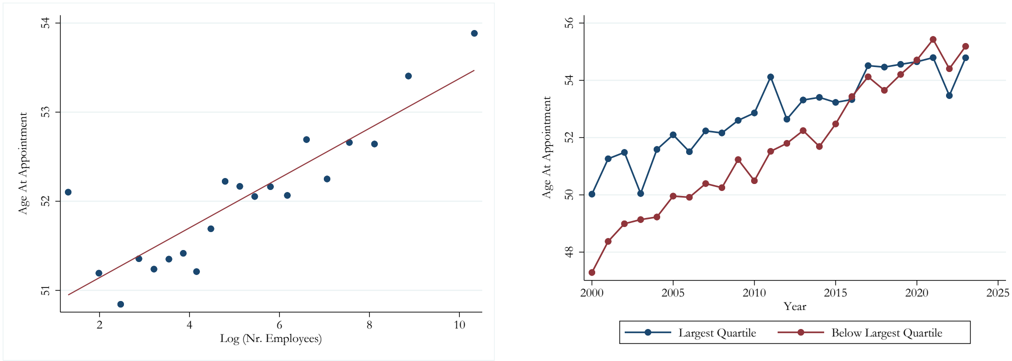

Notes: Panel (a) shows a binscatter plot of age at appointment against the natural logarithm of the number of employees. Panel (b) plots age at appointment separately for those in the largest firm quartile and for those in the other three quartiles. Size quartiles are generated separately by year.

We present several pieces of evidence that suggest that demographics can only account for a small fraction of the CEO aging trend. First, as we report in Figure 1.B.3, Panel (a), the age increase among CEOs in the United States is more than three times that of the overall labor force. This gap is even more pronounced when comparing to college-educated workers, who arguably constitute the relevant baseline population. Second, we document a similar increase in CEO age across a sample of countries in Europe (Figure 1.B.4), despite large differences in demographic trajectories and the timing of cohort sizes such as the baby boomers. Relatedly, our long-run analysis suggests that average CEO age has remained stable for most of the 20$^{\text{th}}$ century, in spite of the substantial demographic changes during this period. Third, we further examine the role of the age distribution in driving the aging trend by leveraging the cross-country variation in our data, with results reported in Figure 1.B.5. Specifically, we regress average CEO age by year and country on the share of the labor force within each 10-year age bin, controlling for country and year fixed effects (Panel b). Multiplying the estimated coefficients by the change in each bin's share over time (Panel a) and summing across bins, we find that shifts in the labor force age distribution account for 1.9 years of the overall increase in CEO age.

Optimal Response to Managerial Entrenchment: A second competing explanation refers to possible increases in rent seeking or entrenchment by CEOs. While entrenchment may directly raise the age of sitting CEOs by extending their tenure, our focus is on age at appointment, which is unaffected by this mechanism. However, entrenchment may also impact age at appointment indirectly through an anticipation channel. Specifically, boards of directors may attempt to prevent entrenchment by appointing older CEOs, who are potentially easier to replace after a shorter tenure. An empirical implication of this conjecture is that age at appointment and tenure should become more (negatively) correlated over time. Figure 1.B.6 examines the joint distribution of age at appointment and tenure by plotting their time-varying correlation. As expected, the two variables are negatively correlated, but the association attenuates throughout our sample period, which contradicts the hypothesis.[15]

Rising Market Concentration: A third potential explanation is that rising market concentration may have influenced CEO appointments. As shown by Bao, De Loecker, and Eeckhout (2022), increasing market power accounts for a substantial share of the growth in top executive compensation. Relatedly, Fee and Hadlock (2000) document a positive link between product market concentration and management turnover. In our context, higher market concentration could also imply a reduced supply of CEO positions at smaller firms. To test this idea, we correlate CEO age at appointment with market concentration measures from the Economic Census. The results, presented in Table 1.C.2, reveal a statistically insignificant association (with a contradictory sign).

Changes in Education Paths: A fourth set of hypotheses centers on the role of human capital in the labor market for CEOs, either due to prolonged educational paths or due to changes in skill accumulation on the job. We examine the former using education data from BoardEx, with results shown in Figure 1.B.7. Overall, we observe an increase of approximately one year in the age at highest degree over the sample period (Panel a). This shift is primarily driven by higher age at graduation among postgraduate degree recipients and by a rising share of CEOs holding such degrees (Panels b and c). Given the modest size of these changes relative to the aging magnitudes we have documented, we conclude that extended education alone cannot account for the rise in CEO age. In contrast, the observed increase in the length and diversity of prior experience supports the view that on-the-job human capital accumulation is the key driver of our results.

Another line of reasoning in this context suggests that firms may have reduced their investment in executive training since the beginning of our sample period. Consequently, only older candidates would have the necessary experience to be appointed CEO. However, the evidence points in the opposite direction: the employment share in professional and management development training has more than doubled (Figure 1.B.8), alongside a broader expansion in business education (Acemoglu, He, and Le Maire 2022).

Sample Composition: A fifth concern is that BoardEx oversamples larger and older firms relative to the population of US firms. To address this consideration,

15. These results are consistent with the overhaul in corporate-governance regulations since the 1980s (e.g., Edmans, Gabaix, and Jenter 2017; Ma and Shleifer 2025), which has increased scrutiny over managerial actions and reduced the potential for entrenchment. Moreover, Dow (2013) suggests that entrenchment plays a limited role in private firms, which are the focus of this paper.

we compare our sample to the Business Dynamics Statistics and reweight CEO age at appointment by the inverse frequency of firms in each category (firm size, firm age, NAICS-2 or NAICS-4 industries) relative to the population. Figure 1.B.9 shows that the 7–8 year increase in CEO age is robust across all reweighting schemes, confirming that compositional shifts cannot explain the age increase at appointment. In Appendix Section 1.E, we also document that the more pronounced age trend among smaller firms is not a mere artifact of the sample composition.[16]

Changes in Firm- and Individual-Level Observables: Sixth, we assess to what extent the age trend can be explained by observable factors. To this end, we regress CEO age at appointment on a linear time trend starting in 2000 and normalized to range between 0 and 1. We sequentially introduce an array of covariates. The estimated coefficients on the time trend are shown in Figure 1.B.10. Absent controls, the predicted age difference between 2000 and 2023 is more than 7 years. This estimate remains stable after accounting for fixed effects at the NAICS-4 level, a firm's listing status, and whether the CEO was hired internally or is female. Next, we consider firm size and age. Controlling for the former is essential given the interplay between age, accumulated skills, and firm size. Firm age is also important, as firm formation tends to be concentrated among younger individuals (Liang, Wang, and Lazear 2018), though to a lesser extent than commonly assumed (Azoulay et al. 2020). An increase in firm age could thus entail a mechanical increase in CEO age. Notably, when adjusting for the natural logarithm of employment and firm age, the estimate shrinks only slightly. However, the coefficient drops markedly to less than 4 years once we restrict the data to the firm sample contained in Compustat, and remains unchanged thereafter. Hence, while changes in firm fundamentals cannot mechanically account for the rise in CEO age, the results suggest a crucial role of smaller and unlisted firms in driving the aggregate patterns.

Reversion to Older CEOs after the Dot-com Crisis: A final possible concern relates to the timing of the starting point of our sample. Specifically, the bursting dot-com bubble in the late 1990s may have disproportionately affected young firms led by young CEOs, thereby mechanically increasing average CEO age afterwards. To assess this, we combine data from Execucomp on executives in the S&P 1500 between 1992 and 2000 with delisting information from CRSP. We then test whether firms delisted for involuntary reasons or through mergers and acquisitions were managed by younger CEOs. The estimates are shown in Figure 1.B.11. For sitting CEOs, we find that firms delisted for involuntary reasons had in fact significantly

16. This result also speaks against the role of superstar firms being the driving force behind the overall age increase. Notably, the narrowing gap in CEO age between smaller and larger firms is consistent with broader structural trends, including the declining contribution of superstar firms to aggregate US productivity growth since 2000 (Gutiérrez and Philippon 2019).

younger CEOs, whereas this pattern does not hold for M&A-related delistings. Crucially, however, we find no such difference in CEO age at appointment. If anything, CEOs appointed toward the end of the sample window at delisted firms were older than those appointed at firms that did not delist. An interpretation of these patterns is that delisted firms were disproportionately managed by their founders, who entered office at relatively young ages before the beginning of the sample and remained in place throughout the 1990s, thereby aging beyond the sample mean. Still, the evidence does not support the notion that our results are driven by a rebound effect following the late 1990s.

1.3.3 Changes in Career Paths and Generalist Skills

We now examine how prospective CEOs allocate the additional time prior to reaching top positions and the extent to which this increase in experience reflects the accumulation of generalist human capital. In Figure 1.3, we highlight the increased importance of external work experience. Panel (a) displays the average number of years spent outside and inside the firm over time. External work experience has increased strongly since the early 2000s, fully accounting for the upward trend in CEO age, whereas internal experience—measured as the time between joining the firm and becoming CEO—has remained relatively stable throughout our sample period. As we discuss below, the share of internal appointments also shows little variation over time. Zooming in on external mobility, Panel (b) reveals a sharp increase in job mobility and the number of external positions held before entry into the firm. Before assuming the CEO role, individuals nowadays spend more time in multiple roles across a wider range of firms and sectors. At the same time, the number of years spent in each position, firm, and sector has decreased substantially since 2000 (Panel c).[17] We interpret these patterns as evidence in favor of the idea that prospective CEOs transition across different positions, firms, and sectors to gather a broader skill set—pointing to a shift toward a *boundaryless career* in this segment of the labor market (Arthur and Rousseau 1996).[18]

Even though our results highlight the importance of generalist skills, the broadening of career paths has not unfolded uniformly across backgrounds. In Figure 1.B.12, we show that an increasing share of CEOs previously held positions

17. These results may be surprising, given that they stand in contrast to a strand of literature that shows a slight increase in tenure among the overall workforce during our sample period (Farber 2010; Hyatt and Spletzer 2016).

18. In addition to reflecting skill accumulation, years of work experience can be an important factor on their own, as they may indirectly reveal a candidate's age during the interview process. In the United States, the *Age Discrimination in Employment Act* prohibits age-based discrimination in the workplace for individuals over 40. Although it is not illegal, asking applicants directly about their age is uncommon and generally considered inappropriate. This lack of explicit age information during interview processes may be less pertinent in recent years, as birth dates have become publicly accessible in many cases.

Figure 1.3. Work Experience Prior to CEO Appointment over Time

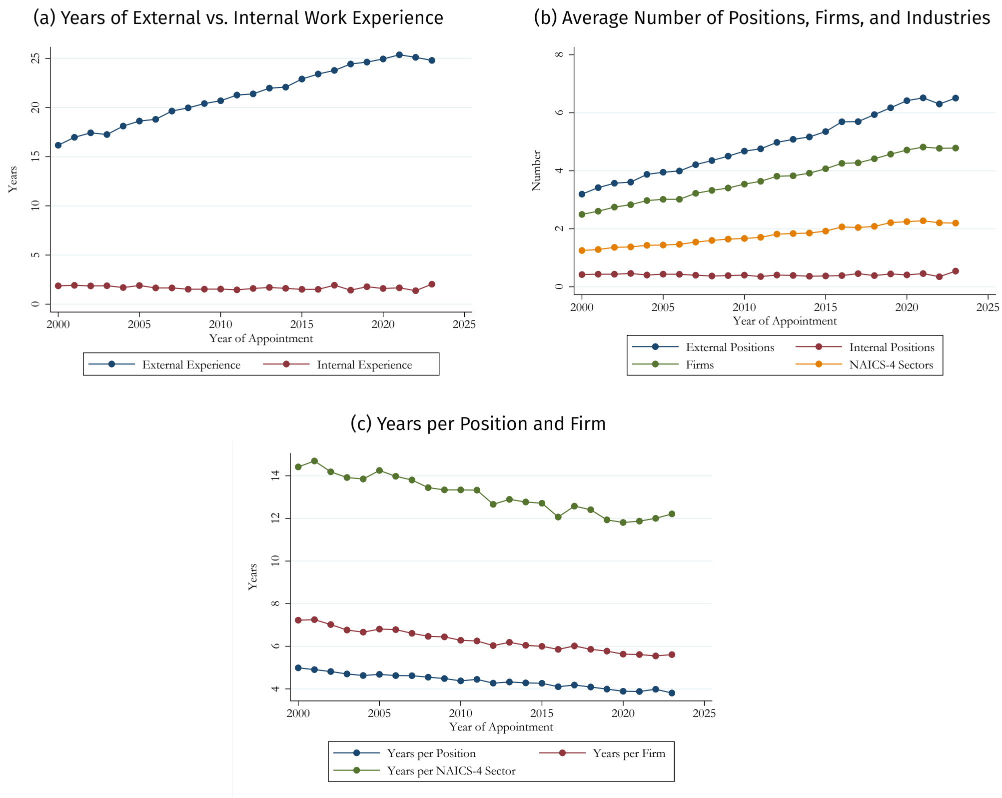

Notes: The figure plots work experience of CEOs over time. Panel (a) plots average external and internal work experience prior to appointment. Panel (b) shows the average number of positions, separately for external and internal positions, as well as the number of firms and industries a prospective CEO worked in. Panel (c) shows the average number of years per position and firm.

as CFO or COO, but this trend is confined to those who served in these roles at other companies. Furthermore, Decressin, Kaplan, and Sorensen (2025) document that CEOs appointed after the Great Financial Crisis score more highly in analytical skills and execution orientation, at the expense of lower levels of interpersonal skills and charisma.

As another way to highlight the rising importance of external experience, we distinguish between externally and internally appointed CEOs, further stratifying the latter into those entering the firm at the managerial level or below. Within this classification, three factors may account for the rise in CEO age: increases in (*i*) external or (*ii*) internal experience within each group, or (*iii*) level differences in CEO age across groups, combined with shifts in group composition. As Figure 1.B.13 shows, a combination of (*i*) and (*iii*) drives the age trend. Average age rose across all groups but more so for externally appointed CEOs, while the level was highest for those entering at the managerial level (Panel a). Panel (b) shows that since

2000, the share of external appointments has been stable, while managerial-level entries nearly doubled as lower-level entries declined. Meanwhile, both internally and externally appointed CEOs join firms at later career stages with more external experience (Panel c), whereas internal experience grew only for the shrinking group of those entering at less senior positions (Panel d). Overall, late entry, both within and across ports of entry, has thus become increasingly common.

In our sample, the share of internally appointed CEOs is relatively low.[19] As we show in Table 1.C.3, this is due to the comprehensive coverage of smaller firms, which are more likely to hire CEOs externally, whereas larger firms tend to favor internal candidates.[20] Among other factors, this heterogeneity across the firm size distribution has been ascribed to differences in internal talent pools among potential applicants and corporate governance practices (Kaplan and Sorensen 2021).[21] Given that prospective CEOs are gaining increasing amounts of external experience (Figure 1.B.15), the prevalence of external hiring among smaller firms is consistent with our finding that these firms account for most of the upward trend in CEO age.

Lastly, we examine the extent to which these time trends among CEOs differ from those of other senior employees who did not become CEOs. If executive skill requirements have evolved, these changes should be most pronounced among CEOs relative to less senior executives and increasingly reflected in their career paths over time. To draw this comparison, we classify all non-supervisory directors in the BoardEx data based on the most senior position they reached throughout the sample and analyze their career paths leading up to the first appointment at that seniority level. In Figures 1.B.16 and 1.B.17, we replicate Figure 1.3, showing the differential effects for CEOs compared to non-CEO executives and non-executive managers, respectively. In terms of levels, CEOs accumulate more external experience than both other groups, spanning a greater number of positions, firms, and industries. These differences are stronger for non-executive managers, who are more distant from the CEO position. Over time, most of these gaps become larger. Notably, the increasing accumulation of external experience is driven by the number of external positions for CEOs relative to non-CEO executives. For CEOs compared to non-executive managers, this development is to a greater extent driven by years of experience across different firms and industries. In the latter comparison, internal experience also exhibits a widening gap, owing to a decline among non-executive managers (Figure

19. With the data processing approach used in the main paper, around 20% of the appointments are internal. Under the alternative preparation method described in Appendix Section 1.E, the share is at 32%.

20. This result holds when firm size is measured by employment and becomes even more pronounced when stratifying by a firm's listing status. Figure 1.B.14 illustrates that these two effects are cumulative, as a firm's listing status remains relevant even after accounting for the number of employees.

21. In turn, these hiring practices directly affect economic outcomes, as more productive firms recruit more talented trainees and promote managers internally (Friedrich 2023).

1.B.17, Panel a). Taken together, this suggests not only that broad external experience is an important factor in the selection of CEOs, but also that its relevance has grown over time.

1.4 Model

To provide a framework for understanding and interpreting the patterns we have documented, we study a frictionless many-to-one matching model with transfers between workers and firms. Workers are heterogeneous along two dimensions: ability and experience at appointment.[22] Firms differ in productivity and feature multiple hierarchical positions with associated productivity requirements.[23] Assortative matching implies that more productive firms are matched with more productive CEOs, which we capture empirically by proxying firm productivity with firm size.[24] The model predicts that CEOs in higher-productivity (larger) firms are older on average, generates a negative correlation between CEO ability and age, and produces an aging of CEOs in response to an increase in the importance of experience for CEO productivity.

Framework: There is a continuum of workers and firms. Workers are indexed by ability $a \in [a_L, a_H]$ and age $t \in [0, T]$. Firms have productivity $z \in [\underline{z}, \bar{z}]$ with distribution G and *capacity* (mass of job positions) $m(z) \geq 0$. Total capacity in the entire economy is $M := \int m(z)\, dG(z)$. Positions or job slots are a continuum, with the *top slot* in each firm being a measure-zero point. Within each firm, positions are indexed by k, with productivity

$$z(k) \;:=\; z - \beta(k-1), \qquad \beta \;>\; 0,$$

so that higher-ranked positions are associated with higher productivity. For the formal analysis in the appendix, it is convenient to work with the induced one-dimensional distribution of slot productivity generated by (z, k). In particular, CEO slots correspond to $k = 1$, implying that the productivity of the CEO equals firm productivity z.

Worker productivity has two components: an *age-adjusted ability* component $x(a, t)$ and an *experience* component $y(t)$. We adopt the additive specification

$$x(a,t) \;\equiv\; a + h(t),$$

22. This distinction parallels a large literature in psychology that differentiates between two forms of cognitive functioning (Cattell 1971): fluid intelligence, capturing the capacity to develop new skills independently from prior knowledge, and crystallized intelligence, reflecting abilities derived from previous learning and experience (Baltes and Mayer 1999; Schaie 2005; Salthouse 2009).

23. See Corblet (2025) and Choné, Gozlan, and Kramarz (2024) for recent empirical contributions on many-to-one matching with multidimensional skills.

24. We therefore use firm productivity and size interchangeably throughout this section.

and define the scalar *worker score*

$$\phi(a,t;\alpha) \quad := \quad (1-\alpha)[a+h(t)] + \alpha y(t), \qquad \alpha \in [0,1).$$

We assume that h is twice continuously differentiable, strictly concave, and single-peaked with maximizer t^*; y is twice continuously differentiable, strictly increasing, and concave.[25] We assume throughout that a and t are independent with continuous densities.

Firm-worker output at position k is multiplicatively separable in firm-side productivity:

$$\pi\big(z(k),a,t;\alpha\big) \quad = \quad z(k)\,\phi(a,t;\alpha).$$

This functional form is consistent with the idea that changes in CEO talent scale with firm productivity, or analogously, firm size (Gabaix and Landier 2008; Terviö 2008; Edmans, Gabaix, and Landier 2009).

Remark 1. Because utility is transferable and the surplus takes the form $\pi(s,a,t;\alpha) = s\,\phi(a,t;\alpha)$, which has increasing differences in (s,ϕ), the efficient and stable assignment is positively assortative between slot productivity and worker score: more productive slots are filled by higher-score workers.

Main Theoretical Results: Appendix 1.A includes a formal analysis of this environment and provides sufficient conditions for three core implications:

(i) *Larger (higher-z) firms have older CEOs on average.* Larger firms have more productive workers. If higher productivity makes higher ages more likely (a stochastic monotonicity property), then selecting workers with higher productivity leads to a stochastically older group. Positive assortative matching transfers this conclusion to firms: more productive firms, which match with higher-score workers, have stochastically older workers.

25. These functional forms are motivated by findings from psychology. The eventual decline of the function h reflects the evolution of energy levels and fluid intelligence. The latter typically increases during youth, peaks in early adulthood, and subsequently declines with age. Cross-sectional studies tend to identify this peak in individuals' twenties, followed by a rapid decline. Longitudinal studies suggest a later peak, occurring up to the forties or fifties, as well as a slower decrease afterwards (for a review, see Desjardins and Warnke 2012). The decline in fluid intelligence correlates with a general slowing of cognitive processes (Schaie 1989) as well as diminished attention, memory capacity, and motivation (Ng and Feldman 2008). The function y captures a variety of gains from age, including broadening networks, having faced a longer history of challenges to learn from, as well as crystallized intelligence. The latter tends to increase steadily until around the mid-fifties and then either remains stable or shows only a modest decline thereafter (Schaie 2005). This decline is mostly confined to individuals with lower skill-usage, while white-collar professionals and those with higher education levels continue accumulating skills beyond their forties (Hanushek et al. 2025).

(ii) *Age and ability of CEOs are negatively correlated.* As older workers benefit from the positive contribution of age to the score, they can clear a given threshold with lower ability, whereas younger workers must compensate with higher ability. This selection induces a negative relationship between ability and age among appointed CEOs.

(iii) *Increasing the weight on experience raises average CEO age.* Higher α shifts selection toward older workers at all ranks, which lowers the relative ability threshold required for older workers and increases the probability that older workers qualify for top positions. This generates a first-order stochastic dominance increase in the age distribution of CEOs.

To understand the intuition for part (iii), consider two simplifications of the model as benchmark cases. If there were only one firm, it is natural to expect that this firm would shift older workers to more senior positions when experience becomes more valuable. However, if there were many firms and only one worker per firm, such a reallocation cannot hold for every firm because of aggregation: if workers in some firms get older, others must get younger. In particular, as experience becomes more highly valued by all firms, the equilibrium shadow value of experience must also adjust, potentially leading some firms to economize on experience, especially in less senior positions. Thus, while the partial-equilibrium logic is straightforward, the general-equilibrium channel introduces conflicting forces. There are a variety of conditions one can impose to guarantee that the distribution of workers in top positions stochastically increases when α increases. In the Appendix we provide sufficient conditions in terms of the induced conditional tail densities.[26]

Figure 1.4 illustrates the theoretical predictions using a simulation and plots the predicted CEO age profiles across the firm productivity distribution for different values of α.[27] Given the link between firm productivity and size, the age profiles can be compared to the empirical age patterns over time across the firm size distribution. For a given α_1 (blue line), CEO age increases with firm size. When $\alpha_2 > \alpha_1$ (red line), CEOs are generally older than under α_1. In the given parametrization, the gradient in CEO age over the firm size distribution is also less pronounced than for α_1. This heterogeneity is driven by two forces. First, due to the concavity of y, any increase in α produces stronger reordering among younger workers. As smaller firms draw from that part of the distribution, their assigned CEO age shifts more. Second, larger firms already select high-ability workers due to assortativity, which constrains how much their CEO age can adjust when α increases. In the numerical example, the average correlation between experience and ability for $\alpha = \alpha_1$ is negative at -0.981, which drops to -0.997 for $\alpha = \alpha_2$. This discrepancy exists

26. A related option is to ensure that CEO positions are in a thin upper tail of the most productive workers in the economy, a condition that is difficult to describe qualitatively.

27. See the figure notes for details on the parametric specification.

Figure 1.4. Illustrative Age Profiles at the Top Position by Firm Productivity

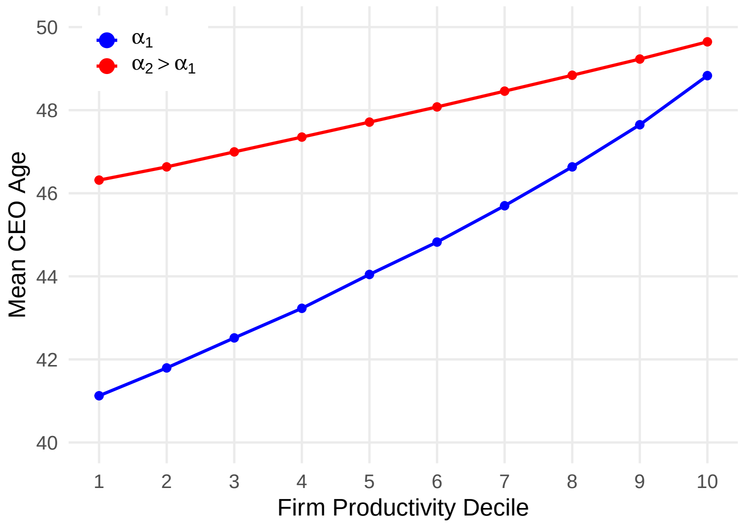

Notes: The figure plots illustrative CEO age profiles across firm productivity deciles for two values of α: $\alpha_1 = 0.4$ (blue line) and $\alpha_2 = 0.8$ (red line). We simulate 100,000 workers with ability $a \sim U(0, 200)$ and age $t \sim U(0, 50)$. Skills are defined as $h(t) = -0.04t(t-50)$ and $y(t) = -0.1t(t-150)$. On the firm side, we draw 1,000 firms with productivity $z \sim U(0, 1)$, capacity $m = 100$, and hierarchical decay parameter $\beta = 0.08$. Within each firm, the CEO corresponds to the top slot ($k = 1$). We exploit the assortative matching implied by the model and compute the average CEO age within each firm productivity decile. The simulation is repeated 500 times, and results are averaged across all replications.

across the entire firm size distribution, but is more pronounced for smaller firms (Figure 1.B.18).

Turnover: We conclude the theory discussion by outlining how the model speaks to the increasing turnover that we observe in our data. Although the model is static, it suggests a simple benchmark for turnover pressure when the importance of experience changes over time. Starting from a baseline environment with weight α_1, consider the workers who initially occupy the top θ positions. As α rises, these workers may move up or down in the frictionless ranking because experience becomes more heavily weighted. The Appendix summarizes this by measuring the cumulative rank movement of these baseline top workers as the economy transitions from α_1 to α_2. While this does not model realized turnover directly, it provides a simple measure of how much the target assignment would need to be reorganized in response to changes in α. A key implication is that larger increases in α generate greater reshuffling pressure, that is, they require more reallocation of workers across positions.

Empirical Evidence on Model Predictions: First, the model predicts that the age distribution of CEOs in larger firms first-order stochastically dominates that in smaller firms. Panel (a) of Figure 1.5 plots the cumulative distribution functions

Figure 1.5. Empirical Evidence for Theoretical Predictions

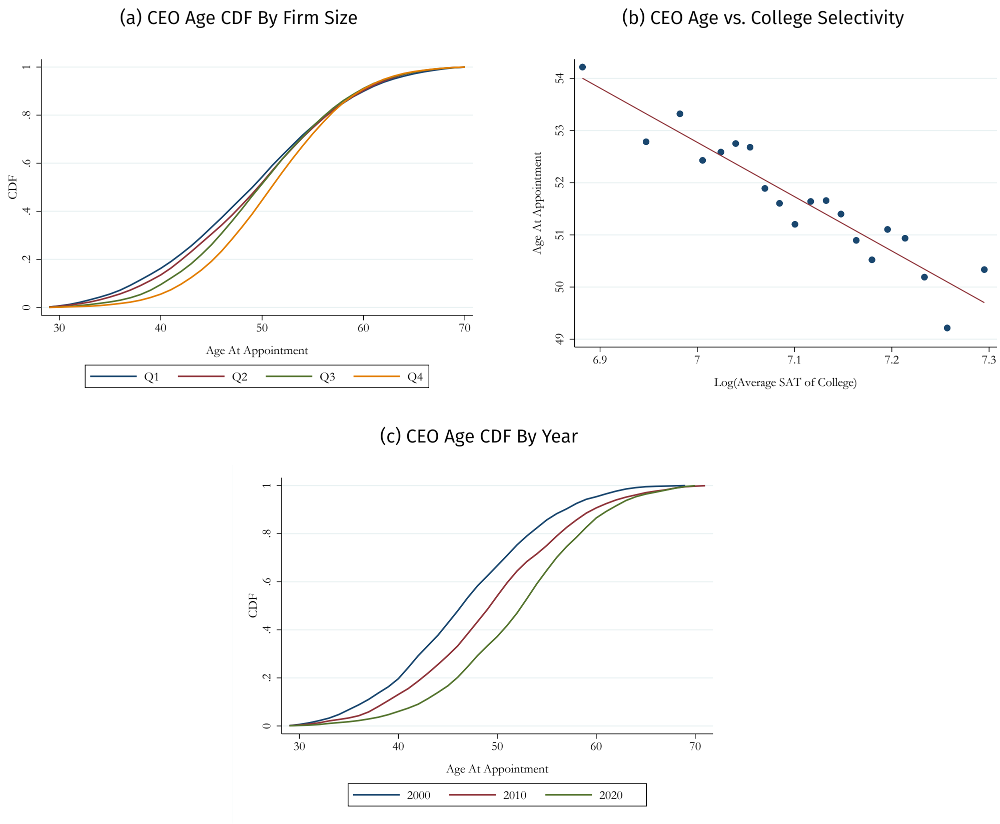

Notes: Panel (a) shows the cumulative distribution function of CEO age at appointment by employment quartile. CEO age and the natural logarithm of firm employment are residualized with respect to NAICS-4 and year fixed effects, with the sample mean added back. Employment quartiles are constructed from these residuals, and the top and bottom percentiles of the age distribution are trimmed. Panel (b) focuses on a CEO's first postsecondary degree, identified either by the stated degree name or, if missing, by the first degree recorded in BoardEx. We exclude first postsecondary degrees obtained from institutions outside of the United States. The plot binscatters age at appointment against the natural logarithm of the institution's average SAT admission score in 2001, taken from Chetty et al. (2020). Controls include the natural logarithm of employment and firm age. Panel (c) shows the cumulative distribution function of CEO age at appointment by year. CEO age is residualized with respect to the natural logarithm of firm employment and NAICS-4 fixed effects, with the sample mean added back, and the top and bottom percentiles of the age distribution are trimmed.

of CEO age at appointment by firm size quartile, after residualizing with respect to NAICS-4 industry and year fixed effects. The plots closely match the expected pattern: relative to smaller firms, larger firms appoint substantially fewer CEOs at younger ages, in particular below age 45.[28]

28. Consistent with this, Adams, Keloharju, and Knüpfer (2018) show that both cognitive and non-cognitive skills matter for CEO selection, with a stronger role for non-cognitive skills. These skills

Second, selection generates a negative association between CEO age at appointment and ability in the model. Even though age and ability are independently distributed in the population, conditioning on CEO appointment creates a collider bias, predicting a negative correlation between the two variables. To test this in the absence of a direct measure of ability, we use information on educational backgrounds to construct different proxies for college selectivity.[29] Consistent with our theory, college selectivity is negatively correlated with CEO age at appointment (Panel b of Figure 1.5 and Table 1.C.4), indicating that age and ability act as substitutes in the selection of CEOs.

Third, Panel (c) of Figure 1.5 plots the cumulative distribution of CEO age at appointment for the years 2000, 2010, and 2020, conditional on employment and NAICS-4 fixed effects. Assuming that the weight on experience has increased over time, we expect the entire distribution of CEO age at appointment to shift to the right throughout the sample period. Indeed, the age distribution in more recent years first-order stochastically dominates that in previous periods.[30]

Potential Extensions: First, for the concept of a CEO position to be meaningful, firms must have multiple, hierarchically ordered positions. In the model, CEO productivity is exogenous, and the remainder of the hierarchy plays only a limited role. However, the model could be extended to allow productivity to depend on the team composition, which would potentially generate additional richness without changing its main qualitative predictions. Second, the model considers firm heterogeneity in terms of productivity, which we empirically proxy by firm size. A natural extension with equivalent results would be to allow firm size $m(z)$ to depend on productivity, with $m(z)$ increasing in z.[31] Third, the model assumes that all workers are matched to some firm. Allowing for potential unemployment would imply that the least productive and disproportionately young workers remain unmatched, which would strengthen our existing results.

include traits such as social maturity, emotional stability, and persistence, which may partly reflect accumulated experience.

29. As college attendance is correlated with SAT performance and, in turn, with intelligence (for a review, see Frey 2019), these measures closely mirror the ability component used in the theoretical analysis.

30. Across the firm size distribution, the heterogeneous age trends (Figure 1.2) closely mirror the simulated response to changes in α (Figure 1.4).

31. We could further allow for heterogeneity in α. If larger firms tend to place higher value on experience in general and if smaller firms become more similar to larger ones over time, our results should hold under less stringent assumptions.

1.5 Identifying Demand and Supply of Generalist Skills

1.5.1 Demand for Generalists in the Face of Greater Uncertainty and Complexity

So far, we have documented that CEOs have become more mobile between positions and firms prior to their appointment, possibly in order to gather generalist skills across a range of functional areas. In this section, we propose a demand-side channel: heightened uncertainty and complexity may require executives with more generalist skill sets,[32] resulting in firms appointing executives who have taken longer career paths before reaching CEO positions and are, thus, older upon appointment. We seek to identify the causal effect of changes in firms' economic environment on the characteristics of appointed CEOs, in particular their level of generalist skills. We start with uncertainty and later extend our empirical framework to a range of measures capturing economic complexity. The rationale for this approach is that while uncertainty and complexity have generally increased over time (cf. Figure 1.B.19), different components thereof may matter more at different points in time, about which we remain agnostic.

Empirical Strategy: The ideal experiment would induce random variation in firm-level uncertainty or complexity, bolstering demand for generalist skills among some firms. However, finding exogenous variation in uncertainty or complexity that operates solely through this channel is challenging. Moreover, even then, comparing CEO age across high- and low-uncertainty (complexity) firms would recover the estimand of interest only if the supply of potential CEOs were sufficiently elastic.

To address these issues, our empirical approach exploits information on both the demand and the supply side. Specifically, we use plausibly exogenous variation in the availability of generalist skills and study its *differential* effect on CEO appointments in high- vs. low-uncertainty (complexity) environments. As a supply shifter, we rely on spatial variation in firms' access to elite strategy consultancies across the United States. These consultancies, with a broad outlook across firms and industries, provide firms with additional manpower in their management teams, allowing them to quickly adjust to changing market environments (Ernst and Kieser 2002; Glückler and Armbrüster 2003).[33] In other words, the presence of external strategy consultants reduces firms' dependence on hiring older CEOs as a means to navigate uncertainty. Consultants effectively substitute for these experienced executives by supplying the generalist skills that become particularly valuable in high-

32. This reasoning is supported by management studies arguing that firm diversification and industry complexity increase the value of generalist skills (see, e.g., Mueller et al. 2021).

33. The empirical evidence highlights that consultancies improve firm performance through productivity-enhancing channels, in particular restructuring, with little support for rent-seeking effects (Bruhn, Karlan, and Schoar 2018; Bijnens, Jäger, and Schoefer 2025).

uncertainty or high-complexity environments. Therefore, we expect that firms in areas with better access to elite consultancies are less likely to appoint older CEOs when faced with higher uncertainty (complexity). This reasoning builds on Yonker (2017), who documents the prevalence of CEO appointments from local labor markets, and on Knyazeva, Knyazeva, and Masulis (2013) and Sauvagnat and Schivardi (2024), who emphasize the role of the local supply of managerial talent, specifically the thickness of the market for executive skills, in shaping firm-level outcomes.

On the supply side, the empirical strategy is based on the idea that the consulting industry is highly concentrated, but to a heterogeneous extent across time and space. We focus on the three firms commonly considered the most prestigious: *McKinsey & Company*, *Boston Consulting Group*, and *Bain & Company*. Throughout our sample period, these companies (henceforth MBB) have seen a substantial expansion, increasing their number of offices from 25 in 2000 to 47 in 2022 (Figure 1.6, Panel a). Despite this growth, their offices remain clustered in relatively few cities. As illustrated in Panel (b) of Figure 1.6, while the number of MSAs with any MBB presence has increased substantially, so has the number of MSAs hosting all three MBB offices. In 2022, this latter group accounted for roughly 80% of all offices, concentrated in merely 12 cities. This spatial variation will constitute one ingredient in capturing firms' exposure to consultancy services in our empirical approach.

We hypothesize that increasing the geographical proximity between consultancies and their clients enhances the dominance of the consultancies, as consultants typically work on-site, arriving from their office location and staying from Monday until Thursday. In other words, having an MBB office nearby implies an increase in the availability of generalist skills. However, the timing and location of new office openings may be correlated with existing firm networks (Glückler 2006) or local economic conditions, since consultancies tend to expand into areas with high demand for their services. To address this potential endogeneity issue, we do not use direct exposure to an MBB office (i.e., a consultancy being located in the same MSA) but, instead, construct a proxy for indirect exposure by exploiting variation in flight times. This approach, inspired by Giroud (2013), creates additional variation that we exploit in our empirical analysis.

Starting with uncertainty, we can spell out the regression equation of interest. Specifically, considering a CEO appointment in firm j, MSA m, industry i and year t, we estimate the following specification:

$$Age_{jt} = \beta\, Uncertainty_{it} \times Travel\ Time_{mt} + \gamma_{mt} + \phi_{it} + X_{jt} + \epsilon_{jt}, \tag{1.1}$$

where Age_{jt} denotes CEO age at appointment obtained from BoardEx, $Uncertainty_{it}$ is an uncertainty (or complexity) measure at the industry-time level[34], and

34. For the baseline specification, we use the cross-sectional standard deviation of the symmetric growth rate of firm revenue within each NAICS-4 industry.

Figure 1.6. Expansion of MBB Offices Across the United States

(a) Number of Offices (b) Concentration of MBB Offices

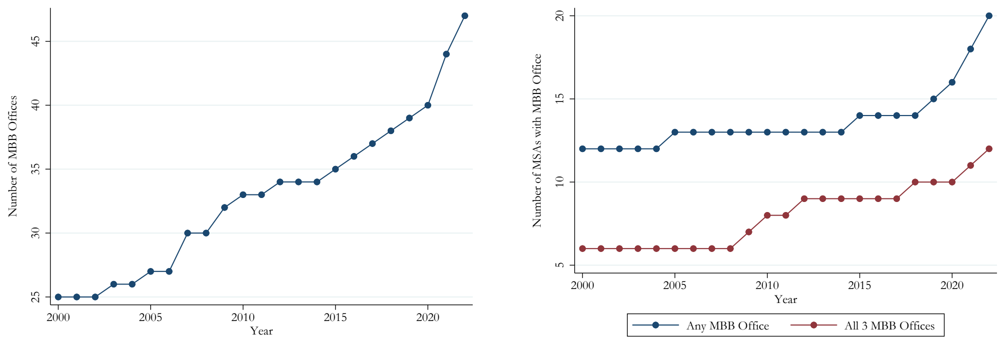

Notes: The figure shows the evolution of MBB offices (*McKinsey & Company*, *Boston Consulting Group*, and *Bain & Company*) in the United States over time. Panel (a) shows the total number of offices, and Panel (b) plots the number of MSAs with at least one or all three offices, respectively. The data are obtained from *archive.org* and cross-checked with information from Weinstein (2025).

$\textit{Travel Time}_{mt}$ is the flight time to the closest MBB office. The fixed effects γ_{mt} and ϕ_{it} are at the MSA-year and NAICS-4-year level, respectively, eliminating the standalone coefficients on $\textit{Uncertainty}_{it}$ and $\textit{Travel Time}_{mt}$. Finally, the vector of controls, X_{jt}, includes the natural logarithm of the number of employees and of firm age.

The coefficient of interest is β, which captures the differential effect of higher uncertainty on CEO age upon appointment, conditional on the availability of young generalists (i.e., MBB consultants). This interpretation hinges on two assumptions. First, proximity to an MBB office tilts the age composition of the available generalist talent pool, expanding firms' access to younger generalists more than to older ones. Second, the part of the variation in uncertainty that is associated with CEO age responds to travel time due to changes in firms' demand for generalist skills. The fixed effects control for all time-varying characteristics at the levels of MSAs and industries, implying that our empirical approach relies solely on the fluctuations generated by the interaction of uncertainty and travel time. We expect firms to appoint older CEOs if uncertainty is high and the available supply of other generalists is low due to exogenous reasons, as experienced executives typically provide the generalist expertise that firms seek in such situations. Since travel time is negatively related to the availability of young generalists, this implies a positive sign for the estimated β as firms in high-uncertainty industries are then left with an older pool of available generalists.

We emphasize that this exercise is *not* aimed at suggesting that increased MBB exposure causally contributed to the selection of older CEOs. If anything, one would expect the opposite effect, given the expansion of offices and flight routes in conjunction with the substitutive relationship between consultants and more experienced executives. Instead, we exploit the supply of consultants as a source of variation to probe our mechanism of interest.

To measure firms' travel time to the closest MBB office, we rely on time-varying flight times that we compute as described in Section 1.2, considering only direct connections. Combining these data with information on the expansion of MBB offices across the United States, we calculate the minimal flight time to any MBB office at the MSA-year level. The final sample covers all observations between 2000 and 2022, excluding those in which an MBB office is present in a given year.[35] As we show in Appendix Table 1.C.5, reductions in travel time translate into higher inflows of employees from these firms. Conditional on time-varying firm-level controls as well as fixed effects at the firm and industry-year level, a one-hour reduction in flight time increases inflows by about 19% (Column 1).

Results on Uncertainty: We use the cross-sectional standard deviation of the symmetric growth rate of firm revenue within each NAICS-4 industry as our main measure of uncertainty. This variable has the advantage that it captures broader, industry-wide fluctuations and is unaffected by idiosyncratic shocks experienced by individual firms. Appendix Figure 1.B.20 plots the measure over time, alongside uncertainty indices from Baker, Bloom, and Davis (2016) and Baker et al. (2019). All three series reveal a secular increase in uncertainty since around 2000 compared to the 1980s and 1990s.

The results obtained from estimating Equation (1.1) are shown in Table 1.1. Column (1) includes the results from a regression with controls only, showing a positive and significant coefficient on the interaction term. By consecutively adding fixed effects for MSAs, NAICS-4 industries, years, and combinations thereof, the estimated coefficient remains consistently significant at the 5%- or 1%-level. Quantitatively, the estimated coefficient from our preferred specification (Column 4) is large and lends itself to an intuitive interpretation: One additional hour of travel time, evaluated at one standard deviation above mean uncertainty, translates into an increase in CEO age of 0.6 years.

We provide additional evidence for our mechanism in four ways. First, we corroborate the result from Table 1.1, Column (4) by conducting a number of robustness checks, presented in Table 1.2. In Column (1), we use the standard deviation of revenue growth at the NAICS-2 level, rather than the NAICS-4 level, as a proxy for uncertainty. The estimated coefficient remains comparable, supporting the idea that broader economic turbulences shift the demand for older CEOs. Similarly, in the remaining regressions, the results remain highly significant and quantitatively

35. We impose this restriction because direct exposure to office entry is more likely to reflect strategic considerations, whereas exposure through flight connections is less prone to such endogeneity concerns. As we show in Column (5) of Table 1.2, however, our main result is insensitive to this choice.

Table 1.1. MBB Availability, Uncertainty, and CEO Age

	CEO Age			
	(1)	(2)	(3)	(4)
Travel Time	0.216	-0.710		
	(0.271)	(0.893)		
Industry-Level SD	-0.471	-0.657**	-0.732*	
	(0.338)	(0.307)	(0.378)	
Travel Time × Industry-Level SD	0.454**	0.625***	0.476**	0.632***
	(0.186)	(0.232)	(0.212)	(0.212)
R^2	0.033	0.088	0.230	0.360
Observations	15,679	15,676	15,235	14,132
MSA fixed effects		✓		
Year fixed effects		✓		
MSA-Year fixed effects			✓	✓
Naics4 fixed effects			✓	
Naics4-Year fixed effects				✓

Notes: The dependent variable in all regressions is CEO age upon appointment, obtained from BoardEx. *Travel Time* denotes the flight time (in hours) to the closest office, which we construct from MBB office openings combined with flight data from the Bureau of Transportation Statistics. *Industry-Level SD*, calculated as the cross-sectional standard deviation of the symmetric growth rate of firm revenue within a firm's NAICS-4 industry, is standardized to mean zero and standard deviation one. The sample is restricted to those MSAs without MBB presence. All regressions control for firm size by including the natural logarithm of employment and firm age (extracted from the employment history data). Standard errors are double-clustered at the NAICS-4 industry and MSA level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

consistent across the board. In Column (2), we replace the standard deviation of revenue growth with stock return volatility. In Column (3), we use lagged values of travel time. In Column (4), we discretize the linear measure of travel time to the nearest MBB office by constructing a dummy for above-median values of travel time. In Column (5), we examine the influence of the current age of the firm's other executives, a factor that could play a significant role in CEO selection. In the case of internal appointment, other executives are part of the pool of potential candidates, while for external appointments older board members may favor older CEOs due to similarity attraction (Davidson, Nemec, and Worrell 2006). However, the estimated coefficients remain qualitatively unaffected. In Column (6), we use the full sample, which includes metro areas with an MBB office and, thus, a travel time of zero. Lastly, in Column (7), we restrict the sample to firms in metro areas that do not experience any office entry throughout the entire sample period.

Second, we provide more direct evidence that the treatment shifts CEO appointments in smaller firms toward older candidates with more generalist skills, while leaving CEO appointments in larger firms essentially unchanged. We do so by splitting the sample at the median employment level and reporting the corresponding

Table 1.2. MBB Availability, Uncertainty, and CEO Age—Robustness

	CEO Age						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Travel Time × Industry-Level SD	0.527**	0.702**	0.640***	0.366**	0.528*	0.628***	0.572*
	(0.219)	(0.315)	(0.205)	(0.149)	(0.284)	(0.148)	(0.327)
R^2	0.365	0.371	0.360	0.360	0.381	0.252	0.397
Observations	14,954	13,856	14,132	14,132	12,289	33,006	10,912
MSA-Year fixed effects	✓	✓	✓	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓	✓	✓	✓
Change Compared to Baseline	NAICS-2 SD	Return Volatility	Lagged	Discretized	Board Age	Full Sample	Never Office

Notes: The dependent variable in all regressions is CEO age upon appointment, obtained from BoardEx. *Travel Time* denotes the flight time (in hours) to the closest office, which we construct from MBB office openings combined with flight data from the Bureau of Transportation Statistics. *Industry-Level SD*, calculated as the cross-sectional standard deviation of the symmetric growth rate of firm revenue within a firm's NAICS-4 industry, is standardized to mean zero and standard deviation one. The sample is restricted to those MSAs without MBB presence. All regressions control for firm size by including the natural logarithm of employment and firm age (extracted from the employment history data). We deviate from this setup in the following ways: Column (1) uses the standard deviation of revenue growth within NAICS-2 sectors; Column (2) uses stock return volatility, calculated as the within-year volatility of monthly stock returns aggregated up by taking the NAICS-4-specific median, with data taken from the Center for Research in Security Prices; Column (3) uses the first lag of travel time; Column (4) replaces the linear measure of travel time to the nearest MBB office with a dummy for above-median travel time; Column (5) controls for average executive age constructed from the employment history data; Column (6) uses the entire sample including those metro areas with an MBB office; Column (7) restricts the sample to firms in MSAs that do not experience any office entry throughout the entire sample. Standard errors double-clustered at the NAICS-4 industry and MSA level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

estimates in Panels A (smaller firms) and B (larger firms) of Table 1.3. Column (1) shows that the effect on CEO age is entirely driven by smaller firms. The remaining columns further probe our main mechanism. Columns (2) and (3) document that treated smaller firms appoint CEOs with significantly more experience outside of the current industry, measured either as years spent in other NAICS-4 or NAICS-2 industries. We also consider years within the same NAICS-4 industry (Column 4) and firm (Column 5) as placebos, both with insignificant responses. Columns (6) and (7) reveal an increase in the breadth of CEOs' prior industry experience among smaller firms, captured by the number of distinct NAICS-4 industries worked in, with very similar conclusions drawn from OLS and quasi-Poisson estimations. By contrast, the estimated coefficients for larger firms are consistently insignificant.

Third, we scrutinize our empirical approach through the lens of a difference-in-differences framework, for which we show event-study estimates in Figure 1.7. To this end, we estimate the effect of MBB entry in MSAs linked by a direct flight connection, improving access to the closest consultancy compared to the counterfactual of no MBB entry. Separate estimations for sectors with high and low levels of uncertainty reveal no differential pre-trend between the two groups and a positive response in low-uncertainty sectors. In contrast, CEO age in high-uncertainty sec-

Table 1.3. MBB Availability, Uncertainty, and Generalists by Firm Size

	CEO Age (1)	Yrs. Oth. NAICS-4 (2)	Yrs. Oth NAICS-2 (3)	Yrs. Same NAICS-4 (4)	Yrs. Internal (5)	No. NAICS-4 (6)	No. NAICS-4 (7)
Panel A: Smaller Firms							
Travel Time × Industry-Level SD	0.897**	0.573**	0.665***	-0.255	0.183	0.124***	0.127***
	(0.437)	(0.219)	(0.189)	(0.427)	(0.358)	(0.037)	(0.035)
Observations	5,932	5,932	5,932	5,932	5,932	5,932	5,696
MSA-Year fixed effects	✓	✓	✓	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓	✓	✓	✓
Estimation	OLS	OLS	OLS	OLS	OLS	OLS	Poisson
Panel B: Larger Firms							
Travel Time × Industry-Level SD	0.287	0.105	0.496	-0.564	-0.135	-0.099	-0.061
	(0.686)	(0.595)	(0.509)	(0.550)	(0.525)	(0.063)	(0.076)
Observations	6,055	6,055	6,055	6,055	6,055	6,055	5,539
MSA-Year fixed effects	✓	✓	✓	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓	✓	✓	✓
Estimation	OLS	OLS	OLS	OLS	OLS	OLS	Poisson

Notes: The dependent variables refer to the time, firm, or industry of CEO appointment: CEO age (Column 1), years of experience in other NAICS-4 industries (Column 2), years of experience in other NAICS-2 sectors (Column 3), years of experience in the same NAICS-4 industry (Column 4), years of international experience within the same firm (Column 5), and the number of NAICS-4 industries the CEO has worked in (Column 6 and 7). *Smaller Firms* are defined as those with below-median employment, while *Larger Firms* comprise those above the median. *Travel Time* denotes the flight time (in hours) to the closest office, which we construct from MBB office openings combined with flight data from the Bureau of Transportation Statistics. *Industry-Level SD*, calculated as the cross-sectional standard deviation of the symmetric growth rate of firm revenue within a firm's NAICS-4 industry, is standardized to mean zero and standard deviation one. The sample is restricted to those MSAs without MBB presence. All regressions control for firm size by including the natural logarithm of employment and firm age (extracted from the employment history data). Columns (1) to (6) are estimated using OLS, while Column (7) uses quasi-Poisson estimation. Standard errors are double-clustered at the NAICS-4 industry and MSA level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

tors declines gradually and persistently following the treatment. As an additional test for pre-trends, we conduct a placebo exercise, in which we replace travel time in levels by future changes relative to the current period. The estimated coefficients, shown in Figure 1.B.21, are statistically indistinguishable from zero across all horizons.

Fourth, we plot the coefficient on the interaction term separately for 5-year subperiods. The estimated coefficients, shown in Figure 1.B.22, can be interpreted as the elasticity of CEO age with respect to variation in uncertainty, evaluated at a given level of travel time. The estimates are consistently positive and statistically significant for the periods 2010–2014 and 2020–2023. Together with the increase since 2005–2009, this provides tentative evidence of a gradual rise in the value of generalist human capital in the aftermath of the financial crisis.

Results on Complexity: Beyond the secular shift in uncertainty, firms' simultaneous expansion across geographies and business areas, coupled with the need to

Figure 1.7. MBB Market Entry, Uncertainty, and CEO Age—Event Study

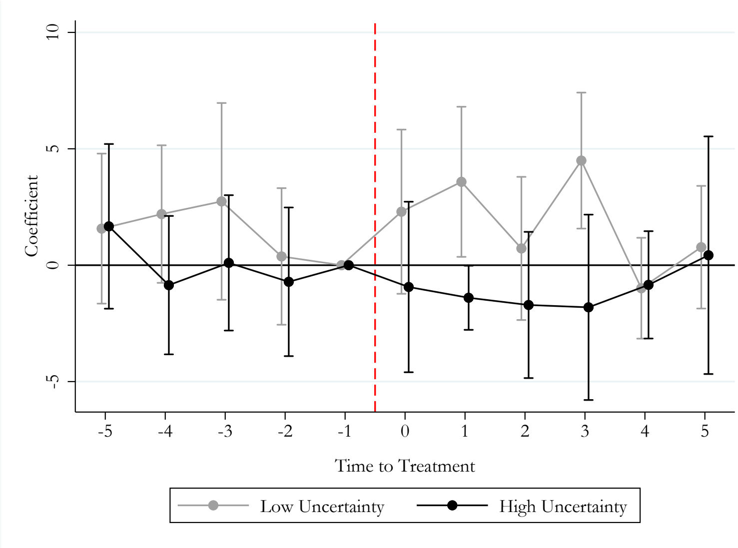

Notes: The plot shows event-study estimates obtained from regressing CEO age at appointment (obtained from BoardEx) on indicators for time relative to treatment, separately for low- and high-uncertainty NAICS-4 industries. The treatment is defined as exposure to MBB market entry in an MSA linked by a direct flight connection. An MSA is classified as treated if the minimum travel time to an MBB office is reduced by at least 5% in response to that entry. To ensure sufficient pre-treatment periods, only treatments that occurred in 2005 or later are considered. Uncertainty is measured as the cross-sectional standard deviation of the symmetric growth rate of firm revenue within a firm's NAICS-4 industry. We take the median of this measure across all periods and define high-uncertainty sectors as those with above-median levels of this variable. The sample is restricted to those MSAs without MBB presence. The regression controls for fixed effects at the NAICS-4-MSA and NAICS-4-year level as well as the natural logarithm of employment and firm age (extracted from the employment history data). The endpoints are binned at $t = -6$ and $t = 6$, while $t = -1$ is omitted as the reference category. Point estimates are shown alongside 95% confidence intervals, which are obtained from standard errors double-clustered at the NAICS-4 industry and MSA level.

operate in multiple regulatory environments, has created new layers of complexity. These changes along several, possibly correlated and compounding, dimensions have potentially contributed to the rising value of generalist human capital. To assess our mechanism along these different dimensions, we introduce a set of proxies for economic complexity to our identifying framework. As before, we estimate the differential effect of these industry-level measures across MSAs with varying local supply of generalist skills, captured by flight time to the closest MBB office. The results are summarized in Table 1.4.

First, we construct a measure of business diversification from the employment history data by counting the number of vice presidents (VPs) in distinct job roles.[36]

36. To identify job roles, we rely on the taxonomy developed by Revelio Labs, which assigns activities to job titles by matching descriptions from resumes and online profiles to responsibilities in job postings. It initially classifies 1,500 roles, which are then hierarchically clustered.

VPs are transparently identifiable in the employment history data, with their titles revealing the activities, product lines, or markets they are responsible for. As firms typically appoint VPs to oversee specific business functions, the number of VPs captures the extent of diversification at the operational level. As shown in Panel (a) of Figure 1.B.19, the average firm-level number of VPs in distinct roles has increased substantially over time. In Column (1) of Table 1.4, we use this count variable as our main explanatory variable, aggregated to the industry-specific median and transformed to its natural logarithm. To rule out that the respective coefficient is driven by structural time-varying differences across firms related to, e.g., size, we control for the interaction of travel time with the overall number of VPs.[37] Consistent with our hypothesis, the estimated interaction of travel time with the number of distinct VPs is positive and highly significant. Our interpretation is that firms tend to appoint older CEOs when demand for generalists stemming from business diversification increases and the supply of other generalist managers is low. In the baseline specification, we employ the job taxonomy at a granularity of 500 distinct roles. As reported in Table 1.C.6, very similar results obtain when using the taxonomy at higher or lower levels of disaggregation.

Next, we examine the role of spatial diversification within the United States. As documented in Rossi-Hansberg, Sarte, and Trachter (2021), large firms have expanded into an increasing number of local markets, with diverging implications for national and local concentration. We use the employment history data to construct a Herfindahl-Hirschman Index (HHI) based on employee counts across offices, defining an office as having at least 5 employees in a given MSA and year. This measure has declined substantially since 2010, suggesting a marked rise in spatial diversification (see Panel (b) of Figure 1.B.19, where we invert the measure). Motivated by these observations, we construct an industry-level measure of high spatial diversification by assigning indicators to firms in the lower tail of the HHI distribution and computing the industry-specific median of this indicator by year. Column (2) of Table 1.4 presents the results for industries where at least half of the firms fall into the bottom decile of the spatial HHI, controlling for travel time interacted with the number of offices. Consistent with expectations, the estimated interaction for spatial diversification is positive and highly significant. Table 1.C.7 suggests that the results are robust to alternative definitions of spatial diversification, including the use of indicators for the bottom quartile of the HHI distribution, a threshold of 10 employees to define offices, and an indicator based on the share of individuals employed outside of a firm's headquarters.

Third, we are interested in trade-induced variation in economic complexity. We focus on sectoral exports, given that exporting is a relatively rare activity, involv-

37. Equivalently, our results are robust to using the ratio of the number of different VP roles to the total number of VPs.

Table 1.4. MBB Availability, Complexity Measures, and CEO Age

	CEO Age				
	(1)	(2)	(3)	(4)	(5)
Travel Time × No. VP Roles	4.683***				
	(1.549)				
Travel Time × No. VPs	-4.281***			-0.688**	1.641
	(1.401)			(0.346)	(1.093)
Travel Time × Spatial HHI, Bottom Decile		4.244***			
		(1.262)			
Travel Time × No. Offices		-0.568		-0.290	-3.050***
		(0.375)		(0.336)	(0.758)
Travel Time × Export Complexity, Above Median			3.374***		
			(0.544)		
Travel Time × No. Countries Export			-9.049**		-9.878***
			(3.487)		(1.569)
Travel Time × Complexity Index (3 Components)				1.071***	
				(0.224)	
Travel Time × Complexity Index (4 Components)					1.098***
					(0.262)
R^2	0.365	0.365	0.408	0.365	0.403
Observations	14,933	14,195	5,136	14,173	4,946
MSA-Year fixed effects	✓	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓

Notes: The dependent variable in all regressions is CEO age upon appointment, obtained from BoardEx. *Travel Time* denotes the flight time (in hours) to the closest office, which we construct from MBB office openings combined with flight data from the Bureau of Transportation Statistics. In Column (1), *No. VP Roles* is a count variable measuring the number of distinct job roles (out of 500) in which a firm employs vice presidents, while *No. VPs* is a count variable of the number of VPs. We aggregate these measures to the year-specific median within NAICS-4 industries and use their natural logarithm. Further details are provided in Table 1.C.6. In Column (2), spatial diversification is measured as the Herfindahl-Hirschman Index (HHI) of employee counts across offices, which we define as having at least 5 employees in an MSA-year observation. The variable *Spatial HHI, Bottom Decile* is derived from an indicator for firms in the lowest decile, aggregated to the NAICS-4-year level by assigning a value of one whenever at least half of the underlying indicators equal one. *No. Offices* denotes the natural logarithm of the median number of offices of each firm within a year and NAICS-4 industry. Further details are provided in Table 1.C.7. In Column (3), *Export Complexity, Above Median* is an indicator for above-median levels of export complexity, which we define as the weighted average economic complexity (Hidalgo and Hausmann 2009) across export destinations, where the weights reflect the export shares of a NAICS-4 industry in a given year. Further details are provided in Table 1.C.8. In Columns (4) and (5), we create a composite index of complexity by transforming all variables to standard normal (if necessary), and summing over the corresponding variables: *Industry-level SD*, *No. VP Roles*, *Spatial HHI, Bottom Decile* (3 components) and additionally *Export Complexity, Above Median* (4 components). The sample is restricted to those MSAs without MBB presence. All regressions control for firm size by including the natural logarithm of employment and firm age, both extracted from the employment history data. Standard errors are double-clustered at the NAICS-4 industry and MSA level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ing fewer than 3% of US firms in 2021,[38] which requires high levels of strategic skill. Specifically, export participation is strongly shaped by managerial experience and associated with a wage premium (Mion and Opromolla 2014). These dynamics are further amplified by globalization: higher economic integration tends to increase managerial slack, in turn raising the importance of generalist human capital (Schymik 2018). We operationalize these considerations by combining trade flows with the country-level economic complexity index (ECI) from Hidalgo and Hausmann (2009) into a measure of export complexity. The ECI is based on product-level trade data and captures both the diversity of a country's exports and the ubiquity of the products it exports. These dimensions are synthesized into a single continuous index that reflects the complexity of a country's product space relative to others. Based on these data, we construct the export-related economic complexity index (EECI) as the weighted average of the ECI of a sector's export destinations, with weights corresponding to the export shares of each NAICS-4 industry in a given year.[39] Column (3) of Table 1.4 shows the results for industries with above-median levels of EECI, interacted with travel time. The estimated coefficient is again positive and highly significant, even after controlling for the natural logarithm of the number of export destinations. We conduct a battery of robustness checks in Table 1.C.8. The findings remain unchanged when using the raw EECI, an indicator for the upper quartile, or when constructing the EECI based on 1999 export shares, the last year before the start of the sample period. The results are also robust to using EECI values derived from HS code data, and to controlling for the number of export destinations that account for at least 1% of sectoral exports. As a placebo test, Table 1.C.9 replicates these regressions using import shares, yielding small and insignificant coefficients across the board.

To summarize these results, we re-estimate the specification using the sum of all four normalized drivers for generalists. We show specifications with and without the EECI, since this variable is available only for manufacturing. Using the index constructed from industry-level uncertainty, business diversification, and spatial diversification, we find a strong positive coefficient (Column 4). The results are very similar when we augment the index with EECI (Column 5).

Substitutability of MBB Experience and Age: Our identification exploits spatial variation in access to consultancies. Such exposure may operate through the usage of their services, which is not captured by our data, or through hiring from these firms, which we document. The latter channel is relevant because current or former consultants can transition directly into CEO positions or enter firms in any other

38. Source: 2022 Annual Business Survey. Summary tables are available under www.census.gov.

39. By construction, the ECI exhibits no trend on average. To be able to interpret the time series, we deduct the ECI of the US economy from the EECI. This difference can be interpreted as the distance of the export destinations to the US economy and has declined in absolute value, suggesting a convergence process since the late 1990s (Figure 1.B.19, Panel c).

capacity. Consequently, our measure of the local supply of young generalists may partly reflect firms' hiring at lower organizational levels, so that the age of the subsequently appointed CEO does not purely reflect firms' demand for such skills at the level of the CEO.

To test how hiring from consultancies at different organizational levels affects CEO age, we examine labor market flows from these firms to other employers. In particular, we document that CEOs with prior MBB experience are younger on average, consistent with the idea that work experience in these firms fast-tracks the accumulation of generalist skills. Moreover, the inflow of former consultants into a firm is positively associated with CEO age, but only for executives without a consulting background. This suggests that consulting firms offer their employees an accelerated path to management-relevant human capital, whereas executives who rise through other channels take longer to acquire the same skills, which the firm nonetheless demands of its CEO despite having hired former consultants into non-CEO positions. Together, these findings indicate that consulting experience provides individuals with generalist skills that substitute for age and interact with executives' internal career paths to affect CEO age at appointment.

We establish these patterns by combining the CEO information from BoardEx with the employment history data. For a CEO w, appointed in firm j, industry i, MSA m, and year t, we estimate regressions of the following form:

$$Age_{wt} = \alpha \, MBB\ Exp_{wt} + \beta \, Inflows_{jt} + \delta \, MBB\ Exp_{wt} \times Inflows_{jt} + \gamma_{mt} + \phi_{it} + X_{jt} + \epsilon_{wt}, \tag{1.2}$$

where Age_{wt} denotes the CEO's age at appointment, $MBB\ Exp_{wt}$ indicates individual w's experience at one of the MBB firms, and $Inflows_{jt}$ captures the share of employees with MBB experience. Specifically, the latter is defined as the firm-level sum of employees with any MBB experience hired over the previous 10 years, divided by the total number of new employees hired in the firm over the same period. As before, the regression includes fixed effects at the MSA-year and the NAICS-4-year level. We also control for the natural logarithm of employment and firm age. The main coefficient of interest is δ, capturing the differential effect of MBB inflows on individuals with MBB experience, compared to those without.

The results are reported in Table 1.C.10. Inflows from MBB firms are consistently and positively associated with CEO age, which is inconsistent with the view that the CEO age trend is primarily driven by supply-side scarcity of executives. Individuals with experience in one of these firms tend to be appointed as CEOs at substantially younger ages, by more than 3 years. The estimated coefficient on the interaction term is negative, indicating that MBB experience can serve as a substitute for age. These findings are robust across all specifications, conditional on controls and across varying sets of fixed effects.[40] Using the coefficient from our preferred

40. The results from a number of additional robustness checks are shown in Table 1.C.11.

specification, which includes fixed effects at the NAICS-4-year and MSA-year level, a one standard deviation increase in MBB inflows translates into a CEO age increase of around 0.4 years for those without MBB experience. This suggests that if the local supply of young generalists in Tables 1.1 to 1.4 operates, among other margins of adjustment, through hiring former consultants at lower organizational levels, then our estimated effects on CEO age constitute lower bounds.

CEO Age and Firm-Level Outcomes: Finally, in Appendix Section 1.D.3, we scrutinize how these shifts in the demand for certain skills lead to the appointment of systematically different types of CEOs, thereby causally affecting firm-level policies and outcomes. Our results reveal that the appointment of older CEOs, as a result of higher demand for generalists, leads to reduced R&D expenditure and innovation in computing-related technologies.

1.5.2 Supply-Side Factors

We next shift the focus to a potential supply-side response in the managerial labor market. In the descriptive analysis, we established that part of the increase in the aging of CEOs goes hand in hand with increased mobility across jobs prior to CEO appointments. In this section, we explore whether the rise in labor market turnover prior to CEO appointments is driven by executives' voluntary job switches aimed at acquiring generalist skills. Upward, or possibly lateral, moves are typically associated with career progression and compensation. In contrast, downward moves, i.e., job transitions to less senior positions across firms, tend to be associated with factors such as employer-employee mismatches, wage disparities, or shocks to product market competition that induce firm exit and worker reallocation.[41] Against this background, *voluntary* downward moves are unlikely to be motivated by short-term financial considerations. Instead, they suggest that executives may be willing to forego short-term wage growth for the sake of accumulating generalist skills by switching to less senior positions in new domains, e.g., to different functions or even industries.

Using employment biographies, we identify transitions between firms that involve moving to lower seniority levels ("downward switches") or maintaining the same level ("lateral switches") as well as across job categories.[42] Figure 1.8 plots the evolution of the share of CEOs who ever moved downwards or laterally prior to

41. As documented in Fee and Hadlock (2004), forced turnover among non-CEO executives occurs at rates at least as high as those for CEOs.

42. The use of seniority levels as a measure of career attainment (e.g., as documented in Amornsiripanitch et al. 2025) is motivated by a large literature highlighting the role of firm-level hierarchies in organizing production processes (for a survey, see Garicano and Rossi-Hansberg 2015), with implications for individual career progression (Caicedo, Lucas, and Rossi-Hansberg 2019), wage growth (Bayer and Kuhn 2023), and firm-level outcomes (Ewens and Giroud 2025).

their appointment, across positions within or across firms (Panel a) or only across firms (Panel b). While all variables exhibit a striking and similar upward trend in absolute terms, the relative increase in the share of those CEOs who ever moved downward is much larger. Consistent with these findings, Figure 1.B.23 plots the shares of downward and lateral switches over total switches—as above—across positions (Panel a) and firms (Panel b). The probability of moving downward conditional on switching has increased substantially over time, while the conditional probability of switching laterally has only seen modest increases.

We examine whether these labor market mobility patterns reflect voluntary switches by testing whether they are partly caused by heterogeneous information shocks on the returns to job switching. Our labor market data are particularly well-suited to answer this question, given that they capture individuals' information sets about their network connections.

For workers to capitalize on opportunities for on-the-job learning across firms, they must be aware of such opportunities and believe that increased mobility will be beneficial for their careers. We therefore expect exposure to information supporting this idea to increase job turnover in subsequent periods. We operationalize this idea by leveraging the role of information flows via coworker networks. The importance of social networks in shaping labor market outcomes is well-documented, both theoretically (Montgomery 1991; Dustmann et al. 2016) and empirically in the overall labor market (Bramoullé and Saint-Paul 2010; Hensvik and Skans 2016) and in the labor market for CEOs (Liu 2014; Kim, Fahlenbrach, and Low 2023). Network ties can emerge for various reasons, including prior employment at the same firm (Cingano and Rosolia 2012; Glitz 2017; Saygin, Weber, and Weynandt 2021; Cornelissen, Dustmann, and Schönberg 2025), shared educational backgrounds (Oyer and Schaefer 2009), or residential proximity (Bayer, Ross, and Topa 2008). Importantly, networks also affect firm policies (Shue 2013) and misconduct among financial advisors (Dimmock, Gerken, and Graham 2018).

The positive information shock that we explore is the event of a former coworker leaving the firm and later becoming CEO. Consistent with the idea of observational learning, individuals, witnessing their former colleagues' career progress, should become more likely to emulate these career moves and become more mobile across firms and industries. That is, we test whether individuals respond to such information shocks by increasing their own job mobility.

Specifically, consider a worker w in year t, who worked alongside a future CEO in firm j, MSA m, and seniority level s. We implement a difference-in-differences design by estimating equations of the following form:

$$y_{wt} = \beta \, Treat_{jms} \times Post_{jst} + \gamma_w + \delta_{jst} + \epsilon_{wt}, \tag{1.3}$$

where y_{wt} is the outcome of interest, $Treat_{jms}$ is an indicator for treated individuals, and $Post_{jst}$ is an indicator for all periods following the CEO appointment of a

Figure 1.8. Share of CEOs That Ever Moved Downwards or Laterally

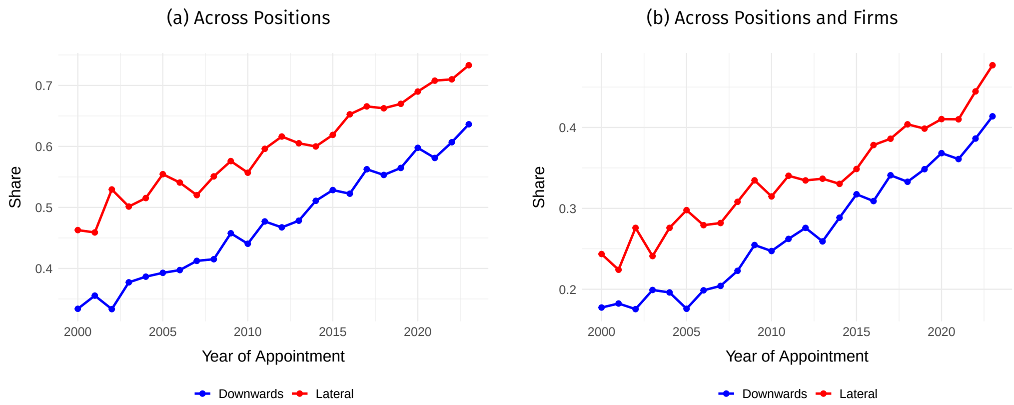

Notes: Plots the share of CEOs by year of appointment that have ever switched downwards or laterally prior to their first CEO appointment. Panel (a) includes all switches across positions, and Panel (b) focuses on those across positions and firms. Switches are identified from a panel constructed from employment biographies from Revelio Labs, prioritizing more senior positions and those that started earlier in the case of overlapping spells.

former colleague. The fixed effects included, γ_w and δ_{jst}, are at the worker and firm-seniority-year level, respectively. The coefficient of interest is β, capturing the increase in job mobility following the CEO promotion event for workers associated with the CEO compared to the counterfactual.

To construct the sample, we start by identifying a list of CEO employment spells, considering only the first CEO appointment for each individual.[43] We then extract the employment histories of these CEOs and retain all spells that meet the following criteria: they ended during the 10 years prior to the CEO appointment; they are at least at seniority level 4; and after these spells, the future CEO left the firm, but did not immediately become CEO. We label these as *connection spells*. The sample includes all individuals who had an employment spell in the same company and at the same seniority level, overlapping with a *connection spell* by at least 90 days. Thus, the sample encompasses all individuals who worked contemporaneously for at least 3 months in the same firm and at the same seniority level as a future CEO.

Among those individuals, the variable $Treat_{jms}$ is set equal to one for all those who were employed in the same MSA as their associated future CEO. In other words, the treatment group is composed of all individuals who have worked alongside a future CEO, that is, at the same company, the same seniority level, and in the same MSA. The control group comprises individuals that worked at the same company

43. We focus on employment spells with the following attributes: (*i*) job title includes "CEO" or "Chief Executive Officer," (*ii*) seniority level of 7, (*iii*) O*NET-code in the major group "11 - Management Occupations," (*iv*) the firm has at least 100 active employment spells in a given year, and (*v*) it is located in the United States.

and seniority level but in other MSAs. The variable $Post_{jst}$ is an indicator for all periods following the former colleague's appointment to CEO, which typically takes place years after leaving the firm in question. Each individual enters the sample with the beginning of the spell that overlaps with the *connection spell*, but at the earliest in 1990, and exits the sample after at most 10 years following the CEO appointment, but at the latest at the end of the sample in 2023. In Table 1.C.12, we show summary statistics for the last year prior to the treatment, separately for the treatment and the control group. On average, individuals in the treatment group are slightly older, more experienced, and work at higher seniority levels, with somewhat higher imputed salaries, although these differences are small.

Our main outcomes comprise a set of indicators for job switches across firms between two consecutive observations.[44] With all fixed effects in place, this implies that the regressions estimate the change in the propensity to switch firms in response to the CEO appointment, comparing those who were working alongside the future CEO to those who were not. The causal effect is identified under a standard parallel trend assumption, which holds if—conditional on all fixed effects and controls—the change in the probability of switching firms would have been the same in both groups before and after the treatment. We are confident that this assumption is justified in our setting, given that the timing of the treatment is arguably quasi-random from the perspective of the treated individuals.

A burgeoning literature discusses potential biases arising from estimating difference-in-differences models with heterogeneous treatment timing using two-way fixed effects regressions (see, e.g., Goodman-Bacon 2021; Roth et al. 2023). A key problem is that these designs generate flawed comparisons between a treatment group and a comparison group that has already been treated. To alleviate concerns about the validity of our approach, we exclude all individuals who were first assigned to a control group and then to a treatment group. If an individual is treated multiple times, we focus on the first treatment event. Finally, we only consider the first CEO appointment for each CEO. These restrictions ensure that the treatment is binary and that the sample includes only individuals who were either treated simultaneously with all other units in their firm-seniority group or who were never treated. In our preferred specification, we include fixed effects for all firm-seniority-year combinations. Consequently, our final estimate is a weighted average of canonical 2×2 difference-in-differences estimates generated by each of these firm-seniority-year combinations, following the logic of a stacked design (Wing et al. 2024). We also highlight that the control group is relatively large compared to the treatment group, accounting for 79% of all post-treatment observations

44. One potential concern is that a CEO may improve job opportunities for connected individuals in firms they have worked in or currently work in. To mitigate this, we exclude transitions to firms where the CEO was employed between the *connection spell* and the CEO appointment.

and 77% of all individuals, further mitigating potential issues arising from negative weights (De Chaisemartin and d'Haultfoeuille 2020).

We proceed by estimating Equation (1.3) on the employment history data. The baseline results are presented in Table 1.5, where the outcome variable is an indicator of switching between firms. Column (1) displays the coefficients from regressions that include the treatment and post indicators, their interaction, and year fixed effects. While the coefficient on the $Treat_{jms}$ is insignificant, the coefficient on the $Post_{jst}$ is negative, suggesting that job mobility during the post-treatment period tends to decrease—possibly because people are at later stages of their careers and less likely to transition between firms. Finally, the estimated coefficient on the interaction term is positive but insignificant.

In the subsequent columns, we augment the regression with different sets of fixed effects, which eliminate the coefficients on $Treat_{jms}$ and $Post_{jst}$. The estimated coefficient on the interaction term becomes highly significant and increases in terms of magnitude. Given that the treatment is at the seniority-firm-year level, the specification in Column (4) is our preferred one. Quantitatively, the probability of switching firms increases by about 0.5 percentage points, amounting to a 5% increase compared to the baseline.

One challenge to our identification strategy involves potential latent differences between individuals working alongside a future CEO and their respective control groups; for instance, the former may work at headquarters while the latter are based in regional offices. The environment in the treatment group may be more competitive and ambitious, or it may offer better outside options due to the clustering of similar industries nearby. While we do not observe any baseline differences across the two groups in Column (1), we explicitly address this issue by including fixed effects at the MSA-NAICS-4-year level, effectively controlling for any time-varying differences across narrow industries and geographic locations in the level of job mobility. The results are shown in Column (5). While the estimated coefficient slightly decreases in size, this change is neither statistically nor economically significant.

We repeat the analysis for switches across NAICS-4 industries, with the results included in Table 1.C.13. While the estimates are naturally smaller in absolute size, the signs and relative magnitudes are closely in line with the previous results. For instance, the estimate from Column (4) suggests a 7% increase in mobility across industries compared to the baseline.

Moving beyond the overall increase in worker mobility in response to the information shock, we turn to examining the underlying dynamics. To this end, we further differentiate the main outcome based on whether an individual moved to a more senior position, a less senior position, or remained at the same seniority level. The results show that the rise in overall mobility is driven entirely by an increase in the frequency of downward transitions across firms (Table 1.6), while upward and lateral transitions remain practically unaffected (Tables 1.C.14 and 1.C.15).

Table 1.5. CEO Appointment of Former Coworker and Job Mobility (Any Firm Switch)

	Firm Switch				
	(1)	(2)	(3)	(4)	(5)
Treat	-0.002				
	(0.001)				
Post	-0.010***	-0.016***	-0.019***		
	(0.001)	(0.001)	(0.001)		
Treat × Post	0.002	0.002	0.002	0.005***	0.004***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.012	0.139	0.140	0.166	0.197
Observations	22,777,381	22,777,381	22,777,381	22,777,381	16,139,755
Pre-Treatment Control Mean	0.094	0.094	0.094	0.094	0.095
Year fixed effects	✓	✓	✓		
Individual fixed effects		✓	✓	✓	✓
Firm-Seniority fixed effects			✓		
Firm-Seniority-Year fixed effects				✓	✓
MSA-Naics4-Year fixed effects					✓

Notes: The dependent variable in all regressions is a dummy equal to one in the last year before an individual switches firms. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors clustered at the firm level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Next, we assess whether the transitions between firms in response to the information shock are accompanied by changes in job categories, distinguishing between seven functional areas.[45] The estimates in Tables 1.C.16 and 1.C.17 indicate that the increase in job mobility is mainly due to a higher probability of switching across job categories. Even though the coefficients from our preferred specification are similar in both cases, the relative increase in cross-category switches is nearly twice as large, once taking the baseline differences into account. These dynamics are potentially self-reinforcing, as individuals with broader experience tend to be more willing to move across positions (Cappelli and Hamori 2014).

In sum, these findings support the idea that workers emulate future CEOs' employment trajectories. While upward or lateral moves may be motivated by financial incentives, the greater prevalence of downward transitions suggests that employees are seeking to acquire a broader range of generalist skills across firms and functional areas—skills that are highly valued for future CEO roles.

As a first robustness check, we modify our preferred specification to obtain event-study estimates, presented in Figure 1.B.24. The estimated coefficients are generally small and insignificant during the pre-treatment periods, followed by a gradual increase in the post-treatment periods for overall and downward switches.

45. The job categories are Admin, Engineering, Finance, Marketing, Operations, Sales, and Science.

Table 1.6. CEO Appointment of Former Coworker and Job Mobility (Downward Firm Switch)

	Firm and Downward Switch				
	(1)	(2)	(3)	(4)	(5)
Treat	-0.005***				
	(0.001)				
Post	-0.010***	-0.012***	-0.013***		
	(0.001)	(0.001)	(0.001)		
Treat × Post	0.003***	0.004***	0.004***	0.005***	0.005***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.005	0.090	0.090	0.111	0.144
Observations	22,777,381	22,777,381	22,777,381	22,777,381	16,139,755
Pre-Treatment Control Mean	0.041	0.041	0.041	0.041	0.041
Year fixed effects	✓	✓	✓		
Individual fixed effects		✓	✓	✓	✓
Firm-Seniority fixed effects			✓		
Firm-Seniority-Year fixed effects				✓	✓
MSA-Naics4-Year fixed effects					✓

Notes: The dependent variable in all regressions is a dummy equal to one in the last year before an individual switches firms and transitions to a lower seniority level. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors clustered at the firm level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In contrast, we observe no discernible response for upward or lateral switches. We also show the results from a placebo test, where we focus on individuals who left the initial firm between one and five years before the prospective CEO was hired at that firm. The estimates, included in Table 1.C.18, are close to zero and insignificant across the board.

To bolster the validity of our mechanism, we provide a battery of consistency checks, focusing particularly on downward mobility. As our argument hinges on network connections, we conjecture that the response in job mobility should be more pronounced among those who have closer ties to the future CEO. One potential measure of proximity is shared tenure, i.e., the time spent together at the same seniority level within a firm during their former job. We test this in Table 1.C.19 by including an interaction with an indicator equal to one for those whose shared tenure is in the highest quartile. The estimated coefficients on the triple interaction are large and significant, suggesting that this group reacts more strongly to the information shock. Figure 1.B.25 shows the coefficients on the marginal effect for downward mobility (Column 4 of Table 1.C.19) separately for all quartiles of shared tenure, revealing a strong gradient in the response across the entire distribution.

In a similar vein, we hypothesize that the information shock spills over to individuals in other seniority levels at the initial firm. To check this, Tables 1.C.20 and 1.C.21 focus on individuals in the two levels above and below the future CEO, respectively. We find that employees who worked just above the seniority level of the future CEO exhibit a similar reaction to the treatment as those who worked at the same seniority level as the future CEO, whereas the response in seniority levels below the future CEOs is more muted.

We also examine whether the response to the information treatment has become stronger over time. In Table 1.C.22, we test for time variation in the treatment effect by interacting the variable of interest with indicators for 5-year bins, excluding the final two bins to ensure a balanced 10-year post-period. The results indicate a structural break during the period 2000 to 2004, which is omitted as a baseline. In the early 1990s, treated individuals reacted to the information shock by increasing their lateral job mobility. During the late 2000s and early 2010s, the response shifted toward downward job mobility, further prolonging the career trajectory of prospective CEOs. Given the increasing coefficients on downward mobility after 2000, this result speaks against a one-off shock but in favor of a gradual increase in the weight attributed to generalist human capital.

Next, we consider heterogeneity with respect to the future CEOs' job mobility between leaving the firm and being promoted to CEO. In Table 1.C.23, we build on the idea that the information shock is stronger for those who have seen their former colleague transition across a higher number of positions. To evaluate this hypothesis, we introduce an interaction term with an indicator equal to one if the number of CEO transitions is in the highest quartile. In line with our expectations, the estimated coefficients on the triple interactions are positive for both overall and downward transitions. Moreover, we test for heterogeneity depending on the direction of the CEO switch in terms of seniority. Assuming individuals imitate successful career paths, one would anticipate a higher probability of upward, lateral, or downward switches after having witnessed a former coworker leaving in the respective direction (the omitted category is upward, i.e., a future CEO leaving the company for a more senior position). The results shown in Table 1.C.24 point in this direction, showing a significant coefficient for downward moves in Column (4), the category of greatest interest to us.

In addition to CEOs' career paths, we hypothesize that the treated individual's own characteristics and career considerations will play a role in the response of job mobility. Downward and industry switches should be more prevalent responses among individuals working in smaller firms where it is harder to acquire generalist skills, possibly due to the limited scope of the firm's operations. We scrutinize this idea by including interactions with size quartiles of a worker's current firm, with results shown in Table 1.C.25 and Figure 1.B.26. As hypothesized, we find that the positive response to the information shock is concentrated among individuals working in smaller firms. This result complements Lazear (2009), who argues that quits should be less common among more idiosyncratic, i.e., smaller, firms, as the market for the skills of their employees is thinner. Importantly, the increase in career mobility is not primarily driven by switches to larger firms. To demonstrate this, we distinguish downward moves by the size of the destination firm. The results in Table 1.C.26 reveal that the rise in turnover is driven by switches to smaller firms, both in absolute terms and compared to the current employer.[46]

As a direct implication of the higher likelihood of moving to less senior positions, we anticipate that individuals forego income in the medium run in order to improve their long-term career prospects. To test this hypothesis, we repeat the analysis with the growth rate of imputed salary as the dependent variable. The results included in Table 1.C.27 are in line with our expectations, indicating a reduction in salary growth of around 11% relative to the counterfactual. As before, this result is robust to adding fixed effects at the MSA-NAICS-4-year level, suggesting that we can rule out differences in local labor market conditions as the driver of these results.

Taken together, these estimates support the hypothesis that the observed trends are partly influenced by an increase in generalist skills on the labor supply side. They provide evidence in favor of the idea that positive information about the returns to job mobility among connected individuals increases employees' probability of switching across firms with the intention of gathering generalist skills across a wide range of functional areas. Consequently, the career paths of prospective CEOs become longer, leading to higher age at CEO appointment.

1.5.3 Decomposing Variation in the Aging Patterns

In this section, we decompose the aging trend into a sectoral and a geographical dimension. Our results in Section 1.5.1 suggest that sectoral shocks, specifically shifts in uncertainty and complexity in the economic environment, have contributed to higher demand for generalist skills. Yet, the implications of these shocks for CEO appointments are heterogeneous across space due to frictions in the availability of generalist skills. As such, one can interpret sectoral variation over time as primarily reflecting changes in the demand for generalist skills, while variation across space tends to capture changes in the supply thereof (in line with our identifying variation in Section 1.5.2).

To quantify the role of these two dimensions, we assess the explanatory power of different combinations of fixed effects for CEO age at appointment. We consider a CEO appointment in firm j, MSA m, industry i and year t, and estimate a series of regressions to obtain the following coefficients of determination

46. This does not exclude the possibility that ambitious and successful individuals are generally willing to trade off working in a larger firm with a lower seniority level. In Figure 1.B.27, we focus on firm switches of future CEOs and regress the differences in firm size against differences in seniority. We find that such a trade-off indeed exists and that it has become stronger over time.

$$
\begin{aligned}
R_i^2 &: \quad Age_{jt} = \phi_i \times (\delta_t + X_{jt}) + \epsilon_{jt} && (1.4)\\
R_m^2 &: \quad Age_{jt} = \gamma_m \times (\delta_t + X_{jt}) + \epsilon_{jt} && (1.5)\\
R_{i+m}^2 &: \quad Age_{jt} = (\phi_i + \gamma_m) \times (\delta_t + X_{jt}) + \epsilon_{jt} && (1.6)\\
R_{i\times m}^2 &: \quad Age_{jt} = \phi_i \times \gamma_m \times (\delta_t + X_{jt}) + \epsilon_{jt}. && (1.7)
\end{aligned}
$$

As before, Age_{jt} denotes CEO age at appointment. The fixed effects ϕ_i and γ_m are at the NAICS-4 and MSA level, respectively, and interacted with year fixed effects, δ_t, and a vector of control variables, X_{jt}, which includes the number of employees, firm age (both in terms of their natural logarithm) and average age of other executives. As we add all lower-level interactions, we estimate to what extent sectoral and geographical factors can explain variation in CEO age at appointment, both differentially across years and firm-level characteristics.

We then use the estimated R^2s to calculate the contribution of each dimension to $R_{i\times m}^2$, which is generated by the full model. In doing so, the key complication is that $R_i^2 + R_m^2 \neq R_{i+m}^2$ if the regressors are non-orthogonal. We address this by calculating Shapley values, which provide the unique decomposition of $R_{i\times m}^2$ that is both additive and symmetric, i.e., it is invariant to the order in which regressors are added to the model (Grömping 2015; Sharapov et al. 2021). Due to the nested fixed effects structure, the analytical expressions for the three components simplify as follows:

$$
\begin{aligned}
\text{Sectoral contribution} &= \frac{1}{2}(R_i^2 + R_{i+m}^2 - R_m^2) && (1.8)\\
\text{Geographical contribution} &= \frac{1}{2}(R_m^2 + R_{i+m}^2 - R_i^2) && (1.9)\\
\text{Interaction} &= R_{i\times m}^2 - R_{i+m}^2. && (1.10)
\end{aligned}
$$

In the final step, we rescale the contributions by $R_{i\times m}^2$ to represent the share of the total explained variance that we can trace back to one of the factors.

Table 1.7 shows the results. In the first row, we focus on interactions with year fixed effects. The estimated $R_{i\times m}^2$ is at 36%, of which approximately 38% and 28% can be attributed to sectoral and geographical factors, respectively, while the interaction term accounts for the remaining one-third of the explained variance. In the subsequent rows, we successively introduce additional interactions with the number of employees, firm age, and average executive age. While we are now able to explain 57% of the variation in age, the relative contributions of the two dimensions remain remarkably stable across specifications. We also report bootstrapped standard errors, which confirm that the components are statistically significantly different from one another.[47]

47. One limitation of this approach is that it does not account for general demographic changes, which are absorbed by both the time-varying sectoral and geographical variation. To gauge the magnitude of this concern, we repeat the exercise using CEO age after residualizing for year fixed effects. Even with this overly conservative adjustment, our preferred specification attributes 35.8% and 26.4%

Table 1.7. Decomposition of CEO Aging Trend (in %)

Independent Variables	Sectoral	Geographical	Interaction	$R^2_{i\times m}$
$(\phi_i + \gamma_m) \times$ Year	38.1	28.2	33.7	36.2
	(0.9)	(0.9)	(0.9)	(0.9)
+ $(\phi_i + \gamma_m) \times$ Log(Employment)	36.5	25.7	37.8	44.6
	(0.8)	(0.7)	(0.9)	(0.8)
+ $(\phi_i + \gamma_m) \times$ Log(Firm Age)	35.4	25.6	39.0	49.4
	(0.9)	(0.8)	(1.0)	(0.9)
+ $(\phi_i + \gamma_m) \times$ Executive Age	37.0	28.4	34.6	56.7
	(0.8)	(0.7)	(0.9)	(0.8)

Notes: Shows the estimated contributions (in %) of sectoral and geographical factors, and their interaction to the CEO aging trend as described in Equations (1.8), (1.9), and (1.10). $R^2_{i\times m}$ denotes the R^2 obtained from estimating Equation (1.7). The first row shows the results obtained from interacting ϕ_i, γ_m, and $\phi_i \times \gamma_m$ with δ_t. The following rows add further interactions with the natural logarithm of both the number of employees and firm age, as well as average age of other executives (obtained from the employment history data). Standard errors obtained from bootstrapping (100 replications) contained in parentheses.

Taken together, the three components can be interpreted as capturing distinct forces behind the rising CEO age at appointment. Given our results from Section 1.5.1, the sectoral dimension maps most closely to demand-side factors that increase the value of a diverse set of labor market experiences. Holding the supply of generalist candidates constant, such demand-side factors, which include higher uncertainty and greater business diversification, jointly account for approximately 37% of the explained variation in CEO age. Analogously, the geographic component, as captured by the local supply of generalists across metro areas, explains around 28% of the variation. This availability is shaped by factors such as labor market thickness, the responsiveness of prospective executives to changes in firm-level demand, and the presence of executive training pipelines.[48] The third term reflects the degree to which these two forces interact with each other, in particular the extent to which the effects of increases in demand depend on the local supply of generalist skills. For instance, firms in sectors experiencing a surge in the demand for generalists can only adjust fully if they are located in labor markets in which these skills are available. The contribution of 35% highlights that such frictions or complementarities are sizable.

to sectoral and geographical factors, respectively ($R^2 = 54.6\%$). As CEOs at appointment have aged more rapidly than the population at large, these estimates can be interpreted as lower bounds.

48. In our sample, about 56% of CEO appointments occur within the same metro area. Yonker (2017) reports that firms are over five times more likely to appoint a local CEO than predicted if geography did not matter in the matching process. This pattern persists when focusing on external appointments from other firms or industries. After considering several competing explanations, he argues that CEOs' geographic preferences are the most plausible driver of the prevalence of local hiring. Amore, Bennedsen, and Larsen (2022) complement this view by documenting that closer physical proximity of CEOs contributes positively to the working environment within firms.

1.6 Conclusion

This paper contributes to our understanding of the evolving market for corporate leadership by documenting and explaining the substantial increase in CEO age upon appointment over the past two and a half decades. Our findings reveal that this trend reflects fundamental changes in how executives build careers and how firms select their leaders. The rise in CEO age is primarily driven by longer external career paths prior to appointment, with prospective CEOs accumulating increasingly diverse experiences across positions, firms, and sectors.

Using plausibly exogenous variation in access to strategy consulting firms, we demonstrate that demand for generalist skills in response to greater industry-level uncertainty and complexity has causally contributed to the appointment of older CEOs. Simultaneously, executives strategically respond to these evolving skill requirements by increasing their job mobility through professional networks, on average at the cost of reduced short-term compensation.

This mechanism explains distinct CEO hiring patterns by firm size. CEOs in large firms are predominantly promoted internally (e.g., Cziraki and Jenter 2021) and have experienced a decline in executive mobility since 2000 (Barry et al. 2025). In smaller firms, by contrast, the majority of CEOs are appointed externally, typically following career paths that involve an increasing number of transitions across other firms. The rising demand for generalist skills rationalizes these patterns: large firms can cultivate generalists internally through diverse assignments across divisions and functions, while smaller firms tend to recruit executives who accumulated generalist experience externally. Consequently, the lengthening of external career paths translates into older CEO appointments primarily at smaller firms.

The implications of this trend are substantial, given that older CEOs tend to manage firms with slower growth rates and less radical innovation. While these patterns might appear concerning for long-term economic dynamism, they may represent a rational response to changing business environments characterized by heightened uncertainty and complexity. Understanding the underlying mechanisms driving this transformation provides important insights into evolving corporate governance and the firm-level origins of aggregate fluctuations.

More broadly, the shift toward generalist human capital reflects the growing demand for skills that enable coordination, adaptation, and decision-making under uncertainty. While central to firm organization (Garicano 2000), these skills are difficult to codify and automate. As technology, including AI, increasingly substitutes for routine tasks, it may disrupt the pathways through which expertise has traditionally been acquired, while simultaneously increasing the value of experienced decision-makers who can effectively leverage these technologies (Garicano and Rayo 2025; Garicano, Li, and Wu 2026). The rising importance of generalist human capital thus not only reflects its relative resilience to automation, but also suggests that the patterns we document may prove persistent.

References

Acemoglu, Daron, Ufuk Akcigit, and Murat Alp Celik. 2022. "Radical and Incremental Innovation: The Roles of Firms, Managers, and Innovators." *American Economic Journal: Macroeconomics* 14 (3): 199–249. [11, 105, 107]

Acemoglu, Daron, Alex He, and Daniel Le Maire. 2022. "Eclipse of Rent-Sharing: The Effects of Managers' Business Education on Wages and the Labor Share in the US and Denmark." Working Paper, 29874, National Bureau of Economic Research. [17]

Adams, Renée, Ali Akyol, and Patrick Verwijmeren. 2018. "Director Skill Sets." *Journal of Financial Economics* 130 (3): 641–62. [9]

Adams, Renée, Matti Keloharju, and Samuli Knüpfer. 2018. "Are CEOs Born Leaders? Lessons from Traits of a Million Individuals." *Journal of Financial Economics* 130 (2): 392–408. [26]

Adams, Robin, Michael Aldous, Philip Fliers, and John Turner. 2023. "Aristocratic Amateurs to Fat Cats? British CEOs in the Twentieth Century." Working Paper, 8, Queen's University Belfast. [11]

Adda, Jerome, and Christian Dustmann. 2023. "Sources of Wage Growth." *Journal of Political Economy* 131 (2): 456–503. [10]

Amore, Mario Daniele, Morten Bennedsen, and Birthe Larsen. 2022. "Neighborhood CEOs." Working Paper, Bocconi University. [50]

Amornsiripanitch, Natee, Paul Gompers, George Hu, Will Levinson, Vladimir Mukharlyamov, and Sachin Srivastava. 2025. "Assessing Career Attainment via a Non-Wage Measure." Working Paper, 34292, National Bureau of Economic Research. [40]

Antia, Murad, Christos Pantzalis, and Jung Chul Park. 2010. "CEO Decision Horizon and Firm Performance: An Empirical Investigation." *Journal of Corporate Finance* 16 (3): 288–301. [10]

Arellano-Bover, Jaime. 2024. "Career Consequences of Firm Heterogeneity for Young Workers: First Job and Firm Size." *Journal of Labor Economics* 42 (2): 549–89. [10]

Arellano-Bover, Jaime, Nicola Bianchi, Salvatore Lattanzio, and Matteo Paradisi. 2024. "One Cohort at a Time: A New Perspective on the Declining Gender Pay Gap." Working Paper, 32612, National Bureau of Economic Research. [12]

Arthur, Michael, and Denise Rousseau. 1996. *The Boundaryless Career as a New Employment Principle.* Oxford University Press. [19]

Ash, Elliott, and Bentley MacLeod. 2024. "Mandatory Retirement for Judges Improved the Performance of US State Supreme Courts." *American Economic Journal: Economic Policy* 16 (1): 518–48. [11]

Azoulay, Pierre, Benjamin Jones, Daniel Kim, and Javier Miranda. 2020. "Age and High-Growth Entrepreneurship." *American Economic Review: Insights* 2 (1): 65–82. [18]

Bai, John Jianqiu, and Anahit Mkrtchyan. 2023. "What Do Outside CEOs Really Do? Evidence from Plant-Level Data." *Journal of Financial Economics* 147 (1): 27–48. [9]

Baker, Scott, Nicholas Bloom, Steven Davis, and Kyle Kost. 2019. "Policy News and Stock Market Volatility." Working Paper, 25720, National Bureau of Economic Research. [31, 81]

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. "Measuring Economic Policy Uncertainty." *Quarterly Journal of Economics* 131 (4): 1593–636. [7, 31, 81]

Baltes, Paul, and Karl Ulrich Mayer, eds. 1999. *The Berlin Aging Study: Aging from 70 to 100.* Edited by Paul Baltes and Karl Ulrich Mayer. Cambridge University Press. [22]

Bao, Renjie, Jan De Loecker, and Jan Eeckhout. 2022. "Are Managers Paid for Market Power?" Working Paper, 29918, National Bureau of Economic Research. [17]

Barry, John, John Graham, Dawoon Kim, and Hyunseob Kim. 2025. “Executive Mobility in the United States, 1920 to 2023.” Working Paper, Nanyang Business School. [51]

Bayer, Christian, and Moritz Kuhn. 2023. “Job Levels and Wages.” Working Paper, 16177, IZA – Institute of Labor Economics. [40]

Bayer, Patrick, Stephen Ross, and Giorgio Topa. 2008. “Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes.” *Journal of Political Economy* 116 (6): 1150–96. [41]

Becker, Gary S. 1964. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education.* New York: Columbia University Press. [10]

Bennedsen, Morten, Francisco Pérez-González, and Daniel Wolfenzon. 2020. “Do CEOs Matter? Evidence from Hospitalization Events.” *Journal of Finance* 75 (4): 1877–911. [9]

Bertrand, Marianne, Simon Johnson, Krislert Samphantharak, and Antoinette Schoar. 2008. “Mixing Family with Business: A Study of Thai Business Groups and the Families Behind Them.” *Journal of Financial Economics* 88 (3): 466–98. [9]

Bertrand, Marianne, and Antoinette Schoar. 2003. “Managing with Style: The Effect of Managers on Firm Policies.” *Quarterly Journal of Economics* 118 (4): 1169–208. [6, 9, 105]

Bertrand, Marianne, and Antoinette Schoar. 2006. “The Role of Family in Family Firms.” *Journal of Economic Perspectives* 20 (2): 73–96. [9]

Bianchi, Nicola, Giulia Bovini, Jin Li, Matteo Paradisi, and Michael Powell. 2023. “Career Spillovers in Internal Labour Markets.” *Review of Economic Studies* 90 (4): 1800–1831. [12]

Bianchi, Nicola, and Matteo Paradisi. 2024. “Countries for Old Men: An Analysis of the Age Pay Gap.” Working Paper, 32340, National Bureau of Economic Research. [12]

Bijnens, Gert, Simon Jäger, and Benjamin Schoefer. 2025. “What Does Consulting Do?” Working Paper, 24272, National Bureau of Economic Research. [28]

Bramoullé, Yann, and Gilles Saint-Paul. 2010. “Social Networks and Labor Market Transitions.” *Labour Economics* 17 (1): 188–95. [41]

Bruhn, Miriam, Dean Karlan, and Antoinette Schoar. 2018. “The Impact of Consulting Services on Small and Medium Enterprises: Evidence from a Randomized Trial in Mexico.” *Journal of Political Economy* 126 (2): 635–87. [28]

Buhai, Sebastian, Miguel Portela, Coen Teulings, and Aico Van Vuuren. 2014. “Returns to Tenure or Seniority?” *Econometrica* 82 (2): 705–30. [12]

Caicedo, Santiago, Robert Lucas, and Esteban Rossi-Hansberg. 2019. “Learning, Career Paths, and the Distribution of Wages.” *American Economic Journal: Macroeconomics* 11 (1): 49–88. [40]

Cappelli, Peter. 2010. “The Rise and Decline of Managerial Development.” *Industrial and Corporate Change* 19 (2): 509–48. [9]

Cappelli, Peter, Rocio Bonet, and Monika Hamori. 2024. “The Changing Ranks of Corporate Leaders.” *California Management Review* 66 (2): 5–29. [11]

Cappelli, Peter, and Monika Hamori. 2004. “The Path to the Top: Changes in the Attributes and Careers of Corporate Executives, 1980–2001.” Working Paper, 10507, National Bureau of Economic Research. [9]

Cappelli, Peter, and Monika Hamori. 2014. “Understanding Executive Job Search.” *Organization Science* 25 (5): 1511–29. [45]

Cattell, Raymond. 1971. *Abilities: Their Structure, Growth, and Action.* Houghton Mifflin. [22]

Chetty, Raj, John Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan. 2020. “Income Segregation and Intergenerational Mobility Across Colleges in the United States.” *Quarterly Journal of Economics* 135 (3): 1567–633. [26, 88]

Choné, Philippe, Nathael Gozlan, and Francis Kramarz. 2024. “Workers-to-Firms Matching When Skills Become Unbundled.” Working Paper, 19438, Centre for Economic Policy Research. [22]

Cingano, Federico, and Alfonso Rosolia. 2012. “People I Know: Job Search and Social Networks.” *Journal of Labor Economics* 30 (2): 291–332. [41]

Cline, Brandon, and Adam Yore. 2016. “Silverback CEOs: Age, Experience, and Firm Value.” *Journal of Empirical Finance* 35: 169–88. [11, 105]

Corblet, Pauline. 2025. “Education, Sorting and Wages: A Structural Matching Approach.” Working Paper, NYU Abu Dhabi. [22]

Cornelissen, Thomas, Christian Dustmann, and Uta Schönberg. 2025. “Knowledge Spillovers, Competition, and Individual Careers.” Working Paper, 80/25, RFBerlin. [41]

Custódio, Cláudia, Miguel A Ferreira, and Pedro Matos. 2013. “Generalists Versus Specialists: Lifetime Work Experience and Chief Executive Officer Pay.” *Journal of Financial Economics* 108 (2): 471–92. [9, 10]

Cziraki, Peter, and Dirk Jenter. 2021. “The Market for CEOs.” Working Paper, 16281, Centre for Economic Policy Research. [9, 12, 13, 51]

d’Astous, Philippe, Thomas Geelen, and Jakub Hajda. 2025. “Old Workers, New Capital.” Working Paper, HEC Montréal. [12]

Daniels, Ronald. 2015. “A Generation at Risk: Young Investigators and the Future of the Biomedical Workforce.” *Proceedings of the National Academy of Sciences* 112 (2): 313–18. [11]

Davidson, Wallace, Carol Nemec, and Dan Worrell. 2006. “Determinants of CEO Age at Succession.” *Journal of Management & Governance* 10 (1): 35–57. [32]

De Chaisemartin, Clément, and Xavier d’Haultfoeuille. 2020. “Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects.” *American Economic Review* 110 (9): 2964–96. [44]

Dechow, Patricia, and Richard Sloan. 1991. “Executive Incentives and the Horizon Problem: An Empirical Investigation.” *Journal of Accounting and Economics* 14 (1): 51–89. [10]

Decressin, Yann, Steven N. Kaplan, and Morten Sorensen. 2025. “Have CEOs Changed?” *Journal of Financial Economics* 173: 104169. [9, 20]

Del Prato, Francesco. 2023. “The Importance of Working for Earnest: Climbing the Job Ladder through Firms’ Connectivity.” Working Paper, Aarhus University. [10]

Deming, David. 2017. “The Growing Importance of Social Skills in the Labor Market.” *Quarterly Journal of Economics* 132 (4): 1593–640. [9]

Deming, David. 2021. “The Growing Importance of Decision-Making on the Job.” Working Paper, 28733, National Bureau of Economic Research. [9]

Deming, David, and Mikko Silliman. 2025. “Skills and Human Capital in the Labor Market.” In *Handbook of Labor Economics,* edited by Christian Dustmann and Thomas Lemieux, 6: 115–57. Elsevier. [9]

Desjardins, Richard, and Arne Jonas Warnke. 2012. “Ageing and Skills.” Working Paper, 72, OECD Education. [23]

Dessein, Wouter, and Tano Santos. 2006. “Adaptive Organizations.” *Journal of Political Economy* 114 (5): 956–95. [7, 10]

Dimmock, Stephen, William Gerken, and Nathaniel Graham. 2018. “Is Fraud Contagious? Coworker Influence on Misconduct by Financial Advisors.” *Journal of Finance* 73 (3): 1417–50. [41]

Dow, James. 2013. “Boards, CEO Entrenchment, and the Cost of Capital.” *Journal of Financial Economics* 110 (3): 680–95. [17]

Dow, James, and Clara C. Raposo. 2005. “CEO Compensation, Change, and Corporate Strategy.” *Journal of Finance* 60 (6): 2701–27. [10]

Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Brücker. 2016. "Referral-Based Job Search Networks." *Review of Economic Studies* 83 (2): 514–46. [41]

Ebner, Natalie, Alexandra Freund, and Paul Baltes. 2006. "Developmental Changes in Personal Goal Orientation from Young to Late Adulthood: From Striving for Gains to Maintenance and Prevention of Losses." *Psychology and Aging* 21 (4): 664. [11]

Edmans, Alex, Xavier Gabaix, and Dirk Jenter. 2017. "Executive Compensation: A Survey of Theory and Evidence." *Handbook of the Economics of Corporate Governance* 1: 383–539. [17]

Edmans, Alex, Xavier Gabaix, and Augustin Landier. 2009. "A Multiplicative Model of Optimal CEO Incentives in Market Equilibrium." *Review of Financial Studies* 22 (12): 4881–917. [9, 23]

Ernst, Berit, and Alfred Kieser. 2002. "In Search of Explanations for the Consulting Explosion." In *The Expansion of Management Knowledge: Carriers, Flows, and Sources,* 47–73. Stanford University Press. [28]

Ewens, Michael, and Xavier Giroud. 2025. "Corporate Hierarchy." Working Paper, 34162, National Bureau of Economic Research. [40]

Farber, Henry S. 2010. "Job Loss and the Decline in Job Security in the United States." In *Labor in the New Economy,* 223–62. University of Chicago Press. [19]

Fee, Edward, and Charles Hadlock. 2000. "Management Turnover and Product Market Competition: Empirical Evidence from the U.S. Newspaper Industry." *Journal of Business* 73 (2): 205–43. [17]

Fee, Edward, and Charles Hadlock. 2004. "Management Turnover Across the Corporate Hierarchy." *Journal of Accounting and Economics* 37 (1): 3–38. [40]

Fee, Edward, Charles Hadlock, and Joshua Pierce. 2013. "Managers With and Without Style: Evidence Using Exogenous Variation." *Review of Financial Studies* 26 (3): 567–601. [9, 108]

Fons-Rosen, Christian, Patrick Gaule, and Taras Hrendash. 2023. "Why Has Science Become an Old Man's Game?" Working Paper, 16365, IZA – Institute of Labor Economics. [11]

Frey, Meredith. 2019. "What We Know, Are Still Getting Wrong, and Have Yet To Learn About the Relationships Among the SAT, Intelligence and Achievement." *Journal of Intelligence* 7 (4): 26. [27]

Freye, Saskia. 2010. "Germany's New Top Managers? The Corporate Elite in Flux, 1960–2005." Working Paper, 10, Max Planck Institute for the Study of Societies. [11]

Friedrich, Benjamin. 2023. "Information Frictions in the Market for Managerial Talent: Theory and Evidence." Working Paper, Northwestern University. [21]

Frydman, Carola. 2019. "Rising through the Ranks: The Evolution of the Market for Corporate Executives, 1936–2003." *Management Science* 65 (11): 4951–79. [8, 11, 15]

Fuller, Joseph, Stephen Hansen, Tejas Ramdas, and Raffaella Sadun. 2021. "The Demand for Executive Skills." Working Paper, 16280, Centre for Economic Policy Research. [9]

Gabaix, Xavier, and Augustin Landier. 2008. "Why Has CEO Pay Increased So Much?" *Quarterly Journal of Economics* 123 (1): 49–100. [9, 23]

Gao, Huasheng, Juan Luo, and Tilan Tang. 2015. "Effects of Managerial Labor Market on Executive Compensation: Evidence from Job-Hopping." *Journal of Accounting and Economics* 59 (2-3): 203–20. [9]

Garicano, Luis. 2000. "Hierarchies and the Organization of Knowledge in Production." *Journal of Political Economy* 108 (5): 874–904. [51]

Garicano, Luis, Jin Li, and Yanhui Wu. 2026. "Weak Bundle, Strong Bundle: How AI Redraws Job Boundaries." Working Paper, London School of Economics. [51]

Garicano, Luis, and Luis Rayo. 2025. "Training in the Age of AI: A Theory of Apprenticeship Viability." Working Paper, 20634, Centre for Economic Policy Research. [51]

Garicano, Luis, and Esteban Rossi-Hansberg. 2015. "Knowledge-Based Hierarchies: Using Organizations to Understand the Economy." *Annual Review of Economics* 7 (1): 1–30. [40]

Gayle, George-Levi, Limor Golan, and Robert Miller. 2015. "Promotion, Turnover, and Compensation in the Executive Labor Market." *Econometrica* 83 (6): 2293–369. [9]

Gibbons, Robert, and Michael Waldman. 2004. "Task-Specific Human Capital." *American Economic Review* 94 (2): 203–7. [10]

Giroud, Xavier. 2013. "Proximity and Investment: Evidence from Plant-Level Data." *Quarterly Journal of Economics* 128 (2): 861–915. [29]

Glitz, Albrecht. 2017. "Coworker Networks in the Labour Market." *Labour Economics* 44: 218–30. [41]

Glückler, Johannes. 2006. "A Relational Assessment of International Market Entry in Management Consulting." *Journal of Economic Geography* 6 (3): 369–93. [29]

Glückler, Johannes, and Thomas Armbrüster. 2003. "Bridging Uncertainty in Management Consulting: The Mechanisms of Trust and Networked Reputation." *Organization Studies* 24 (2): 269–97. [28]

Gompers, Paul, Steven Kaplan, and Vladimir Mukharlyamov. 2023. "The Market for CEOs: Evidence from Private Equity." Discussion Paper, 30899, National Bureau of Economic Research. [9]

Gong, Jie, Ang Sun, and Zhichao Wei. 2018. "Choosing the Pond: On-the-Job Experience and Long-Run Career Outcomes." *Management Science* 64 (2): 860–72. [10]

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics* 225 (2): 254–77. [43]

Groes, Fane, Philipp Kircher, and Iourii Manovskii. 2015. "The U-Shapes of Occupational Mobility." *Review of Economic Studies* 82 (2): 659–92. [10]

Grömping, Ulrike. 2015. "Variable Importance in Regression Models." *Wiley Interdisciplinary Reviews: Computational Statistics* 7 (2): 137–52. [49]

Gutiérrez, Germán, and Thomas Philippon. 2019. "Fading Stars." *AEA Papers and Proceedings* 109: 312–16. [18]

Haltiwanger, John, Ron Jarmin, and Javier Miranda. 2013. "Who Creates Jobs? Small versus Large versus Young." *Review of Economics and Statistics* 95 (2): 347–61. [7]

Hamori, Monika, and Maria Kakarika. 2009. "External Labor Market Strategy and Career Success: CEO Careers in Europe and the United States." *Human Resource Management* 48 (3): 355–78. [9]

Hanushek, Eric, Lavinia Kinne, Frauke Witthöft, and Ludger Woessmann. 2025. "Age and Cognitive Skills: Use It or Lose It." *Science Advances* 11 (10): eads1560. [23]

Hensvik, Lena, and Oskar Nordström Skans. 2016. "Social Networks, Employee Selection, and Labor Market Outcomes." *Journal of Labor Economics* 34 (4): 825–67. [41]

Hidalgo, César, and Ricardo Hausmann. 2009. "The Building Blocks of Economic Complexity." *Proceedings of The National Academy of Sciences* 106 (26): 10570–75. [37, 38, 92, 93]

Hopenhayn, Hugo, Julian Neira, and Rish Singhania. 2022. "From Population Growth to Firm Demographics: Implications for Concentration, Entrepreneurship and the Labor Share." *Econometrica* 90 (4): 1879–914. [7]

Hyatt, Henry, and James Spletzer. 2016. "The Shifting Job Tenure Distribution." *Labour Economics* 41: 363–77. [19]

Jarosch, Gregor, Ezra Oberfield, and Esteban Rossi-Hansberg. 2021. "Learning from Coworkers." *Econometrica* 89 (2): 647–76. [10]

Jedwab, Remi, Paul Romer, Asif Islam, and Roberto Samaniego. 2023. "Human Capital Accumulation at Work: Estimates for the World and Implications for Development." *American Economic Journal: Macroeconomics* 15 (3): 191–223. [10]

Jenter, Dirk, and Katharina Lewellen. 2015. "CEO Preferences and Acquisitions." *Journal of Finance* 70 (6): 2813–52. [6, 105]

Jones, Benjamin. 2010. "Age and Great Invention." *Review of Economics and Statistics* 92 (1): 1–14. [11]

Jones, Benjamin, and Bruce Weinberg. 2011. "Age Dynamics in Scientific Creativity." *Proceedings of the National Academy of Sciences* 108 (47): 18910–14. [11]

Kaplan, Steven, and Morten Sorensen. 2021. "Are CEOs Different?" *Journal of Finance* 76 (4): 1773–811. [9, 21]

Kim, Hyemin, Rüdiger Fahlenbrach, and Angie Low. 2023. "CEO Networks and the Labor Market for Directors." *Journal of Empirical Finance* 70: 1–21. [41]

Knyazeva, Anzhela, Diana Knyazeva, and Ronald Masulis. 2013. "The Supply of Corporate Directors and Board Independence." *Review of Financial Studies* 26 (6): 1561–605. [29]

Koch, Michael, Bernard Forgues, and Vanessa Monties. 2017. "The Way to the Top: Career Patterns of Fortune 100 CEOs." *Human Resource Management* 56 (2): 267–85. [9]

Kozeniauskas, Nicholas. 2025. "What's Driving the Decline in Entrepreneurship?" *Journal of Monetary Economics,* 103812. [7]

Lazear, Edward. 2009. "Firm-Specific Human Capital: A Skill-Weights Approach." *Journal of Political Economy* 117 (5): 914–40. [10, 47]

Levi, Maurice, Kai Li, and Feng Zhang. 2010. "Deal or No Deal: Hormones and the Mergers and Acquisitions Game." *Management Science* 56 (9): 1462–83. [11, 107]

Li, Xiaoyang, Angie Low, and Anil Makhija. 2017. "Career Concerns and the Busy Life of the Young CEO." *Journal of Corporate Finance* 47: 88–109. [10, 11, 105, 107]

Liang, James, Hui Wang, and Edward P Lazear. 2018. "Demographics and Entrepreneurship." *Journal of Political Economy* 126 (S1): 140–96. [18]

Liu, Yun. 2014. "Outside Options and CEO Turnover: The Network Effect." *Journal of Corporate Finance* 28: 201–17. [41]

Lyman, Noah, John Barry, and Lin Zhao. 2025. "Human Capital, Competition and Mobility in the Managerial Labor Market." Working Paper, University of Warwick. [9]

Ma, Yueran, and Andrei Shleifer. 2025. "The Invention of Corporate Governance." *Journal of Financial Economics,* 104115. [17]

Macis, Mario, and Fabiano Schivardi. 2016. "Exports and Wages: Rent Sharing, Workforce Composition, or Returns to Skills?" *Journal of Labor Economics* 34 (4): 945–78. [10]

MacLeod, W Bentley. 2016. "Human Capital: Linking Behavior to Rational Choice via Dual Process Theory." *Labour Economics* 41: 20–31. [11]

Mion, Giordano, and Luca David Opromolla. 2014. "Managers' Mobility, Trade Performance, and Wages." *Journal of International Economics* 94 (1): 85–101. [38]

Mohnen, Paul. 2025. "The Impact of the Retirement Slowdown on the US Youth Labor Market." *Journal of Labor Economics* 43 (1): 203–46. [12]

Montgomery, James. 1991. "Social Networks and Labor-Market Outcomes: Toward an Economic Analysis." *American Economic Review* 81 (5): 1408–18. [41]

Mueller, Philipp, Dimitrios Georgakakis, Peder Greve, Simon Peck, and Winfried Ruigrok. 2021. "The Curse of Extremes: Generalist Career Experience and CEO Initial Compensation." *Journal of Management* 47 (8): 1977–2007. [28]

Mullins, William, and Antoinette Schoar. 2016. “How Do CEOs See Their Roles? Management Philosophies and Styles in Family and Non-Family Firms.” *Journal of Financial Economics* 119 (1): 24–43. [9]

Murphy, Kevin, and Jan Zabojnik. 2004. “CEO Pay and Appointments: A Market-Based Explanation for Recent Trends.” *American Economic Review* 94 (2): 192–96. [8]

Murphy, Kevin, and Jan Zabojnik. 2007. “Managerial Capital and the Market for CEOs.” Working Paper, University of Southern California. [8]

Neal, Derek. 1995. “Industry-Specific Human Capital: Evidence from Displaced Workers.” *Journal of Labor Economics* 13 (4): 653–77. [10]

Ng, Thomas WH, and Daniel Feldman. 2008. “The Relationship of Age to Ten Dimensions of Job Performance.” *Journal of Applied Psychology* 93 (2): 392. [23]

Nguyen, Kieu-Trang. 2025. “Trust and Innovation Within the Firm: Evidence from Matched CEO-Firm Data.” *Quarterly Journal of Economics,* forthcoming. [9]

Nix, Emily. 2020. “Learning Spillovers in the Firm.” Working Paper, 14, Institute for Evaluation of Labour Market and Education Policy (IFAU), Uppsala. [10]

Oyer, Paul, and Scott Schaefer. 2009. “The Personnel-Economic Geography of US Law Firms and Law Schools.” Working Paper, Stanford University. [41]

Pan, Yihui. 2017. “The Determinants and Impact of Executive-Firm Matches.” *Management Science* 63 (1): 185–200. [7, 10]

Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Nicholas Trachter. 2021. “Diverging Trends in National and Local Concentration.” *NBER Macroeconomics Annual* 35 (1): 115–50. [36]

Roth, Jonathan, Pedro Sant’Anna, Alyssa Bilinski, and John Poe. 2023. “What’s Trending in Difference-in-Differences? A Synthesis of the Recent Econometrics Literature.” *Journal of Econometrics* 235 (2): 2218–44. [43]

Salthouse, Timothy. 2009. “When Does Age-Related Cognitive Decline Begin?” *Neurobiology of Aging* 30 (4): 507–14. [22]

Salthouse, Timothy. 2012. “Consequences of Age-Related Cognitive Declines.” *Annual Review of Psychology* 63: 201–26. [11]

Sauvagnat, Julien, and Fabiano Schivardi. 2024. “Are Executives in Short Supply? Evidence from Death Events.” *Review of Economic Studies* 91 (1): 519–59. [10, 29]

Saygin, Perihan Ozge, Andrea Weber, and Michele Weynandt. 2021. “Coworkers, Networks, and Job-Search Outcomes Among Displaced Workers.” *ILR Review* 74 (1): 95–130. [41]

Schaie, Warner. 1989. “The Hazards of Cognitive Aging.” *Gerontologist* 29 (4): 484–93. [23]

Schaie, Warner. 2005. *Developmental Influences on Adult Intelligence: The Seattle Longitudinal Study.* Oxford University Press. [22, 23]

Schoar, Antoinette, Kelvin Yeung, and Luo Zuo. 2024. “The Effect of Managers on Systematic Risk.” *Management Science* 70 (2): 815–33. [11, 108]

Schoar, Antoinette, and Luo Zuo. 2016. “Does the Market Value CEO Styles?” *American Economic Review* 106 (5): 262–66. [6, 9]

Schoar, Antoinette, and Luo Zuo. 2017. “Shaped by Booms and Busts: How the Economy Impacts CEO Careers and Management Styles.” *Review of Financial Studies* 30 (5): 1425–56. [9]

Schymik, Jan. 2018. “Globalization and The Evolution of Corporate Governance.” *European Economic Review* 102: 39–61. [38]

Serfling, Matthew. 2014. “CEO Age and the Riskiness of Corporate Policies.” *Journal of Corporate Finance* 25: 251–73. [11, 105, 107]

Sharapov, Dmitry, Paul Kattuman, Diego Rodriguez, and Javier Velazquez. 2021. “Using the Shapley Value Approach to Variance Decomposition in Strategy Research: Diversification, Internationalization, and Corporate Group Effects on Affiliate Profitability." *Strategic Management Journal* 42 (3): 608–23. [49]

Shue, Kelly. 2013. "Executive Networks and Firm Policies: Evidence from the Random Assignment of MBA Peers." *Review of Financial Studies* 26 (6): 1401–42. [41]

Singh-Manoux, Archana, Mika Kivimaki, Maria Glymour, Alexis Elbaz, Claudine Berr, Klaus Ebmeier, Jane Ferrie, and Aline Dugravot. 2012. "Timing of Onset of Cognitive Decline: Results from Whitehall II Prospective Cohort Study." *BMJ* 344. [11]

Sluyterman, Keetie, and Gerarda Westerhuis. 2022. "The Changing Role of CEOs in Dutch Listed Companies, 1957–2007." *Enterprise & Society* 23 (3): 711–45. [11]

Sullivan, Paul. 2010. "Empirical Evidence on Occupation and Industry Specific Human Capital." *Labour Economics* 17 (3): 567–80. [10]

Sundaram, Rangarajan, and David Yermack. 2007. "Pay Me Later: Inside Debt and Its Role in Managerial Compensation." *Journal of Finance* 62 (4): 1551–88. [11]

Taussig, Frank, and Carl Joslyn. 1932. *American Business Leaders: A Study in Social Origins and Social Stratification.* Literary Licensing, LLC. [11]

Terviö, Marko. 2008. "The Difference That CEOs Make: An Assignment Model Approach." *American Economic Review* 98 (3): 642–68. [9, 23]

Van Dijcke, David, Marcus Buckmann, Arthur Turrell, and Tomas Key. 2023. "Vacancy Posting, Firm Balance Sheets, and Pandemic Policy." Working Paper, 1033, Bank of England. [14]

Waelchli, Urs, and Jonas Zeller. 2013. "Old Captains at the Helm: Chairman Age and Firm Performance." *Journal of Banking & Finance* 37 (5): 1612–28. [11, 105]

Warner, William Lloyd, Darab Unwalla, John Trimm, and June Collins. 1967. *The Emergent American Society.* Yale University Press. [11]

Weinstein, Russell. 2025. "The Great Recession and the Widening Income Gap Between Alumni of Elite and Less Selective Universities." *American Economic Journal: Economic Policy* 17 (1): 216–41. [14, 30]

Wing, Coady, Madeline Yozwiak, Alex Hollingsworth, Seth Freedman, and Kosali Simon. 2024. "Designing Difference-in-Difference Studies with Staggered Treatment Adoption: Key Concepts and Practical Guidelines." *Annual Review of Public Health* 45. [43]

Yim, Soojin. 2013. "The Acquisitiveness of Youth: CEO Age and Acquisition Behavior." *Journal of Financial Economics* 108 (1): 250–73. [11, 107]

Yonker, Scott. 2017. "Geography and the Market for CEOs." *Management Science* 63 (3): 609–30. [29, 50]

Yotzov, Ivan, Jose Maria Barrero, Nicholas Bloom, Philip Bunn, Steven Davis, Kevin Foster, Aaron Jalca, et al. 2026. "Firm Data on AI." Working Paper, 34836, National Bureau of Economic Research. [11]

Appendix 1.A Model Appendix

1.A.1 Overview

We formalize the analysis underlying the model section. Job slots are indexed by a scalar productivity s; in the CEO application $s = z$, but the analysis extends to any scalar slot index generated by a hierarchical firm structure. Under transferable utility and a multiplicative surplus, equilibrium matching is positively assortative in slot productivity and worker score, so age patterns across the productivity distribution reduce to the distribution of worker age conditional on score. The analysis delivers three results.

(1) **Age and slot productivity at fixed α (Theorem 1.1).** Higher score tails contain stochastically older workers; under positive assortative matching (PAM), more productive slots are filled by stochastically older workers.

(2) **Within-tail age–ability selection (Theorem 1.2).** Selection into a score tail induces a weakly (strictly under mild regularity) negative age–ability relationship.

(3) **Comparative statics in α at fixed selectivity (Theorem 1.3).** Holding the selected fraction fixed, raising α shifts the tail age distribution upward in first-order stochastic dominance.

A parametric example with exponential ability illustrates how the key condition for the third result reduces to monotonicity of the age-output profile $y(\cdot)$.

1.A.2 Environment, Surplus, and Equilibrium Sorting

Slots have productivity $S \sim G_S$ on $[\underline{s}, \bar{s}]$ with $\underline{s} > 0$ (continuous CDF). Workers are characterized by ability A and age T. Match surplus is

$$\pi(s, a, t; \alpha) = s \cdot \phi(a, t; \alpha), \tag{1.A.1}$$

where the worker score is

$$\phi(a, t; \alpha) = (1 - \alpha)[a + h(t)] + \alpha y(t), \qquad \alpha \in [0, 1). \tag{1.A.2}$$

The restriction $\alpha < 1$ ensures ability enters with positive weight. Writing $\Phi_\alpha := \phi(A, T; \alpha)$, it is convenient to decompose the score as

$$\Phi_\alpha = (1 - \alpha)A + g_\alpha(T), \qquad g_\alpha(t) := (1 - \alpha)h(t) + \alpha y(t). \tag{1.A.3}$$

We maintain the following assumptions throughout.

Assumption 1.

(1) *Continuous score.* For each α, Φ_α has a continuous distribution.

(2) *Independence.* $A \perp T$ in the population.

(3) *Log-concave ability.* A has a log-concave density f_A.

(4) *Monotone composite age component.* $g_\alpha(\cdot)$ is weakly increasing on $[\underline{t}, \bar{t}]$.

Remark 2 (On Assumption 1.(4)). The composite term $g_\alpha = (1-\alpha)h + \alpha y$ can be increasing even when h is single-peaked, provided $\alpha y'(t) \geq -(1-\alpha)h'(t)$ for all t, which is satisfied when α is large enough relative to the rate of decline of h, or y' is sufficiently steep.

Remark 3 (Positive assortative matching). Under transferable utility, stable matchings maximize total surplus. Because $(s, \varphi) \mapsto s\varphi$ has strictly increasing differences on $\mathbb{R}^2_+$, equilibrium matching is positively assortative in (S, Φ_α): higher-productivity slots are filled by higher-score workers. Assumption 1.(1) ensures score cutoffs at fixed tail fractions are unique.

1.A.3 Fixed-α Tail Results

Fix α and write $\Phi := \Phi_\alpha$. For a tail fraction $\theta \in (0, 1)$, let $\kappa(\theta)$ denote the unique cutoff with $\mathbb{P}(\Phi \geq \kappa(\theta)) = \theta$, and write $C(\kappa) := \{\Phi \geq \kappa\}$.

Age increases stochastically with score

Under Assumption 1, the joint density of (Φ_α, T) is totally positive of order two (TP2). To see this, the change of variables $(a, t) \mapsto (x, t)$ with $x = (1-\alpha)a + g_\alpha(t)$ yields

$$f_{\Phi_\alpha, T}(x, t) \;=\; \frac{f_T(t)}{1-\alpha} f_A\!\left(\frac{x - g_\alpha(t)}{1-\alpha}\right). \tag{1.A.4}$$

Log-concavity of f_A implies the kernel $f_A(u - v)$ is TP2 in (u, v) (Karlin 1968); monotonicity of g_α preserves TP2 under the change of variables; and TP2 of $f_{\Phi_\alpha, T}$ implies that T is stochastically increasing in Φ_α—i.e. $\mathbb{P}(T > t \mid \Phi_\alpha = x)$ is nondecreasing in x for every t (Shaked and Shanthikumar 2007, Theorem 1.C.1). We denote this property $\mathrm{SI}(T \mid \Phi_\alpha)$.

Theorem 1.1 (Higher score tails select older workers; more productive slots filled by older workers). *Under Assumption 1,* $\mathrm{SI}(T \mid \Phi)$ *holds. Consequently:*

(i) (Tail selectivity.) *If* $\kappa_2 > \kappa_1$, *then* $T \mid \{\Phi \geq \kappa_2\} \geq_{st} T \mid \{\Phi \geq \kappa_1\}$. *Equivalently, for* $0 < \theta_2 < \theta_1 < 1$, $T \mid C(\kappa(\theta_2)) \geq_{st} T \mid C(\kappa(\theta_1))$.

(ii) (Slot productivity.) *Under PAM, the age distribution of the assigned worker is stochastically increasing in slot productivity. In particular, CEOs at larger firms are stochastically older.*

Proof. $\mathrm{SI}(T \mid \Phi)$ follows from the TP2 argument above.

Part (i). Fix t and let $q_t(x) := \mathbb{P}(T > t \mid \Phi = x)$, nondecreasing by SI. For $\kappa_2 > \kappa_1$, $\Phi \mid \{\Phi \geq \kappa_2\}$ FOSD-dominates $\Phi \mid \{\Phi \geq \kappa_1\}$, so $\mathbb{E}[q_t(\Phi) \mid \Phi \geq \kappa_2] \geq \mathbb{E}[q_t(\Phi) \mid \Phi \geq \kappa_1]$, i.e. $\mathbb{P}(T > t \mid \Phi \geq \kappa_2) \geq \mathbb{P}(T > t \mid \Phi \geq \kappa_1)$ for all t.

Part (ii). Under PAM with continuous distributions, the equilibrium assignment is an increasing map $s \mapsto \varphi(s)$ from slot productivity to worker score. The conditional age distribution at slot s is $T \mid \Phi = \varphi(s)$, which is stochastically increasing in s by SI and monotonicity of φ. □

Selection-induced age–ability correlation within a score tail

Membership in $C(\kappa) = \{\Phi \geq \kappa\}$ can be written as an age-dependent ability truncation:

$$C(\kappa) = \{A \geq c_\kappa(T)\}, \qquad c_\kappa(t) := \frac{\kappa - g_\alpha(t)}{1-\alpha}. \tag{1.A.5}$$

By Assumption 1.(4), $c_\kappa(\cdot)$ is weakly decreasing: older workers face a (weakly) lower ability threshold to enter the same tail.

Theorem 1.2 (Selection induces a negative age–ability relationship within a top tail). *Under Assumption 1, for any cutoff κ with $\mathbb{P}(C(\kappa)) > 0$:*

$$\mathrm{Cov}(A, T \mid C(\kappa)) \leq 0.$$

If moreover g_α is strictly increasing on a set of positive $T \mid C(\kappa)$-measure and the truncated-mean $m(x) := \mathbb{E}[A \mid A \geq x]$ is strictly increasing on the essential range of $c_\kappa(T)$ under $T \mid C(\kappa)$, then the covariance is strictly negative, and if additionally $\mathrm{Var}(A \mid C(\kappa)) > 0$ *then* $\mathrm{Corr}(A, T \mid C(\kappa)) < 0$.

Proof. For any t with $\mathbb{P}(C \mid T = t) > 0$, independence gives $\mathbb{E}[A \mid T = t, C] = m(c_\kappa(t))$. Hence $\mathrm{Cov}(A, T \mid C) = \mathrm{Cov}(m(c_\kappa(T)), T \mid C)$. Since $m(\cdot)$ is weakly increasing[49] and $c_\kappa(\cdot)$ is weakly decreasing, the composite $U := m \circ c_\kappa$ is weakly decreasing. For $\widetilde{T}$ an independent copy of $T \mid C$,

$$2\,\mathrm{Cov}(U(T), T \mid C) = \mathbb{E}[(U(T) - U(\widetilde{T}))(T - \widetilde{T}) \mid C] \leq 0,$$

since the factors have opposite signs. Strict inequality follows under the additional conditions. □

Remark 4 (Intuition). To clear a fixed score cutoff, an older worker can qualify with lower ability because age raises $g_\alpha(t)$. Hence within a high-score tail, older workers tend to have lower ability than younger workers, generating a negative relationship that is absent in the overall population.

49. $m'(x) = h_A(x)[m(x) - x] \geq 0$ since $m(x) \geq x$ and $h_A \geq 0$; strict when f_A is strictly positive near x.

1.A.4 Comparative Statics in α

Fix $\theta \in (0,1)$ and $\alpha_1 < \alpha_2$. For each α_i, let $\kappa_i := \kappa_{\alpha_i}(\theta)$ be the unique cutoff with $\mathbb{P}(\Phi_{\alpha_i} \geq \kappa_i) = \theta$, and define

$$C_i := \{\Phi_{\alpha_i} \geq \kappa_i\} = \{A \geq c_i(T)\}, \qquad c_i(t) := \frac{\kappa_i - g_{\alpha_i}(t)}{1-\alpha_i}. \tag{1.A.6}$$

By Assumption 1.(4) each $c_i(\cdot)$ is weakly decreasing (with slope $-g'_{\alpha_i}(t)/(1-\alpha_i) \leq 0$).

Remark 5 (Mapping to the CEO application). Each firm's rank in the slot-productivity distribution G_S is independent of α, so the tail fraction θ corresponding to a given firm is fixed across economies. Theorem 1.3 therefore compares CEO age distributions at a given firm-size quantile across environments with different α.

Cutoff schedules across α

A direct calculation yields

$$c_2(t) - c_1(t) = \underbrace{\left(\frac{\kappa_2}{1-\alpha_2} - \frac{\kappa_1}{1-\alpha_1}\right)}_{\text{constant in } t} - \underbrace{\left(\frac{\alpha_2}{1-\alpha_2} - \frac{\alpha_1}{1-\alpha_1}\right)}_{>0} \cdot y(t). \tag{1.A.7}$$

If $y(\cdot)$ is weakly increasing, then $c_2(t) - c_1(t)$ is weakly decreasing in t: raising α relatively lowers the ability threshold for older workers. Because θ is held fixed, the schedules generically cross—a uniform ordering $c_2 \leq c_1$ would typically violate $\mathbb{P}(C_1) = \mathbb{P}(C_2) = \theta$.

By Bayes' rule and independence, the age density conditional on belonging to the top-θ tail is

$$f_{T|C_i}(t) \propto f_T(t)\,\bar{F}_A(c_i(t)), \tag{1.A.8}$$

so the likelihood ratio of the two tail densities satisfies $f_{T|C_2}(t)/f_{T|C_1}(t) \propto \bar{F}_A(c_2(t))/\bar{F}_A(c_1(t))$.

Main comparative-statics result

Assumption 2 (Tail likelihood-ratio monotonicity (TLR))**.** The ratio $t \mapsto \bar{F}_A(c_2(t))/\bar{F}_A(c_1(t))$ is weakly increasing on $[\underline{t}, \bar{t}]$.

Theorem 1.3 (Greater weight on experience shifts the selected age distribution upward). *Fix $\theta \in (0,1)$ and $\alpha_1 < \alpha_2$. Under Assumptions 1.(1), 1.(2), and 2, increasing the weight on experience shifts selection toward older workers. Formally,*

$$T \mid C_2 \geq_{st} T \mid C_1.$$

Proof. By (1.A.8), the likelihood ratio $f_{T|C_2}/f_{T|C_1}$ is proportional to $\bar{F}_A(c_2(t))/\bar{F}_A(c_1(t))$, which is nondecreasing by Assumption 2. MLR dominance implies FOSD (Whitt 1982). □

Sufficient conditions for TLR

Exponential ability. A revealing example uses an exponential distribution of ability. If $\bar{F}_A(x) = e^{-\lambda x}$ on the relevant cutoff range, then $\bar{F}_A(c_2(t))/\bar{F}_A(c_1(t)) = \exp\{-\lambda(c_2(t) - c_1(t))\}$. When $y(\cdot)$ is weakly increasing, $c_2(t) - c_1(t)$ is weakly decreasing by (1.A.7), so the ratio is weakly increasing—TLR holds. The exponential case is especially transparent because the constant hazard rate means selection probabilities depend only on cutoff *differences*, eliminating any dependence on levels: monotonicity of y is both necessary and sufficient.

General hazard-rate condition. For non-exponential ability, differentiating the log of the likelihood ratio gives

$$\frac{d}{dt}\log\frac{\bar{F}_A(c_2(t))}{\bar{F}_A(c_1(t))} = h_A(c_2(t)) \cdot \frac{g'_{\alpha_2}(t)}{1-\alpha_2} - h_A(c_1(t)) \cdot \frac{g'_{\alpha_1}(t)}{1-\alpha_1},$$

where $h_A(x) := f_A(x)/\bar{F}_A(x)$ is the ability hazard rate. TLR holds whenever this expression is nonnegative for all t. Since $g'_\alpha(t)/(1-\alpha) = h'(t) + [\alpha/(1-\alpha)]y'(t)$ is increasing in α when $y'(t) \geq 0$, the slope term favors TLR; whether this dominates depends on how the hazard varies across the two endogenous cutoff levels $c_1(t)$ and $c_2(t)$.

1.A.5 Reshuffling Pressure from Changes in α

The model is static but quantifies how much the frictionless target ranking changes when α rises. For worker $w = (a, t)$, define her upper-tail rank under α by $q_\alpha(w) := \mu\{w' : \Phi_\alpha(w') \geq \Phi_\alpha(w)\}$, where μ is the population measure. Let $S_{\alpha_1,\theta} := \{w : q_{\alpha_1}(w) \leq \theta\}$ be the baseline top-θ workers. Cumulative reshuffling pressure is

$$V_\theta(\alpha_1, \alpha_2) := \int_{S_{\alpha_1,\theta}} TV_{[\alpha_1,\alpha_2]}\big(q_\cdot(w)\big)\, d\mu(w), \tag{1.A.9}$$

where TV denotes total variation. Using cumulative rather than net rank change ensures intermediate movements do not cancel.

Proposition 1.4 (Larger α-changes generate more reshuffling pressure). *For fixed θ and α_1, the map $\alpha_2 \mapsto V_\theta(\alpha_1, \alpha_2)$ is weakly increasing on $[\alpha_1, 1)$.*

Proof. For each pair (w, w'), $\Phi_\alpha(w') - \Phi_\alpha(w)$ is affine in α, so $\mathbf{1}\{\Phi_\alpha(w') \geq \Phi_\alpha(w)\}$ changes sign at most once, bounding $TV_{[\alpha_1,\alpha_2]}(q_\cdot(w)) \leq 1$ and ensuring (1.A.9) is well-defined. Monotonicity follows because total variation is monotone in the interval: $[\alpha_1, \alpha_2] \subseteq [\alpha_1, \alpha'_2]$ implies $TV_{[\alpha_1,\alpha_2]}(q_\cdot(w)) \leq TV_{[\alpha_1,\alpha'_2]}(q_\cdot(w))$ for every w; integrating over $S_{\alpha_1,\theta}$ gives the result. □

Remark 6. This measures *reshuffling pressure*, not realized turnover: it captures how much the frictionless assignment must be reorganized when the importance of experience rises. Larger changes in α require more extensive re-ranking of initially top-ranked workers.

Appendix 1.B Additional Figures

Figure 1.B.1. CEO Age at Appointment

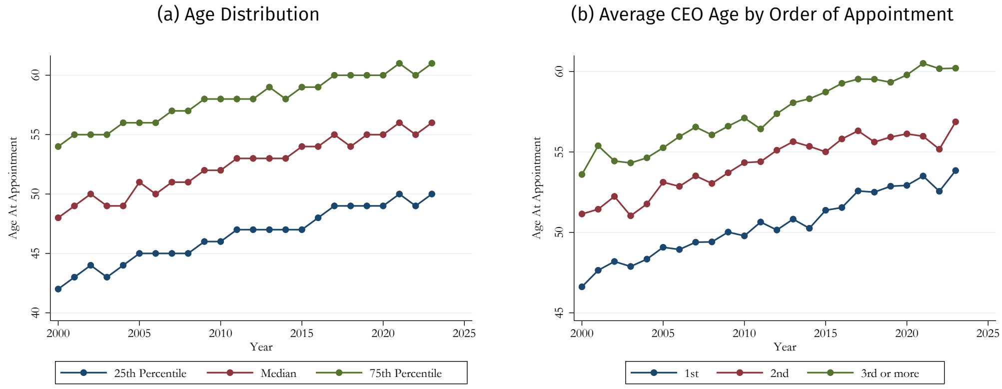

Notes: The figure shows CEO age at appointment over time. While Panel (a) illustrates the age distribution by plotting the quartiles of CEO age, Panel (b) shows average CEO age by order of appointment, distinguishing between first, second, and third or higher appointments.

Figure 1.B.2. Long-Run CEO Age Patterns

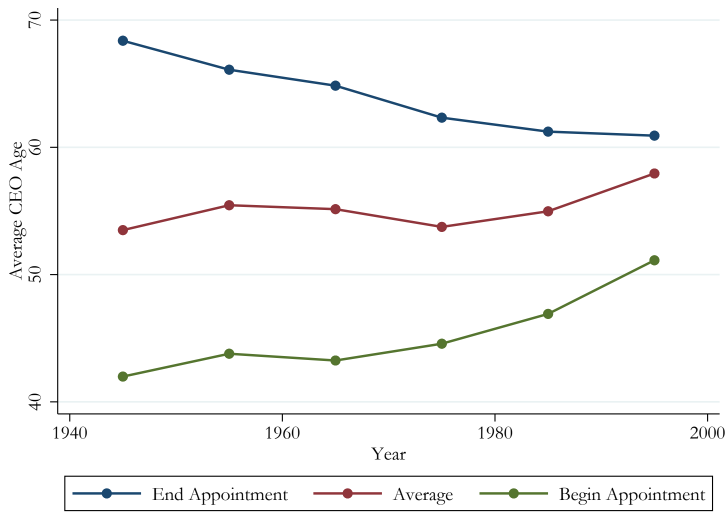

Notes: The figure plots long-run developments in average age as well as age at the beginning and the end of the appointment of American executives, obtained from the *Great American Business Leaders of the 20th Century database* (https://www.hbs.edu/leadership/20th-century-leaders). We show aggregates by decade to reduce noise.

Figure 1.B.3. Change in Average Age Since 2000 for CEOs vs. Labor Force

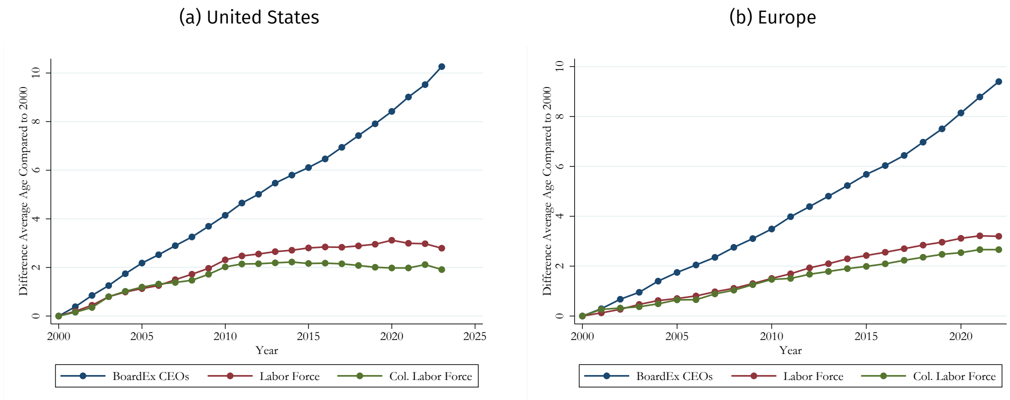

Notes: The figure compares the average age of sitting CEOs from BoardEx with that of the overall and the college-educated labor force, calculated from survey data. Panel (a) uses data from the *Current Population Survey*, showing the average age of all respondents in the labor force, calculated using ASEC weights. We define individuals holding at least a bachelor's degree as college-educated. The data have been obtained via *IPUMS* (Flood et al., IPUMS CPS: Version 12.0 [dataset]). Panel (b) uses data from the *European Labour Force Survey*. We compute average labor force age at the country level and then aggregate across countries, weighting by the number of BoardEx CEOs in each country-year cell. The countries included are Austria, Belgium, Switzerland, Cyprus, Germany, Denmark, Estonia, Greece, Spain, Finland, France, Hungary, Italy, Lithuania, Luxembourg, Norway, Poland, Portugal, and Sweden. We define individuals with some tertiary education as college-educated (ISCED-level 5 or above).

Figure 1.B.4. Average CEO Age Over Time—Europe

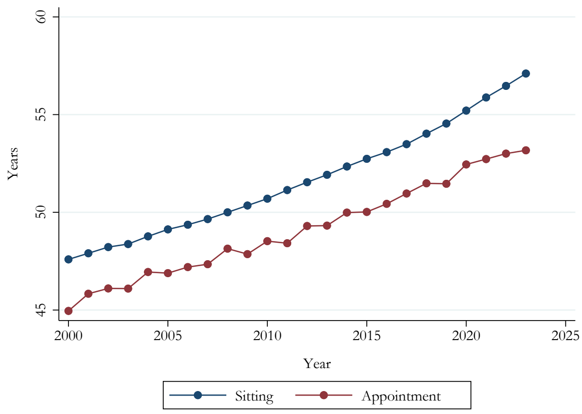

Notes: The plot shows the average age of CEOs in Europe over time, separately for sitting CEOs and CEOs at appointment. The countries included are Austria, Belgium, Switzerland, Cyprus, Germany, Denmark, Estonia, Greece, Spain, Finland, France, Hungary, Italy, Lithuania, Luxembourg, Norway, Poland, Portugal, and Sweden.

Figure 1.B.5. Contribution of Labor Force Aging to CEO Age Increase

(a) Change in Share of Labor Force by Age Bin

(b) Coefficient of CEO Age on Labor Force Age Bins

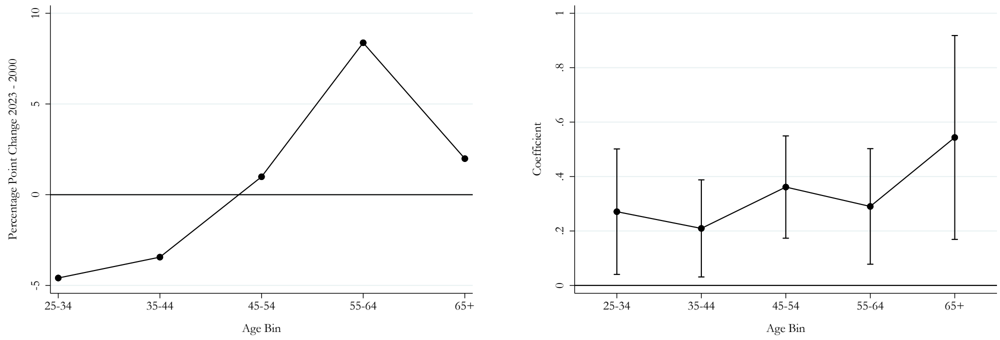

Notes: The figure shows results used to quantify the contribution of labor force aging to the CEO age increase. Panel (a) shows the average percentage point change in the share of the labor force by 10-year age bin between 2000 and 2023, using data from the ILO (*Labour force by sex and age (thousands))*. The countries included are Austria, Belgium, Switzerland, Cyprus, Germany, Denmark, Estonia, Greece, Spain, Finland, France, Hungary, Italy, Lithuania, Luxembourg, Norway, Poland, Portugal, Sweden, and the United States. Panel (b) shows coefficients obtained from regressing average CEO age by year and country on the share of the labor force by 10-year age bin, with individuals 15–24 as the omitted category. The regression controls for country and year fixed effects. Point estimates are plotted alongside 95% confidence intervals, based on robust standard errors.

Figure 1.B.6. Time-Varying Correlation of CEO Tenure and Age at Appointment

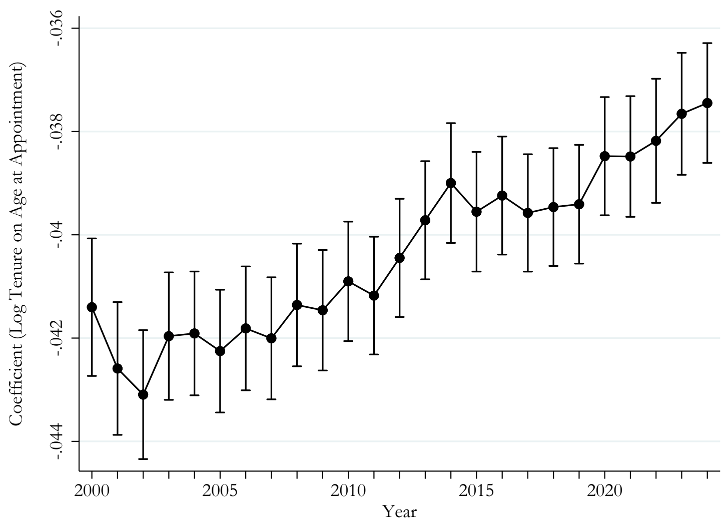

Notes: The figure shows the time-varying correlation of CEO tenure with age at appointment. The data are at the CEO-year level. The coefficients are obtained from regressions of the natural logarithm of tenure on year dummies, both alone and interacted with CEO age at appointment. Point estimates are plotted alongside 95% confidence intervals, based on standard errors clustered at the firm level.

Figure 1.B.7. CEO Education by Year of Appointment

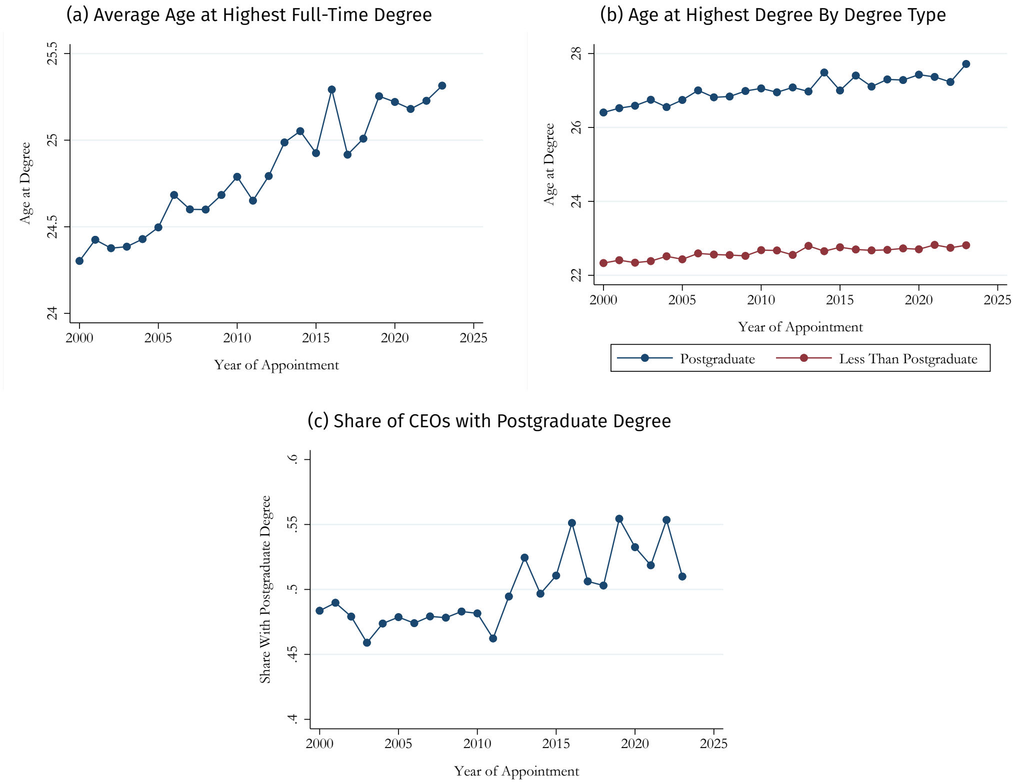

Notes: The figure shows CEO education by year of appointment, focusing on full-time degrees obtained after at most 5 years post labor market entry. The data are from BoardEx. Panel (a) plots average age at the highest full-time degree. Panel (b) shows average age at the highest full-time degree, distinguishing between postgraduate and less than postgraduate degrees. Panel (c) plots the share of CEOs with a post-graduate degree.

Figure 1.B.8. Employment Share of “Professional and Management Development Training”

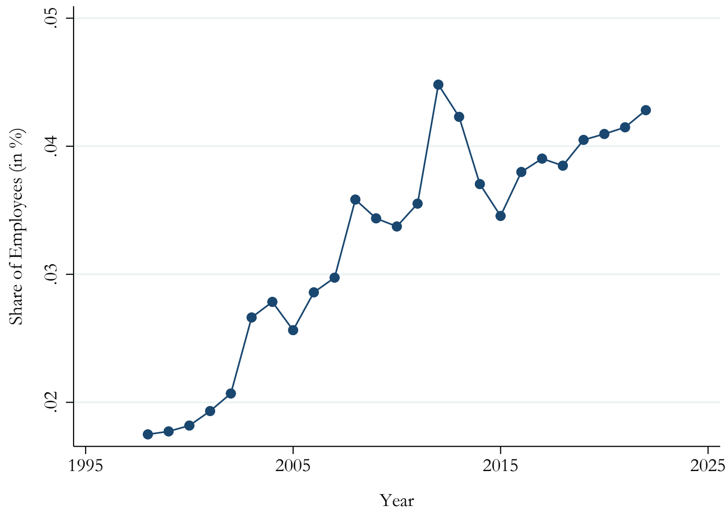

Notes: The figure shows the number of employees employed in the NAICS sector “611430: Professional and Management Development Training”, divided by the overall number of employees in the economy. The data are obtained from the US flat files of County Business Patterns.

Figure 1.B.9. CEO Age at Appointment—Robustness to Population Reweighting

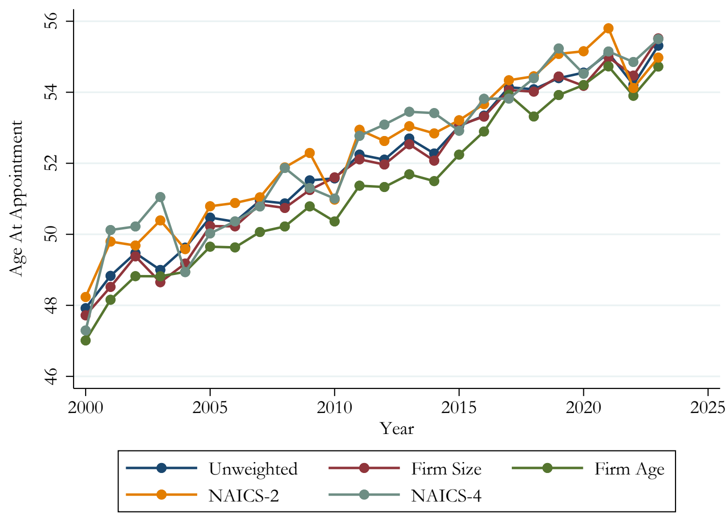

Notes: The figure shows CEO age at appointment over time. The blue line plots the unweighted average, whereas the other lines compare our sample to the Business Dynamics Statistics (BDS) and reweight CEO age at appointment by the inverse frequency of firms in each category (firm size, firm age, NAICS-2 or NAICS-4 industries) relative to the population in the respective year. The BDS data have been drawn from the US Census Bureau. Firm size (file *BDSFSIZE*) is binned to the following categories: at least 1, 20, 100, 500, 1,000, 2,500, 5,000, and 10,000 employees. Firm age (file *BDSFAGE*) is binned to the following categories: at least 1, 6, 11, 16, and 21 years. Sectoral counts at the NAICS-2 and NAICS-4 level have been obtained from file *BDSNAICS*.

Figure 1.B.10. CEO Age at Appointment—Predicted Age Difference 2000–2023 Conditional on Controls

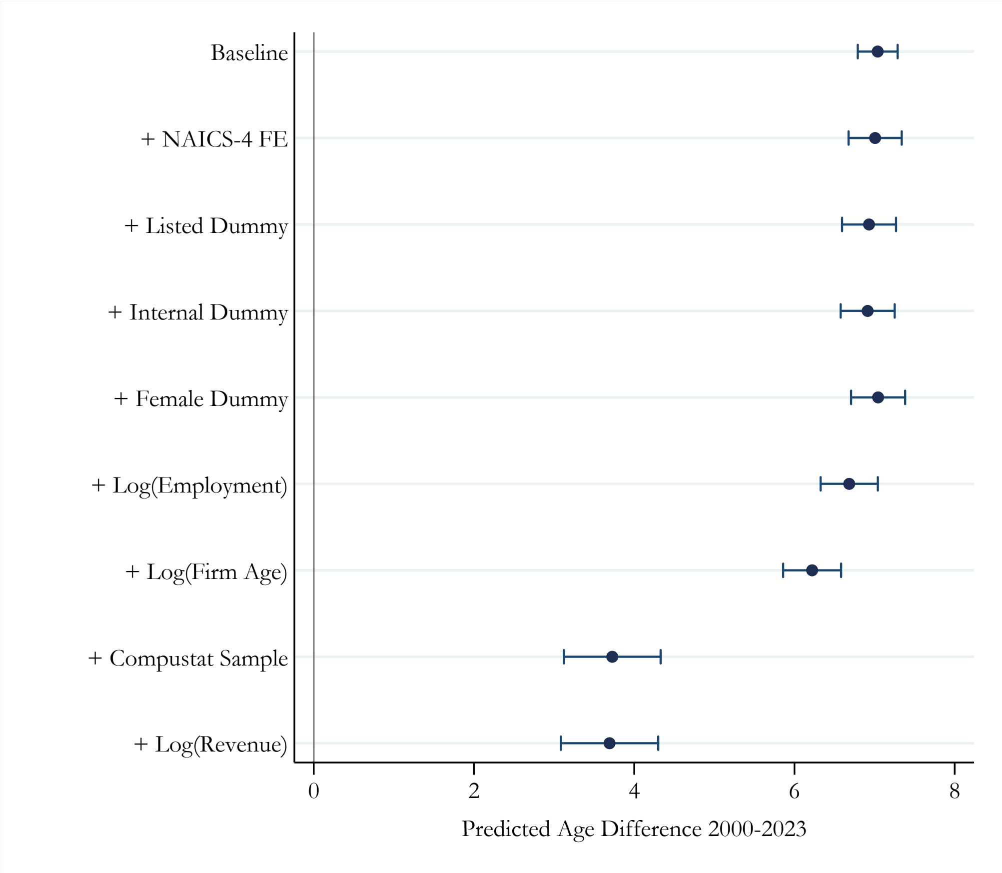

Notes: The figure shows coefficients obtained from regressing CEO age at appointment on a linear time trend (normalized to range between 0 and 1 from 2000 to 2023). While the baseline specification includes no controls, we then consecutively control for NAICS-4 fixed effects, dummies for whether a firm is listed, or whether the CEO is promoted internally or is female, and the natural logarithm of employment and firm age (both derived from employment biographies). We then restrict the sample to the set of firms contained in Compustat. Finally, we control for the natural logarithm of revenue. Point estimates are plotted alongside 95% confidence intervals, based on standard errors clustered at the firm level.

Figure 1.B.11. CEO Age of Firms Delisted 1997–2002—Differential Effects

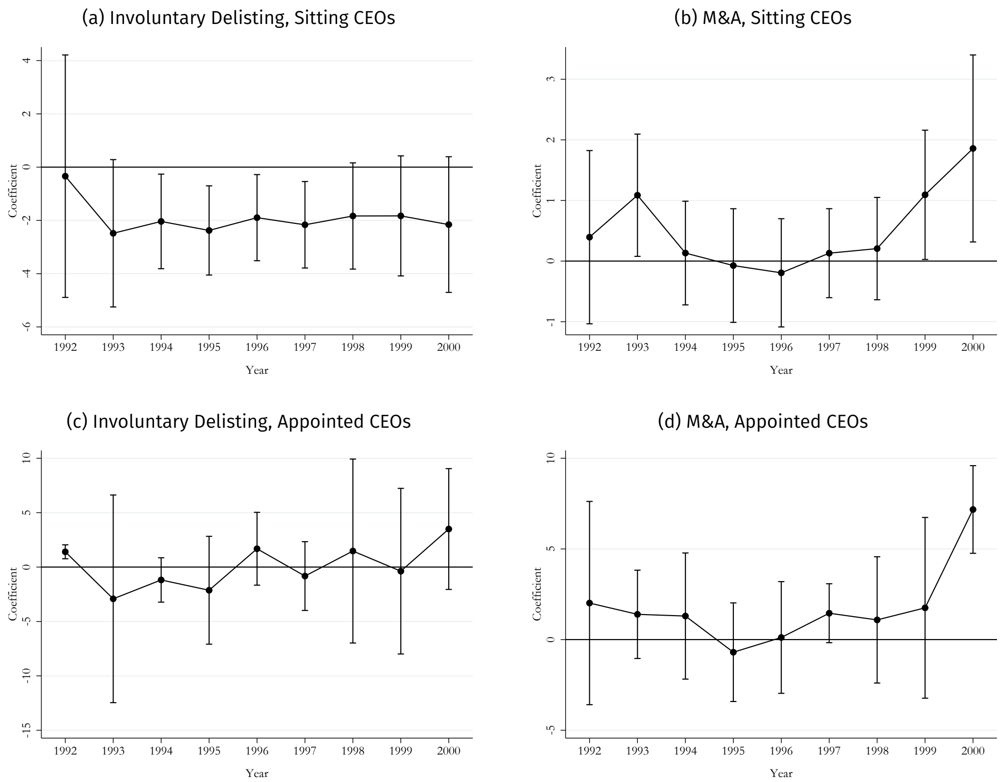

Notes: The figure shows coefficients from regressions of CEO age on year fixed effects interacted with indicators for different types of delistings (involuntary delistings in Panels (a) and (c) and delistings due to M&A activity in Panels (b) and (d)). The CEO data are from Execucomp and cover the S&P1500 between 1992 and 2000. The data have been combined with information on delistings between 1997 and 2002, obtained from CRSP. We follow the literature (e.g., Bharath and Dittmar 2010) in our definition of involuntary delistings due to bankruptcy or liquidation (DLSTCD codes $\geq$ 400 excluding 570 and 573) and mergers and acquisitions (DLSTCD codes 200–399 excluding 332). All regressions include NAICS-4-year fixed effects and control for the natural logarithm of firm age, employment, and revenue, obtained from Compustat. Point estimates are plotted alongside 95% confidence intervals, based on standard errors clustered at the NAICS-4 level.

Figure 1.B.12. Share of CEO Appointments with Experience as CFO or COO Over Time

(a) CFO (b) COO

Notes: This figure shows the share of appointed CEOs who previously worked as CFO (Panel a) or COO (Panel b), distinguishing between those who served in this role at the same firm or a different firm from the one that promoted them to CEO. CFOs are identified by role names containing any of the following keywords: “CFO”, “Chief Financial Officer”, “General Manager - Finance”, “Head of Finance”, and “Director - Finance”. For COOs, the corresponding keywords are “COO”, “Chief Operations Officer”, “General Manager - Operations”, “Head of Operations”, and “Director - Operations”.

Figure 1.B.13. CEO Appointments by Seniority of First Position in Firm

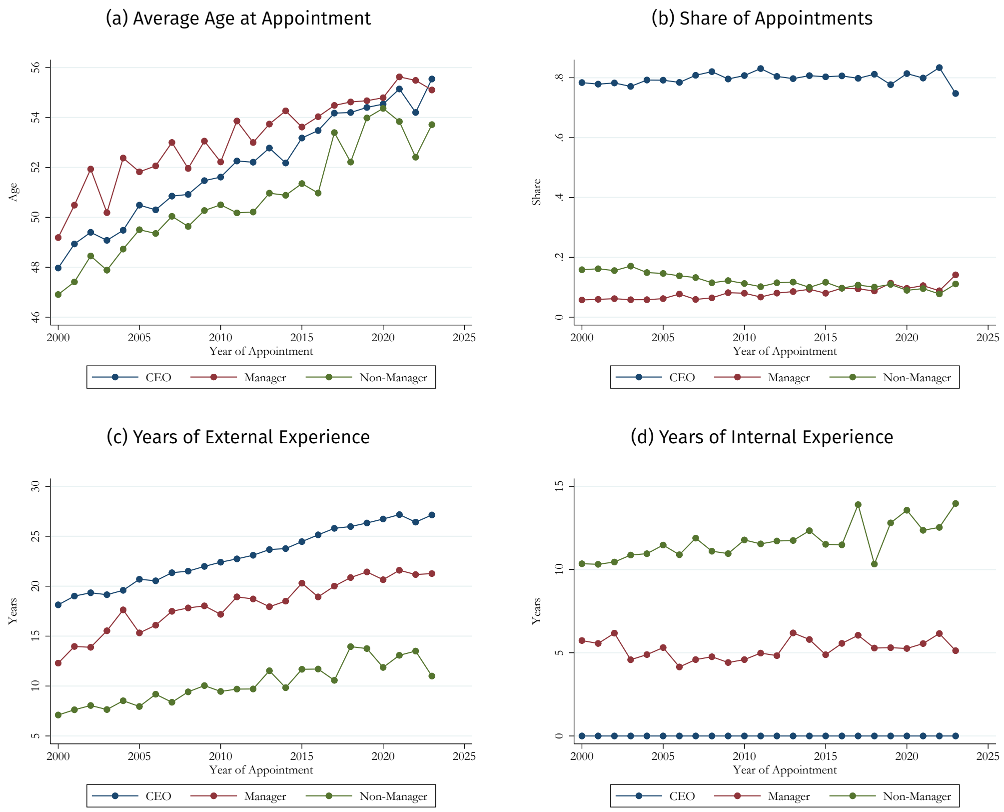

Notes: The figure plots descriptives on CEO appointments over time, stratified by the seniority of their first position in the firm. Panel (a) displays the average age at appointment, while Panel (b) shows the share of appointments for each group. Panels (c) and (d) present the average years of external and internal experience, respectively.

Figure 1.B.14. Share of Internal Appointments by Firm Size and Listing Status

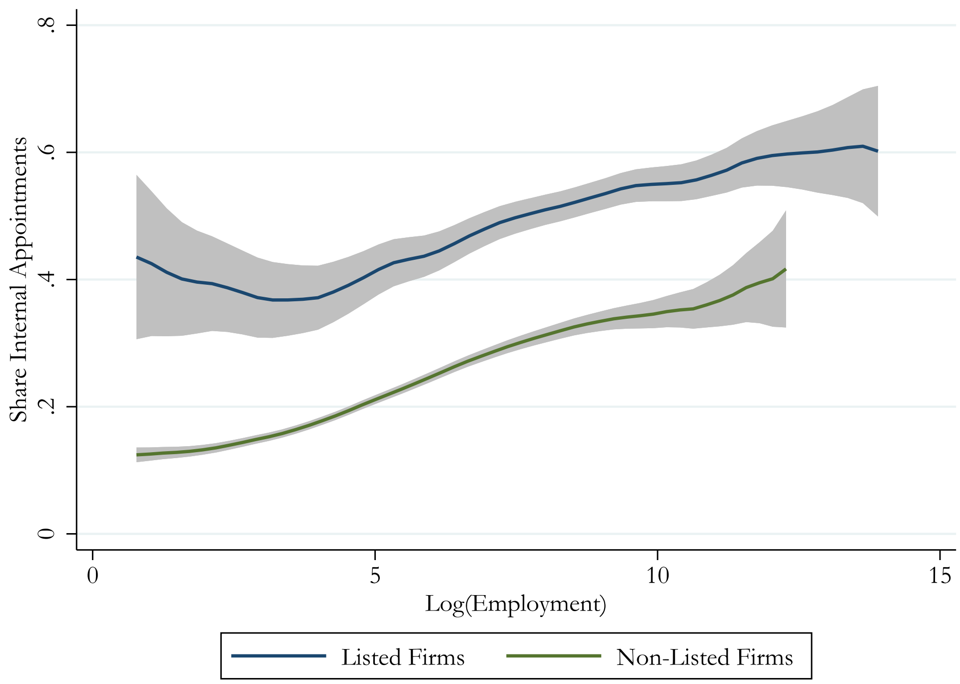

Notes: The figure plots the share of internal appointments against firm size measured by the natural logarithm of employment (obtained from employment histories). The data are smoothed with a local polynomial with Epanechnikov kernel and bandwidths of 0.91 and 0.78, respectively, which have been obtained via optimal bandwidth selection. Point estimates are shown alongside 95% confidence intervals.

Figure 1.B.15. Years of External vs. Internal Experience by Firm Size

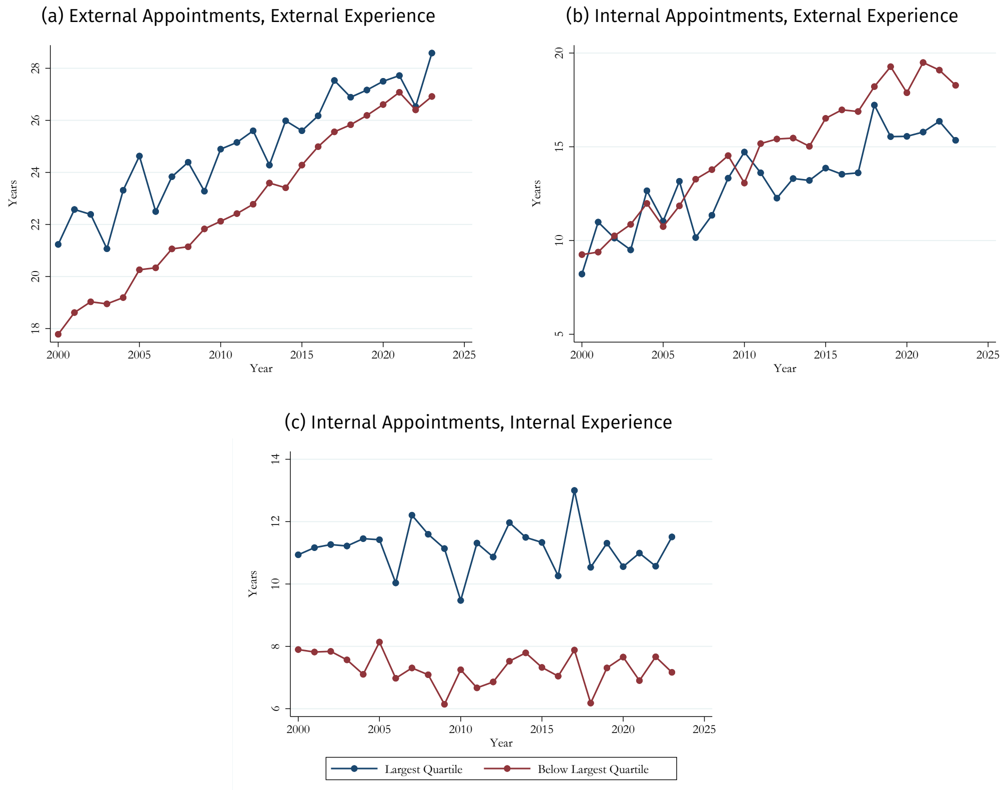

Notes: The figure plots years of experience, separately for those in the largest firm quartile and for those in the other three quartiles. Panel (a) shows years of external experience for externally appointed CEOs, Panel (b) shows years of external experience for internally appointed CEOs, and Panel (c) shows years of internal experience for internally appointed CEOs.

Figure 1.B.16. Work Experience at First Appointment to Most Senior Position: CEOs vs. Non-CEO Executives

(a) Years of External vs. Internal Work Experience

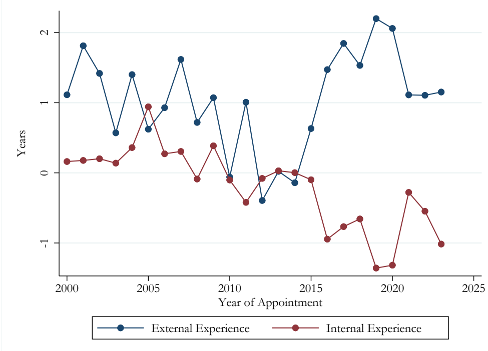

(b) Average Number of Positions, Firms, and Sectors

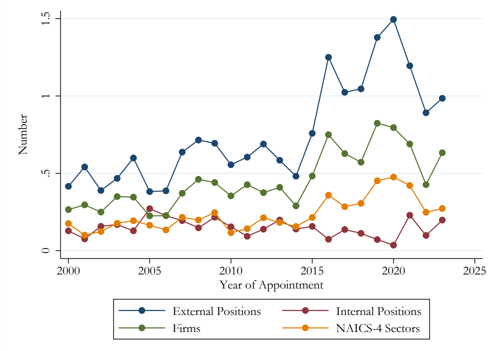

(c) Years per Position and Firm

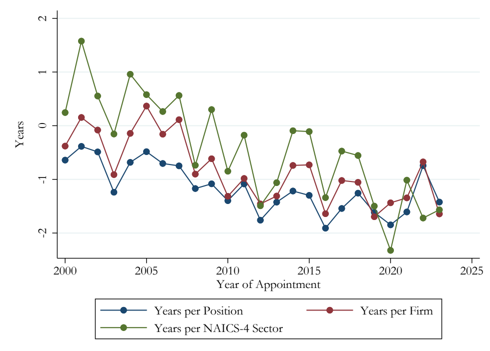

Notes: The figure plots work experience of CEOs compared to non-CEO executives over time, computed as the CEO average minus the non-CEO executive average in each year. The groups are defined based on the most senior position reached by the individual throughout the sample, focusing on the first appointment. Panel (a) plots the average external and internal work experience prior to appointment. Panel (b) shows the average number of positions (separately for external and internal positions) as well as the number of firms and sectors the executive worked in. Panel (c) shows the average number of years per position, firm, and NAICS-4 industry.

Figure 1.B.17. Work Experience at First Appointment to Most Senior Position: CEOs vs. Non-Executive Managers

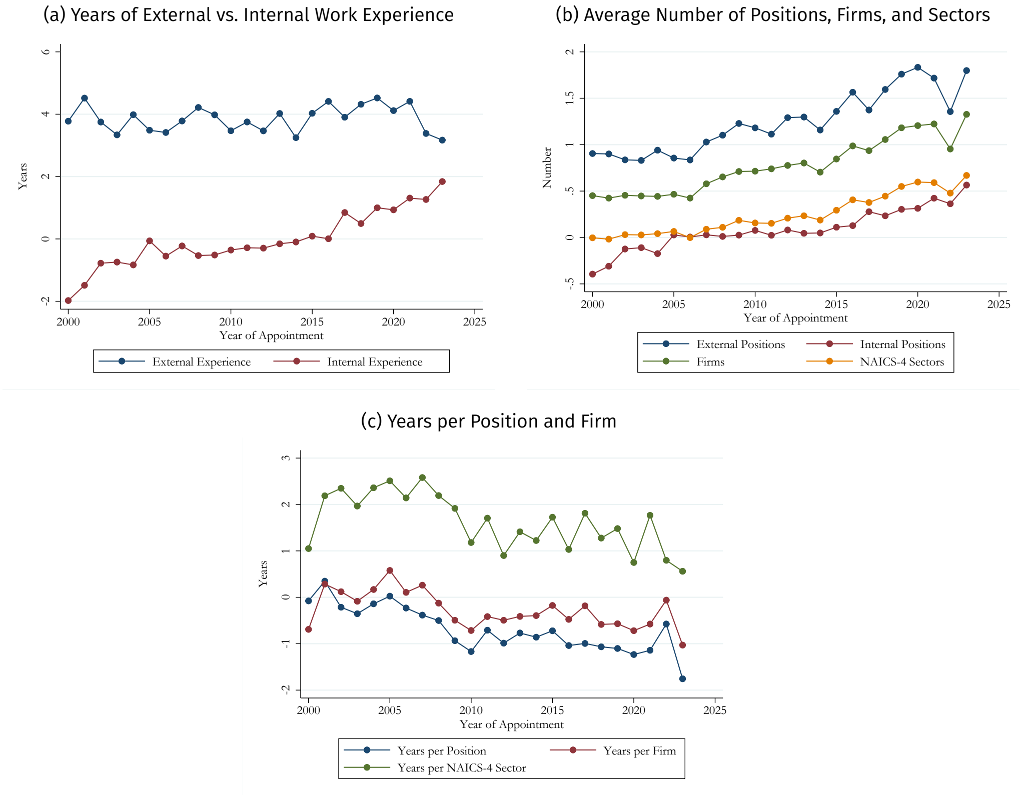

Notes: The figure plots work experience of CEOs compared to non-executive managers over time, computed as the CEO average minus the non-executive manager average in each year. The groups are defined based on the most senior position reached by the individual throughout the sample, focusing on the first appointment. Panel (a) plots the average external and internal work experience prior to appointment. Panel (b) shows the average number of positions (separately for external and internal positions) as well as the number of firms and sectors the CEO worked in. Panel (c) shows the average number of years per position, firm, and NAICS-4 industry.

Figure 1.B.18. Illustrative Correlation of Age and Ability at the Top Position by Firm Productivity

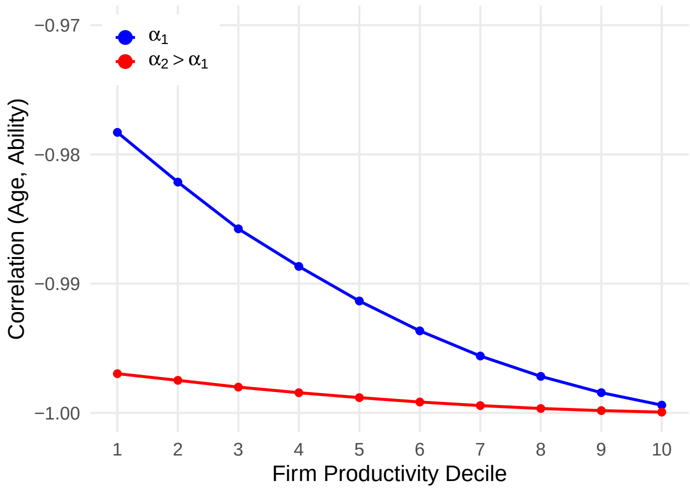

Notes: The figure plots the correlation of age and ability across firm productivity deciles for two values of α: $\alpha_1 = 0.4$ (blue line) and $\alpha_2 = 0.8$ (red line) and an illustrative parametrization. We simulate 100,000 workers with ability $a \sim U(0, 200)$ and age $t \sim U(0, 50)$. Skills are defined as $h(t) = -0.04t(t - 50)$ and $y(t) = -0.1t(t - 150)$. On the firm side, we draw 1,000 firms with productivity $z \sim U(0, 1)$, capacity $m = 100$, and hierarchical decay parameter $\beta = 0.08$. Within each firm, the CEO corresponds to the top slot ($k = 1$). We exploit the assortative matching implied by the model and compute the correlation of a and t within each firm productivity decile. The simulation is repeated 500 times, and results are averaged across all replications.

Figure 1.B.19. Complexity Measures Over Time

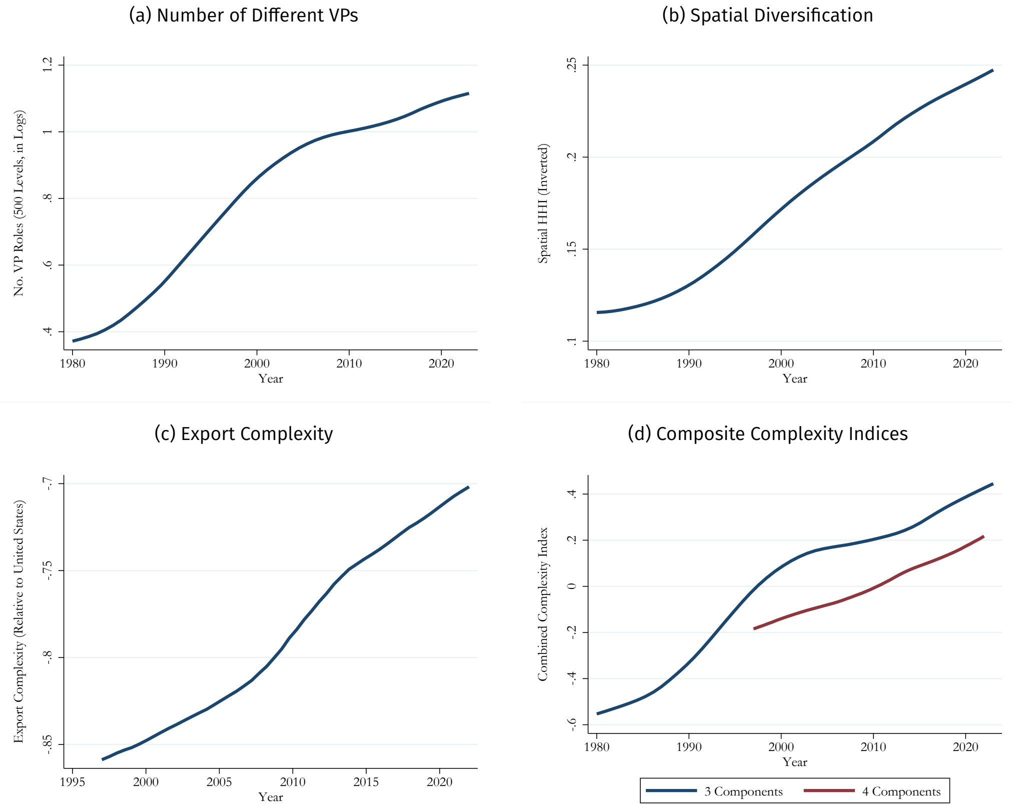

Notes: This figure shows different measures of sectoral complexity over time, smoothed using a local polynomial. Panel (a) plots the average number of distinct VPs per firm, taking the median at the NAICS-4 level and the natural logarithm thereof. Panel (b) shows the average HHI of employees across offices, where an office is defined as having at least five employees in a given MSA-year. We take the median at the NAICS-4 level. For exposition, we plot the inverse of the measure so that higher values correspond to greater diversification. Panel (c) shows average export complexity, defined as the economic complexity of trade partners weighted by the NAICS-4 industry-specific share of exports to that country. We express this as the distance from the complexity frontier by subtracting the average complexity of the US economy from the EECI. Panel (d) shows composite indices of complexity, constructed by standardizing all variables to mean zero and a standard deviation of one (if necessary), and summing across components. The 3-component index includes the industry-level standard deviation of revenue growth (the uncertainty measure used in the main analysis) together with the measures from Panels (a) and (b); the 4-component index additionally uses the export complexity measure from Panel (c).

Figure 1.B.20. Uncertainty Indices Over Time

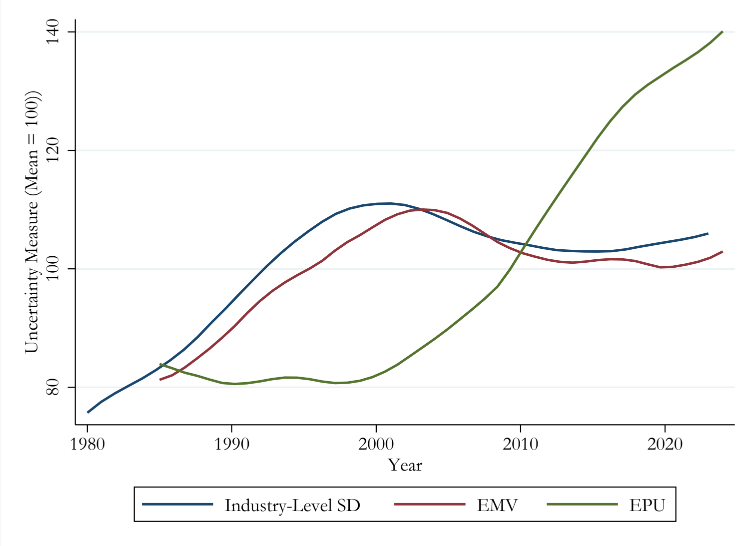

Notes: The figure plots different measures of aggregate uncertainty over time. *Industry-Level SD* is the cross-sectional (NAICS-4-level) standard deviation of revenue growth rate (obtained from Compustat), which we aggregate by weighting with revenue. *EMV* is a newspaper-based Equity Market Volatility index constructed in Baker et al. (2019). *EPU* is the Economic-Policy Uncertainty index for the US, obtained from Baker, Bloom, and Davis (2016). We annualize the variables by computing their arithmetic means and normalize them by adjusting their average over the sample window to 100. The data are smoothed with a local polynomial with Epanechnikov kernel and bandwidths of 3.78, 5.26, and 5.08, respectively, which have been obtained via optimal bandwidth selection.

Figure 1.B.21. MBB Availability, Uncertainty, and CEO Age—Placebo Test

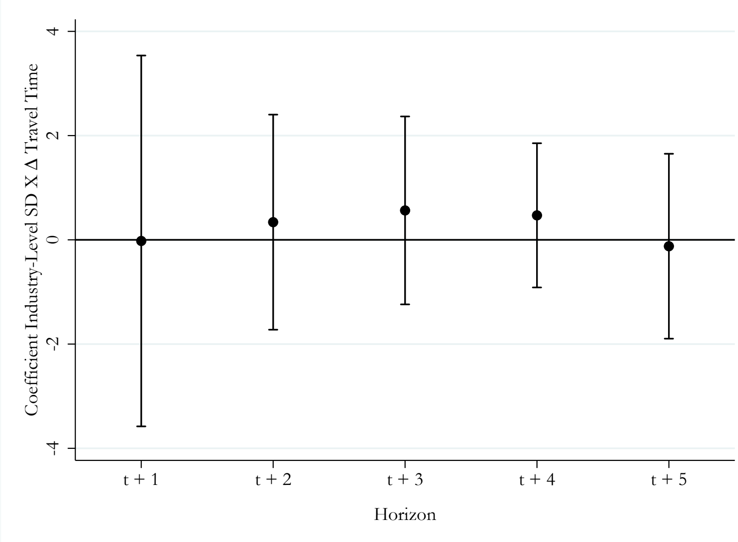

Notes: The plot shows coefficients from estimating the specification from Table 1.1, Column (4), while replacing *Travel Time*$_{mt}$ with *ΔTravel Time*$_{m\,\tau-t}$, $\tau = \{1 \ldots 5\}$. Point estimates are plotted alongside 95% confidence intervals, based on standard errors double-clustered at the NAICS-4 industry and MSA level.

Figure 1.B.22. MBB Availability, Uncertainty, and CEO Age—Time Variation

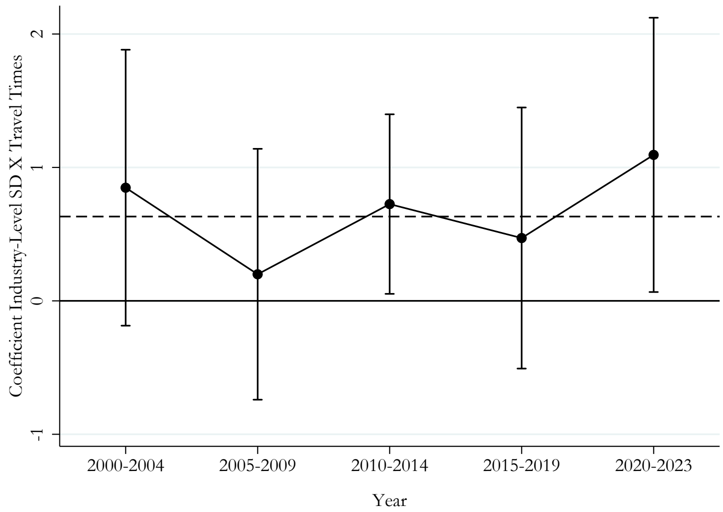

Notes: The plot shows coefficients from time-varying estimates, obtained by augmenting the specification from Table 1.1, Column (4) with interactions for 5-year bins. Uncertainty, calculated as the cross-sectional standard deviation of the symmetric growth rate of firm revenue within a firm's NAICS-4 industry, is standardized within each 5-year bin to mean zero and standard deviation one. Point estimates are plotted alongside 95% confidence intervals, based on standard errors double-clustered at the NAICS-4 industry and MSA level. The dashed horizontal line shows the point estimate from the baseline regression.

Figure 1.B.23. Share of Downward/Lateral Switches of CEOs Prior to Appointment

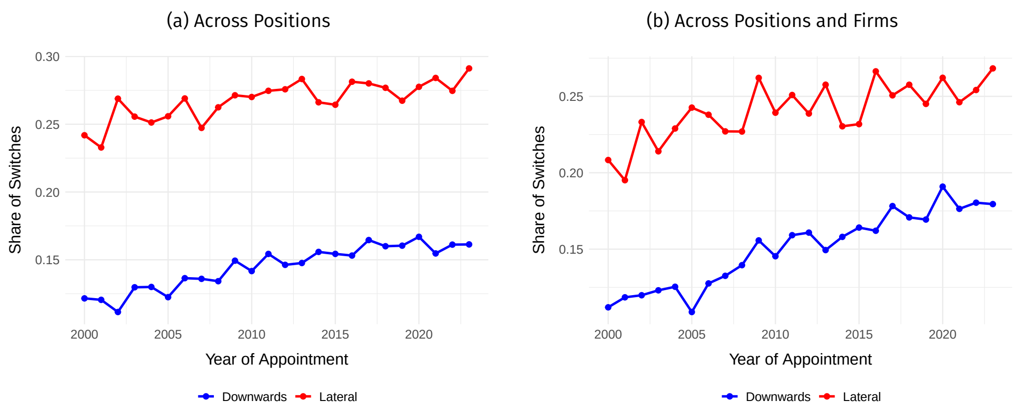

Notes: The plot shows the share of downward/lateral switches of CEOs prior to appointment by year of appointment, calculated as the ratio of the number of downward or lateral switches to the number of total switches. Panel (a) includes all switches across positions, Panel (b) focuses on those across positions and firms. Switches are identified from a panel constructed from employment biographies from Revelio Labs, prioritizing more senior positions and those that started earlier in the case of overlapping spells.

Figure 1.B.24. CEO Appointment of Former Coworker and Job Mobility—Event Study

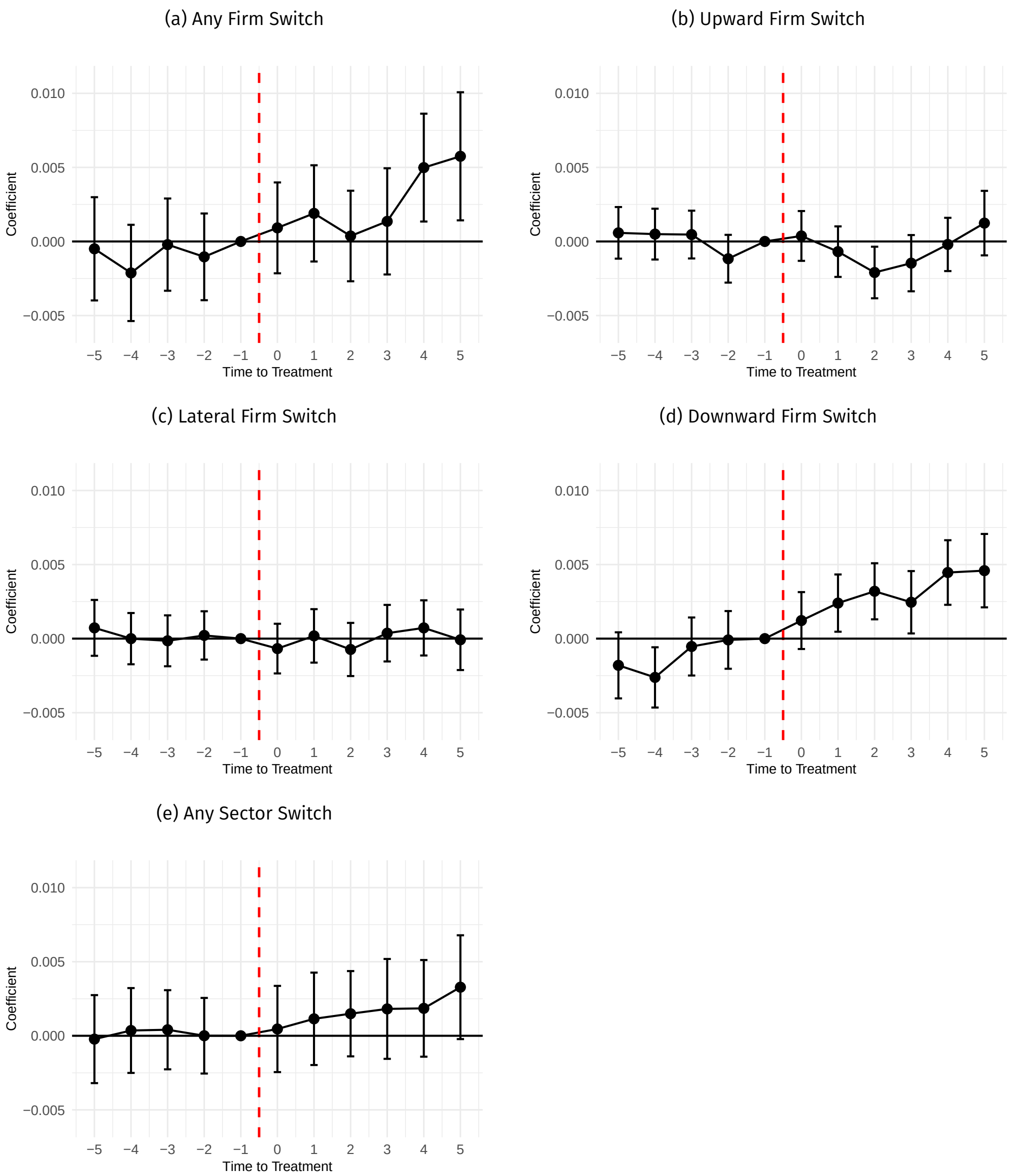

Notes: The figure shows event-study estimates, obtained by estimating the specification in Column (4) from Table 1.5 (Panel a), Table 1.C.14 (Panel b), Table 1.C.15 (Panel c), Table 1.6 (Panel d), and Table 1.C.13 (Panel e), respectively, replacing *Post* by indicators for the number of years until treatment. The endpoints are binned at $t = -6$ and $t = 6$, while $t = -1$ is omitted as the reference category. Point estimates are shown alongside 95% confidence intervals, which are obtained from standard errors clustered at the firm level.

Figure 1.B.25. Heterogeneity for Downward Mobility Response by Quartile of Shared Tenure

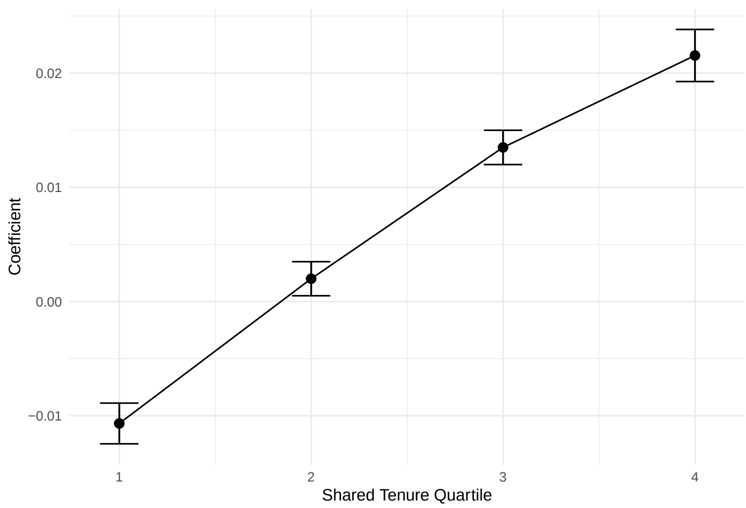

Notes: The specification equals Column (4) from Table 1.C.19 and shows the point estimates for the coefficients on $Treat \times Post \times Quartile = k \ \forall \ k = 1, 2, 3, 4$ alongside 95% confidence intervals, which are obtained from standard errors clustered at the firm level.

Figure 1.B.26. Heterogeneity for Downward Mobility Response by Size Quartile of Current Firm

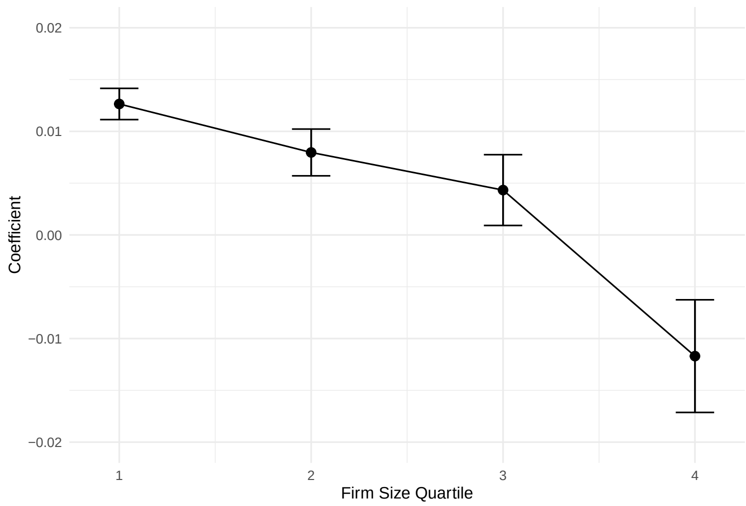

Notes: The specification equals Column (4) from Table 1.C.25 and shows the point estimates for the coefficients on $Treat \times Post \times Quartile = k \ \forall \ k = 1, 2, 3, 4$ alongside 95% confidence intervals, which are obtained from standard errors clustered at the firm level.

Figure 1.B.27. Time-Varying Correlation of Firm Size and Seniority Differences from CEO Employment Histories

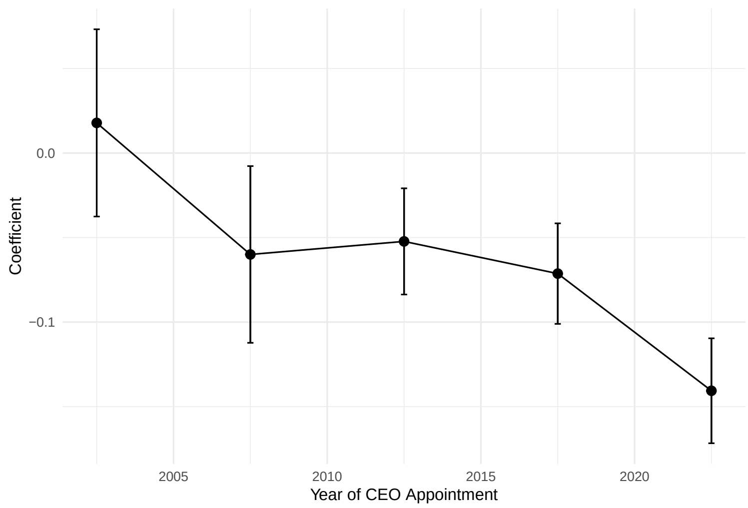

Notes: The figure plots the time-varying correlation of differences in firm employment and differences in seniority levels for job switches, using data from CEO employment histories. The coefficients are obtained from regressions of first differences of the natural logarithm of firm employment on 5-year bins for year of appointment, both alone and interacted with first differences in seniority. Point estimates are plotted alongside 95% confidence intervals, which are obtained from standard errors clustered at the CEO level.

Appendix 1.C Additional Tables

Table 1.C.1. Summary Statistics—Descriptives and Demand-Side Results

	Mean	SD	P25	P50	P75	N
Panel A: Descriptives						
CEO Age at Appointment	51.89	9.11	46.00	52.00	58.00	50,510
Internal Appointment	0.20	0.40	0.00	0.00	0.00	50,510
Internal Appointment, Firm Entry Manager	0.10	0.30	0.00	0.00	0.00	50,510
Years of External Experience	21.12	10.76	14.00	21.00	29.00	50,510
Years of Internal Experience	1.67	5.15	0.00	0.00	0.00	50,510
Number External Positions	4.86	3.01	3.00	5.00	7.00	50,510
Number Internal Positions	0.41	1.11	0.00	0.00	0.00	50,510
Number of Firms	3.91	2.21	2.00	4.00	5.00	50,510
Number of NAICS-4 Sectors	1.47	1.20	1.00	1.00	2.00	50,510
Panel B: Demand Side, Baseline						
CEO Age	53.18	8.58	48.00	53.00	59.00	15,679
Age Firm Entry	47.25	10.92	41.00	48.00	55.00	15,679
Internal Experience	5.93	9.10	0.00	1.00	8.00	15,679
Travel Time to Closest Office	1.17	0.40	0.94	1.13	1.25	15,679
Employment, Revelio Labs	4,212.22	20,782.49	44.17	266.82	1,649.44	15,679
Firm Age, Revelio Labs	27.82	17.83	12.00	26.00	42.00	15,679
Panel C: Demand Side, External Validity						
No. VP Roles (50 Levels)	7.08	7.12	2.00	5.00	10.00	11,033
No. VP Roles (150 Levels)	9.64	12.42	2.00	5.00	12.00	11,033
No. VP Roles (300 Levels)	11.01	16.68	2.00	5.00	13.00	11,033
No. VP Roles (500 Levels)	12.09	20.99	2.00	6.00	14.00	11,033
No. VP Roles (1,000 Levels)	13.85	29.74	2.00	6.00	15.00	11,033
No. VPs	43.77	573.95	2.00	6.00	19.00	11,033
Spatial HHI	0.58	0.34	0.27	0.54	1.00	11,279
Spatial HHI (Office = 10 Individuals)	0.61	0.35	0.29	0.58	1.00	9,907
Share Employees Non-HQ	0.40	0.34	0.00	0.37	0.69	11,279
No. Offices	11.30	24.67	1.00	2.00	8.00	14,609
Export-Related Economic Complexity Index	-0.12	1.15	-0.85	0.06	0.71	6,070
Import-Related Economic Complexity Index	-0.12	1.05	-0.41	0.14	0.58	6,070
No. Export Countries	96.14	16.20	98.00	102.00	104.00	6,070
No. Import Countries	77.64	18.27	72.00	82.00	90.00	6,070
Panel D: MBB Experience vs. Age						
MBB Inflows	0.01	0.05	0.00	0.00	0.00	18,115
MBB Experience	0.05	0.23	0.00	0.00	0.00	18,115

Notes: The table shows summary statistics for the results in Sections 1.3 and 1.5.1.

Table 1.C.2. CEO Age and Industry-Level Market Concentration

	CEO Age					
	(1)	(2)	(3)	(4)	(5)	(6)
4-firm concentration ratio	-1.158					
	(0.888)					
8-firm concentration ratio		-1.058				
		(0.865)				
20-firm concentration ratio			-1.096			
			(0.920)			
HHI-Floor Sales				-7.883		
				(4.952)		
HHI-Floor Employment					-4.500	
					(6.379)	
HHI-Floor Payroll						-4.605
						(5.924)
R^2	0.119	0.119	0.119	0.119	0.131	0.131
Observations	45,370	45,370	45,370	45,326	33,092	33,092
Year fixed effects	✓	✓	✓	✓	✓	✓
Naics4 fixed effects	✓	✓	✓	✓	✓	✓

Notes: The table shows results from regressing CEO age at appointment on different measures of market concentration at the NAICS4-level, using data from the most recent wave of the Economic Census (1997, 2002, 2007, 2012, and 2017). Data from the Economic Census have been obtained from Keil (2017) for years up to 2012 and via the US Census Bureau for 2017. In Columns (1), (2), and (3), the concentration measures are the combined percentage of the total industry sales accounted for by the largest 4, 8, and 20 firms in the industry, respectively. Columns (4), (5), and (6) use the HHI-floor based on sales, employment, and payroll, respectively. The HHI-floor is a lower-bound estimate of the Herfindahl–Hirschman Index (HHI), constructed from concentration ratios, as described in Keil (2017). All regressions control for NAICS-4 and year fixed effects as well as the natural logarithm of employment and firm age, both obtained from the employment history data. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.3. Share of Appointments by Firm Size and Listing Status

Panel A: By Firm Size Quintile					
	Q5	Q4	Q3	Q2	Q1
External	0.60	0.71	0.79	0.83	0.87
Internal	0.40	0.29	0.21	0.17	0.13
Manager	0.19	0.14	0.11	0.09	0.07
Non-Manager	0.21	0.15	0.10	0.08	0.06
Panel B: By Listing Status					
	S&P 500	S&P 400	S&P 600	Other Listed	Private
External	0.44	0.49	0.55	0.71	0.85
Internal	0.56	0.51	0.45	0.29	0.15
Manager	0.30	0.30	0.29	0.19	0.06
Non-Manager	0.26	0.21	0.16	0.10	0.09

Notes: The table includes the share of appointments by seniority at firm entry against firm size quintile by employment (Panel a) and listing status (Panel b).

Table 1.C.4. CEO Age and College Selectivity

	CEO Age					
	(1)	(2)	(3)	(4)	(5)	(6)
College Selectivity	-2.200***	-2.406***	-1.890***	-2.020***	-9.904***	-10.812***
	(0.212)	(0.231)	(0.157)	(0.175)	(0.722)	(0.825)
R^2	0.252	0.347	0.253	0.348	0.265	0.369
Observations	16,182	15,137	16,182	15,137	13,107	12,113
Controls	✓	✓	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓	✓	✓
MSA-Year fixed effects		✓		✓		✓
Selectivity Measure	Ivy Plus		Highly Selective		Ln(SAT 2001)	

Notes: The table regresses CEO age on different measures of school selectivity. We focus on a CEO's first postsecondary degree, identified either by the stated degree name or, if missing, by the first degree recorded in BoardEx. We exclude first postsecondary degrees obtained from institutions outside of the US. *Ivy Plus* is an indicator taken from Chetty et al. (2020), comprising Brown, Columbia, Cornell, Dartmouth, Duke, Harvard, MIT, Princeton, Stanford, Chicago, UPenn, and Yale. *Selective* is an indicator for institutions classified as highly selective in Barron's Selectivity Index (in 1982, 1992, and 2002 in order to avoid selection issues over time). *Ln(SAT)* is the natural logarithm of the institution's average SAT admission score in 2001, taken from Chetty et al. (2020). Controls include the natural logarithm of employment and firm age. Standard errors are clustered at the firm level.

Table 1.C.5. Variation in MBB Availability and Employee Inflows from MBB

	Inflows from MBB			
	(1)	(2)	(3)	(4)
Δ(Travel Time)	-0.071***	-0.067**	-0.072***	-0.076**
	(0.026)	(0.026)	(0.027)	(0.034)
Lag(Inflows from MBB)		0.281***		
		(0.041)		
R^2	0.472	0.523	0.479	0.501
Observations	707,068	682,272	707,068	707,068
Outcome Mean	0.370	0.377	0.399	0.498
Naics4-Year fixed effects	✓	✓	✓	✓
Firm fixed effects	✓	✓	✓	✓
Time Between Spells	1 Month	1 Month	3 Months	1 Year

Notes: The dependent variable is the natural logarithm of one plus the number of employees who transition from an MBB firm to the firm of interest (multiplied by 100), constructed from the employment history data. In Columns (1) and (2), we restrict the sample to job-to-job transitions with at most one month between spells; Column (3) relaxes this window to three months, and Column (4) to one year. The main independent variable is the first difference in flight time (in hours) to the nearest MBB office, constructed using MBB office openings combined with flight data from the Bureau of Transportation Statistics. The sample is restricted to MSAs without MBB presence. All regressions control for firm size by including the natural logarithm of employment and firm age (extracted from the employment history data). Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.6. MBB Availability, Number of Distinct VPs, and CEO Age

	CEO Age				
	(1)	(2)	(3)	(4)	(5)
Travel Time × No. VP Roles (50 Levels)	1.748*				
	(0.905)				
Travel Time × No. VP Roles (150 Levels)		3.281**			
		(1.262)			
Travel Time × No. VP Roles (300 Levels)			5.057***		
			(1.600)		
Travel Time × No. VP Roles (500 Levels)				4.683***	
				(1.549)	
Travel Time × No. VP Roles (1,000 Levels)					3.832**
					(1.744)
Travel Time × No. VPs	-1.659***	-2.987***	-4.494***	-4.281***	-3.743**
	(0.610)	(0.984)	(1.377)	(1.401)	(1.623)
R^2	0.364	0.364	0.365	0.365	0.364
Observations	14,933	14,933	14,933	14,933	14,933
MSA-Year fixed effects	✓	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓

Notes: The dependent variable in all regressions is CEO age upon appointment, obtained from BoardEx. *Travel Time* denotes the flight time (in hours) to the closest office, which we construct from MBB office openings combined with flight data from the Bureau of Transportation Statistics. The variable *No. VP Roles* is defined as the number of distinct roles in which a firm employs vice presidents. *No. VPs* is a count variable measuring the number of vice presidents a firm employs. We aggregate these measures to the year-specific median within NAICS-4 industries and use their natural logarithm. We identify all vice presidents (VPs) as individuals with job titles including “vp” or “vice president” (case-insensitive), excluding any titles that contain “avp” or “assistant”. The job taxonomy assigns activities to job titles by matching descriptions from resumes and online profiles to responsibilities in job postings. It initially classifies 1,500 distinct roles, which are then hierarchically clustered. We utilize this taxonomy at various levels of role granularity (50, 150, 300, 500, and 1,000 roles). The sample is restricted to those MSAs without MBB presence. All regressions control for firm size by including the natural logarithm of employment and firm age, both extracted from the employment history data. Standard errors are double-clustered at the NAICS-4 industry and MSA level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.7. MBB Availability, Spatial Diversification, and CEO Age

	CEO Age			
	(1)	(2)	(3)	(4)
Travel Time × HHI, Bottom Decile	4.244***			
	(1.262)			
Travel Time × HHI, Bottom Quartile		2.409***		
		(0.723)		
Travel Time × HHI, Bottom Decile (Office = 10 Individuals)			3.784**	
			(1.478)	
Travel Time × Share Employees Non-HQ, Top Decile				2.967**
				(1.269)
Travel Time × No. Offices	-0.568	-0.856**	-0.374	-0.425
	(0.375)	(0.367)	(0.412)	(0.345)
R^2	0.365	0.365	0.368	0.364
Observations	14,195	14,195	12,527	14,185
MSA-Year fixed effects	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓
Controls	✓	✓	✓	✓

Notes: The dependent variable in all regressions is CEO age upon appointment, obtained from BoardEx. *Travel Time* denotes the flight time (in hours) to the closest office, which we construct from MBB office openings combined with flight data from the Bureau of Transportation Statistics. Spatial diversification is measured as the distribution of employees across offices in the LinkedIn data, which we define as having at least 5 employees in an MSA-year observation. We then construct the following indicators: *HHI, Bottom Decile* and *HHI, Bottom Quartile* for firms in the lowest decile and quartile, respectively; *HHI, Bottom Decile (Office = 10 Individuals)* for firms in the lowest decile, but defining an office as having at least 10 employees in an MSA-year observation; *Share Employees Non-HQ, Top Decile* for firms in the top decile of the share of employees outside of the headquarters. All measures are then aggregated to the NAICS-4-year level by assigning a value of one whenever at least half of the underlying indicators equal one. *No. Offices* denotes the natural logarithm of the median number of offices of each firm within a year and NAICS-4 industry. The sample is restricted to those MSAs without MBB presence. All regressions control for firm size by including the natural logarithm of employment and firm age, both extracted from the employment history data. Standard errors are double-clustered at the NAICS-4 industry and MSA level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.8. MBB Availability, Export-Related Economic Complexity (EECI), and CEO Age

	CEO Age								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Travel Time × ECI, Above Median	2.327***	3.374***					5.648**	17.730***	2.910***
	(0.682)	(0.544)					(2.237)	(3.287)	(0.769)
Travel Time × ECI			1.294**						
			(0.518)						
Travel Time × ECI, Upper Quartile				2.376**					
				(1.045)					
Travel Time × ECI (1999 Weights)					1.922***				
					(0.698)				
Travel Time × ECI, Above Median (1999 Weights)						2.780***			
						(0.908)			
Travel Time × No. Countries Export ≥ 1%									-4.328***
									(1.490)
Travel Time × No. Countries Export		-9.049**	-8.759**	-6.873**	-6.262**	-5.460**	-8.078**	-16.638***	
		(3.487)	(3.955)	(3.332)	(2.694)	(2.620)	(3.642)	(5.858)	
R^2	0.406	0.408	0.407	0.407	0.408	0.407	0.407	0.388	0.408
Observations	5,136	5,136	5,136	5,136	5,136	5,136	5,136	2,235	5,136
MSA-Year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Measure	SITC	SITC	SITC	SITC	SITC	SITC	HS 92	HS 2012	SITC

Notes: The dependent variable in all regressions is CEO age upon appointment, obtained from BoardEx. *Travel Time* denotes the flight time (in hours) to the closest office, which we construct from MBB office openings combined with flight data from the Bureau of Transportation Statistics. The complexity measures are based on the country-level economic complexity index (ECI) from Hidalgo and Hausmann (2009), weighted by the corresponding exports from the United States and the respective country. The complexity data have been obtained from the Atlas of Economic Complexity. Annual data on exports (FAS value) at the NAICS-4 level have been obtained from the United States International Trade Commission. In Column (1), travel time is interacted with an indicator for above-median values of economic complexity. Column (2) additionally controls for the natural logarithm of the number of countries the US exports to. Column (3) uses the EECI in linear terms, standardized to mean zero and a standard deviation of one. Column (4) uses an indicator for the upper quartile of the EECI. Columns (5) and (6) use weights based on exports in 1999. Columns (7) and (8) rely on the EECI derived from HS 92 and HS 12 product codes, respectively. Column (9) controls for the natural logarithm of all countries to which a NAICS-4 industry exports at least 1% of value in a given year. The sample is restricted to those MSAs without MBB presence. All regressions control for firm size by including the natural logarithm of employment and firm age, both extracted from the employment history data. Standard errors are double-clustered at the NAICS-4 industry and MSA level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.9. MBB Availability, Import-Related Economic Complexity (IECI), and CEO Age

	CEO Age								
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Travel Time × ECI, Above Median	0.932	1.101					-0.373	1.753	0.696
	(0.838)	(0.701)					(1.254)	(2.524)	(0.767)
Travel Time × ECI			-0.100						
			(0.533)						
Travel Time × ECI, Upper Quartile				-0.086					
				(0.881)					
Travel Time × ECI (1999 Weights)					0.087				
					(0.531)				
Travel Time × ECI, Above Median (1999 Weights)						0.331			
						(0.617)			
Travel Time × No. Countries Import ≥ 1%									1.581
									(1.617)
Travel Time × No. Countries Import		-2.355	-1.855	-2.090	-2.307	-2.207	-1.762	-9.675***	
		(2.459)	(2.350)	(2.435)	(2.285)	(2.432)	(2.316)	(3.212)	
R^2	0.405	0.406	0.406	0.406	0.406	0.406	0.406	0.384	0.406
Observations	5,136	5,136	5,136	5,136	5,136	5,136	5,136	2,235	5,136
MSA-Year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓	✓
Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Measure	SITC	SITC	SITC	SITC	SITC	SITC	HS 92	HS 2012	SITC

Notes: The dependent variable in all regressions is CEO age upon appointment, obtained from BoardEx. *Travel Time* denotes the flight time (in hours) to the closest office, which we construct from MBB office openings combined with flight data from the Bureau of Transportation Statistics. The complexity measures are based on the country-level economic complexity index (ECI) from Hidalgo and Hausmann (2009), weighted by the corresponding imports of the United States from the respective country. The complexity data have been obtained from the Atlas of Economic Complexity. Annual data on imports (general customs value) at the NAICS-4 level have been obtained from the United States International Trade Commission. In Column (1), travel time is interacted with an indicator for above-median values of economic complexity. Column (2) additionally controls for the natural logarithm of the number of countries the US imports from. Column (3) uses the IECI in linear terms, standardized to mean zero and a standard deviation of one. Column (4) uses an indicator for the upper quartile of the IECI. Columns (5) and (6) use weights based on imports in 1999. Columns (7) and (8) rely on the IECI derived from HS 92 and HS 12 product codes, respectively. Column (9) controls for the natural logarithm of all countries from which a NAICS-4 industry imports at least 1% of value in a given year. The sample is restricted to those MSAs without MBB presence. All regressions control for firm size by including the natural logarithm of employment and firm age, both extracted from the employment history data. Standard errors are double-clustered at the NAICS-4 industry and MSA level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.10. Substitutability of MBB Experience and Age

	CEO Age					
	(1)	(2)	(3)	(4)	(5)	(6)
MBB Inflows	9.744***	8.332***	7.509***	7.741***	7.882***	3.315
	(2.055)	(1.910)	(2.073)	(1.975)	(2.359)	(5.131)
MBB Experience	-3.040***	-3.529***	-3.190***	-3.447***	-3.367***	-2.494***
	(0.304)	(0.297)	(0.304)	(0.344)	(0.375)	(0.573)
MBB Inflows × MBB Experience	-9.447***	-7.376**	-9.729***	-8.789**	-9.490**	-17.73*
	(3.337)	(3.102)	(3.373)	(3.540)	(4.093)	(9.455)
R^2	0.086	0.121	0.177	0.296	0.384	0.753
Observations	18,115	18,115	17,844	16,154	15,305	10,454
Controls	✓	✓	✓	✓	✓	✓
Year fixed effects		✓	✓			
Naics4 fixed effects			✓			
Naics4-Year fixed effects				✓	✓	✓
MSA-Year fixed effects					✓	✓
Firm fixed effects						✓

Notes: The dependent variable in all regressions is CEO age at appointment among CEOs. *MBB Experience* is a dummy equal to one if the individual had ever worked in an MBB firm. *MBB Inflows* is defined as the company-level sum of new employees with MBB experience during the previous 10 years, divided by the total number of new employees in a company summed over the previous 10 years. Controls are the natural logarithm of employment and firm age, both obtained from the employment biographies. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.11. Substitutability of MBB Experience and Age—Robustness

	CEO Age			
	(1)	(2)	(3)	(4)
MBB Inflows	1.183***	7.390***	7.222***	13.87***
	(0.275)	(1.901)	(2.531)	(3.478)
MBB Experience	-3.122***	-3.354***	-3.313***	-3.377***
	(0.426)	(0.382)	(0.375)	(0.635)
MBB Inflows × MBB Experience	-1.512**	-7.630**	-10.36**	-79.99**
	(0.734)	(3.146)	(4.051)	(36.49)
R^2	0.384	0.384	0.384	0.460
Observations	15,305	15,305	15,278	5,897
Controls	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓
MSA-Year fixed effects	✓	✓	✓	✓
Change Compared to Baseline	Discretized	Seniority ≥ 4	5 Year	Controls

Notes: The dependent variable in all regressions is CEO age at appointment among CEOs. *MBB Experience* is a dummy equal to one if the individual had ever worked in an MBB firm. *MBB Inflows* is defined as the company-level sum of new employees with MBB experience during the previous 10 years, divided by the total number of new employees in a company summed over the previous 10 years. Controls are the natural logarithm of employment and firm age, both obtained from the employment biographies. We deviate from this setup in the following ways: Column (1) discretizes MBB inflows, assigning a value of 1 to all values in the highest quartile (conditional on being strictly positive); Column (2) normalizes the MBB inflows by dividing by inflows at seniority level 4 or above; Column (3) restricts inflows to the past 5 years prior to appointment; Column (4) replaces the control variables by the natural logarithm of employment, revenue, and firm age, obtained from Compustat. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.12. Summary Statistics—Supply-Side Regressions

	Mean	SD	P25	P50	P75	N
Panel A: Treatment Group						
Male	0.64	0.48	0.00	1.00	1.00	320,522
White	0.78	0.41	1.00	1.00	1.00	321,699
Age	40.41	8.44	35.00	40.00	45.00	304,686
Bachelor	0.44	0.50	0.00	0.00	1.00	264,045
More than Bachelor	0.53	0.50	0.00	1.00	1.00	264,045
Salary (in 1,000 USD)	126.05	62.35	85.02	111.97	151.40	321,762
Seniority	4.47	1.26	4.00	5.00	5.00	321,781
Tenure at Firm	8.29	6.94	3.00	6.00	11.00	321,776
External Experience	10.11	8.80	2.00	9.00	16.00	321,700
Panel B: Control Group						
Male	0.65	0.48	0.00	1.00	1.00	1,058,186
White	0.81	0.39	1.00	1.00	1.00	1,061,239
Age	40.15	8.67	34.00	39.00	45.00	1,012,153
Bachelor	0.48	0.50	0.00	0.00	1.00	787,624
More than Bachelor	0.48	0.50	0.00	0.00	1.00	787,624
Salary (in 1,000 USD)	114.00	59.60	75.67	100.31	136.18	1,061,264
Seniority	4.33	1.34	4.00	5.00	5.00	1,061,494
Tenure at Firm	8.08	6.87	3.00	6.00	11.00	1,061,478
External Experience	10.06	8.91	2.00	9.00	16.00	1,061,267

Notes: The table shows summary statistics for the sample used to study network effects (Section 1.5.2), separately for the treatment (Panel A) and control (Panel B) group, and restricted to observations in the year prior to treatment. We weight the treated (control) individuals by the share of the control (treatment) group within each firm-seniority-year cell.

Table 1.C.13. CEO Appointment of Former Coworker and Job Mobility (Industry Switch)

	Industry Switch				
	(1)	(2)	(3)	(4)	(5)
Treat	0.002*				
	(0.001)				
Post	-0.004***	-0.007***	-0.008***		
	(0.001)	(0.001)	(0.001)		
Treat × Post	-0.001	0.001	0.001	0.003**	0.002**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.006	0.161	0.163	0.186	0.215
Observations	14,144,297	14,144,297	14,144,297	14,144,297	12,898,708
Pre-Treatment Control Mean	0.042	0.042	0.042	0.042	0.041
Year fixed effects	✓	✓	✓		
Individual fixed effects		✓	✓	✓	✓
Firm-Seniority fixed effects			✓		
Firm-Seniority-Year fixed effects				✓	✓
MSA-Naics4-Year fixed effects					✓

Notes: The dependent variable in all regressions is a dummy equal to one in the last year before an individual switches across NAICS-4 industries. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.14. CEO Appointment of Former Coworker and Job Mobility (Upward Firm Switch)

	Firm and Upward Switch				
	(1)	(2)	(3)	(4)	(5)
Treat	0.001**				
	(0.001)				
Post	0.002***	0.000	-0.001**		
	(0.000)	(0.000)	(0.000)		
Treat × Post	-0.000	-0.000	-0.000	-0.001	-0.001
	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)
R^2	0.004	0.090	0.091	0.105	0.134
Observations	22,777,381	22,777,381	22,777,381	22,777,381	16,139,755
Pre-Treatment Control Mean	0.025	0.025	0.025	0.025	0.025
Year fixed effects	✓	✓	✓		
Individual fixed effects		✓	✓	✓	✓
Firm-Seniority fixed effects			✓		
Firm-Seniority-Year fixed effects				✓	✓
MSA-Naics4-Year fixed effects					✓

Notes: The dependent variable in all regressions is a dummy equal to one in the last year before an individual switches firms and rises to a higher seniority level. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.15. CEO Appointment of Former Coworker and Job Mobility (Lateral Firm Switch)

	Firm and Lateral Switch				
	(1)	(2)	(3)	(4)	(5)
Treat	0.002***				
	(0.001)				
Post	-0.002***	-0.004***	-0.005***		
	(0.000)	(0.000)	(0.000)		
Treat × Post	-0.001*	-0.002***	-0.002***	0.000	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.004	0.110	0.111	0.129	0.160
Observations	22,777,381	22,777,381	22,777,381	22,777,381	16,139,755
Pre-Treatment Control Mean	0.029	0.029	0.029	0.029	0.029
Year fixed effects	✓	✓	✓		
Individual fixed effects		✓	✓	✓	✓
Firm-Seniority fixed effects			✓		
Firm-Seniority-Year fixed effects				✓	✓
MSA-Naics4-Year fixed effects					✓

Notes: The dependent variable in all regressions is a dummy equal to one in the last year before an individual switches firms but remains at the same seniority level. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.16. CEO Appointment of Former Coworker and Job Mobility (Across Firms and Job Categories)

	Switch Across Firms and Job Categories				
	(1)	(2)	(3)	(4)	(5)
Treat	-0.002***				
	(0.001)				
Post	-0.003***	-0.007***	-0.008***		
	(0.001)	(0.001)	(0.001)		
Treat × Post	0.001**	0.002***	0.002***	0.003***	0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.003	0.127	0.127	0.144	0.176
Observations	22,777,381	22,777,381	22,777,381	22,777,381	16,139,755
Pre-Treatment Control Mean	0.034	0.034	0.034	0.034	0.033
Year fixed effects	✓	✓	✓		
Individual fixed effects		✓	✓	✓	✓
Firm-Seniority fixed effects			✓		
Firm-Seniority-Year fixed effects				✓	✓
MSA-Naics4-Year fixed effects					✓

Notes: The dependent variable in all regressions is a dummy equal to one in the last year before an individual switches firms and job category. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.17. CEO Appointment of Former Coworker and Job Mobility (Across Firms, Within Job Categories)

	Switch Across Firms Within Job Categories				
	(1)	(2)	(3)	(4)	(5)
Treat	0.000				
	(0.001)				
Post	-0.007***	-0.010***	-0.011***		
	(0.001)	(0.001)	(0.001)		
Treat × Post	0.001	0.000	0.000	0.002***	0.001*
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.009	0.139	0.140	0.160	0.191
Observations	22,777,381	22,777,381	22,777,381	22,777,381	16,139,755
Pre-Treatment Control Mean	0.060	0.060	0.060	0.060	0.061
Year fixed effects	✓	✓	✓		
Individual fixed effects		✓	✓	✓	✓
Firm-Seniority fixed effects			✓		
Firm-Seniority-Year fixed effects				✓	✓
MSA-Naics4-Year fixed effects					✓

Notes: The dependent variable in all regressions is a dummy equal to one in the last year before an individual switches firms but stays in the same job category. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.18. CEO Appointment of Former Coworker and Job Mobility—Placebo Test

	Any Firm	Upward	Lateral	Downward	Any Sector
	(1)	(2)	(3)	(4)	(5)
Treat × Post	-0.0019	-0.0005	-0.0012	-0.0002	-0.0036
	(0.0022)	(0.0012)	(0.0011)	(0.0013)	(0.0026)
R^2	0.21794	0.15915	0.18694	0.15936	0.25047
Observations	2,130,135	2,130,135	2,130,135	2,130,135	1,009,521
Pre-Treatment Control Mean	0.1462	0.0518	0.0467	0.0478	0.0721
Individual fixed effects	✓	✓	✓	✓	✓
Firm-Seniority-Year fixed effects	✓	✓	✓	✓	✓

Notes: The specification equals Column (4) from the baseline tables. The sample includes individuals who left the initial firm between one and five years before the prospective CEO was hired at that firm. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.19. Heterogeneity by Length of Shared Tenure

	Any Firm (1)	Upward (2)	Lateral (3)	Downward (4)	Any Sector (5)
Treat × Post	-0.003**	-0.002***	-0.003***	0.002**	-0.001
	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)
Treat × Post × High Overlap	0.037***	0.006***	0.013***	0.019***	0.016***
	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.166	0.105	0.129	0.111	0.186
Observations	22,777,381	22,777,381	22,777,381	22,777,381	14,144,297
Pre-Treatment Control Mean	0.094	0.025	0.029	0.041	0.042
Individual fixed effects	✓	✓	✓	✓	✓
Firm-Seniority-Year fixed effects	✓	✓	✓	✓	✓

Notes: The specification equals Column (4) from the baseline tables. *High Overlap* is a dummy equal to 1 if the length of the overlapping period of the individual and the future CEO working in the same company is in the highest quartile, and zero otherwise. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.20. Effects on Individuals at Seniority Level below Future CEO

	Any Firm (1)	Upward (2)	Lateral (3)	Downward (4)	Any Sector (5)
Treat × Post	0.002*	-0.001***	-0.000	0.003***	0.000
	(0.001)	(0.000)	(0.000)	(0.000)	(0.001)
Treat × Post × Two Levels Below	0.000	-0.000	0.000	-0.000	0.000
	(0.001)	(0.000)	(0.000)	(0.001)	(0.001)
R^2	0.162	0.103	0.127	0.106	0.181
Observations	53,842,096	53,842,096	53,842,096	53,842,096	33,975,831
Pre-Treatment Control Mean	0.102	0.033	0.030	0.038	0.045
Individual fixed effects	✓	✓	✓	✓	✓
Firm-Seniority-Year fixed effects	✓	✓	✓	✓	✓

Notes: The specification equals Column (4) from the baseline tables, with the sample restricted to those working at the two seniority levels directly below the future CEO. For the workers two levels below the future CEO, we focus on a random 50% subsample of all individuals to maintain computational feasibility. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.21. Effects on Individuals at Seniority Level above Future CEO

	Any Firm (1)	Upward (2)	Lateral (3)	Downward (4)	Any Sector (5)
Treat × Post	0.005**	-0.001	0.001	0.004***	0.003*
	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)
Treat × Post × Two Levels Above	0.001	0.003***	-0.000	-0.002	0.002
	(0.002)	(0.001)	(0.001)	(0.002)	(0.002)
R^2	0.170	0.112	0.135	0.120	0.189
Observations	13,119,386	13,119,386	13,119,386	13,119,386	8,465,394
Pre-Treatment Control Mean	0.091	0.019	0.029	0.042	0.041
Individual fixed effects	✓	✓	✓	✓	✓
Firm-Seniority-Year fixed effects	✓	✓	✓	✓	✓

Notes: The specification equals Column (4) from the baseline tables, with the sample restricted to those working at the two seniority levels directly above the future CEO. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.22. Time Variation in Effect of CEO Appointment of Former Coworker and Job Mobility (5-Year Bins)

	Any Firm (1)	Upward (2)	Lateral (3)	Downward (4)	Any Sector (5)
Treat × Post	-0.003*	-0.001	-0.002**	-0.000	-0.000
	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)
Treat × Post × 1990–1994	0.009**	-0.000	0.005**	0.004	0.008**
	(0.004)	(0.002)	(0.003)	(0.003)	(0.004)
Treat × Post × 1995–1999	0.003	0.001	0.003	-0.001	0.002
	(0.004)	(0.002)	(0.002)	(0.003)	(0.003)
Treat × Post × 2005–2009	0.006**	-0.000	0.002*	0.004**	0.003
	(0.003)	(0.001)	(0.001)	(0.001)	(0.002)
Treat × Post × 2010–2014	0.015***	0.002*	0.004***	0.009***	0.008**
	(0.003)	(0.001)	(0.001)	(0.002)	(0.003)
R^2	0.147	0.092	0.114	0.093	0.164
Observations	13,829,579	13,829,579	13,829,579	13,829,579	8,685,741
Pre-Treatment Control Mean	0.079	0.021	0.024	0.035	0.035
Individual fixed effects	✓	✓	✓	✓	✓
Firm-Seniority-Year fixed effects	✓	✓	✓	✓	✓

Notes: The specification equals Column (4) from the baseline tables. The indicators for the 5-year bins refer to the year of the CEO appointment. The sample is restricted to CEO appointments between 1990 and 2014. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The baseline period includes the years 2000 to 2004. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.23. Heterogeneity by Mobility of CEO

	Any Firm (1)	Upward (2)	Lateral (3)	Downward (4)	Any Sector (5)
Treat × Post	0.003***	-0.001**	-0.000	0.005***	0.002*
	(0.001)	(0.000)	(0.001)	(0.001)	(0.001)
Treat × Post × High Mobility CEO	0.004*	0.001	0.001	0.002*	0.002
	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)
R^2	0.166	0.105	0.129	0.111	0.186
Observations	22,777,381	22,777,381	22,777,381	22,777,381	14,144,297
Pre-Treatment Control Mean	0.094	0.025	0.029	0.041	0.042
Individual fixed effects	✓	✓	✓	✓	✓
Firm-Seniority-Year fixed effects	✓	✓	✓	✓	✓

Notes: The specification equals Column (4) from the baseline tables. *High Mobility CEO* is a dummy equal to 1 if the number of positions held by the CEO between leaving the firm and being appointed to CEO is in the highest quartile, and zero otherwise. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.24. Heterogeneity by Direction of CEO Switch

	Any Firm (1)	Upward (2)	Lateral (3)	Downward (4)	Any Sector (5)
Treat × Post	0.005***	0.000	0.000	0.005***	0.003***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Treat × Post × Lateral Switch	-0.002	-0.002*	-0.001	0.001	-0.004**
	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)
Treat × Post × Downward Switch	0.003	-0.001	0.001	0.003*	0.003
	(0.003)	(0.001)	(0.001)	(0.002)	(0.003)
R^2	0.166	0.105	0.129	0.111	0.186
Observations	22,777,381	22,777,381	22,777,381	22,777,381	14,144,297
Pre-Treatment Control Mean	0.094	0.024	0.028	0.041	0.041
Individual fixed effects	✓	✓	✓	✓	✓
Firm-Seniority-Year fixed effects	✓	✓	✓	✓	✓

Notes: The specification equals Column (4) from the baseline tables. “Lateral Switch” (“Downward Switch”) are dummies equal to 1 if the future CEO leaves the company while staying at the same seniority level (switching to a lower seniority level), and zero otherwise. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.25. Heterogeneity by Size Quartile of Current Firm

	Any Firm (1)	Upward (2)	Lateral (3)	Downward (4)	Any Sector (5)
Quartile = 2	0.036***	0.007***	0.010***	0.019***	0.018***
	(0.002)	(0.0006)	(0.0007)	(0.001)	(0.002)
Quartile = 3	0.051***	0.008***	0.011***	0.032***	0.032***
	(0.002)	(0.0009)	(0.0009)	(0.002)	(0.002)
Quartile = 4	0.066***	0.009***	0.014***	0.042***	0.045***
	(0.004)	(0.002)	(0.002)	(0.004)	(0.004)
Treat × Post	0.007***	-0.004***	-0.0010	0.013***	0.010***
	(0.001)	(0.0007)	(0.0007)	(0.0008)	(0.002)
Treat × Post × Quartile = 2	0.007***	0.007***	0.004***	-0.005***	-0.002
	(0.002)	(0.0008)	(0.0010)	(0.001)	(0.002)
Treat × Post × Quartile = 3	0.003	0.008***	0.003***	-0.008***	-0.004
	(0.003)	(0.001)	(0.001)	(0.002)	(0.003)
Treat × Post × Quartile = 4	-0.029***	0.0002	-0.005***	-0.024***	-0.024***
	(0.004)	(0.002)	(0.002)	(0.003)	(0.004)
R^2	0.168	0.105	0.129	0.112	0.187
Observations	22,777,381	22,777,381	22,777,381	22,777,381	14,144,297
Individual fixed effects	✓	✓	✓	✓	✓
Firm-Seniority-Year fixed effects	✓	✓	✓	✓	✓

Notes: The specification equals Column (4) from the baseline tables. The quartiles refer to the firm the individual works for in a given year. Firm size is measured by number of employees. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.26. CEO Appointment of Former Coworker and Job Mobility (Downward Firm Switch) by Firm Size of Current vs. Destination Firm

	Larger Firm (1)	Smaller Firm (2)	Higher Quartile (3)	Same Quartile (4)	Smaller Quartile (5)	Highest Quartile (6)
Treat × Post	0.001***	0.005***	0.001***	0.001***	0.004***	0.000**
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
R^2	0.112	0.101	0.109	0.117	0.096	0.102
Observations	22,745,916	22,745,916	22,745,916	22,745,916	22,745,916	22,745,916
Pre-Treatment Control Mean	0.013	0.029	0.008	0.011	0.022	0.005
Individual fixed effects	✓	✓	✓	✓	✓	✓
Firm-Seniority-Year fixed effects	✓	✓	✓	✓	✓	✓

Notes: The dependent variable in all regressions is a dummy equal to one in the final year before an individual switches firms and moves to a position with lower seniority. We differentiate transitions by the relative size of the destination firm, measured by the number of employees. We distinguish between switches to firms that are larger (Column 1) or smaller (Column 2) than the individual's current firm, as well as moves to a higher (Column 3), the same (Column 4), or a lower (Column 5) firm size quartile. In Column (6), we focus on switches toward firms in the highest size quartile. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 1.C.27. CEO Appointment of Former Coworker and Salary Growth

	Log Salary in First Differences				
	(1)	(2)	(3)	(4)	(5)
Treat	0.002***				
	(0.000)				
Post	0.001**	-0.002***	-0.003***		
	(0.000)	(0.000)	(0.000)		
Treat × Post	-0.001	-0.001	-0.001	-0.002***	-0.002***
	(0.000)	(0.000)	(0.000)	(0.000)	(0.001)
R^2	0.000	0.052	0.052	0.070	0.098
Observations	22,739,139	22,739,139	22,739,139	22,739,139	16,111,774
Pre-Treatment Control Mean	0.016	0.016	0.016	0.016	0.016
Year fixed effects	✓	✓	✓		
Individual fixed effects		✓	✓	✓	✓
Firm-Seniority fixed effects			✓		
Firm-Seniority-Year fixed effects				✓	✓
MSA-Naics4-Year fixed effects					✓

Notes: The dependent variable in all regressions is log salary in first differences. *Treat* is an indicator for having worked alongside a future CEO. *Post* is an indicator for periods following the CEO appointment. The fixed effects and standard errors refer to the position the individual worked in alongside the future CEO. Standard errors are clustered at the firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix 1.D CEO Age and Firm-Level Outcomes

This section summarizes the literature on the link between CEO age and firm-level outcomes, complemented by a number of our own analyses. Summary statistics for all outcome variables are provided in Table 1.D.3.

1.D.1 Correlational Evidence

Motivated by upper echelons theory (Hambrick and Mason 1984), a vast body of research explores the role of executive characteristics in determining company outcomes (Bertrand and Schoar 2003). This literature has established an association between higher CEO age and reduced levels of business dynamism and firm risk. Our analysis revisits this relationship, yielding results in line with previous research. Figure 1.D.1 summarizes the estimated coefficients, obtained by regressing firm-level outcomes on CEO age alongside NAICS-4-year fixed effects and lagged controls.

We find that CEO age is negatively associated with growth rates of revenue, employment, and total assets. A large strand of literature, starting with Child (1974) and Hart and Mellors (1970), has established a negative correlation between CEO age and firm growth: older CEOs are associated with lower investment and sales growth (Acemoglu, Akcigit, and Celik 2022); asset growth (Barba Navaretti, Castellani, and Pieri 2022); profitability (Belenzon, Shamshur, and Zarutskie 2019); firm value, operating performance, and corporate deal-making activity (Cline and Yore 2016); return on assets (Waelchli and Zeller 2013); employment growth (Li, Low, and Makhija 2017); firm valuation (Nguyen, Rahman, and Zhao 2018), and a higher propensity to receive a takeover bid (Jenter and Lewellen 2015). Furthermore, firms that undergo strategic change are characterized by management teams with a lower average age (Wiersema and Bantel 1992; Datta, Rajagopalan, and Zhang 2003). Recently, return-to-office policies have been found to be more prevalent among firms with older CEOs (Flynn, Ghent, and Nair 2025), reflecting a more traditional management approach within this group.

Our analysis also reveals that CEO age is systematically related to innovative activity. This pattern is evident across several dimensions: research inputs, such as R&D expenses (consistent with Serfling 2014, Li, Low, and Makhija 2017 and Barker and Mueller 2002); the extensive margin of research output, measured as changes in a firm's product mix; and the intensive margin of research output as captured by the share of computer-related patents. Notably, overall patenting exhibits no clear association with CEO age.[50] In line with this, Acemoglu, Akcigit, and

50. To document these patterns, we draw on data from the *United States Patent and Trademark Office* provided via *PatentsView*. We merge assignee organizations with company names from BoardEx using the same fuzzy merge algorithm as described above. From these data, we construct an indicator

Figure 1.D.1. Correlations of CEO Age and Firm-Level Outcomes

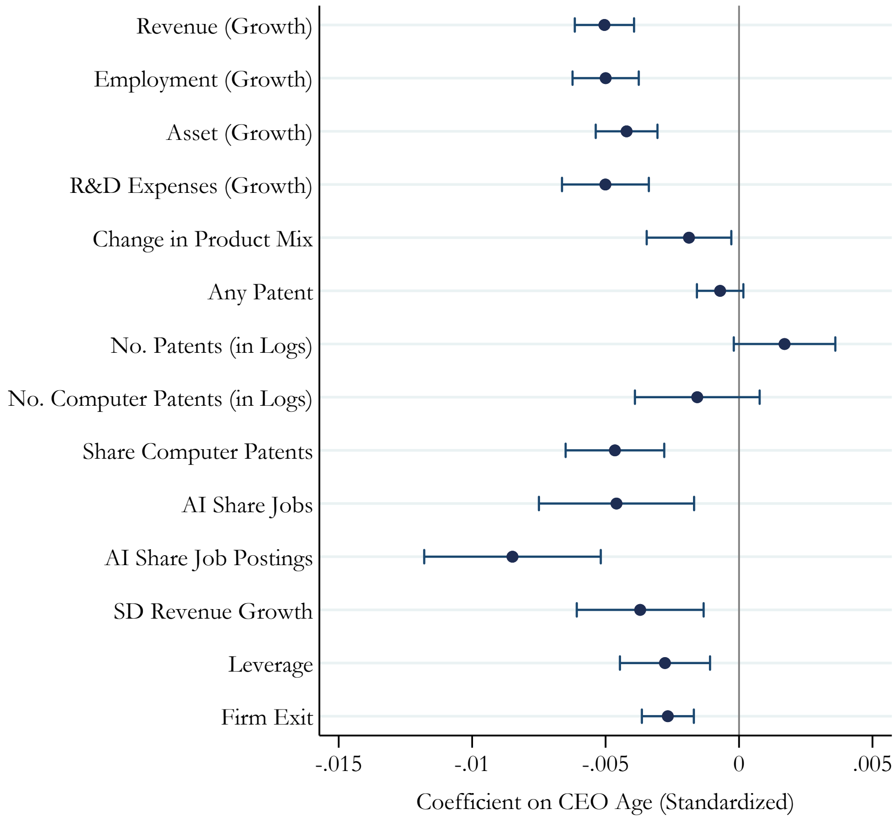

Notes: The figure plots correlations of firm-level outcomes with current CEO age, controlling for fixed effects at the NAICS-4-year level, as well as the natural logarithm of employment and firm age (both obtained from employment history data). Each coefficient stems from a separate regression. For the sake of exposition, the coefficients are standardized by dividing by the standard deviation of the outcome variable. The sample starts in 2000 and ends in 2022. The outcomes are the symmetric growth rates of revenue, employment, assets, and R&D expenses (all obtained from Compustat), changes in the product mix (Hoberg, Phillips, and Prabhala 2014), an indicator for filing any patent, the natural logarithm of the number of patents, the natural logarithm of the number of computer-related patents (share of inventional patents of product class G06), and the share of computer patents (derived from USPTO data), the shares of AI jobs and job postings (Babina et al. 2024), the 5-year ahead standard deviation of revenue growth, leverage, and firm exit (derived from Compustat). Standard errors are clustered at the NAICS-4 level.

Celik (2022) document lower levels of radical innovation among firms with older CEOs, while the same pattern does not hold for incremental innovation. Consistent with evidence from other contexts (Kitchell 1997), we observe reduced technology adoption in firms with older CEOs, specifically in the creation of AI-related jobs.

---

for whether a firm applied for a patent in a given year and calculate the natural logarithm of the number of patents. We define computer-related patents as those invention patents in product class

Given these drawbacks, why would firms appoint older CEOs? One reason may be that CEO age is systematically linked to lower firm risk. In our sample, older CEOs tend to manage firms with lower variation in 5-year ahead revenue growth, lower leverage, and a smaller probability of firm exit. Accordingly, the literature has found CEO age to be associated with lower idiosyncratic risk, stock return volatility (Peltomäki et al. 2021), stock price crash risk (Andreou, Louca, and Petrou 2017), and higher firm survival rates (Belenzon, Shamshur, and Zarutskie 2019). This finding of reduced risk-taking extends to the banking sector (Ahmed, Sihvonen, and Vähämaa 2023). One explanation for this pattern is that older managers adopt more conservative investment strategies (Vroom and Pahl 1971). Serfling (2014) documents that older CEOs take fewer risks as they age, showing that firms managed by older CEOs are more diversified across business segments. Efforts to reduce risk also manifest in that firms with older CEOs tend to have higher financial reporting quality (Huang, Rose-Green, and Lee 2012) and take longer to IPO (Yang, Zimmerman, and Jiang 2011). Conversely, younger CEOs are more likely to pursue acquisitions (Yim 2013), aggressive takeovers (Levi, Li, and Zhang 2010), and strategic restructuring (Li, Low, and Makhija 2017).

Desir et al. (2024) report lower levels of managerial ability among older CEOs. This correlation is particularly strong in high-tech firms, but it is reversed in mature and more regulated firms. An obvious but rather understudied question is whether these correlations merely reflect assortative sorting of CEOs and firms, or whether they can be interpreted as causal. Younger CEOs are more common in smaller, high-growth firms, while larger, complex firms tend to select older CEOs (Joos, Leone, and Zimmerman 2003). Acemoglu, Akcigit, and Celik (2022) find that radical innovation often precedes the appointment of younger CEOs, suggesting a sorting effect, while the causal effects of younger CEOs on innovation are comparatively small.

1.D.2 CEO Age and Adjustments to Macroeconomic Shocks

We now scrutinize whether these associations are also reflected in how firms react to aggregate shocks. To this end, we estimate the impact of business cycle fluctuations on firm-level outcomes over a longer time horizon, focusing on heterogeneity by CEO age. Figure 1.D.2 presents results from local projections using the shock series from Angeletos, Collard, and Dellas (2020). Panel (a) shows the differential responses of revenue, employment, total assets, and R&D spending to an exogenous one standard deviation decrease in the cyclical component of real GDP per capita. Across all outcomes, the coefficients are positive, highly significant, and persistent over several years, even conditional on fixed effects at the NAICS-4-year and firm level. Quantitatively, the effects are sizeable: with an estimated coefficient of

G06. Based on this, we calculate both the natural logarithm of the number of computer-related patents and their share among all invention patents.

Figure 1.D.2. Cyclical Fluctuations of Firm-Level Outcomes—Differential Effect by CEO Age

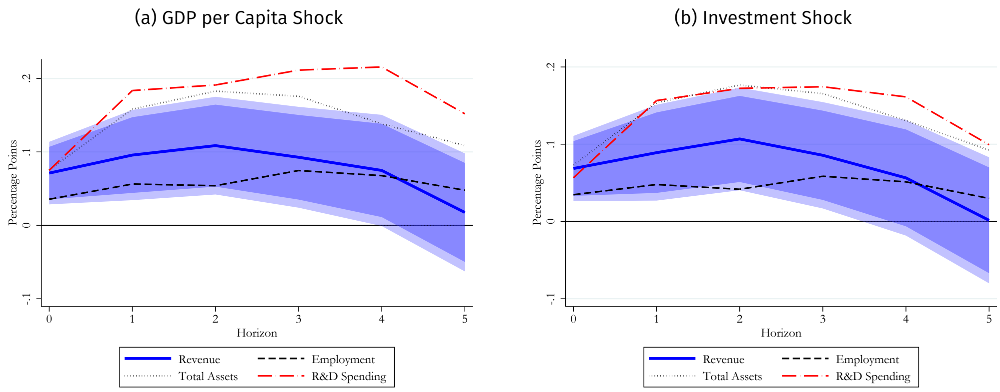

Notes: The figure shows results from local projections, obtained from h-step-ahead symmetric growth rates of the outcome variables compared to $t-1$ on CEO age, alone and interacted with different shock series. The plots show the estimated coefficients on this interaction term. Panels (a) and (b) show the results for the real GDP per capita shock and the investment shock series, respectively, which have been obtained from Angeletos, Collard, and Dellas (2020). The shock series have been averaged by year, standardized to mean zero and a standard deviation of one, and included in first lags. We control for the natural logarithm of employment and firm age (both obtained from employment history data) as well as CEO tenure. We also include fixed effects at the NAICS-4-year and firm level. The sample runs from 1980 to 2016. The plot shows point estimates alongside 90% and 95% confidence bands, obtained from standard errors clustered at the firm level. Confidence intervals are based on the estimated standard errors for the result on firm revenue.

around 0.1, a ten-year older CEO increases revenue growth by about 1 percentage point following a negative one standard deviation business cycle shock. Panel (b) presents analogous estimates using the investment shock series, revealing a similar pattern. These results are consistent with Schoar, Yeung, and Zuo (2024), who document systematic variation in the role of managers in shaping firms' risk exposure. Together, these results suggest that older CEOs enhance firms' resilience to aggregate fluctuations.

1.D.3 IV Regressions

We establish causality through a two-stage estimation approach that leverages our finding that greater industry-level uncertainty leads to older CEOs in the absence of other generalists (Section 1.5.1). First, we estimate firm-by-CEO fixed effects for various time-varying firm-level outcomes to capture the average performance of each CEO throughout their tenure. Second, we regress these fixed effects on the CEO's age at appointment, instrumenting age at appointment with firms' demand for generalists. This identification strategy is motivated by the idea that shifts in the demand for certain skills lead to the appointment of systematically different types of CEOs, with causal consequences for firm-level policies and outcomes (Fee, Hadlock, and Pierce 2013).

Table 1.D.1. CEO Age and Firm-Level Outcomes—Results from IV Regressions

	Emp. Gr.	R&D Gr.	Any Patent	Log(Patent)	Log(Comp. Patent)	Share Comp. Patent
	(1)	(2)	(3)	(4)	(5)	(6)
CEO Age	-0.016	-0.036***	-0.017	0.038	-0.057**	-0.022***
	(0.010)	(0.011)	(0.017)	(0.038)	(0.023)	(0.008)
R^2	-0.206	-0.853	-0.196	-0.053	-0.364	-0.441
Observations	2,505	1,138	9,919	1,951	1,951	1,951
MSA-Year fixed effects	✓	✓	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓	✓	✓
Kleibergen-Paap F-Statistic	11.159	33.357	14.694	10.049	10.049	10.049

Notes: The table shows results from estimating Equation (1.D.2). The dependent variables are the firm-by-CEO fixed effects $\hat{\nu}_{wj}$, which we estimate as described in Equation (1.D.1) for the following variables: (1) symmetric growth rates of employment, (2) symmetric growth rates of R&D expenditure, (3) an indicator for filing any patent, (4) the natural logarithm of the number of patents, (5) the natural logarithm of the number of computer-related patents, and (6) the share of computer-related patents. The independent variable is CEO age, which we instrument with the interaction of the flight time (in hours) to the closest MBB office with the cross-sectional standard deviation of the symmetric growth rate of firm revenue within a firm's NAICS-4 industry. The sample is restricted to those MSAs without MBB presence. All regression controls for fixed effects at the MSA-year and NAICS-4-year level. Standard errors are double-clustered at the NAICS-4 industry and MSA level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In the first step, consider a CEO w in firm j, industry i, and year τ. We then estimate the following equation for a time-varying firm outcome $\nu_{j\tau}$

$$\nu_{j\tau} \;=\; \nu_{wj} + \phi_{i\tau} + X_{j\tau} + \epsilon_{j\tau}, \tag{1.D.1}$$

where ν_{wj} are firm-by-CEO fixed effects and $\phi_{i\tau}$ denote NAICS-4-year fixed effects. The vector of controls, $X_{j\tau}$, includes the natural logarithm of the number of employees and firm age, both obtained from employment history data, as well as the natural logarithm of the CEO's tenure at the firm. The coefficients of interest are the estimated firm-by-CEO fixed effects $\hat{\nu}_{wj}$, which capture a firm's performance throughout a CEO's tenure, net of time-varying industry conditions and firm-level characteristics.

In the second step, we use this measure to estimate the causal effect of CEO age at appointment t on the match-specific firm outcomes $\hat{\nu}_{wj}$, leveraging our plausibly exogenous measure of demand for more experienced CEOs as an instrument. Specifically, we run the following regression

$$\hat{\nu}_{wj} \;=\; \beta\widehat{Age}_j + \gamma_{mt} + \phi_{it} + \epsilon_{jt}, \tag{1.D.2}$$

where $\widehat{Age}_{jt}$ is the predicted age at appointment based on Equation (1.1). That is, we exploit variation in the demand for generalist skills, as captured by the interaction of industry-level uncertainty and travel time to the offices of top strategy consulting firms, to instrument for CEO age at appointment and identify its causal effect on our

Table 1.D.2. CEO Age and Firm-Level Outcomes—Results from IV Regressions (Adjusted Standard Errors)

	Emp. Gr.	R&D Gr.	Any Patent	Log(Patent)	Log(Comp. Patent)	Share Comp. Patent
	(1)	(2)	(3)	(4)	(5)	(6)
CEO Age	-0.016	-0.036*	-0.016	0.038	-0.057**	-0.022**
	(0.014)	(0.020)	(0.018)	(0.042)	(0.025)	(0.010)
R^2	-0.206	-0.853	-0.196	-0.053	-0.364	-0.441
Observations	2,505	1,138	9,919	1,951	1,951	1,951
MSA-Year fixed effects	✓	✓	✓	✓	✓	✓
Naics4-Year fixed effects	✓	✓	✓	✓	✓	✓
Kleibergen-Paap F-Statistic	11.159	33.357	14.694	10.049	10.049	10.049

Notes: The table shows results from estimating (1.D.2). The dependent variables are the firm-by-CEO fixed effects $\hat{v}_{wj}$, which we estimate as described in (1.D.1) for the following variables: (1) symmetric growth rates of employment, (2) symmetric growth rates of R&D expenditure, (3) an indicator for filing any patent, (4) the natural logarithm of the number of patents, (5) the natural logarithm of the number of computer-related patents, and (6) the share of computer-related patents. The independent variable is CEO age, which we instrument with the interaction of the flight time (in hours) to the closest MBB office with the cross-sectional standard deviation of the symmetric growth rate of firm revenue within a firm's NAICS-4 industry. The sample is restricted to those MSAs without MBB presence. All regressions control for fixed effects at the MSA-year and NAICS-4-year level. Standard errors are generated using a bootstrapping procedure with $R = 500$ repetitions, in which we stratify draws with replacement within CEO spells in the first step to obtain a distribution of estimated second-step coefficients $\hat{\beta}_r$, where r indexes the estimation. The corresponding analytical standard errors $\hat{\sigma}_r$ are obtained from double-clustering at the NAICS-4 industry and MSA level. We compute the between-variance $\hat{\sigma}^2_\beta$ as the variance over all $\hat{\beta}_r$, and the within-variance $\hat{\sigma}^2$ by averaging over all squares of $\hat{\sigma}_r$. The reported standard errors are then obtained by combining both components using the formula for multiple imputation described in Rubin (1987) as $\sqrt{\hat{\sigma}^2 + (1 + \frac{1}{R})\hat{\sigma}^2_\beta}$. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

measures of CEO performance. Analogously to our first-stage regression described by Equation (1.1), we control for time-varying characteristics at the level of metro areas and NAICS-4 industries.

We conduct this analysis for two different types of outcomes. On the one hand, we consider balance sheet variables, specifically the symmetric growth rates of employment and R&D expenditures. On the other hand, we exploit information on firms' patenting activity, both overall and restricted to computer-related patents. The results are shown in Table 1.D.1. While we find no significant coefficient for employment growth (Column 1), we detect a negative effect for R&D expenditure growth. Further, we find no evidence of a causal effect of CEO age on overall patenting behavior, whether measured by an indicator for filing any patent in a given year (Column 3) or the number of patents (Column 4). Lastly, our results reveal a significantly negative effect on the number of computer patents (Column 5) and on the share of computer patents among total patents (Column 6). The instrument is relevant throughout, with an F-statistic above 10 in all specifications. In terms of magnitude, having a CEO who is one year older is associated with a decrease in

Table 1.D.3. Summary Statistics—Firm-Level Outcomes

	Mean	SD	P25	P50	P75	N
Panel A: Firm-Level Outcomes						
Revenue (Growth)	10.57	38.70	-1.61	8.19	20.78	86,458
Employment (Growth)	5.51	26.71	-3.59	3.29	13.33	84,701
Asset (Growth)	9.83	31.63	-2.82	6.14	18.34	88,672
R&D Expenses (Growth)	9.17	39.74	-5.74	8.08	24.27	40,131
Change in Product Mix	17.48	15.46	7.32	13.76	23.15	79,281
Any Patent	0.20	0.40	0.00	0.00	0.00	328,690
Number of Patents	17.35	108.54	1.00	2.00	6.00	64,665
Number of Computer Patents	2.86	31.11	0.00	0.00	0.00	64,665
Share Computer Patents	0.15	0.32	0.00	0.00	0.00	64,665
AI Share Jobs (in %)	0.10	0.36	0.00	0.00	0.06	33,103
AI Share Job Postings (in %)	0.28	1.55	0.00	0.00	0.00	18,509
SD Revenue Growth	21.77	27.47	7.06	13.18	24.94	71,720
Leverage	0.89	1.62	0.06	0.41	1.01	86,711
Firm Exit	0.03	0.17	0.00	0.00	0.00	92,654
Panel B: Firm-Level Outcomes, Fixed Effects Estimates						
Employment (Growth)	-0.01	0.29	-0.17	0.02	0.18	4,420
R&D Expenses (Growth)	-0.02	0.38	-0.21	0.02	0.20	2,153
Any Patent	0.05	0.36	-0.21	-0.09	0.26	11,120
Number of Patents (in Logs)	-0.10	1.03	-0.76	-0.28	0.35	3,209
Number of Computer Patents (in Logs)	0.02	0.82	-0.41	-0.24	0.12	3,198
Share Computer Patents	0.01	0.32	-0.16	-0.14	-0.02	3,209

Notes: The table shows summary statistics for outcome variables used to analyze the link between firm-level effects of CEO age.

employment and R&D expenditure growth of 1.7 and 3.6 percentage points, respectively, as well as a 2.2% reduction in the share of computer-related patents.[51]

51. One potential concern is that the reported standard errors may be understated, as they do not explicitly factor in the uncertainty stemming from the first-step estimation. In Table 1.D.2, we address this by bootstrapping the first step to obtain a distribution of point estimates, which we then combine with the analytical standard errors from the second-step estimation using Rubin's rules for multiple imputation. The results remain qualitatively unchanged, with all coefficients except that on employment growth remaining statistically significant.

Appendix 1.E Details on BoardEx CEO Sample

1.E.1 Alternative Data Preparation

For the descriptive analysis, we transform the data into a panel to allow for a meaningful decomposition of employment histories and job turnover. One drawback of removing parallel spells is that it decreases the probability of classifying a CEO appointment as internal. Since our analysis focuses on changes over time rather than levels, this issue is practically irrelevant for our purposes. Nevertheless, we show below that our main conclusions do not hinge on this classification. We proceed as follows: starting from the spell-level data, we retain each individual's first CEO appointment per firm and classify it as internal if the individual held any other position in the same firm during the preceding year. For individuals with overlapping CEO spells, we remove all spells but the one that began earliest.

First, as shown in Figure 1.E.1, CEO age at appointment increases by almost eight years over the sample period (Panel a). Second, when decomposing total work experience into tenure accumulated before and after joining the firm where the individual became CEO, we find that the increase is entirely driven by a rise in external experience, while internal experience remains stable over time (Panel b). Third, external experience becomes more important, irrespective of the seniority level at which a CEO joined the firm. Both internally and externally appointed CEOs have been entering at later career stages with more external experience (Panel c), whereas internal experience has remained stable (Panel d).

The key difference that emerges relative to our baseline analysis is the frequency of internal appointments. In comparison to our baseline probability of 20%, the corresponding figure is now at 32%. To document heterogeneity by firm size and listing status, Figure 1.E.2 reproduces Figure 1.B.14, which plots the share of internal appointments, stratified by these two variables. Consistent with our prior analysis, firms in the S&P 500 are far more likely than unlisted firms to hire internally (74% vs. 25%). A very similar pattern is evident when comparing firms in the largest with those in the smallest employment quartile (60% vs. 25%).

1.E.2 Sample Composition

Another potential concern relates to BoardEx's inclusion criteria. In particular, once an individual becomes CEO of a large firm, BoardEx might retrospectively add their entire employment history, including prior CEO appointments at smaller firms that would not appear in the data otherwise. This could mechanically generate a pattern whereby younger CEOs are more likely to appear in smaller firms in earlier years, while a similar bias would not be present for larger firms. Thus, the inclusion criteria could spuriously produce the gradient in CEO age with respect to firm size that we document in the main analysis. To address this concern, we reproduce Panel (b) of

Figure 1.E.1. CEO Age at Appointment Over Time, Alternative Data Preparation

(a) Age

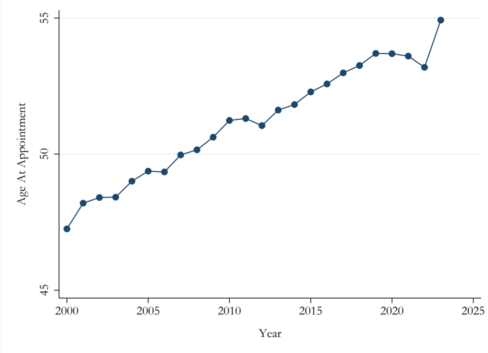

(b) External vs. Internal Work Experience

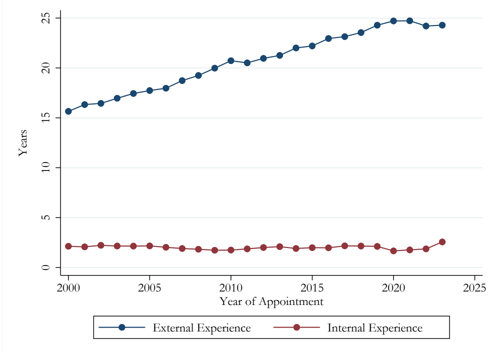

(c) External Experience, by Seniority of First Position

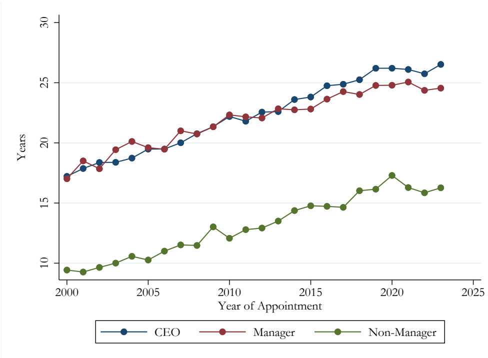

(d) Internal Experience, by Seniority of First Position

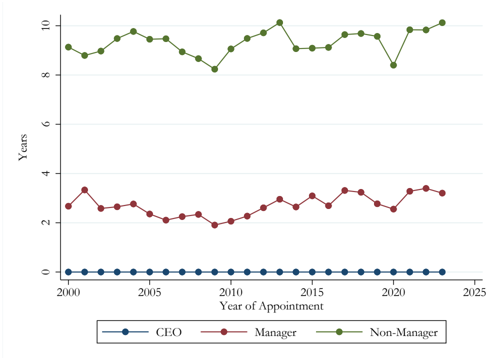

Notes: Panel (a) shows the average age at appointment of CEOs in the United States over time. Panel (b) plots the average external vs. internal work experience prior to appointment. Panels (c) and (d) show the average years of external and internal experience, respectively, stratified by the seniority of their first position in the firm.

Figure 1.2 excluding all individuals who ever held a CEO position at a smaller firm before later becoming CEO of a larger firm. The results are shown in Figure 1.E.3. Reassuringly, albeit overly conservative, excluding this group of CEOs leaves the pattern unchanged.

Figure 1.E.2. Share of Internal Appointments by Firm Size and Listing Status, Alternative Data Preparation

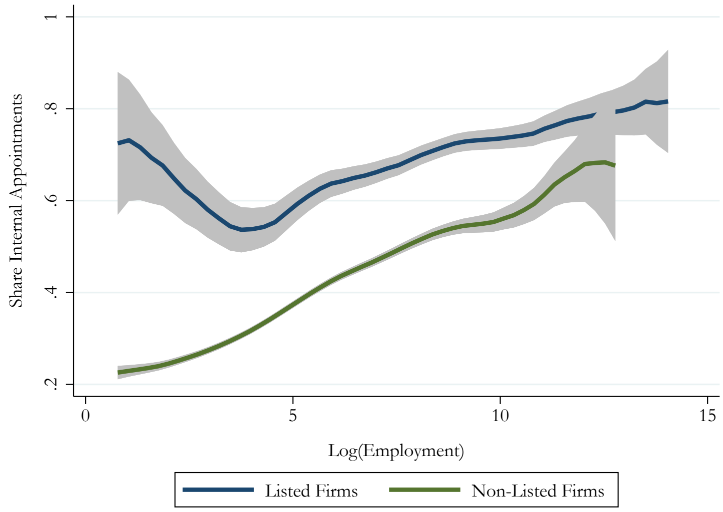

Notes: The figure plots the share of internal appointments against firm size measured by the natural logarithm of employment (obtained from employment histories). The data are smoothed with a local polynomial with Epanechnikov kernel and bandwidths of 0.91 and 0.78, respectively, consistent with Figure 1.B.14. Point estimates are shown alongside 95% confidence intervals.

Figure 1.E.3. Average CEO Age Over Time by Firm Size (Number of Employees), Restricted Sample

Notes: The plot shows age at appointment separately for those in the largest firm quartile and for those in the other three quartiles. Size quartiles are generated separately by year. The sample excludes all individuals who ever held a CEO position at a smaller firm before later becoming CEO of a larger firm.

References

Acemoglu, Daron, Ufuk Akcigit, and Murat Alp Celik. 2022. "Radical and Incremental Innovation: The Roles of Firms, Managers, and Innovators." *American Economic Journal: Macroeconomics* 14 (3): 199–249. [11, 105, 107]

Ahmed, Shaker, Jukka Sihvonen, and Sami Vähämaa. 2023. "Executive Age and Bank Risk-Taking." Working Paper, Vaasa University. [107]

Andreou, Panayiotis, Christodoulos Louca, and Andreas Petrou. 2017. "CEO Age and Stock Price Crash Risk." *Review of Finance* 21 (3): 1287–325. [107]

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas. 2020. "Business-Cycle Anatomy." *American Economic Review* 110 (10): 3030–70. [107, 108]

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. "Artificial Intelligence, Firm Growth, and Product Innovation." *Journal of Financial Economics* 151: 103745. [106]

Baker, Scott, Nicholas Bloom, Steven Davis, and Kyle Kost. 2019. "Policy News and Stock Market Volatility." Working Paper, 25720, National Bureau of Economic Research. [31, 81]

Baker, Scott R., Nicholas Bloom, and Steven J. Davis. 2016. "Measuring Economic Policy Uncertainty." *Quarterly Journal of Economics* 131 (4): 1593–636. [7, 31, 81]

Barba Navaretti, Giorgio, Davide Castellani, and Fabio Pieri. 2022. "CEO Age, Shareholder Monitoring, and the Organic Growth of European Firms." *Small Business Economics* 59 (1): 361–82. [105]

Barker, Vincent, and George Mueller. 2002. "CEO Characteristics and Firm R&D Spending." *Management Science* 48 (6): 782–801. [105]

Belenzon, Sharon, Anastasiya Shamshur, and Rebecca Zarutskie. 2019. "CEO's Age and the Performance of Closely Held Firms." *Strategic Management Journal* 40 (6): 917–44. [105, 107]

Bertrand, Marianne, and Antoinette Schoar. 2003. "Managing with Style: The Effect of Managers on Firm Policies." *Quarterly Journal of Economics* 118 (4): 1169–208. [6, 9, 105]

Bharath, Sreedhar, and Amy Dittmar. 2010. "Why do Firms Use Private Equity to Opt Out of Public Markets?" *Review of Financial Studies* 23 (5): 1771–818. [72]

Chetty, Raj, John Friedman, Emmanuel Saez, Nicholas Turner, and Danny Yagan. 2020. "Income Segregation and Intergenerational Mobility Across Colleges in the United States." *Quarterly Journal of Economics* 135 (3): 1567–633. [26, 88]

Child, John. 1974. "Managerial and Organizational Factors Associated with Company Performance Part I." *Journal of Management Studies* 11 (3): 175–89. [105]

Cline, Brandon, and Adam Yore. 2016. "Silverback CEOs: Age, Experience, and Firm Value." *Journal of Empirical Finance* 35: 169–88. [11, 105]

Datta, Deepak, Nandini Rajagopalan, and Yan Zhang. 2003. "New CEO Openness to Change and Strategic Persistence: The Moderating Role of Industry Characteristics." *British Journal of Management* 14 (2): 101–14. [105]

Desir, Rosemond, Joseph Rakestraw, Scott Seavey, James Wainberg, and George Young. 2024. "Managerial Ability, CEO Age and the Moderating Effect of Firm Characteristics." *Journal of Business Finance & Accounting* 51 (1-2): 148–79. [107]

Fee, Edward, Charles Hadlock, and Joshua Pierce. 2013. "Managers With and Without Style: Evidence Using Exogenous Variation." *Review of Financial Studies* 26 (3): 567–601. [9, 108]

Flood, Sarah, Miriam King, Renae Rodgers, Steven Ruggles, J. Robert Warren, Daniel Backman, Annie Chen, et al. 2024. (IPUMS CPS: Version 12.0 [dataset]. Minneapolis, MN). https://doi.org/10.18128/D030.V12.0. [67]

Flynn, Sean, Andra C Ghent, and Vasudha Nair. 2025. "Determinants and Consequences of Return to Office Policies." Working Paper, Cornell University. [105]

Hambrick, Donald, and Phyllis Mason. 1984. "Upper Echelons: The Organization as a Reflection of Its Top Managers." *Academy of Management Review* 9 (2): 193–206. [105]

Hart, Peter, and John Mellors. 1970. "Management Youth and Company Growth: A Correlation?" *Management Decision* 4 (1): 50–53. [105]

Hidalgo, César, and Ricardo Hausmann. 2009. "The Building Blocks of Economic Complexity." *Proceedings of The National Academy of Sciences* 106 (26): 10570–75. [37, 38, 92, 93]

Hoberg, Gerard, Gordon Phillips, and Nagpurnanand Prabhala. 2014. "Product Market Threats, Payouts, and Financial Flexibility." *Journal of Finance* 69 (1): 293–324. [106]

Huang, Hua-Wei, Ena Rose-Green, and Chih-Chen Lee. 2012. "CEO Age and Financial Reporting Quality." *Accounting Horizons* 26 (4): 725–40. [107]

Jenter, Dirk, and Katharina Lewellen. 2015. "CEO Preferences and Acquisitions." *Journal of Finance* 70 (6): 2813–52. [6, 105]

Joos, Philip, Andrew Leone, and Jerold Zimmerman. 2003. "Selecting CEOs: Matching the Person to the Job." Working Paper, Tilburg University. [107]

Karlin, Samuel. 1968. *Total Positivity.* Vol. 1. Stanford University Press. [61]

Keil, Jan. 2017. "The Trouble with Approximating Industry Concentration from Compustat." *Journal of Corporate Finance* 45: 467–79. [87]

Kitchell, Susan. 1997. "CEO Characteristics and Technological Innovativeness: A Canadian Perspective." *Canadian Journal of Administrative Sciences/Revue Canadienne des Sciences de l'Administration* 14 (2): 111–21. [106]

Levi, Maurice, Kai Li, and Feng Zhang. 2010. "Deal or No Deal: Hormones and the Mergers and Acquisitions Game." *Management Science* 56 (9): 1462–83. [11, 107]

Li, Xiaoyang, Angie Low, and Anil Makhija. 2017. "Career Concerns and the Busy Life of the Young CEO." *Journal of Corporate Finance* 47: 88–109. [10, 11, 105, 107]

Nguyen, Pascal, Nahid Rahman, and Ruoyun Zhao. 2018. "CEO Characteristics and Firm Valuation: A Quantile Regression Analysis." *Journal of Management & Governance* 22: 133–51. [105]

Peltomäki, Jarkko, Jukka Sihvonen, Steve Swidler, and Sami Vähämaa. 2021. "Age, Gender, and Risk-Taking: Evidence from the S&P 1500 Executives and Market-Based Measures of Firm Risk." *Journal of Business Finance & Accounting* 48 (9-10): 1988–2014. [107]

Rubin, Donald. 1987. *Multiple Imputation for Nonresponse in Surveys.* John Wiley & Sons. [110]

Schoar, Antoinette, Kelvin Yeung, and Luo Zuo. 2024. "The Effect of Managers on Systematic Risk." *Management Science* 70 (2): 815–33. [11, 108]

Serfling, Matthew. 2014. "CEO Age and the Riskiness of Corporate Policies." *Journal of Corporate Finance* 25: 251–73. [11, 105, 107]

Shaked, Moshe, and J George Shanthikumar. 2007. *Stochastic Orders.* Springer. [61]

Vroom, Victor, and Bernd Pahl. 1971. "Relationship Between Age and Risk Taking Among Managers." *Journal of Applied Psychology* 55 (5): 399. [107]

Waelchli, Urs, and Jonas Zeller. 2013. "Old Captains at the Helm: Chairman Age and Firm Performance." *Journal of Banking & Finance* 37 (5): 1612–28. [11, 105]

Whitt, Ward. 1982. "Multivariate Monotone Likelihood Ratio and Uniform Conditional Stochastic Order." *Journal of Applied Probability* 19 (3): 695–701. [63]

Wiersema, Margarethe, and Karen Bantel. 1992. "Top Management Team Demography and Corporate Strategic Change." *Academy of Management Journal* 35 (1): 91–121. [105]

Yang, Qin, Monica Zimmerman, and Crystal Jiang. 2011. "An Empirical Study of the Impact of CEO Characteristics on New Firms' Time to IPO." *Journal of Small Business Management* 49 (2): 163–84. [107]

Yim, Soojin. 2013. "The Acquisitiveness of Youth: CEO Age and Acquisition Behavior." *Journal of Financial Economics* 108 (1): 250–73. [11, 107]

Chapter 2

Bankers' Outside Options and Financial Risk

Joint with Georg Schneider

2.1 Introduction

In imperfect labor markets, workers' outside options play a central role in determining wages and job mobility. Improvements in the quality or quantity of alternative employment opportunities strengthen workers' bargaining positions, translating into higher pay and more frequent transitions across employers (Caldwell and Harmon 2019; Caldwell and Danieli 2024; Johnson, Lavetti, and Lipsitz 2025). However, outside options may not only affect workers' matching to firms, but also their actions within firms. We investigate the link between outside options and on-the-job behavior in the context of the financial industry, a setting well suited to study this question for two reasons. First, job-to-job transitions represent a relatively high share of new hires in the financial sector (Azzopardi et al. 2020). Bankers should therefore attend closely to the on-the-job arrival of offers, making their behavior inside the firm sensitive to the state of the external labor market. Second, employees have considerable discretion in making risky, high-stakes choices on behalf of their firms, which have limited capacity to monitor these decisions.

We present empirical evidence that bankers with better labor market outside options make riskier lending decisions and argue that an insurance channel drives these results: improved outside options reduce the alignment of interests between workers and firms, as they undercut the disciplinary role of dismissal. Because employment opportunities are procyclical, this effect varies systematically over the business cycle and can contribute to fluctuations in credit supply. Given the large, negative spillovers of financial crises, these dynamics in a narrow segment of the labor market can carry consequences well beyond the firms involved.

To formalize this mechanism, we develop a stylized model in which a risk-averse banker chooses the risk level of a loan portfolio, balancing higher expected returns against greater payoff volatility. The banker faces a stochastic labor market outside option that is realized after the risk choice. The presence of this outside option truncates the distribution of potential payoffs from below, implying that the banker is partly insured against adverse portfolio realizations. As a result, the outside option reduces effective risk aversion and distorts incentives by weakening the downside consequences of failure. The main prediction of the model is that exogenous improvements in outside options lead bankers to make riskier lending decisions.[1]

In testing this prediction empirically, the key challenge is to identify shocks to outside options that are exogenous to contemporaneous bank-level dynamics. For instance, a natural proxy for outside options is the rate at which workers are departing for other banks, as it plausibly reflects their current employment opportunities. However, as outflows may also respond to the bank's own financial conditions, a negative shock to a bank could raise outflows while depressing credit supply, generating a spurious negative correlation between the two variables. To address this concern, we construct a measure of outside options that is orthogonal to a bank's own contemporaneous dynamics. Specifically, we use individual employment biographies of senior US investment bankers to create a network-based measure, capturing labor demand at other banks. This approach rests on the idea that information about hiring dynamics is transmitted through coworker networks. By exploiting variation in the strength of ties between banks, we measure the differential exposure of workers across banks to labor demand at other institutions. In a series of validation exercises, we show that our measure of outside options is orthogonal to past employee outflows as well as to contemporaneous and past bank-level financial fundamentals. Consistent with prior literature, the measure is also predictive of both worker transitions to other firms and compensation gains for the remaining workers.

In our main empirical analysis, we examine how bankers' outside options affect their lending decisions and the resulting financial risk. With the inclusion of fixed effects at the bank and year-quarter level, we isolate within-bank variation from latent bank-level differences as well as aggregate shifts in risk attitude over time. Drawing on transaction-level data on syndicated loans, our results indicate that banks expand lending volumes in response to improved outside options by arranging more loans. This credit expansion is accompanied by higher value at risk, a bank-level risk measure, with effects concentrated in the tail of the loss distribution.

1. This mechanism parallels the logic of the literature on state guarantees, which has documented increased risk-taking among banks as a response to state-provided insurance (Dam and Koetter 2012; Brandao-Marques, Correa, and Sapriza 2013; Bianchi 2016), with negative implications for financial stability (Demirgüç-Kunt and Detragiache 2002; Diamond and Rajan 2002). However, in our context, the insurance mechanism operates through markets at the level of individual bankers rather than through governments at the level of banks.

Importantly, these effects are not confined to individual banks, as improved outside options also raise a bank's contribution to systemic risk at high levels of distress.

We scrutinize the loan-level dynamics underlying these patterns and find that the credit expansion reflects a shift toward riskier borrowers. Focusing on the likelihood that a borrower is rated investment-grade, a one standard-deviation increase in outside options reduces the share of prime borrowers by 1.7% of its mean. This deterioration in borrower quality is not compensated by higher spreads, which is difficult to reconcile with yield-seeking behavior and rather points to an increase in risk-taking. We also demonstrate that this result is robust to controlling for information flows between employers. Finally, we corroborate these findings by leveraging differences in the enforceability of non-compete agreements (NCAs) in the bank's headquarters state, showing that the response in loan-level risk to outside options is concentrated in states where NCA enforceability has fallen. As weaker enforceability makes outside options more readily exercisable, this pattern suggests that bankers respond to outside options only when these map into feasible opportunities.

Turning to the underlying mechanism, we provide two sets of results showing that improved outside options increase bankers' willingness to take risk by offering implicit insurance against adverse outcomes. First, we link individual loans to underwriting bankers and study heterogeneity across bankers' demographics. Conditional on banker and bank-year fixed effects, the effects are concentrated among younger bankers, who are most exposed to dismissal risk. This result is consistent with the insurance channel and suggests that the effect we identify is not driven by senior poaching or team lift-outs, which would predict the inverse age gradient. Second, we extend the analysis to the market for financial advisors, where we find that misconduct becomes more prevalent when outside options improve. Furthermore, better outside options partly offset the disciplinary effects of layoffs. Following misconduct, they reduce the probability of job separation and, even conditional on separation, increase the likelihood of reemployment in the same sector. This state-dependence indicates that stronger labor markets weaken the disciplinary force of dismissal risk, the mechanism at the heart of the insurance channel.

These results also speak to aggregate dynamics. As documented in Figure 2.A.1, labor mobility in the financial sector is closely associated with lending standards over time. By offering a microfoundation for this relationship, we identify a worker-level channel through which labor market conditions translate into systematic variation in credit supply. Because outside options are generally considered procyclical in both theoretical and empirical accounts (Mortensen and Pissarides 1999; Nakamura et al. 2019; Figueiredo 2022), real economic shocks that expand both credit demand and competition over banking talent can reinforce one another, contributing to excess lending during expansions. Due to the role of credit growth in increasing both the likelihood and the severity of financial crises (Schularick and Taylor 2012; Jordà, Schularick, and Taylor 2013), our findings point to a neglected role of labor markets in macroprudential policy.

Literature: Our main contribution is to provide empirical evidence on a novel mechanism through which labor markets affect credit supply, adding to four strands of literature. First, a small theoretical literature has studied potential interactions between agency problems within banks, labor market dynamics in the financial sector, and the quality and quantity of credit (Myerson 2012; Thanassoulis 2012; Myerson 2014; Axelson and Bond 2015; Acharya, Pagano, and Volpin 2016; Van Boxtel 2017). These papers analyze different channels through which the labor market for bankers can impact credit supply, including frictions in the supply of bankers, externalities generated by competition for talent, and bankers’ job opportunities outside the financial sector. We add to this literature by providing empirical evidence on this interaction, and specifically on the channel through which hiring at one bank improves other bankers’ outside options and reduces alignment with their employer.

Second, we inform the discussion about the role of career concerns in mitigating risky behavior in the financial sector, which can be traced back to the seminal work of Holmström (1999). The evidence on whether bankers are punished for adverse choices is mixed: some studies report no effect in response to corporate misconduct (Griffin, Kruger, and Maturana 2019; Hamdi, Kalda, and Pal 2023), while others detect such effects in the context of syndicated loans and hedge funds (Ellul, Pagano, and Scognamiglio 2020; Gao, Kleiner, and Pacelli 2020). Our study complements these findings by showing that the career consequences of risky choices vary with outside options, and that bankers internalize this state-dependence.

Third, a related empirical literature has examined different aspects of the labor market in the financial sector, such as the high labor mobility of investment bankers (Gao, Wang, and Wu 2024; Gao, Wang, and Yu 2024) and the effect of labor mobility restrictions on financial intermediation (Bonelli 2019; Cici, Hendriock, and Kempf 2021; Agarwal et al. 2024; Norden, Van Doornik, and Wang 2025). Other papers have highlighted the discretionary role of individual bankers in setting lending terms for syndicated loans (Bushman et al. 2021; Herpfer 2021; Carvalho, Gao, and Ma 2023). We contribute to this literature by showing that labor mobility across firms can lead to the origination of more credit, more risky credit, and more financial risk in the context of corporate lending.

Finally, our paper speaks to a recent literature examining the role of outside options in labor markets (Schubert, Stansbury, and Taska 2022; Caldwell and Danieli 2024; Caldwell, Haegele, and Heining 2024). Jäger et al. (2024) demonstrate that information from coworkers shifts worker beliefs about outside options, in turn affecting search and bargaining behavior. Caldwell and Harmon (2019) study the role of personal networks in transmitting information about outside options across firms and show that news transmitted through these networks promotes wage growth and labor mobility. Ahammer, Fahn, and Stiftinger (2023) find a negative relationship between outside options and worker effort by exploiting variation in unemployment insurance. We add to this work by documenting that outside options shape the job-related decisions of high-skilled workers through a risk-taking channel.

The structure of the paper is as follows: Section 2.2 lays out the conceptual framework and main predictions, followed by Section 2.3, which details the data. Section 2.4 describes our outside options measure and identification strategy, and Section 2.5 presents the results from the syndicated loan market. We examine the underlying insurance mechanism in Section 2.6, and Section 2.7 concludes.

2.2 Conceptual Framework

To guide our analysis of how outside options affect risk-taking, we present a stylized framework featuring a banker who chooses optimal portfolio risk in the presence of a random outside option. The key prediction is that better outside options induce higher levels of risk-taking. The intuition behind this result builds on Curello, Sinander, and Whitmeyer (2025), who show that improvements in outside options are equivalent to reductions in effective risk aversion. The latter is defined as the curvature of the utility function that rationalizes a banker's choices once the outside option is taken into account, which differs from their underlying preferences.

While this setup resembles a conventional portfolio-choice problem, it is nonstandard. This is because the presence of an outside option changes the selection rule by replacing poor portfolio outcomes with a fallback.[2] The transformation works through the quality of the outside option, summarized by the reverse hazard rate of its distribution, and leaves both the mean and the variance of payoffs unchanged. The standard wealth and variance channels therefore play no role. The higher the reverse hazard rate, the better the option insures against bad realizations, and the lower the banker's effective risk aversion. In the following, we parameterize an application of the framework of Curello, Sinander, and Whitmeyer (2025) to generate closed-form comparative statics that we bring to the data.

Setup: We consider a banker with CARA utility choosing a portfolio risk measure $r \in [0, \bar{r}]$. Both the expected return and payoff risk are increasing in r. Specifically, we model portfolio payoffs by a normal approximation,

$$x(r) \;=\; \mu_0 + a\,r + \varepsilon, \qquad \varepsilon \;\sim\; \mathcal{N}\big(0,\, v_0 + b\,r^2\big), \tag{2.1}$$

with $a > 0$ and $b > 0$. The banker receives a share $y(r) = \alpha x(r)$ of the portfolio returns, with the bank retaining $(1-\alpha)x(r)$. Let baseline utility v be CARA with baseline risk aversion $\sigma > 0$, $v(x) = -\frac{1}{\sigma}e^{-\sigma x}$.

We assume that the outside option is reversed exponential on $(-\infty, 0]$ with rate $\lambda \geq 0$, implying a constant reverse hazard: $G'(x)/G(x) = \lambda$.[3] The labor market outside option is a random payoff K with cumulative distribution G, which is continuous

2. This sets it apart from ordinary background risk, which adds mean-zero noise and typically raises effective risk aversion.

3. Define $K = -E$ with $E \sim \mathrm{Exp}(\lambda)$. Then K follows a mirrored exponential distribution bounded above by 0, with cumulative distribution function $G(x) = e^{\lambda x}$ for $x \leq 0$ and $G(x) = 1$ for $x > 0$. Thus,

and has positive density on its support. That is, the banker knows the distribution of outside option payoffs but is unaware of the realization during the decision-making process. A higher λ shifts probability mass of the outside-option distribution toward its upper bound 0, thinning the left tail and making low realizations of K less likely. Intuitively, outside options improve in expectation while becoming less dispersed. To maintain tractability, we assume K is independent of the payoff shock in $x(r)$, ruling out outside options that co-move with the banker's own portfolio outcomes.

The banker chooses r to maximize

$$\max_r \mathbb{E}[v(\max\{y(r), K\})]. \tag{2.2}$$

The timing is as follows: the banker chooses r, then the payoff shock ε and the outside option K realize, determining consumption $Y^{\text{eff}} = \max\{y(r), K\}$. As the banker chooses r before K realizes but with full knowledge of its distribution, risk-taking responds to λ, which captures expectations about the outside option.

Decomposition: We show that the outside option generates a CARA representation with effective risk aversion $\sigma - \lambda$. Let $W(y) \equiv \mathbb{E}_K[v(\max\{y, K\})]$ denote the objective function induced by the outside option. Given that v is increasing and twice-differentiable and G has positive density g on its support, differentiation yields

$$W'(y) = \mathbb{E}\big[v'(\max\{y, K\}) \cdot \mathbf{1}\{K \leq y\}\big] = v'(y)\, G(y), \tag{2.3}$$

because an incremental increase in y matters only in states where y binds, that is, where $K \leq y$. Differentiating again and dividing by $W'(y) = v'(y)G(y)$ yields

$$\frac{W''(y)}{W'(y)} = \frac{v''(y)}{v'(y)} + \frac{G'(y)}{G(y)}. \tag{2.4}$$

Since choices are invariant to positive affine transformations of the objective function, we can take $u(y) = aW(y) + b$ with $a > 0$, which is also CARA with coefficient of absolute risk aversion

$$-\frac{u''(y)}{u'(y)} = \sigma - \lambda. \tag{2.5}$$

Effective risk aversion $\sigma - \lambda$ is thus strictly below σ, confirming the intuition above. Only a fraction $G(y)$ of states use y rather than K, so marginal utility is effectively scaled by $G(y)$. The second differentiation converts this scaling into an additive term, the reverse hazard rate $\frac{G'(y)}{G(y)}$, which lowers the agent's absolute risk aversion.

---

when λ is large, more probability mass lies just below 0. The exponential is the case in which the reverse hazard is constant, which delivers closed-form comparative statics, while more general distributions yield a state-dependent effective risk aversion.

Comparative Statics: The decomposition shows that effective utility u is also CARA with effective absolute risk aversion $\rho = \sigma - \lambda$. That is, improvements in the outside option (higher λ) translate into a one-for-one reduction of ρ. Under CARA–normal, expected utility is summarized by the certainty equivalent, which takes the form of

$$\mathrm{CE}(r) \;=\; \mathbb{E}[y(r)] - \frac{\rho}{2}\mathrm{Var}[y(r)] \;=\; \alpha(\mu_0 + ar) - \frac{\rho}{2}\alpha^2\big(v_0 + br^2\big). \tag{2.6}$$

Assuming that $\sigma > \lambda$, maximizing CE(r) yields

$$\frac{d\,\mathrm{CE}}{dr} \;=\; \alpha a - \rho\,\alpha^2 b\,r \;=\; 0 \;\Rightarrow\; r^* \;=\; \frac{a}{\rho\,\alpha b} \;=\; \frac{a}{(\sigma-\lambda)\,\alpha b}$$
$$\frac{\partial r^*}{\partial \lambda} \;=\; \frac{a}{b}\frac{1}{\alpha(\sigma-\lambda)^2} \;>\; 0 \;\;(\lambda \;<\; \sigma). \tag{2.7}$$

Consistent with standard models, risk-taking increases with the risk sensitivity of returns a and decreases with both the convexity of risk b and the share α accruing to the banker.

Crucially, the framework also predicts that higher levels of λ increase the optimal level of risk-taking.[4] A more favorable outside option effectively provides insurance, as the joint probability of bad realizations for the portfolio return and the outside option becomes very small. This truncation of downside risk reduces the effective risk aversion of the banker ($\rho = \sigma - \lambda$). Empirically, we thus expect better outside options to induce greater risk-taking, reflected in higher loan volumes and a lower share of investment-grade borrowers.

While our framework treats outside options as exogenously given, they may in turn depend on the chosen risk level r. We abstract from this feedback loop, as our empirical strategy exploits variation in outside options that is orthogonal to individual behavior. The framework also takes λ as fixed, whereas outside options are in practice tightly linked to overall labor market conditions. Because outside options are procyclical, risk-taking through this channel co-moves with the credit cycle. As the labor market strengthens during expansions, improved outside options in finance add to the surge in credit demand by further relaxing lending standards. In downturns, the mechanism works in reverse: weaker outside options tighten the implicit insurance, prompting bankers to reduce risk, thereby reinforcing credit contractions. The model thus provides a micro-level channel linking labor market conditions to credit supply over the cycle.

Several additional assumptions are required to arrive at the prediction that better labor market outside options increase credit risk. First, we assume that changes

4. If the interior solution fails, i.e., if $\sigma \leq \lambda$ or $\rho \leq 0$, the solution is at the upper bound $\bar{r}$.

in other banks’ hiring behavior lead to variation in bankers’ outside options. This assumption is relatively weak, as the banking sector experiences limited entry by workers, particularly at mid-to-late career stages (Philippon and Reshef 2012; Bell and Van Reenen 2014; Böhm, Metzger, and Strömberg 2023). Moreover, the supply of qualified graduates is likely relatively inelastic in the short run, and senior roles cannot easily be filled by junior bankers. As a result, job-to-job transitions are common in the financial industry. Second, our results hinge on the assumption that individual bankers influence loan decisions. Although loan sizes are substantial in the syndicated loan market, making this effect non-trivial, there is extensive evidence documenting the discretionary role of individual bankers in determining lending terms for syndicated loans (Bushman et al. 2021; Herpfer 2021; Carvalho, Gao, and Ma 2023). Third, we implicitly assume that banks do not immediately and fully adjust compensation to current labor market conditions in order to offset the adverse effects of improved outside options. To induce less risky lending, a bank could increase the share α of portfolio returns that go to the banker. We treat α as fixed, keeping the determination of compensation contracts outside the model.

While our model focuses on the insurance effect of outside options as alternative employment opportunities, a number of complementary mechanisms could be at play. First, better outside options are likely to increase turnover. As a result, expected tenure decreases, reducing the probability of a loan default occurring during the current employment spell and hence of being punished for excessive risk-taking. The diminished threat of punishment may then incentivize bankers to take on greater risk. As this channel operates through earlier departures to other firms, whereas our insurance mechanism operates even for bankers who stay, controlling for outflows allows us to rule out the former. In a robustness exercise, we find no evidence that turnover accounts for our main results. Second, behavioral factors could also be relevant in this setting. For example, the literature on diagnostic expectations suggests that exposure to positive cues leads individuals to form more optimistic expectations (Bordalo et al. 2024). Accordingly, favorable news about outside options could make bankers overly optimistic, leading them to behave less cautiously and engage in riskier lending practices. Our mechanism results in Section 2.6.2, however, suggest that more optimistic beliefs alone cannot explain our main findings. Finally, improved outside options may drive strategic behaviors beyond the scope of our model. For instance, bankers may leverage their enhanced bargaining power to negotiate *golden parachutes*, i.e., more generous severance packages, which lower the cost of dismissal and thereby encourage risk-taking. They may also bargain for greater discretion over lending itself. These channels operate through bankers’ strengthened position in the labor market and reinforce, rather than confound, the relationship we document.

2.3 Data

For our main analysis, we combine data on syndicated lending with employment histories of senior US investment bankers obtained from *Revelio Labs* as well as supplementary data sources on firm performance and wages. We focus on the syndicated lending market for three main reasons. First, both the size of individual loans and the size of the overall loan market make syndicated loans an important factor for financial stability. With a median loan size in our sample of 400 million USD, decisions over these loan contracts constitute high-stakes choices. In 2018, the 2.8 trillion USD syndicated lending market exceeded the 1.3 trillion USD corporate bond market, highlighting its central role in financing US firms. Second, the risk associated with loan default is not contained within a single financial institution but disseminated throughout the financial sector via the syndication process itself as well as through trading and securitization in the secondary market (Mugasha 2004). Finally, lending choices can be attributed to the arranging bank, with prior literature suggesting that individual lead bankers exert substantial influence over the design of these loans (Bushman et al. 2021; Herpfer 2021).

Dealscan: Data on lending activity by banks are downloaded from *Dealscan*. From the universe of loans recorded in the database, we identify the 1,500 largest lead arranging banks in terms of the number of loans issued and include all tranches arranged by these banks in our sample. For each loan, we include information on the borrowing and lending firm identity, tranche amount, and maturity, whether the loan is secured, and whether covenants are attached. Summary statistics on the tranche-level data are shown in Table 2.B.1. To measure credit quantity, we also aggregate the data to the lender-quarter level to obtain both the number of loans and total lending volume.

Revelio Labs: We draw on labor market data from Revelio Labs, which collects employment biographies reported on LinkedIn. The dataset includes employment details for the universe of public profiles, covering nearly 500 million workers across approximately 1.8 billion employment spells. The linked employer-employee data contain employment spell start and end dates, as well as self-reported job titles. From this information, Revelio Labs generates a task-based classification of job titles, which maps activities to each title by analyzing descriptions from resumes and online profiles, alongside responsibilities listed in job postings.[5] Initially, the algorithm identifies 1,500 job roles, which are successively aggregated using a clustering algorithm. The data also include a seniority ranking derived from job titles, ranging from 1 (lowest) to 7 (highest).

5. For additional details, see www.reveliolabs.com.

To construct our sample of senior bankers, we proceed as follows. First, we exclude all employment spells outside the United States. Second, we use the job taxonomy at the hierarchical level that defines 300 distinct roles, allowing us to identify the individuals employed in the relevant roles.[6] Third, to identify senior bankers, we only retain spells with a seniority level at or above 4. Finally, we only retain full-time bankers by excluding overlapping spells. We use all employment spells between 1994 and 2020.[7] Panel A of Table 2.B.2 shows summary statistics for this sample. As expected, the sample predominantly consists of male, white, and highly educated employees with high incomes (annually around 150,000 USD on average). Due to the backfilled nature of the employment histories, the mean age of employees is relatively low, at 37 years. Table 2.B.3 compares key demographic variables between the LinkedIn dataset and administrative data sources. Qualitatively, the sample captures the key characteristics of individuals employed in the financial industry reported in the CPS and ACS.[8] These similarities are even more pronounced when the restriction on seniority is removed.

For our main analysis, we combine the lending data from Dealscan with an outside options proxy calculated from employment histories at the parent company level. To this end, we sequentially merge on exact co-occurrences of parent companies’ gvkey, ticker, and company names, which have been cleaned and normalized beforehand, retaining only unique merges. For unmatched observations, we repeat the procedure at the subsidiary level, linking parent companies through their subsidiaries when matches are identified. In cases where multiple subsidiaries match, we prioritize the match associated with the highest number of spells for each lender parent. We manually verify all matches and identify additional ones. In this way, we are able to merge 1,268 parent companies, comprising two-thirds of subsidiaries in Dealscan. Among the 1,000 largest lenders, this accounts for 99% of individual loans and 91% of total lending volume.

For a limited sample of loans, we are able to identify individual lead arranging bankers by name and affiliation, and match them to their respective employment history in the employment history data. We follow Herpfer (2021) in extracting banker signatures from syndicated loan contracts filed with the Securities and Exchange Commission (SEC) and link these to their employment history information using the combination of reported name, affiliation, and filing date. Panel B of Table 2.B.2 shows summary statistics for this sample.

6. The included roles are “Investment Specialist”, “Investment Banking”, “Banker”, “Credit Specialist”, “Financial Advisor”, “Loan Officer”, and “Financial Analyst”.

7. The data have been collected since 2019 but contain all self-reported employment spells from earlier periods.

8. Both are obtained via IPUMS. For details on the ACS data, we refer to Ruggles et al. (IPUMS USA: Version 16.0).

Compustat: We obtain data on firm-level fundamentals from Compustat. For both borrowing and lending firms, we include data on assets, liabilities, net income, EBIT, short-term debt, long-term debt, employees, stockholder equity, and capital expenditure, with summary statistics shown in Table 2.B.4. We also exploit borrower firm S&P credit ratings as a measure of credit risk and repayment probability. We use the linking tables by Chava and Roberts (2008) and Schwert (2020) to match Dealscan and Compustat borrower and lender firms (Panels A and B of Table 2.B.4, respectively). To ensure a consistent sample of firms throughout the analysis, we only use data on banks for which we observe syndicated loan activity, employment histories, and balance sheet information.

(Conditional) Value at Risk: To measure financial risk, we use data from Adrian and Brunnermeier (2016), which are available at quarterly frequency between 1994 and 2009. The first measure is *Value at Risk* (VaR), which captures the maximum loss for an institution i at the 95% or 99% confidence level, i.e., the financial risk of the lending bank. The second variable is *Conditional Value at Risk* (CoVaR), defined as the VaR of the financial sector conditional on bank i being in distress. In other words, CoVaR quantifies the expected loss to the financial system conditional on a 95^{th} or 99^{th} percentile loss at bank i, capturing the (non-causal) contribution of bank i to systemic risk. Adrian and Brunnermeier (2016) estimate these measures using quantile regressions. Summary statistics are shown in Panel C of Table 2.B.4.

BoardEx: We use information on banker pay from BoardEx, a database provided by *Delinian*. These data include detailed information on compensation for top executives of large, listed companies in the United States since 1997. We construct year-over-year growth rates in the firm-level median of salaries, bonuses, and equity at risk. To better capture wage bargaining, we restrict the sample to roles with operational responsibilities rather than executive oversight.[9] Summary statistics are shown in Panel D of Table 2.B.4.

Investment Advisor Data: We leverage data on misconduct among US investment advisors, covering individuals registered with the SEC, a state authority, or both. The dataset is a panel spanning 2007 to 2015 and includes detailed employment histories alongside annual counts of disclosures related to misconduct or customer complaints. We combine this information with our measure of outside options at the parent company level, matching sequentially on CIK, web domain, and bank name. The final sample includes approximately 140,000 distinct individuals.[10]

9. This includes all role names containing “Director”, “President”, “VP”, or “MD”, excluding those that contain “CEO”, “Chair”, “Trustee”, or “Partner”.

10. Egan, Matvos, and Seru construct the underlying data using publicly available records from the SEC’s Investment Advisor Public Disclosure website. For additional details, we refer to Egan, Matvos, and Seru’s homepage, Egan, Matvos, and Seru (2019, 2022), as well as Egan, Matvos, and Seru (2024) for a recent review of the literature.

2.4 Measuring Outside Options

This section describes our network-based measure of outside options, followed by a number of validation tests. We construct a measure that captures meaningful variation in outside options while remaining orthogonal to bank-level shocks. The identification approach exploits variation in labor demand at other institutions, transmitted to a bank's workers through their former coworkers, which is plausibly unrelated to the bank's current situation. The role of social networks in shaping labor market outcomes has been extensively documented, both theoretically (Montgomery 1991; Dustmann et al. 2016) and empirically (Bramoullé and Saint-Paul 2010; Cingano and Rosolia 2012; Shue 2013; Hensvik and Skans 2016; Gee, Jones, and Burke 2017; Glitz 2017; Saygin, Weber, and Weynandt 2021). Our measure is built from LinkedIn data on employee job mobility, drawn from public profiles.

At its core, our outside options measure counts new hires at other banks, weighted by the share of former coworkers now employed there, capturing differential exposure to signals about labor market dynamics across banks. More formally, for bankers at a given bank i in quarter t, our measure of outside options is defined as follows:

$$InflowsAtTies_{it} = \sum_{j \neq i} w_{ijt} \times Inflows_{-i,jt} \tag{2.8}$$

Thus, $InflowsAtTies_{it}$ is the weighted sum of inflows at other banks j, denoted by $Inflows_{-i,jt}$. The latter term is defined as job-to-job transitions from all other banks $k \neq i$ to bank j in quarter t:

$$Inflows_{-i,jt} = \sum_{k \neq i} Transitions_{kjt}. \tag{2.9}$$

We exclude transitions from bank i to j in t from our calculation in order to ensure orthogonality between our outside options measure and any bank-level dynamics in bank i over time.

The weights measure the share of former coworkers now employed at a given bank. More specifically, the weight bank i attributes to j at time t counts the number of past transitions from the former to the latter, normalized by the number of past transitions to all other banks:

$$w_{ijt} = \frac{\sum_{h=4}^{12} Transitions_{ij,t-h}}{\sum_{m \neq i} \sum_{h=4}^{12} Transitions_{im,t-h}}. \tag{2.10}$$

The weights thus capture the social ties a bank's workers have to other banks, and hence their exposure to hiring at those institutions.[11] In our baseline specification,

11. The LinkedIn dataset does not include information about the profile connections of LinkedIn profiles. We are therefore unable to base our network measure on social ties as indicated by these

we focus on ties that are between 4 and 12 quarters old. We limit the tie duration to 3 years, given that the median spell length in our sample is around 13 quarters. We also include ties only one year after the transition to ensure that outflows from a bank do not contemporaneously affect the tie weight component of our measure. A bank's contemporaneous outflows therefore enter neither the labor-demand nor the tie-weight component of the measure. We demonstrate the robustness of our main results to alternative assumptions in the appendix.

The measure $InflowsAtTies_{it}$ was initially proposed by Caldwell and Harmon (2019) and has several properties that make it well-suited to our setting. First, it captures firm-level variation in signals about labor demand, shifting workers' beliefs about their outside options. Prior work shows that such beliefs affect not only the job search and bargaining intentions of workers (Jäger et al. 2024) but also their search and bargaining outcomes (Caldwell and Harmon 2019). We expect this information transmission channel to be relevant, given the importance of professional (Godechot 2014) and alumni networks (McNamara and McLoughlin 2009) and the widespread use of LinkedIn in the financial industry (Fracassi, Petry, and Tate 2016; Jiang, Wang, and Wang 2018). In many cases, senior positions remain unadvertised, increasing banks' reliance on established network connections and employee referrals. Second, the weighting approach provides an intuitive method of accounting for a bank *i*'s worker placement record in the recent past. Workers should place more weight on labor demand at banks that are achievable for them. Because we weight inflows by the share rather than the number of ties, bank size and the corresponding size of the coworker network do not mechanically inflate the proxy. Third, the measure relies exclusively on variation in employment flows across institutions. Alternative approaches (e.g., Caldwell, Haegele, and Heining 2024) based on worker concentration across industries or occupations are unsuitable in the context of a single occupational group. We also caution against relying solely on spatial variation, given the concentration of senior investment bankers in a small number of financial hubs as well as potential demand-side effects (Ee, Huang, and Cheng 2023; Ee and Huang 2024).

To construct the measure, we identify senior investment bankers as those with at least one employment spell that meets the exclusion criteria outlined in Section 2.3. Next, we retrieve the employment histories of these individuals to trace their connections to other institutions. We also extract all senior investment banking spells at those other banks. For the hiring dynamics, we focus on job-to-job transitions, which we define as those with employment gaps of one month or less. In a given quarter, we observe on average 317 bankers per institution, with 24 connected

connections. Instead, we assume that ties form between senior investment bankers working at the same institution at the same time, and that those ties persist for a limited time after one banker moves to another bank.

Figure 2.1. Outside Options and Outflows to Other Banking Roles

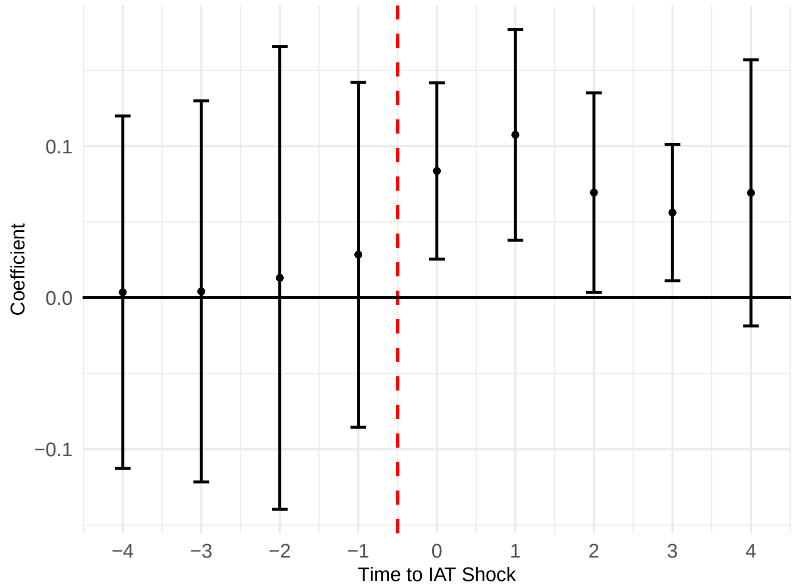

Notes: The figure shows coefficients from regressions of employee outflows to other banks in our sample on inflows at ties, using the Poisson estimator for dispersed count variables with meaningful zeros (Wooldridge 1999). The coefficients come from separate regressions on a series of leads and lags of inflows at ties, all including lender and year-quarter fixed effects. The inflows at ties measure is standardized to a mean of zero and a standard deviation of one. The figure shows 95% intervals calculated using two-way clustered standard errors at the lender and year-quarter level.

banks, each of which hires 7 bankers.[12] Based on this sample of spells, we compute $InflowsAtTies_{it}$ in the way described above and standardize the resulting measure to have a mean of zero and a standard deviation of one. As we plot in Figure 2.A.2, the outside options measure exhibits a procyclical pattern over time, closely co-moving with the credit-to-GDP ratio and the number of employees in credit intermediation.

We validate the relevance of the measure by replicating the findings of Caldwell and Harmon (2019) on the effects of outside options on mobility and wages. To examine the relationship between outside options and employee mobility, we regress employee outflows from a given institution to other banks in our sample on our proxy for outside options, including bank and year-quarter fixed effects. We estimate separate models for each lead and lag of the outside options measure. The results, presented in Figure 2.1, show that outside options significantly increase employee outflows both on impact and in subsequent periods. This persistence is consistent with news shocks shifting beliefs and influencing search behavior over the medium term. Importantly, the coefficients on the leads are insignificant, indicating that elevated past outflows do not predict our outside options proxy. In Fig-

12. We refer to Table 2.B.5 for summary statistics.

ure 2.A.3, we replicate the analysis using all employee outflows as the dependent variable, with results closely mirroring those in Figure 2.1. The pattern also remains unchanged for alternative timing assumptions for the construction of $InflowsAtTies_{it}$ (Figure 2.A.4).

Second, we examine whether better outside options translate into pay increases for employees who remain with their institution. To this end, we regress symmetric growth rates of bank-level compensation components on our outside options proxy, controlling for lender and year fixed effects. Table 2.1 presents the results. Columns (1) and (2) show that improved outside options are associated with substantial increases in median salary growth, with and without lender controls. Columns (3) to (6) consider other key pay components, bonuses and equity at risk, and find no significant response to variation in outside options. This contrast indicates that, conditional on controls, the proxy does not capture cyclical variation in banks' financial position, as we would otherwise expect more volatile pay components to respond more strongly. It also aligns with prior evidence that workers primarily bargain over wages rather than bonuses or amenities (Caldwell, Haegele, and Heining 2024). Overall, the results suggest that banks raise fixed pay to retain talent when outside options improve, confirming that our proxy reflects genuine variation in outside options.

In the main analysis, we regress loan-, banker-, and bank-level outcomes on our outside options measure. We can attribute a causal interpretation to the estimated coefficients if, conditional on covariates and fixed effects, unobserved determinants of lending behavior or bank risk are uncorrelated with variation in labor demand at connected banks. As a shift-share variable, $InflowsAtTies_{it}$ inherits the exogeneity conditions of such designs, where the shares are the weights w_{ijt} and the shocks are labor demand $Inflows_{-i,jt}$ at connected banks. While most of the literature exploits exogenous shares, we follow the exogenous-shifts approach (Borusyak, Hull, and Jaravel 2025), which only imposes orthogonality conditions on the shocks. This distinction is relevant to our setting, since a bank's network may be correlated with its characteristics while identification remains valid as long as the hiring shocks transmitted through it are not. Formally, two conditions have to be satisfied (Borusyak, Hull, and Jaravel 2022). First, the underlying shocks must be as-good-as-randomly assigned, without requiring exogeneity of the shares. Second, identification hinges on sufficient variation in the shocks, which arises when they are many and relatively independent, each affecting only a small number of units.

The first condition is plausibly met in our setting, as labor demand shocks at other institutions are, conditional on controls, unlikely to be correlated with unobservable factors that determine a bank's outcomes. To provide direct evidence, we regress inflows at ties on balance-sheet financials from Compustat and find that bank observables do not predict the measure. Conditional on fixed effects, future inflows at ties are unrelated to a bank's financial position, both in levels (Table 2.B.6) and in growth rates (Table 2.B.7).

Table 2.1. Outside Options and Pay Growth (in Percentage Points)

	Salary		Bonus		Equity At Risk	
	(1)	(2)	(3)	(4)	(5)	(6)
IAT	0.081***	0.078***	-0.177	-0.182	0.082	0.082
	(0.029)	(0.029)	(0.123)	(0.123)	(0.094)	(0.096)
Observations	8,462	8,462	8,462	8,462	8,462	8,462
Lender Controls		✓		✓		✓
Lender fixed effects	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓

Notes: The table shows regressions of symmetric growth rates of median firm-level compensation on $InflowsAtTies_{it}$, lender controls, and fixed effects at the lender and year-quarter level. The dependent variables are salary (Columns 1 and 2), bonus payments (Columns 3 and 4), and equity at risk (Columns 5 and 6). Lender controls are the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, and equity. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

We similarly expect the second condition to hold. The transition matrix across banks is sparse and mobility between institutions is heterogeneous, generating substantial variation in exposure across banks. As shown in Table 2.B.5, the average bank is connected to 24 other institutions in a given quarter. This number drops by about half when adjusting for size heterogeneity or restricting the sample to connected banks with positive inflows. At the same time, the total number of connected banks with non-zero hiring stands at 367 and has increased substantially over time (Figure 2.A.5). Therefore, an average bank is only connected to 7% of all hiring institutions, implying that individual banks' shocks are unlikely to dominate the measure. Next, we examine the extent to which hiring at other banks can be predicted by characteristics of the bank of interest. To do so, we regress hiring at connected banks on various sets of fixed effects and inspect the resulting R^2, with results presented in Table 2.B.8. $Bank_i$ fixed effects account for only a small share of hiring variation at connected banks, and their marginal contribution becomes negligible once fixed effects for the connected bank are included. Moreover, even the most comprehensive specification leaves roughly one-third of the variation in hiring unexplained.

2.5 Main Results

In this section, we examine the effect of our outside options measure on credit supply. We document that improved outside options are associated with higher credit volumes and increased financial risk at the lender level, thereby contributing to elevated systemic risk. We then demonstrate that these aggregate patterns reflect a shift toward riskier borrowers.

2.5.1 Lender-Level Outcomes

We estimate the response of lender-level variables to variation in inflows at ties. As discussed in Section 2.2, improved outside options are expected to result in a credit expansion accompanied by increases in bank risk. To test this empirically, we use the following specification for a bank i at year-quarter t:

$$y_{it} = \beta InflowsAtTies_{it} + \Gamma V_{it} + \alpha_i + \lambda_t + \epsilon_{it} \tag{2.11}$$

where y_{it} is the outcome of interest and V_{it} is a vector of time-varying controls, comprising the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, and equity. The terms α_i and λ_t are fixed effects at the bank and year-quarter level. Standard errors are double-clustered at the lender and year-quarter level.

We assess the effects of outside options on lending volumes constructed from the syndicated loan data, with results presented in Table 2.2. In Columns (1) and (2), we estimate the effect of inflows at ties on loan volume growth. In line with our prediction, the estimated coefficients are positive and significant, both with and without time-varying lender controls. Quantitatively, a one standard-deviation increase in outside options corresponds to a 1.9 percentage point rise in loan volume growth. To further dissect this effect, we distinguish between extensive and intensive margin responses. Columns (3) and (4) show that the likelihood of a bank serving as lead arranger rises when outside options improve, with an average marginal effect of 2.8 percentage points in our preferred specification in Column (4). In contrast, the coefficients on loan size, conditional on issuing any loan, are small and statistically insignificant. The volume response thus operates through banks arranging more loans rather than larger ones. As shown in Tables 2.B.9 and 2.B.10, our results are robust to varying both the activation lag and the duration for which ties remain active.

Next, we examine the effect of bankers' outside options on financial risk at the lending bank and its implications for systemic risk. Because credit risk accounts for a substantial share of the variation in banks' return characteristics (Begenau, Piazzesi, and Schneider 2025), riskier lending should translate into elevated risk at the bank level. To measure financial risk, we draw on the metrics for both types of risk developed in Adrian and Brunnermeier (2016). Financial risk for a bank in isolation is captured by *Value at Risk* (VaR), defined as the maximum loss to bank i at a given confidence level. A bank's contribution to systemic risk is measured by CoVaR, defined as the change in the VaR of the financial sector conditional on a corresponding percentile loss to bank i. In other words, CoVaR reflects the non-causal contribution of a bank to systemic risk. We use the estimates provided by Adrian and Brunnermeier (2016) at the 95% and 99% confidence levels, transformed into symmetric growth rates, and re-estimate Equation (2.11).

Table 2.2. Outside Options and Lending Volumes

	Growth Rate		Extensive Margin		Log Volume	
	(1)	(2)	(3)	(4)	(5)	(6)
IAT	1.880*	1.900*	0.226*	0.243**	0.010	0.007
	(1.010)	(1.020)	(0.119)	(0.119)	(0.028)	(0.027)
Observations	6,277	6,277	3,999	3,999	4,058	4,058
Lender Controls		✓		✓		✓
Lender fixed effects	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓
Estimation	OLS	OLS	Logit	Logit	OLS	OLS

Notes: The table reports results from regressing the outcome of interest on $InflowsAtTies_{it}$, lender controls, and fixed effects at the lender and year-quarter level. The outcome variables are symmetric growth rates of total lending volume in percentage points (Columns 1 and 2), an indicator equal to one if a bank arranged any loan (Columns 3 and 4), and the natural logarithm of the loan amount, conditional on lending a strictly positive amount (Columns 5 and 6). Lender controls are the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, and equity. Columns (1), (2), (5), and (6) are estimated via OLS, while Columns (3) and (4) are obtained from logit regressions. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The results are reported in Table 2.3. Columns (1) to (4) present the estimates for VaR and CoVaR at the 95% confidence level, where the coefficients are positive but statistically insignificant. In Columns (5) to (8), we shift focus to more extreme levels of financial distress by using VaR and CoVaR at the 99% confidence level as outcomes. The estimated coefficients roughly double in magnitude and become significant at the 1% level. A one standard deviation improvement in outside options is associated with increases in VaR and CoVaR of approximately 0.45 and 0.40 percentage points, respectively. This pattern suggests that the additional risk is concentrated in the tail of the loss distribution, where the implications for financial stability are greatest. Results for alternative tie lengths, reported in Tables 2.B.11 and 2.B.12, provide support for these conclusions.

Crucially, these effects are not confined to the originating bank under scrutiny but extend to the broader financial system, pointing to systemic implications of bankers' labor market conditions. To gauge the magnitude of the estimates, consider that between 2003 and 2008 outside options rose by more than two standard deviations, while CoVaR increased by approximately 52%. A back-of-the-envelope calculation, extrapolating our marginal estimate over this range, attributes around 1.8% of the rise in systemic risk to changes in outside options. The channel is thus economically meaningful.

Table 2.3. Outside Options and Financial Risk (Changes in Percentage Points)

	VaR 95%		CoVaR 95%		VaR 99%		CoVaR 99%	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IAT	0.247	0.246	0.190	0.187	0.452***	0.447***	0.408***	0.401***
	(0.254)	(0.253)	(0.195)	(0.192)	(0.162)	(0.161)	(0.129)	(0.126)
Observations	10,635	10,635	10,635	10,635	10,635	10,635	10,635	10,635
Lender Controls		✓		✓		✓		✓
Lender fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓	✓	✓

Notes: The table reports coefficients from OLS regressions of symmetric growth rates in VaR and CoVaR (in percentage points) on $InflowsAtTies_{it}$, lender controls, and fixed effects at the lender and year-quarter level. The outcomes are VaR (Columns 1, 2, 5, and 6) and CoVaR (Columns 3, 4, 7, and 8), computed at the 95% confidence level (Columns 1 to 4) and the 99% confidence level (Columns 5 to 8). Lender controls are the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, and equity. The data are at quarterly frequency and span 1994 to 2009. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

2.5.2 Borrower Risk

We scrutinize the loan-level dynamics underlying these bank-level patterns. Consider a bank i issuing a loan l to borrower b in year-quarter t. The loan-level specification is given by:

$$y_{bt} \; = \; \beta \, InflowsAtTies_{it} + \Gamma V_{it} + \Theta W_{bt} + \Pi Z_l + \alpha_i + \lambda_t + \epsilon_{iblt}, \qquad (2.12)$$

where y_{bt} is a measure of credit risk, proxied by the borrower's credit rating. As before, we regress the outcome on inflows at ties, with lender-level controls and fixed effects at the lender and year-quarter level. The loan-level specification additionally allows us to control for the same balance-sheet variables on the borrower side, W_{bt}, and for loan-level covariates Z_l, which comprise the inverse hyperbolic sine of the tranche amount and maturity as well as indicators for whether the loan is secured and whether covenants are attached.

The results are presented in Table 2.4. In Columns (1) to (3), the outcome is a borrower-level indicator for an investment-grade rating by S&P, while in Columns (4) to (6), it is an ordinal credit rating normalized to range between 0 and 1. The specifications in Columns (1) and (4) mirror those at the bank level, but the estimated coefficients are statistically insignificant. Once we include borrower controls in Columns (2) and (5), the coefficients increase in absolute magnitude and become highly significant, while controlling for loan-level characteristics in Columns (3) and (6) does not substantially alter the estimates. The borrower controls capture firm size, which is in turn correlated with investment-grade status (Table 2.B.13). In the absence of these controls, the coefficient masks two offsetting effects: a shift toward larger firms and, conditional on firm size, toward riskier ones. Controlling

Table 2.4. Outside Options and Borrower Risk

	Borrower Prime			Borrower Rating		
	(1)	(2)	(3)	(4)	(5)	(6)
IAT	-0.002	-0.005***	-0.005***	0.000	-0.002**	-0.002**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	55,327	55,327	55,327	55,327	55,327	55,327
Outcome Mean	0.296	0.296	0.296	0.462	0.462	0.462
Lender Controls	✓	✓	✓	✓	✓	✓
Borrower Controls		✓	✓		✓	✓
Loan Controls			✓			✓
Lender fixed effects	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓

Notes: The table reports results from estimating Equation (2.12) at the loan level. The dependent variables are an indicator for investment-grade borrowers (Columns 1 to 3) and an ordinal borrower rating normalized to range between 0 and 1 (Columns 4 to 6). The independent variable is $InflowsAtTies_{it}$. Borrower and lender controls include the inverse hyperbolic sine transformation of total assets, number of employees, capital expenditure, total liabilities, EBIT, and (for borrowers) fixed assets. Loan controls are the inverse hyperbolic sine transformation of tranche amount and maturity and indicators for whether the loan is secured and whether covenants are attached. All regressions include fixed effects at the lender and year-quarter level, and standard errors are double-clustered at the same levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

for size isolates the latter effect, which is quantitatively meaningful. A one standard deviation increase in outside options reduces the share of investment-grade borrowers by 1.7% of its mean. Together with the bank-level results, this points to a credit expansion accompanied by a deterioration in borrower quality.

We perform a series of robustness checks for our main result in Column (3), reported in Tables 2.B.14 and 2.B.15. First, the coefficient of interest remains stable when we include leads and lags of inflows at ties. The estimated coefficients also remain significant and comparable in magnitude when accounting for time-varying demand-side factors (Huang et al. 2024; Huang et al. 2025) by adding fixed effects for borrowers' SIC sectors and states, each interacted with year-quarters. Similarly, our baseline findings are robust to using only within-lender-year variation, as well as to controlling for outflows or the number of lenders. The results also continue to hold when we restrict the sample to post-2008 loans or to newly originated loans. Lastly, we show that our main results are robust to alternative assumptions about the duration for which ties remain active (Tables 2.B.16 and 2.B.17).

A first-order question is whether this shift reflects a genuine increase in risk-taking rather than yield-seeking behavior. In case of the latter, the decline in ratings should be compensated by higher spreads. We test this by regressing the log all-in-spread-drawn, the standard measure of loan prices, on inflows at ties, conditional on the same controls and fixed effects as in the baseline specification. As reported in Table 2.B.18, the estimated coefficients are small and statistically insignificant, allowing us to rule out spread increases larger than roughly four basis points per standard-deviation shock, which amounts to about 1% of the outcome mean, at the 95% confidence level. Banks thus take on riskier borrowers without demanding significantly higher compensation, which is difficult to reconcile with yield-seeking.

Another concern is that stronger ties between institutions may increase the flow of information about a banker's current on-the-job behavior to other employers. Former colleagues or individuals who have interacted with the banker may be better positioned to assess their performance and provide hiring recommendations. One might therefore worry that bankers behave more cautiously when well-connected firms are hiring, as these firms are better able to observe current risk-taking. From a theoretical standpoint, however, such an effect should make bankers more cautious, meaning the information flow channel would counteract our main mechanism. As a result, any bias would attenuate our estimated coefficients toward zero. We test the channel directly by leveraging the syndication process, through which banks become linked via repeated co-arrangement of loans. We construct a variant of our measure in which ties are based on syndication links rather than coworker networks, capturing inter-bank information flow without the labor-mobility content. Including both measures in Column (7) of Table 2.B.15 leaves the coefficient of interest unaffected, indicating that information flows on the employers' side do not drive our results.

We provide more direct evidence for our main mechanism by investigating whether the response of borrower ratings to outside options depends on the ease of job-to-job transitions. We focus on changes in the enforceability of non-compete agreements (NCAs), which restrict employees' ability to work for direct competitors in the same region and have been shown to diminish worker mobility and wages (Johnson, Lavetti, and Lipsitz 2025). Crucially, their enforceability varies at the state level and is shifted by state court decisions that clarify or overturn precedent, providing plausibly exogenous variation in labor mobility. When enforceability decreases in the state where a bank is headquartered, its bankers become less constrained in exercising their outside options. We therefore expect the effect of outside options on borrower risk to be stronger following such decreases, and conversely weaker following increases. To test this, we interact our measure with the state-level NCA enforceability index of Mueller (2023), assigning each bank to its headquarters state and including state-year fixed effects to absorb time-varying geographic heterogeneity. The results in Table 2.B.19 are broadly in line with this prediction. The interaction of inflows at ties with past decreases in enforceability (Column 3) is negative and highly significant, indicating that states moving toward weaker enforceability drive the effect. By contrast, the interactions for increases in enforceability and a combined measure are insignificant.

The estimated coefficients reported in this section are moderate in magnitude but can be considered lower-bound estimates for several reasons. First, by estimating effects at the bank level, we average across all investment bankers, including those who are unresponsive to outside options, whether owing to low observed productivity, niche specialization, or strong firm-specific attachments (Caldwell, Haegele, and Heining 2025). Second, our identification strategy removes cyclical variation in outside options, which represents a substantial share of the overall variation and likely the component to which bankers respond most.[13] Third, our measure does not capture news about labor demand at other firms that is transmitted outside of personal networks, such as public job postings, positions at non-connected banks, or opportunities outside the financial industry. Taken together, our results reflect the average response to variation in one plausibly exogenous component of outside options.

2.6 Mechanism

So far, we have provided evidence that better outside options increase bank risk by inducing a credit expansion toward riskier borrowers. To assess whether the mechanism introduced in Section 2.2 can account for these effects, we turn to individual-level data. We present several pieces of evidence that, taken together, are difficult to reconcile with alternative explanations, such as behavioral factors (Cohn, Fehr, and Maréchal 2014; Cohn et al. 2015; Bordalo, Gennaioli, and Shleifer 2018) or residual demand-side forces, as the primary drivers of our results. Specifically, we show that individuals more exposed to dismissal risk respond more strongly to variation in outside options. We also document that outside options lead to behavior that harms both employers and clients, and reduce banks' ability to discipline employees for misconduct, consistent with the implicit insurance channel.

2.6.1 Banker-Level Heterogeneity

We assess banker-level heterogeneity in how outside options affect borrower risk, focusing on age for two key reasons. First, age can be reasonably inferred from employment history data. Second, and more importantly, age has a theoretically meaningful interpretation in the context of risk choices and outside options. The classical argument of career concerns is that such considerations discourage risky or shirking behavior, as employees have an incentive to signal desirable traits to employers. Under this framework, labor markets serve as a disciplining device and mitigate agency problems within firms (Shapiro and Stiglitz 1984; Holmström 1999).

In our case, however, the differential response to variation in outside options across age groups is theoretically ambiguous. On the one hand, because improvements in outside options reduce the cost of job loss, younger individuals should be more responsive to the implicit insurance they provide and make riskier choices, as illustrated in Section 2.2. On the other hand, younger individuals also have a longer

13. Year-quarter fixed effects explain around 39% of the overall variation in inflows at ties, which further increases to 47% once controlling for bank fixed effects.

Table 2.5. Outside Options and Borrower Prime Rating: Age Heterogeneity

	Borrower Prime Rating						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IAT	-0.587***	-0.495**	-0.482**				
	(0.179)	(0.189)	(0.200)				
IAT × Age	0.010**	0.009***	0.009**	0.012***	0.012***	0.012***	0.012**
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)
Observations	275	275	275	327	327	327	327
Outcome Mean	0.153	0.153	0.153	0.174	0.174	0.174	0.174
Lender Controls	✓	✓	✓				
Borrower Controls		✓	✓	✓	✓	✓	✓
Loan Controls			✓	✓	✓	✓	✓
IAT × Spell Controls							✓
Lender fixed effects	✓	✓	✓				
Year-Quarter fixed effects	✓	✓	✓				
Banker fixed effects	✓	✓	✓	✓	✓	✓	✓
Lender × Year-Quarter fixed effects				✓	✓		✓
MSA × Year-Quarter fixed effects					✓		✓
Lender × MSA × Year-Quarter fixed effects						✓	

Notes: The table reports results from estimating Equation (2.12) at the loan level. The dependent variable is an indicator for investment-grade borrowers. The independent variables are $InflowsAtTies_{it}$ and its interaction with the banker's age. Borrower and lender controls include the inverse hyperbolic sine transformation of total assets, number of employees, capital expenditure, total liabilities, EBIT, and (for borrowers) fixed assets. Loan controls are the inverse hyperbolic sine transformation of tranche amount and maturity and indicators for whether the loan is secured and whether covenants are attached. *Spell controls* contain the position number, the seniority level, the inverse hyperbolic sine transformation of imputed salary and job tenure, each included interacted with $InflowsAtTies_{it}$ and (if possible) as a standalone variable. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

remaining career over which reputational damage can compound. If improvements in outside options are perceived as persistent, or if excessive risk-taking inflicts lasting reputational harm, these career-long consequences could outweigh the insurance benefit and make younger bankers more cautious instead.

To take these considerations to the data, we extend Equation (2.12) by interacting banker age with inflows at ties, using our sample of loans matched to individual bankers. All regressions include banker fixed effects, which, together with the other fixed effects, absorb the standalone coefficient on age. The results, presented in Table 2.5, show that outside options are consistently negatively associated with the likelihood of lending to prime-rated borrowers, while the interaction with age is large, positive, and highly significant. Despite the small sample size, this pattern is robust across Columns (1) to (3), which replicate the specifications in Table 2.4, and persists under more granular fixed effects. Column (4) adds lender fixed effects interacted with year-quarters, which absorb the standalone coefficient on inflows at ties, and Column (5) further includes time-varying metro-area fixed effects to account for local economic conditions.

A further concern is that bankers' labor market mobility is partly driven by liftouts, whereby senior bankers are poached and take their entire team with them.

Such a pattern would, if anything, imply stronger effects among senior bankers. The age interaction instead places the effect among younger bankers. We also address this concern directly in Column (6) by interacting lender-year-quarter fixed effects with metro areas, thus relying only on variation within the same bank across space, with the estimate remaining essentially unchanged. Finally, Column (7) interacts inflows at ties with spell-level controls, namely seniority, position number, job tenure, and imputed salary, and the age interaction remains stable. Tables 2.B.20 and 2.B.21 report the equivalent regressions with continuous rating and tranche amounts as outcomes, with consistent results.

For all outcomes, the coefficients on the interaction with age are of the opposite sign to the corresponding baseline coefficient, implying that the effects of outside options on lending are concentrated among younger employees. This is consistent with the idea that outside options act as insurance against downside risk. Typically, less experienced bankers have accumulated less firm-specific capital and smaller internal networks, making them more dependent on outside labor market conditions. Similarly, younger individuals may face tighter financial constraints in their private consumption decisions.

2.6.2 Evidence on Advisor Misconduct

Since our results suggest that better outside options lead to riskier behavior, this effect should manifest beyond the domain of loan officers' decisions. We turn to the market for investment advisors for three reasons. First, misconduct provides clear instances of wrongdoing, cases in which employees acted against their employer's interest. Credit decisions offer no such indication, since the optimal level of risk-taking is ex-ante unclear. The advisor setting is therefore a cleaner test of whether outside options induce harmful behavior. Second, as disclosures of misconduct are publicly available, this context allows us to examine how such behavior affects labor market outcomes. Finally, the results generalize our main finding on the effect of outside options on workers' job-related choices to the workforce in the financial industry more broadly. In doing so, we contribute to a growing literature studying the determinants of advisor misconduct (Egan, Matvos, and Seru 2019; Clifford and Gerken 2021; Dimmock, Gerken, and Van Alfen 2021).

We first investigate how variation in outside options influences the frequency of advisor misconduct, building on Equation (2.12). We aggregate inflows at ties to the annual frequency and focus on the first lag of the measure, as misconduct typically takes time to be discovered and disclosed. Our analysis considers three outcome variables: the number of misconduct disclosures an advisor receives in a given year, the number of customer complaints filed against them, and the number of complaints that result in a settlement. Given the within-firm spillovers in misconduct (Egan, Matvos, and Seru 2019), we control for the first lag of the firm-level average of the outcome variable.

Table 2.6. Outside Options and Misconduct Disclosures

	Misconduct		Complaints		Complaints (Settled)	
	(1)	(2)	(3)	(4)	(5)	(6)
IAT_{t-1}	0.003***	0.003***	0.005***	0.006***	0.004***	0.004***
	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)
Observations	582,568	551,079	582,568	551,079	582,568	551,079
Outcome Mean	0.010	0.009	0.018	0.017	0.008	0.008
Lender Controls	✓	✓	✓	✓	✓	✓
Lender fixed effects	✓		✓		✓	
Year fixed effects	✓	✓	✓	✓	✓	✓
Lender × Advisor fixed effects		✓		✓		✓

Notes: The table reports results from estimating Equation (2.12) at annual frequency at the advisor level. The dependent variables are the numbers of misconduct disclosures an advisor had in a given year (Columns 1 and 2), the number of complaints filed against the advisor (Columns 3 and 4), and the number of settled customer complaints (Columns 5 and 6). The main independent variable is $InflowsAtTies_{it}$. Lender controls are the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, and equity, together with the first lag of the firm-level average of the outcome variable. All models are estimated with year fixed effects and lender controls, with the latter replaced by lender-advisor fixed effects in Columns (2), (4), and (6). Standard errors are clustered at the lender level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The results are shown in Table 2.6. Across the board, better outside options correspond to large and significant increases in overall misconduct, customer complaints, and settlements. For each of these outcomes, this holds conditional on the baseline fixed effects and additional lender-advisor fixed effects. In Tables 2.B.22, 2.B.23, and 2.B.24, we provide a battery of robustness checks for these results. The findings remain unaffected by augmenting the regressions with state-year fixed effects, including the contemporaneous levels and the second lag of inflows at ties, controlling for lagged outflows, and binarizing the outcome variable. Quantitatively, these effects are sizeable, as a one standard-deviation improvement in outside options maps into an increase in misconduct disclosures of roughly 30%. The associated costs to the financial sector are substantial. With average settlement costs of 550,000 USD per case (Egan, Matvos, and Seru 2019), a one standard-deviation increase in inflows at ties corresponds to an annual increase in settlement costs of approximately 140 million USD for the firms in our sample.

Egan, Matvos, and Seru (2019) show that firms tend to punish misconduct among investment advisors. Employment separations become much more likely following related disclosures, even though the broader market partly offsets this penalty through rehiring. Our mechanism posits that outside options insure workers against adverse outcomes by lowering both the likelihood and the cost of separation, reducing the deterrent effect of job loss following misconduct. An advisor who anticipates this insurance and takes more risk responds rationally to weakened incentives, which distinguishes our mechanism from purely behavioral explanations such as overoptimism.

We test this conjecture by estimating the differential effect of misconduct on three distinct outcomes, comparing individuals with better versus worse outside options. Specifically, for an advisor a employed at bank i in year t, we define the following indicators for transitions between year t and $t+1$: (*i*) Separation$_{at}$, equal to one if the advisor is no longer employed at i; (*ii*) IndustryExit$_{at}$, equal to one if the advisor is no longer recorded in the sample, which we define as exiting the industry; and (*iii*) Reemployment$_{at}$, equal to one if the advisor is employed at another firm in the sample. We use two empirical specifications. First, we estimate the unconditional effects of misconduct, outside options, and their interaction on Separation$_{at}$ and IndustryExit$_{at}$. Second, conditional on experiencing a job separation, we estimate the differential effects on IndustryExit$_{at}$ and Reemployment$_{at}$.

For the unconditional effects, consider the following specification:

$$\begin{aligned} y_{at} = {} & \beta_1 Misconduct_{at} + \beta_2 Misconduct_{at} \times InflowsAtTies_{it} \\ & + \beta_3 Misconduct_{at} \times InflowsAtTies_{it-1} + \Gamma V_{it} + \alpha_{ai} + \lambda_{it} + \epsilon_{ait}, \end{aligned} \tag{2.13}$$

that is, we regress an outcome y_{at} on an indicator for misconduct, alone and interacted with contemporaneous values and the first lag of inflows at ties. We include fixed effects at the lender-year level, removing the baseline coefficient on outside options, and at the advisor or lender-advisor level. For the conditional effects, we further augment this setup by including all potential interactions with Separation$_{at}$, focusing on IndustryExit$_{at}$ and Reemployment$_{at}$ as outcome variables.

The results are reported in Table 2.7. While we are particularly interested in the interactions with lagged outside options, we control for contemporaneous inflows at ties and all potential interactions in all regressions. In Columns (1) and (2), we estimate the unconditional specification for the probability of a job separation, including advisor and lender-advisor fixed effects, respectively. As expected, misconduct enters with a large, positive, and highly significant coefficient, corroborating the results of Egan, Matvos, and Seru (2019). However, the estimated coefficient on the interaction with lagged inflows at ties is negative, highly significant, and sizeable: the probability of job separation following misconduct is reduced by around 22% relative to the baseline when inflows at ties are one standard deviation above the mean. Thus, better outside options lower the probability of experiencing a job separation following a misconduct disclosure, providing evidence of employers' reduced disciplining power.

In Columns (3) and (4), we repeat the estimation using IndustryExit$_{at}$ as the outcome. Again, better outside options counteract the main effect, reducing the probability of leaving the industry following a misconduct disclosure. This pattern could be driven by both changes in job separations and hiring at other firms. To isolate the latter channel, we re-estimate Equation (2.13) and interact the variables of interest

Table 2.7. Outside Options and Labor Market Disciplining

	Job Separation		Leave Industry				Reemployment Next Year	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Misconduct	0.425***	0.443***	0.237***	0.239***	0.004	0.003	-0.006	-0.004
	(0.052)	(0.050)	(0.029)	(0.029)	(0.009)	(0.007)	(0.011)	(0.010)
Misconduct × IAT_t	0.015	0.018	0.018	0.018	-0.0001	-0.0006	0.002	0.003
	(0.025)	(0.022)	(0.017)	(0.016)	(0.004)	(0.003)	(0.004)	(0.003)
Misconduct × IAT_{t-1}	-0.075***	-0.086***	-0.054***	-0.055***	0.001	0.002	-0.002	-0.003
	(0.023)	(0.019)	(0.013)	(0.013)	(0.004)	(0.004)	(0.005)	(0.005)
Job Separation × Misconduct					0.133**	0.127*	-0.140**	-0.127*
					(0.058)	(0.063)	(0.068)	(0.072)
Job Separation × IAT_t					-0.015	0.002	0.028*	0.014
					(0.016)	(0.017)	(0.016)	(0.015)
Job Separation × IAT_{t-1}					-0.040***	-0.033***	0.032***	0.022**
					(0.008)	(0.007)	(0.010)	(0.009)
Job Separation × Misconduct × IAT_t					0.053*	0.043	-0.057	-0.051
					(0.027)	(0.029)	(0.035)	(0.035)
Job Separation × Misconduct × IAT_{t-1}					-0.068***	-0.061***	0.077***	0.072***
					(0.015)	(0.021)	(0.020)	(0.025)
Observations	598,839	589,763	598,839	589,763	598,839	589,763	598,839	589,763
Outcome Mean	0.119	0.109	0.038	0.037	0.038	0.037	0.079	0.071
Lender × Year fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Advisor fixed effects	✓		✓		✓		✓	
Lender × Advisor fixed effects		✓		✓		✓		✓

Notes: The table reports results from estimating versions of Equation (2.13) at annual frequency at the advisor level. The dependent variables are an indicator equal to one if the advisor is no longer employed at i in year $t+1$ (Columns 1 and 2), an indicator equal to one if the advisor is no longer recorded in the sample (Columns 3 to 6), and an indicator equal to one if the advisor is employed at another firm in the sample in $t+1$ (Columns 7 and 8). The main independent variables are $\mathit{InflowsAtTies}_{it}$, $\mathit{InflowsAtTies}_{it-1}$, and an indicator for having a misconduct disclosure in year t. The specification is given by Equation (2.13) in Columns (1) to (4), with all variables interacted with the indicator for job separations in Columns (5) to (8). All models include fixed effects at the lender-year and at the advisor (Columns 1, 3, 5, and 7) or lender-advisor level (Columns 2, 4, 6, and 8). Standard errors are clustered at the lender level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

with Separation_{at}. The results, included in Columns (5) and (6), point to an important role for hiring. As before, misconduct-related separations raise the probability of exiting the industry, but this effect is muted for individuals with above-average outside options. Lastly, we find consistent results for Reemployment_{at}, an indicator for working at a different employer within the financial sector in the following year, reported in Columns (7) and (8).

Taken together, these results suggest that outside options reduce the probability of separation following misconduct and, even conditional on separation, make reemployment more likely. Outside options thus reduce the downside of risky behavior, thereby offsetting part of the disciplining role of labor markets.

2.7 Conclusion

This paper documents that better labor market outside options for bankers lead to riskier lending behavior and contribute to fluctuations in systemic risk. Using transaction-level data on syndicated loans, we find that banks exposed to improved outside options expand lending volumes by shifting toward riskier borrowers, increasing financial risk at both the bank and system levels. We argue that these effects arise because stronger outside options provide implicit insurance against adverse outcomes, weakening the disciplinary role of dismissal risk. We show that these patterns persist in a related setting, the market for financial advisors, where better outside options raise misconduct and attenuate the labor market penalties of such violations.

These findings have implications for market participants and policymakers alike. Our results suggest that variation in sectoral labor market conditions affects the strength of employers' disciplining power, with consequences for operational risk and settlement costs. During hot labor markets, banks may benefit from strengthening incentive alignment to limit the scope for adverse on-the-job behavior, for instance by increasing the use of deferred compensation and clawback provisions as well as by setting stricter risk limits. Yet an individual bank's ability to do so is limited. The disciplinary value of these tools erodes precisely when outside options improve, since better options shorten expected tenure, and a bank that defers compensation more aggressively becomes less attractive to talent when rivals are hiring. The mechanism we document thus operates through a labor market externality that individual banks do not internalize and cannot offset in isolation. As a result, our findings point to a potential role for coordination. Minimum sector-wide standards on the time horizon of compensation, implemented through deferred pay and clawback provisions, could help address the externality at the level at which it arises.

References

Acharya, Viral, Marco Pagano, and Paolo Volpin. 2016. "Seeking Alpha: Excess Risk Taking and Competition for Managerial Talent." *Review of Financial Studies* 29 (10): 2565–99. [122]

Adrian, Tobias, and Markus K. Brunnermeier. 2016. "CoVaR." *American Economic Review* 106 (7): 1705–41. [129, 135, 157]

Agarwal, Sumit, Yupeng Lin, Yunqi Zhang, and Zilong Zhang. 2024. "Labor Mobility and Loan Origination." *Journal of Financial and Quantitative Analysis* 59 (5): 2099–132. [122]

Ahammer, Alexander, Matthias Fahn, and Flora Stiftinger. 2023. "Outside Options and Worker Motivation." Working Paper, JKU Linz. [122]

Axelson, Ulf, and Philip Bond. 2015. "Wall Street Occupations." *Journal of Finance* 70 (5): 1949–96. [122]

Azzopardi, Damien, Fozan Fareed, Mikkel Hermansen, Patrick Lenain, and Douglas Sutherland. 2020. "The Decline in Labour Mobility in the United States: Insights from New Administrative Data." Working Papers, 1644, OECD Economics Department. [119]

Begenau, Juliane, Monika Piazzesi, and Martin Schneider. 2025. "Banks' Risk Exposures." Working Paper, 21334, National Bureau of Economic Research. [135]

Bell, Brian, and John Van Reenen. 2014. "Bankers and Their Bonuses." *Economic Journal* 124 (574): F1–F21. [126]

Bianchi, Javier. 2016. "Efficient Bailouts?" *American Economic Review* 106 (12): 3607–59. [120]

Böhm, Michael J, Daniel Metzger, and Per Strömberg. 2023. ""Since You're So Rich, You Must Be Really Smart": Talent, Rent Sharing, and the Finance Wage Premium." *Review of Economic Studies* 90 (5): 2215–60. [126]

Bonelli, Maxime. 2019. "Labor Mobility and Capital Misallocation in the Mutual Fund Industry." Working Paper, University of Rochester. [122]

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer. 2024. "Belief Overreaction and Stock Market Puzzles." *Journal of Political Economy* 132 (5): 1450–84. [126]

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2018. "Diagnostic Expectations and Credit Cycles." *Journal of Finance* 73 (1): 199–227. [140]

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2022. "Quasi-Experimental Shift-Share Research Designs." *Review of Economic Studies* 89 (1): 181–213. [133]

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2025. "A Practical Guide to Shift-Share Instruments." *Journal of Economic Perspectives* 39 (1): 181–204. [133, 158]

Bramoullé, Yann, and Gilles Saint-Paul. 2010. "Social Networks and Labor Market Transitions." *Labour Economics* 17 (1): 188–95. [130]

Brandao-Marques, Luis, Ricardo Correa, and Horacio Sapriza. 2013. "International Evidence on Government Support and Risk Taking in the Banking Sector." Working Paper, 1086, Board of Governors of the Federal Reserve System. [120]

Bushman, Robert, Janet Gao, Xiumin Martin, and Joseph Pacelli. 2021. "The Influence of Loan Officers on Loan Contract Design and Performance." *Journal of Accounting and Economics* 71 (2-3): 101384. [122, 126, 127]

Caldwell, Sydnee, and Oren Danieli. 2024. "Outside Options in the Labour Market." *Review of Economic Studies.* [119, 122]

Caldwell, Sydnee, Ingrid Haegele, and Jörg Heining. 2024. "Bargaining and Inequality in the Labor Market." Working Paper, 33396, National Bureau of Economic Research. [122, 131, 133]

Caldwell, Sydnee, Ingrid Haegele, and Jörg Heining. 2025. "Firm Pay and Worker Search." Working Paper, 33445, National Bureau of Economic Research. [140]

Caldwell, Sydnee, and Nikolaj Harmon. 2019. "Outside Options, Bargaining, and Wages: Evidence from Coworker Networks." Working Paper, Copenhagen University. [119, 122, 131, 132]

Carvalho, Daniel, Janet Gao, and Pengfei Ma. 2023. "Loan Spreads and Credit Cycles: The Role of Lenders' Personal Economic Experiences." *Journal of Financial Economics* 148 (2): 118–49. [122, 126]

Chava, Sudheer, and Michael R Roberts. 2008. "How Does Financing Impact Investment? The Role of Debt Covenants." *Journal of Finance* 63 (5): 2085–121. [129]

Cici, Gjergji, Mario Hendriock, and Alexander Kempf. 2021. "The Impact of Labor Mobility Restrictions on Managerial Actions: Evidence from the Mutual Fund Industry." *Journal of Banking & Finance* 122: 105994. [122]

Cingano, Federico, and Alfonso Rosolia. 2012. "People I Know: Job Search and Social Networks." *Journal of Labor Economics* 30 (2): 291–332. [130]

Clifford, Christopher P, and William C Gerken. 2021. "Property Rights to Client Relationships and Financial Advisor Incentives." *Journal of Finance* 76 (5): 2409–45. [142]

Cohn, Alain, Jan Engelmann, Ernst Fehr, and Michel André Maréchal. 2015. "Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals." *American Economic Review* 105 (2): 860–85. [140]

Cohn, Alain, Ernst Fehr, and Michel André Maréchal. 2014. "Business Culture and Dishonesty in the Banking Industry." *Nature* 516 (7529): 86–89. [140]

Curello, Gregorio, Ludvig Sinander, and Mark Whitmeyer. 2025. "Outside Options and Risk Attitude." Working Paper, University of Mannheim. [123]

Dam, Lammertjan, and Michael Koetter. 2012. "Bank Bailouts and Moral Hazard: Evidence from Germany." *Review of Financial Studies* 25 (8): 2343–80. [120]

Demirgüç-Kunt, Asli, and Enrica Detragiache. 2002. "Does Deposit Insurance Increase Banking System Stability? An Empirical Investigation." *Journal of Monetary Economics* 49 (7): 1373–406. [120]

Diamond, Douglas W, and Raghuram G Rajan. 2002. "Bank Bailouts and Aggregate Liquidity." *American Economic Review* 92 (2): 38–41. [120]

Dimmock, Stephen G, William C Gerken, and Tyson Van Alfen. 2021. "Real Estate Shocks and Financial Advisor Misconduct." *Journal of Finance* 76 (6): 3309–46. [142]

Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Brücker. 2016. "Referral-Based Job Search Networks." *Review of Economic Studies* 83 (2): 514–46. [130]

Ee, Mong Shan, and He Huang. 2024. "Employee Mobility Constraints and Debt Choice." *Economics Letters* 234: 111467. [131]

Ee, Mong Shan, He Huang, and Mingying Cheng. 2023. "Do Labor Mobility Restrictions Affect Debt Maturity?" *Journal of Financial Stability* 66: 101121. [131]

Egan, Mark, Gregor Matvos, and Amit Seru. 2019. "The Market for Financial Adviser Misconduct." *Journal of Political Economy* 127 (1): 233–95. [129, 142–144]

Egan, Mark, Gregor Matvos, and Amit Seru. 2022. "When Harry Fired Sally: The Double Standard in Punishing Misconduct." *Journal of Political Economy* 130 (5): 1184–248. [129]

Egan, Mark, Gregor Matvos, and Amit Seru. 2024. "The Problem of Good Conduct among Financial Advisers." *Journal of Economic Perspectives* 38 (4): 193–210. [129]

Ellul, Andrew, Marco Pagano, and Annalisa Scognamiglio. 2020. "Career Risk and Market Discipline in Asset Management." *Review of Financial Studies* 33 (2): 783–828. [122]

Figueiredo, Ana. 2022. "Wage Cyclicality and Labor Market Sorting." *American Economic Review: Insights* 4 (4): 425–42. [121]

Fracassi, Cesare, Stefan Petry, and Geoffrey Tate. 2016. "Does Rating Analyst Subjectivity Affect Corporate Debt Pricing?" *Journal of Financial Economics* 120 (3): 514–38. [131]

Gao, Janet, Kristoph Kleiner, and Joseph Pacelli. 2020. "Credit and Punishment: Are Corporate Bankers Disciplined for Risk-Taking?" *Review of Financial Studies* 33 (12): 5706–49. [122]

Gao, Janet, Wenyu Wang, and Yufeng Wu. 2024. "Human Capital Mobility and Careers of M&A Advisors." *Review of Financial Studies* 37 (9): 2732–78. [122]

Gao, Janet, Wenyu Wang, and Xiaoyun Yu. 2024. "Big Fish in Small Ponds: Human Capital Migration and the Rise of Boutique Banks." *Management Science* 70 (11): 7829–50. [122]

Gee, Laura K, Jason Jones, and Moira Burke. 2017. "Social Networks and Labor Markets: How Strong Ties Relate to Job Finding on Facebook's Social Network." *Journal of Labor Economics* 35 (2): 485–518. [130]

Glitz, Albrecht. 2017. "Coworker Networks in the Labour Market." *Labour Economics* 44: 218–30. [130]

Godechot, Olivier. 2014. "Getting a Job in Finance: The Strength of Collaboration Ties." *European Journal of Sociology/Archives Européennes de Sociologie* 55 (1): 25–56. [131]

Griffin, John M, Samuel Kruger, and Gonzalo Maturana. 2019. "Do Labor Markets Discipline? Evidence from RMBS Bankers." *Journal of Financial Economics* 133 (3): 726–50. [122]

Hamdi, Naser, Ankit Kalda, and Avantika Pal. 2023. "External Labor Market Punishment in Finance." Working Paper, Indiana University. [122]

Hensvik, Lena, and Oskar Nordström Skans. 2016. "Social Networks, Employee Selection, and Labor Market Outcomes." *Journal of Labor Economics* 34 (4): 825–67. [130]

Herpfer, Christoph. 2021. "The Role of Bankers in the US Syndicated Loan Market." *Journal of Accounting and Economics* 71 (2-3): 101383. [122, 126–128]

Holmström, Bengt. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *Review of Economic Studies* 66 (1): 169–82. [122, 140]

Huang, Ruidi, James S Linck, Erik J Mayer, and Chris Parsons. 2024. "Can Human Capital Explain Income-Based Disparities in Financial Services." *Review of Financial Studies,* 1785–822. [138]

Huang, Ruidi, Erik J Mayer, Sheridan Titman, and David X Xu. 2025. "Human Capital and Local Credit Supply: Evidence from the Mortgage Industry." Working Paper, SMU Cox School of Business. [138]

Jäger, Simon, Christopher Roth, Nina Roussille, and Benjamin Schoefer. 2024. "Worker Beliefs about Outside Options." *Quarterly Journal of Economics* 139 (3): 1505–56. [122, 131]

Jiang, John, Isabel Y Wang, and K Philip Wang. 2018. "Revolving Rating Analysts and Ratings of Mortgage-Backed and Asset-Backed Securities: Evidence from LinkedIn." *Management Science* 64 (12): 5832–54. [131]

Johnson, Matthew S, Kurt Lavetti, and Michael Lipsitz. 2025. "The Labor Market Effects of Legal Restrictions on Worker Mobility." *Journal of Political Economy* 133 (9): 2735–93. [119, 139]

Jordà, Òscar, Moritz Schularick, and Alan M. Taylor. 2013. "When Credit Bites Back." *Journal of Money, Credit and Banking* 45: 3–28. [121]

McNamara, Yvonne, and Damien P McLoughlin. 2009. "Corporate Alumni Networks and Knowledge Flows." Working Paper, University College Dublin. [131]

Montgomery, James. 1991. "Social Networks and Labor-Market Outcomes: Toward an Economic Analysis." *American Economic Review* 81 (5): 1408–18. [130]

Mortensen, Dale T, and Christopher A Pissarides. 1999. "New Developments in Models of Search in the Labor Market." *Handbook of Labor Economics* 3: 2567–627. [121]

Mueller, Clemens. 2023. "Non-Compete Agreements and Labor Allocation Across Product Markets." In *Proceedings of the EUROFIDAI-ESSEC Paris December Finance Meeting.* [139]

Mugasha, Agasha. 2004. "The Secondary Market for Syndicated Loans: Loan Trading, Credit Derivatives, and Collateralized Debt Obligations." *Banking and Finance Law Review* 19 (2): 199–236. [127]

Myerson, Roger B. 2012. "A Model of Moral-Hazard Credit Cycles." *Journal of Political Economy* 120 (5): 847–78. [122]

Myerson, Roger B. 2014. "Moral-Hazard Credit Cycles with Risk-Averse Agents." *Journal of Economic Theory* 153: 74–102. [122]

Nakamura, Alice, Emi Nakamura, Kyle Phong, and Jón Steinsson. 2019. "New Evidence on the Cyclicality of Employer-to-Employer Flows from Canada." In *AEA Papers and Proceedings,* 109: 456–60. [121]

Norden, Lars, Bernardus Van Doornik, and Weichao Wang. 2025. "Labor Turnover, Information Production, and Bank Risk." Working Paper, 626, Banco Central do Brasil. [122]

Philippon, Thomas, and Ariell Reshef. 2012. "Wages and Human Capital in the US Finance Industry: 1909–2006." *Quarterly Journal of Economics* 127 (4): 1551–609. [126]

Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, Renae Rodgers, Jonathan Schroeder, and Kari C.W. Williams. 2025. (IPUMS USA: Version 16.0). [128]

Saygin, Perihan Ozge, Andrea Weber, and Michele Weynandt. 2021. "Coworkers, Networks, and Job-Search Outcomes Among Displaced Workers." *ILR Review* 74 (1): 95–130. [130]

Schubert, Gregor, Anna Stansbury, and Bledi Taska. 2022. "Employer Concentration and Outside Options." Working Paper, UCLA. [122]

Schularick, Moritz, and Alan M Taylor. 2012. "Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870–2008." *American Economic Review* 102 (2): 1029–61. [121]

Schwert, Michael. 2020. "Does Borrowing from Banks Cost More than Borrowing from the Market?" *Journal of Finance* 75 (2): 905–47. [129]

Shapiro, Carl, and Joseph E Stiglitz. 1984. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review* 74 (3): 433–44. [140]

Shue, Kelly. 2013. "Executive Networks and Firm Policies: Evidence from the Random Assignment of MBA Peers." *Review of Financial Studies* 26 (6): 1401–42. [130]

Thanassoulis, John. 2012. "The Case for Intervening in Bankers' Pay." *Journal of Finance* 67 (3): 849–95. [122]

Van Boxtel, Anton. 2017. "Trader Compensation, Risk, and the Banking Labour Market." Working Paper, Institute for Advanced Studies Vienna. [122]

Wooldridge, Jeffrey M. 1999. "Distribution-Free Estimation of Some Nonlinear Panel Data Models." *Journal of Econometrics* 90 (1): 77–97. [132, 152, 153]

Appendix 2.A Additional Figures

Figure 2.A.1. Labor Market Turnover in the Financial Sector and Variation in Lending Standards

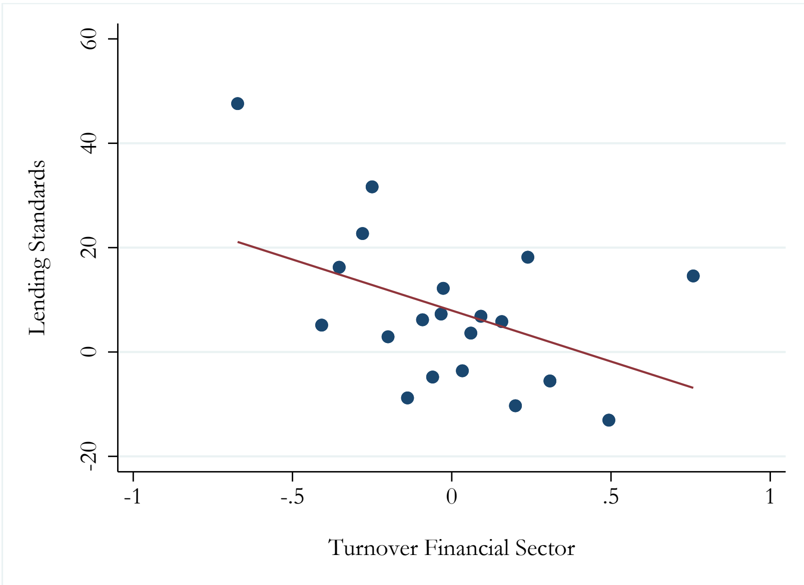

Notes: The figure binscatters voluntary labor market turnover in the financial sector against variation in lending standards. Labor market turnover is defined as the sum of hiring and quit rates in financial activities (FRED series JTU510099HIL and JTU510099QUR), which we transform into first differences. Lending standards are measured as the net percentage of domestic banks tightening their standards across loan categories, which is calculated based on data from Senior Loan Officer Opinion Survey on Bank Lending Practices (SLOOS) and weighted by banks' outstanding loan balances (FRED series SUBLPDMOSXWBNQ).

Figure 2.A.2. Procyclicality of Inflows at Ties

(a) Credit-to-GDP Ratio (b) Employment in Credit Intermediation

Notes: The plot shows average inflows at ties over time against the credit-to-GDP ratio (Panel a) and the number of employees in credit intermediation (Panel b). The aggregate data are obtained via FRED (series QUSPAM770A and CEU5552200001).

Figure 2.A.3. Outside Options and All Employee Outflows

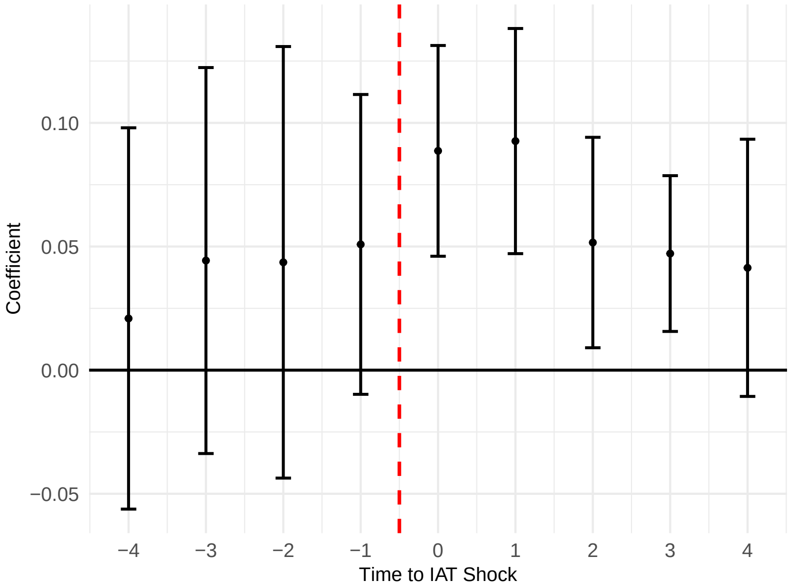

Notes: The figure shows coefficients from regressions of all employee outflows on inflows at ties, using the Poisson estimator for dispersed count variables with meaningful zeros (Wooldridge 1999). The coefficients come from separate models on a series of leads and lags of inflows at ties. The inflows at ties measure is standardized to a mean of zero and a standard deviation of one. The figure shows 95% intervals calculated using two-way clustered standard errors at the lender and year-quarter level.

Figure 2.A.4. Outside Options and Employee Outflows—Robustness to Alternative Tie Lengths

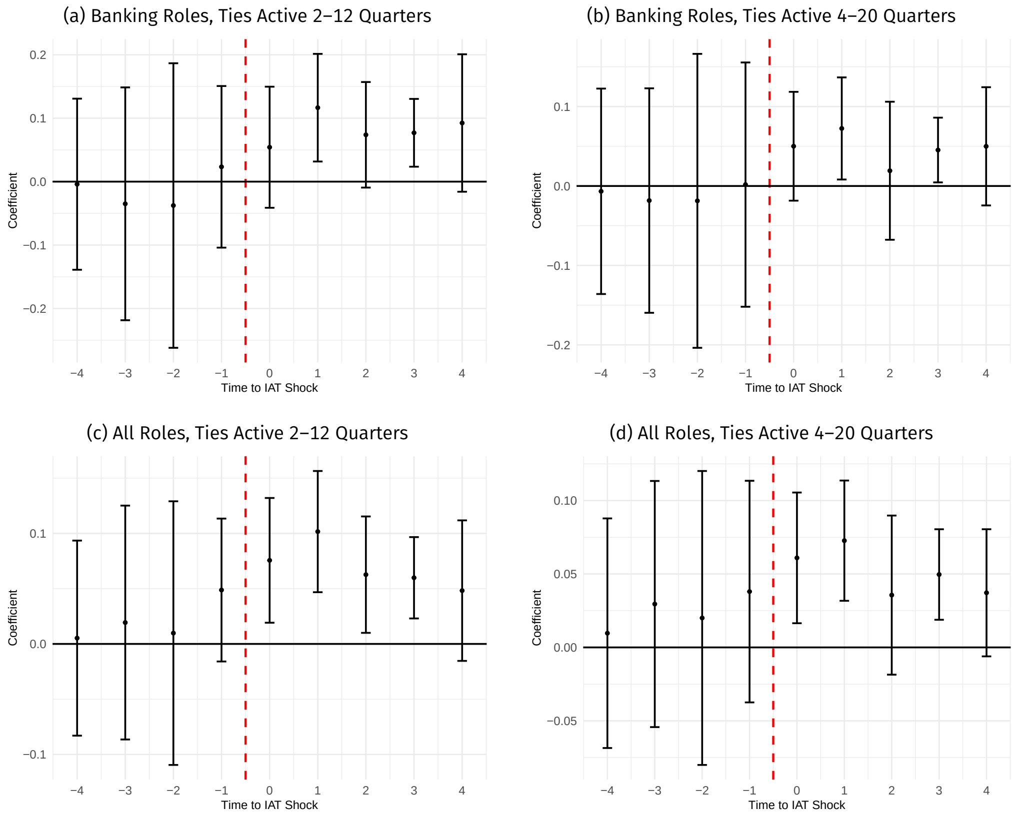

Notes: The figure shows coefficients from regressions of employee outflows on inflows at ties calculated under alternative timing assumptions, using the Poisson estimator for dispersed count variables with meaningful zeros (Wooldridge 1999). The dependent variables are outflows to other banking roles (Panels a and b) and to all other roles (Panels c and d). The main independent variable is *InflowsAtTies*, recalculated so that ties are active between 2 and 12 quarters after a coworker transition (Panels a and c) and between 4 and 20 quarters (Panels b and d), respectively. The coefficients come from separate regressions on a series of leads and lags of inflows at ties. The inflows at ties measure is standardized to a mean of zero and a standard deviation of one. The figure shows 95% intervals calculated using two-way clustered standard errors at the lender and year-quarter level.

Figure 2.A.5. Unique Number of Connected Banks Over Time

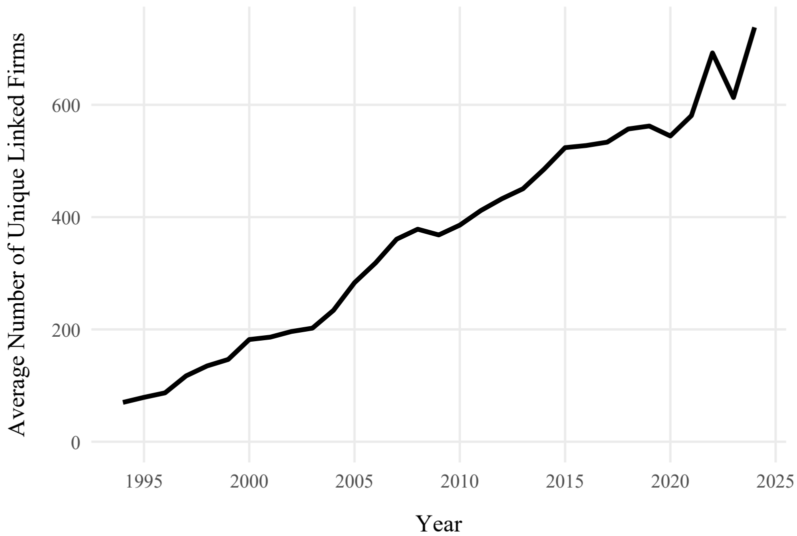

Notes: The figure plots, over time, the number of unique banks with non-zero hiring and at least one active tie to a bank in our main estimation sample.

Appendix 2.B Additional Tables

Table 2.B.1. Summary Statistics: Dealscan Loan-Level Data

	N	Mean	SD	P25	P50	P75
Tranche Amount	55,327	905.00	1,623.00	150.00	400.00	1,000.00
Maturity	55,327	50.00	23.00	36.00	58.00	60.00
Secured	55,327	0.51	0.50	0.00	1.00	1.00
Covenants	55,327	0.48	0.50	0.00	0.00	1.00

Notes: The table reports loan-tranche level summary statistics for our sample of Dealscan loans. *Tranche Amount* is converted to million USD. *Maturity* is indicated in months. *Secured* is an indicator for whether or not the tranche is secured. *Covenants* is an indicator for whether or not financial covenants exist.

Table 2.B.2. Summary Statistics: Employment History Data

	N	Mean	W. Mean	SD	P25	P50	P75
Panel A: Full Sample							
Male	801,049	0.69	0.68	0.46	0.00	1.00	1.00
White	802,908	0.78	0.78	0.42	1.00	1.00	1.00
Bachelor	634,528	0.50	0.50	0.50	0.00	1.00	1.00
More than Bachelor	634,528	0.48	0.48	0.50	0.00	0.00	1.00
Salary (in $1,000)	803,076	150.06	149.55	69.22	102.56	135.27	181.38
Age at Start	771,951	33.55	33.53	8.18	28.20	32.60	38.40
Age at Middle	771,951	36.69	36.66	8.28	30.85	35.68	41.71
Age at End	771,951	39.83	39.80	9.69	32.91	38.31	45.36
Seniority	803,096	4.75	4.75	0.74	4.00	5.00	5.00
No. of Previous Positions	803,096	3.02	3.03	2.62	1.00	3.00	4.00
Panel B: Matched Sample							
Age	327	40.28	40.27	7.93	35.10	40.00	45.34
Seniority	327	4.76	4.76	0.90	5.00	5.00	5.00
Job Tenure	327	9.04	9.05	7.31	3.48	7.20	12.47
Salary (in $1,000)	327	224.62	224.49	89.16	169.11	200.57	269.28

Notes: The table reports spell-level summary statistics on demographic variables for the sample of senior employees working in credit origination described in Section 2.3. The full sample in Panel A includes all employment spells that are reported in the United States, in the respective job roles, and at seniority level 4 or above. *Male*, *White*, *Bachelor*, and *More Than Bachelor* are indicator variables equal to one for the corresponding categories. *Salary* is an imputed variable based on observable characteristics of an employment spell and measured in 1,000 USD. The age variables refer to the beginning, middle, and end of an employment spell. *Seniority* refers to the seniority ranking, taking values between 4 (lowest) and 7 (highest). *No. of Previous Positions* counts the number of positions held prior to the current spell. The sample in Panel B shows demographic summary statistics on the sample of bankers matched to individual loans, focusing on the date of loan signature. *Age* and *Job Tenure* are precise values in years, *Seniority* refers to the seniority ranking. *Salary* is an imputed variable based on observable characteristics of an employment spell and measured in 1,000 USD. The variable *Weighted Mean* exploits sampling weights constructed by Revelio Labs (see www.data-dictionary.reveliolabs.com for details).

Table 2.B.3. Summary Statistics: Employment History vs. Administrative Data

	Estimation Sample	Full Sample	CPS	ACS
Male	0.68	0.63	0.57	0.57
White	0.78	0.74	0.84	0.82
Bachelor	0.50	0.56	0.49	0.46
More than Bachelor	0.48	0.40	0.17	0.17
Salary (in $1,000)	149.55	102.10	98.31	100.91
Age	36.66	32.15	40.84	40.13
Observations (in 1,000)	607.60	1,657.54	20.23	304.14

Notes: The table reports spell-level weighted averages of demographic variables, comparing the employment biography data against demographic sources. *Estimation Sample* includes all employment spells that are reported in the United States, in the respective job roles, and at seniority level 4 or above, as described in Section 2.3. *Full Sample* relaxes the latter restriction and includes all seniority levels of the employment spell data. *CPS* and *ACS* rely on administrative data from the *Current Population Survey* and the *American Community Survey*, respectively. For the CPS and ACS, we retain all observations between 1994 and 2023 of individuals that are in the labor force, at work and in full-time employment, below 66 years, and with hourly income of at least 10 USD (21,600 USD annually). We restrict the sample to employment in the financial industry (1990 industry classification code between 700 and 710) and occupations related to loan origination (2010 occupation codes 0120, 0830, 0840, 0850, 0910, and 4820). *Male*, *White*, *Bachelor*, and *More Than Bachelor* are indicator variables equal to one for the corresponding categories. *Salary* is measured in 1,000 USD. For the employment history data, it is an imputed measure based on observable characteristics of an employment spell. *Age* refers to age at the middle of the employment spell for the spell-level data. *Observations* shows the number of observations that are non-missing for all of the above variables.

Table 2.B.4. Summary Statistics: Borrower and Lender-Level Outcomes and Controls

	N	Mean	SD	P25	P50	P75
Panel A: Borrower Characteristics						
Borrower Prime	55,327	0.30	0.46	0.00	0.00	1.00
Borrower Rating	55,327	0.46	0.20	0.25	0.50	0.62
Total Assets	55,327	29,702.32	126,413.95	1,906.80	6,370.07	18,754.20
Capex	55,327	929.33	2,594.73	32.23	138.41	597.25
Employees	55,327	31.12	96.96	2.33	8.41	25.70
Liabilities	55,327	23,160.37	111,742.11	1,200.10	4,217.55	13,194.50
EBIT	55,327	1,308.43	3,782.17	79.63	349.00	1,097.00
Net Income	55,327	608.39	2,833.28	-1.61	111.60	507.10
Current Liabilities	55,327	2,651.48	19,409.86	10.00	74.20	503.00
Long Term Liabilities	55,327	8,157.54	21,938.02	568.89	2,077.75	6,840.00
Stockholder Equity	55,327	6,162.13	17,864.16	381.04	1,500.63	4,700.43
Panel B: Lender Characteristics						
Total Assets	55,327	1,325,920.98	809,725.31	687,935.00	1,198,942.00	1,930,115.00
Capex	55,327	1,411.17	2,171.85	73.00	934.11	1,497.20
Employees	55,327	112.34	81.66	51.20	80.00	163.50
Liabilities	55,327	1,235,089.55	761,271.58	645,607.00	1,108,646.00	1,757,212.00
EBIT	55,327	20,601.01	14,427.06	10,750.51	17,104.00	26,337.57
Net Income	55,327	7,169.74	7,446.01	2,655.00	5,579.94	10,101.96
Current Liabilities	55,327	176,537.77	130,782.04	73,136.00	161,592.54	262,878.00
Long Term Liabilities	55,327	120,202.79	84,346.79	32,771.00	127,663.64	178,324.00
Stockholder Equity	55,327	88,314.10	62,474.38	44,354.00	75,397.31	110,450.56
Panel C: (Conditional) Value at Risk						
Value At Risk (95^{th} Pctile)	10,635	-0.08	23.00	-12.00	-0.85	11.00
Conditional Value at Risk (95^{th} Pctile)	10,635	-0.07	21.00	-10.00	-0.80	9.60
Value At Risk (99^{th} Pctile)	10,635	-0.05	21.00	-11.00	-0.63	10.00
Conditional Value at Risk (99^{th} Pctile)	10,635	-0.04	21.00	-9.90	-0.45	9.30
Panel D: Compensation Data						
Median Salary (in $1,000)	8,457	784.00	393.00	506.00	767.00	1,000.00
Median Bonus (in $1,000)	8,457	579.00	1,745.00	0.00	0.00	600.00
Median Equity at Risk (in $1,000)	8,333	9,170.00	19,784.00	2,342.00	5,672.00	11,147.00
Median Salary (Growth, in Ppt.)	8,462	3.60	27.00	0.00	2.70	9.90
Median Bonus (Growth, in Ppt.)	8,462	-4.40	85.00	0.00	0.00	1.50
Median Equity at Risk (Growth, in Ppt.)	8,462	2.60	75.00	-36.00	0.00	43.00

Notes: This table reports summary statistics for the borrower firm (Panel A) and lender firm (Panel B) control variables in the loan level sample. Variables are taken from the Compustat Fundamentals Annual North America and Global databases. *Borrower Prime* is an indicator for investment-grade borrowers. *Borrower Rating* is an ordinal borrower rating normalized to range between 0 and 1. Total assets, total liabilities, capital expenditures, earnings before interest and taxes (EBIT), net income, current liabilities, long-term liabilities and stockholders equity are given in million USD. Employee counts are given in thousands. Panel C shows summary statistics of (Conditional) Value at Risk, obtained from Adrian and Brunnermeier (2016). The data in Panel C are at quarterly frequency, span the period between 1994 and 2009, are in symmetric growth rates and have been transformed into percentage points. Panel D shows summary statistics of compensation data from BoardEx, and includes averages of salaries, bonuses, and equity at risk of senior non-executive bankers at the bank-year level. The data in Panel D are at annual frequency, span the period between 1997 and 2020, and (if applicable) have been constructed in terms of symmetric growth rates.

Table 2.B.5. Summary Statistics: Outside Options Measure

	N	Mean	SD	P25	P50	P75
No. of Employees	19,378	316.77	750.68	19.00	82.00	258.00
No. of Connected Bankers	19,378	40.34	105.46	2.00	8.00	29.00
No. of Bank-Level Ties	19,378	23.58	45.27	2.00	7.00	23.00
No. of Bank-Level Ties (Size-corrected)	19,378	13.12	15.68	2.00	6.40	18.28
No. of Bank-Level Ties With Positive Inflows	19,378	11.54	20.49	1.00	3.00	13.00
Average Job-to-Job Hiring At Ties	19,378	6.67	9.26	1.56	4.40	8.12
Inflows At Ties	19,378	6.55	9.28	1.00	4.00	8.83

Notes: This table shows summary statistics related to the outside options measure *InflowsAtTies* by lender at the quarterly level. *Number of Employees* refers to the number of sampled individuals at the bank. *Number of Connected Bankers* is the number of individuals working in other institutions with active ties. *No. of Bank-Level Ties* is, respectively, the number of distinct banks with active ties, reported equally weighted, size-corrected using the inverse Herfindahl index of Borusyak, Hull, and Jaravel (2025), and restricted to ties with positive inflows. *Average Job-to-Job Hiring at Ties* is the mean inflow at tied banks, conditional on hired individuals transitioning between jobs with a gap of one month or less. *Inflows at Ties* is the outside options proxy prior to normalization.

Table 2.B.6. Validation: Outside Options and Bank-Level Covariates in Levels

	IAT Lead 1Q (1)	IAT Lead 2Q (2)	IAT Lead 3Q (3)	IAT Lead 4Q (4)
Log Assets	-0.338	-0.313	-0.319	-0.333
	(0.343)	(0.328)	(0.334)	(0.343)
Log EBIT	0.005	0.005	0.012	0.028
	(0.040)	(0.035)	(0.034)	(0.034)
Log Current Liabilities	0.113*	0.072	0.052	0.049
	(0.064)	(0.075)	(0.084)	(0.089)
Log Long-Term Debt	-0.120	-0.144	-0.133	-0.131
	(0.174)	(0.179)	(0.174)	(0.173)
Log Liabilities	0.473	0.450	0.432	0.422
	(0.368)	(0.357)	(0.351)	(0.346)
Log Equity	0.019	0.014	0.027	0.023
	(0.081)	(0.079)	(0.079)	(0.076)
Log Capital Expenditures	-0.176	-0.124	-0.121	-0.112
	(0.115)	(0.128)	(0.133)	(0.142)
Log Net Income	-0.007	0.001	0.001	0.003
	(0.028)	(0.025)	(0.021)	(0.021)
Observations	11,964	11,964	11,964	11,964
Lender fixed effects	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓

Notes: The table reports results from regressing future values (one to four quarters ahead) of inflows at ties on current bank-level covariates, conditional on fixed effects at the lender and year-quarter level. Standard errors clustered at the lender level are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.7. Validation: Outside Options and Bank-Level Covariates in Growth Rates

	IAT Lead 1Q	IAT Lead 2Q	IAT Lead 3Q	IAT Lead 4Q
	(1)	(2)	(3)	(4)
Asset Growth	-0.217	0.206	0.201	0.180
	(0.573)	(0.666)	(0.657)	(0.658)
EBIT Growth	-0.008	0.002	0.022	0.038**
	(0.022)	(0.019)	(0.018)	(0.017)
Current Liabilities Growth	0.111*	0.076*	0.044	0.036
	(0.059)	(0.041)	(0.034)	(0.038)
Long-Term Debt Growth	0.008	0.006	-0.011	-0.018
	(0.211)	(0.218)	(0.205)	(0.187)
Liabilities Growth	0.170	-0.017	-0.047	0.024
	(0.473)	(0.452)	(0.441)	(0.415)
Equity Growth	0.009	0.004	0.018	0.021
	(0.042)	(0.040)	(0.046)	(0.043)
Capital Expenditures Growth	-0.025	0.124	0.129	0.145
	(0.158)	(0.101)	(0.113)	(0.145)
Net Income Growth	-0.016	-0.009	-0.008	-0.004
	(0.015)	(0.015)	(0.016)	(0.017)
Observations	11,268	11,268	11,268	11,268
Lender fixed effects	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓

Notes: The table reports results from regressing future values (one to four quarters ahead) of inflows at ties on growth rates in current bank-level covariates, conditional on fixed effects at the lender and year-quarter level. Standard errors clustered at the lender level are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.8. Variation in Hiring at Connected Banks

Fixed Effects	R^2
Bank_i	0.025
Bank_i + Year-Quarter	0.064
$\text{Bank}_i \times$ Year-Quarter	0.112
Bank_j	0.598
Bank_j + Year-Quarter	0.648
Bank_j + Bank_i + Year-Quarter	0.649
Bank_j + $\text{Bank}_i \times$ Year-Quarter	0.675

Notes: The table reports R^2 obtained from regressing hiring at connected banks on varying combinations of fixed effects. $Bank_i$ refers to the bank of interest, $Bank_j$ to the connected bank.

Table 2.B.9. Outside Options and Lending Volumes (Ties Active 2–12 Quarters)

	Growth Rate		Extensive Margin		Log Volume	
	(1)	(2)	(3)	(4)	(5)	(6)
IAT	1.690	1.680	0.262**	0.281**	0.019	0.015
	(1.110)	(1.130)	(0.126)	(0.126)	(0.029)	(0.028)
Observations	6,277	6,277	3,999	3,999	4,058	4,058
Lender Controls		✓		✓		✓
Lender fixed effects	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓
Estimation	OLS	OLS	Logit	Logit	OLS	OLS

Notes: The table reports results from regressing the outcome of interest on *InflowsAtTies*$_{it}$, lender controls, and fixed effects at the lender and year-quarter level. The measure *InflowsAtTies* is recalculated so that ties are active between 2 and 12 quarters after a coworker transition, rather than between 4 and 12 quarters as in the baseline. The outcome variables are symmetric growth rates of total lending volume in percentage points (Columns 1 and 2), an indicator equal to one if a bank arranged any loan (Columns 3 and 4), and the natural logarithm of the loan amount, conditional on lending a strictly positive amount (Columns 5 and 6). Lender controls are the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, and equity. Columns (1), (2), (5), and (6) are estimated via OLS, while Columns (3) and (4) are obtained from logit regressions. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.10. Outside Options and Lending Volumes (Ties Active 4–20 Quarters)

	Growth Rate		Extensive Margin		Log Volume	
	(1)	(2)	(3)	(4)	(5)	(6)
IAT	2.250***	2.210**	0.181*	0.192*	0.014	0.005
	(0.843)	(0.893)	(0.103)	(0.104)	(0.032)	(0.031)
Observations	6,277	6,277	3,999	3,999	4,058	4,058
Lender Controls		✓		✓		✓
Lender fixed effects	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓
Estimation	OLS	OLS	Logit	Logit	OLS	OLS

Notes: The table reports results from regressing the outcome of interest on *InflowsAtTies*$_{it}$, lender controls, and fixed effects at the lender and year-quarter level. The measure *InflowsAtTies* is recalculated so that ties are active between 4 and 20 quarters after a coworker transition, rather than between 4 and 12 quarters as in the baseline. The outcome variables are symmetric growth rates of total lending volume in percentage points (Columns 1 and 2), an indicator equal to one if a bank arranged any loan (Columns 3 and 4), and the natural logarithm of the loan amount, conditional on lending a strictly positive amount (Columns 5 and 6). Lender controls are the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, and equity. Columns (1), (2), (5), and (6) are estimated via OLS, while Columns (3) and (4) are obtained from logit regressions. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.11. Outside Options and Changes in Financial Risk (in Percentage Points, Ties Active 2–12 Quarters)

	VaR 95%		CoVaR 95%		VaR 99%		CoVaR 99%	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IAT	0.254	0.255	0.204	0.202	0.425***	0.421***	0.385***	0.379***
	(0.218)	(0.217)	(0.180)	(0.179)	(0.157)	(0.155)	(0.136)	(0.133)
Observations	10,635	10,635	10,635	10,635	10,635	10,635	10,635	10,635
Lender Controls		✓		✓		✓		✓
Lender fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓	✓	✓

Notes: The table reports coefficients from OLS regressions of symmetric growth rates in VaR and CoVaR (in percentage points) on $InflowsAtTies_{it}$, lender controls, and fixed effects at the lender and year-quarter level. The measure *InflowsAtTies* is recalculated so that ties are active between 2 and 12 quarters after a coworker transition, rather than between 4 and 12 quarters as in the baseline. The outcomes are VaR (Columns 1, 2, 5, and 6) and CoVaR (Columns 3, 4, 7, and 8), computed at the 95% confidence level (Columns 1 to 4) and the 99% confidence level (Columns 5 to 8). Lender controls are the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, and equity. The data are at quarterly frequency and span 1994 to 2009. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.12. Outside Options and Changes in Financial Risk (in Percentage Points, Ties Active 4–20 Quarters)

	VaR 95%		CoVaR 95%		VaR 99%		CoVaR 99%	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
IAT	0.406*	0.408*	0.336*	0.333*	0.411**	0.407**	0.355*	0.348**
	(0.217)	(0.214)	(0.174)	(0.171)	(0.201)	(0.199)	(0.179)	(0.175)
Observations	10,749	10,749	10,749	10,749	10,749	10,749	10,749	10,749
Lender Controls		✓		✓		✓		✓
Lender fixed effects	✓	✓	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓	✓	✓

Notes: The table reports coefficients from OLS regressions of symmetric growth rates in VaR and CoVaR (in percentage points) on $InflowsAtTies_{it}$, lender controls, and fixed effects at the lender and year-quarter level. The measure *InflowsAtTies* is recalculated so that ties are active between 4 and 20 quarters after a coworker transition, rather than between 4 and 12 quarters as in the baseline. The outcomes are VaR (Columns 1, 2, 5, and 6) and CoVaR (Columns 3, 4, 7, and 8), computed at the 95% confidence level (Columns 1 to 4) and the 99% confidence level (Columns 5 to 8). Lender controls are the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, and equity. The data are at quarterly frequency and span 1994 to 2009. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.13. Borrower Characteristics and Investment-Grade Status

	Borrower Prime	
	Coef.	SE
Borrower Assets (log)	0.056***	(0.006)
Borrower Capex (log)	0.015***	(0.004)
Borrower Employees (log)	0.041***	(0.006)
Borrower Liabilities (log)	0.049***	(0.005)
Borrower EBIT (log)	0.030***	(0.002)
Borrower Net Income (log)	0.026***	(0.002)
Borrower Current Debt (log)	0.022***	(0.003)
Borrower Total Debt (log)	0.026***	(0.004)
Borrower Equity (log)	0.009***	(0.002)

Notes: The table reports results from regressing an indicator for investment-grade borrower status separately on borrower characteristics, one per row, each entered as its inverse hyperbolic sine transformation. All columns include fixed effects at the lender and year-quarter level. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.14. Outside Options and Borrower Risk—Robustness

	Borrower Prime			
	(1)	(2)	(3)	(4)
IAT	-0.006***	-0.007**	-0.010***	-0.006***
	(0.002)	(0.003)	(0.003)	(0.001)
Observations	55,327	52,475	50,131	55,278
Borrower Controls	✓	✓	✓	✓
Lender Controls	✓	✓	✓	
Loan Controls	✓	✓	✓	✓
IAT Leads and Lags	✓			
Lender fixed effects	✓	✓	✓	
Year-Quarter fixed effects	✓	✓	✓	✓
Borrower SIC Sector × Year-Quarter fixed effects		✓		
Borrower State × Year-Quarter fixed effects			✓	
Lender × Year fixed effects				✓

Notes: The table reports results from robustness exercises for Column (3) of Table 2.4, obtained by estimating Equation (2.12) at the loan level. The dependent variable in all regressions is an indicator for investment-grade borrowers. The independent variable is $InflowsAtTies_{it}$. Borrower and lender controls include the inverse hyperbolic sine transformation of total assets, number of employees, capital expenditure, total liabilities, EBIT, and (for borrowers) fixed assets. Loan controls are the inverse hyperbolic sine transformation of tranche amount and maturity and indicators for whether the loan is secured and whether covenants are attached. All regressions include fixed effects at the lender and year-quarter level. Standard errors are double-clustered at the lender and year-quarter level. Compared to Column (3) of Table 2.4, the following additional variables are included: four lags and leads of $InflowsAtTies_{it}$ (Column 1), borrower-SIC × year-quarter fixed effects (Column 2), borrower-state × year-quarter fixed effects (Column 3), and lender × year fixed effects (Column 4). * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.15. Outside Options and Borrower Risk—Robustness (Part II)

	Borrower Prime				
Sample	Full Sample			Post-2008	Origination
	(1)	(2)	(3)	(4)	(5)
IAT	-0.005***	-0.005***	-0.005***	-0.004***	-0.008***
	(0.001)	(0.001)	(0.002)	(0.001)	(0.001)
Outflows	0.011				
	(0.007)				
Number of Lenders		0.000			
		(0.008)			
IAC			-0.005		
			(0.003)		
Observations	55,327	55,327	49,959	44,523	25,400
Borrower Controls	✓	✓	✓	✓	✓
Lender Controls	✓	✓	✓	✓	✓
Loan Controls	✓	✓	✓	✓	✓
Lender fixed effects	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓

Notes: The table reports results from robustness exercises for Column (3) of Table 2.4, obtained by estimating Equation (2.12) at the loan level. The dependent variable in all regressions is an indicator for investment-grade borrowers. The independent variable is $\mathit{InflowsAtTies}_{it}$. Borrower and lender controls include the inverse hyperbolic sine transformation of total assets, number of employees, capital expenditure, total liabilities, EBIT, and (for borrowers) fixed assets. Loan controls are the inverse hyperbolic sine transformation of tranche amount and maturity and indicators for whether the loan is secured and whether covenants are attached. All regressions include fixed effects at the lender and year-quarter level. Standard errors are double-clustered at the lender and year-quarter level. Compared to Column (3) of Table 2.4, the following additional variables are included: the inverse hyperbolic sine transformation of outflows (Column 1), the inverse hyperbolic sine transformation of the number of lenders (Column 2), and IAC ($\mathit{InflowsAtCosignatories}_{it}$), a variant of $\mathit{InflowsAtTies}_{it}$ where we replace the ties from coworker networks by the ties resulting from co-arranging loans (Column 3). In Column (4), we restrict the sample to loans in 2008 and afterwards. In Column (5), we restrict the sample to new originations. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.16. Outside Options and Borrower Risk (Ties Active 2–12 Quarters)

	Borrower Prime			Borrower Rating		
	(1)	(2)	(3)	(4)	(5)	(6)
IAT	-0.002	-0.005***	-0.005***	0.000	-0.002**	-0.002**
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	55,327	55,327	55,327	55,327	55,327	55,327
Outcome Mean	0.296	0.296	0.296	0.462	0.462	0.462
Lender Controls	✓	✓	✓	✓	✓	✓
Borrower Controls		✓	✓		✓	✓
Loan Controls			✓			✓
Lender fixed effects	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓

Notes: The table reports results from estimating Equation (2.12) at the loan level. The dependent variables are an indicator for investment-grade borrowers (Columns 1 to 3) and an ordinal borrower rating normalized to range between 0 and 1 (Columns 4 to 6). The measure *InflowsAtTies* is recalculated so that ties are active between 2 and 12 quarters after a coworker transition, rather than between 4 and 12 quarters as in the baseline. The independent variable is $InflowsAtTies_{it}$. Borrower and lender controls include the inverse hyperbolic sine transformation of total assets, number of employees, capital expenditure, total liabilities, EBIT, and (for borrowers) fixed assets. Loan controls are the inverse hyperbolic sine transformation of tranche amount and maturity and indicators for whether the loan is secured and whether covenants are attached. All regressions include fixed effects at the lender and year-quarter level, and standard errors are double-clustered at the same levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.17. Outside Options and Borrower Risk (Ties Active 4–20 Quarters)

	Borrower Prime			Borrower Rating		
	(1)	(2)	(3)	(4)	(5)	(6)
IAT	-0.004***	-0.006***	-0.006***	-0.001	-0.002***	-0.002***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Observations	55,327	55,327	55,327	55,327	55,327	55,327
Outcome Mean	0.296	0.296	0.296	0.462	0.462	0.462
Lender Controls	✓	✓	✓	✓	✓	✓
Borrower Controls		✓	✓		✓	✓
Loan Controls			✓			✓
Lender fixed effects	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓

Notes: The table reports results from estimating Equation (2.12) at the loan level. The dependent variables are an indicator for investment-grade borrowers (Columns 1 to 3) and an ordinal borrower rating normalized to range between 0 and 1 (Columns 4 to 6). The measure *InflowsAtTies* is recalculated so that ties are active between 4 and 20 quarters after a coworker transition, rather than between 4 and 12 quarters as in the baseline. The independent variable is $InflowsAtTies_{it}$. Borrower and lender controls include the inverse hyperbolic sine transformation of total assets, number of employees, capital expenditure, total liabilities, EBIT, and (for borrowers) fixed assets. Loan controls are the inverse hyperbolic sine transformation of tranche amount and maturity and indicators for whether the loan is secured and whether covenants are attached. All regressions include fixed effects at the lender and year-quarter level, and standard errors are double-clustered at the same levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.18. Outside Options and Loan Prices

	All-In-Spread-Drawn (log bps)					
	(1)	(2)	(3)	(4)	(5)	(6)
IAT	-0.004	0.000	0.002	-0.005	-0.002	0.000
	(0.006)	(0.005)	(0.003)	(0.005)	(0.005)	(0.003)
Observations	47,760	47,760	47,760	47,760	47,760	47,760
Outcome Mean	5.935	5.935	5.935	5.935	5.935	5.935
Lender Controls	✓	✓	✓	✓	✓	✓
Borrower Controls		✓	✓		✓	✓
Loan Controls			✓			✓
Lender fixed effects	✓	✓	✓	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓	✓	✓	✓
Borrower Rating fixed effects				✓	✓	✓

Notes: The table reports results from estimating Equation (2.12) at the loan level. The dependent variable is the log all-in-spread-drawn (in basis points). The independent variable is $InflowsAtTies_{it}$. Borrower and lender controls include the inverse hyperbolic sine transformation of total assets, number of employees, capital expenditure, total liabilities, EBIT, and (for borrowers) fixed assets. Loan controls are the inverse hyperbolic sine transformation of tranche amount and maturity and indicators for whether the loan is secured and whether covenants are attached. All regressions include fixed effects at the lender and year-quarter level, and standard errors are double-clustered at the same levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.19. Outside Options and Borrower Risk by Non-Compete Agreement Enforceability

	Borrower Prime		
	(1)	(2)	(3)
IAT	-0.007**	-0.004	-0.004
	(0.003)	(0.004)	(0.005)
IAT × Enforcement Change	0.011		
	(0.011)		
IAT × Enforcement Increase		-0.013	
		(0.011)	
IAT × Enforcement Decrease			-0.036***
			(0.011)
Observations	50,490	50,490	50,490
Outcome Mean	0.296	0.296	0.296
Lender Controls	✓	✓	✓
Borrower Controls	✓	✓	✓
Loan Controls	✓	✓	✓
Lender fixed effects	✓	✓	✓
Year-Quarter fixed effects	✓	✓	✓
State-Year fixed effects	✓	✓	✓

Notes: The table reports a robustness exercise for Column (3) of Table 2.4, estimating Equation (2.12) at the loan level. The dependent variable in all Columns is an indicator for investment-grade borrowers. The main independent variable is $InflowsAtTies_{it}$, interacted with measures of changes in Non-Compete Agreement (NCA) enforceability in the state of the bank's headquarters. *Enforcement Change* equals 1 if a state increased its NCA enforceability and −1 if it decreased it, in any past or current period, and 0 otherwise; *Enforcement Increase* and *Enforcement Decrease* are separate indicators for past increases and decreases. Borrower and lender controls include the inverse hyperbolic sine transformation of total assets, number of employees, capital expenditure, total liabilities, EBIT, and (for borrowers) fixed assets. Loan controls are the inverse hyperbolic sine transformation of tranche amount and maturity and indicators for whether the loan is secured and whether covenants are attached. All columns include fixed effects at the lender and year-quarter level and at the state-year level. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.20. Outside Options and Borrower Rating: Age Heterogeneity

	Borrower Prime Rating (Normalized)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IAT	-0.321**	-0.335**	-0.331**				
	(0.124)	(0.143)	(0.152)				
IAT × Age	0.008**	0.008**	0.008**	0.005***	0.005***	0.005***	0.004*
	(0.003)	(0.003)	(0.003)	(0.001)	(0.001)	(0.001)	(0.002)
Observations	275	275	275	327	327	327	327
Outcome Mean	0.495	0.495	0.495	0.488	0.488	0.488	0.488
Lender Controls	✓	✓	✓				
Borrower Controls		✓	✓	✓	✓	✓	✓
Loan Controls			✓	✓	✓	✓	✓
IAT × Spell Controls							✓
Lender fixed effects	✓	✓	✓				
Year-Quarter fixed effects	✓	✓	✓				
Banker fixed effects	✓	✓	✓	✓	✓	✓	✓
Lender × Year-Quarter fixed effects				✓	✓		✓
MSA × Year-Quarter fixed effects					✓		✓
Lender × MSA × Year-Quarter fixed effects						✓	

Notes: The table reports results from estimating Equation (2.12) at the loan level, augmented with an interaction for the underwriting banker's age. The dependent variable is an ordinal scale of borrower rating, normalized to range between 0 and 1. The independent variables are $\mathit{InflowsAtTies}_{it}$ and its interaction with the banker's age. Borrower and lender controls include the inverse hyperbolic sine transformation of total assets, number of employees, capital expenditure, total liabilities, EBIT, and (for borrowers) fixed assets. Loan controls are the inverse hyperbolic sine transformation of tranche amount and maturity and indicators for whether the loan is secured and whether covenants are attached. *Spell controls* contain the position number, the seniority level, the inverse hyperbolic sine transformation of imputed salary and job tenure, each included interacted with $\mathit{InflowsAtTies}_{it}$ and (if possible) as a standalone variable. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.21. Outside Options and Tranche Amounts: Age Heterogeneity

	Asinh (Tranche Amount)						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IAT	-1.340	0.322	0.358				
	(0.770)	(0.569)	(0.449)				
IAT × Age	0.035***	0.008	0.008	-0.031***	-0.030**	-0.030**	-0.044***
	(0.009)	(0.008)	(0.008)	(0.010)	(0.010)	(0.010)	(0.009)
Observations	322	322	322	383	383	383	383
Outcome Mean	6.900	6.900	6.900	6.800	6.800	6.800	6.800
Lender Controls	✓	✓	✓				
Borrower Controls		✓	✓	✓	✓	✓	✓
Loan Controls			✓	✓	✓	✓	✓
IAT × Spell Controls							✓
Lender fixed effects	✓	✓	✓				
Year-Quarter fixed effects	✓	✓	✓				
Banker fixed effects	✓	✓	✓	✓	✓	✓	✓
Lender × Year-Quarter fixed effects				✓	✓		✓
MSA × Year-Quarter fixed effects					✓		✓
Lender × MSA × Year-Quarter fixed effects						✓	

Notes: The table reports results from estimating Equation (2.12) at the loan level, augmented with an interaction for the underwriting banker's age. The dependent variable is the inverse hyperbolic sine transformation of the tranche amount. The independent variables are $InflowsAtTies_{it}$ and its interaction with the banker's age. Borrower and lender controls include the inverse hyperbolic sine transformation of total assets, number of employees, capital expenditure, total liabilities, EBIT, and (for borrowers) fixed assets. Loan controls are the inverse hyperbolic sine transformation of maturity and indicators for whether the loan is secured and whether covenants are attached. *Spell controls* contain the position number, the seniority level, the inverse hyperbolic sine transformation of imputed salary and job tenure, each included interacted with $InflowsAtTies_{it}$ and (if possible) as a standalone variable. Standard errors are double-clustered at the lender and year-quarter level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.22. Outside Options and Misconduct Disclosures

	Number of Misconduct Disclosures						Binarized
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IAT_{t-1}	0.003***	0.003***	0.003***	0.003***	0.002***	0.004***	0.003***
	(0.0009)	(0.0009)	(0.0008)	(0.0008)	(0.0006)	(0.0009)	(0.0008)
IAT_{t}					0.003**		
					(0.001)		
IAT_{t-2}					-0.0008		
					(0.0005)		
$Outflows_{t-1}$						0.003	
						(0.002)	
Observations	582,568	561,018	551,079	551,079	551,079	551,079	551,079
Outcome Mean	0.010	0.010	0.009	0.009	0.009	0.009	0.009
Lender Controls	✓	✓	✓	✓	✓	✓	✓
Lender fixed effects	✓	✓					
Year fixed effects	✓	✓	✓		✓	✓	✓
Advisor fixed effects		✓					
Lender × Advisor fixed effects			✓	✓	✓	✓	✓
State × Year fixed effects				✓			

Notes: The table reports results from estimating Equation (2.12) at annual frequency at the advisor level. The dependent variable is the number of misconduct disclosures an advisor had in a given year, which is binarized in Column (7). The main independent variable is $InflowsAtTies_{it-1}$. Lender controls are the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, equity, and the first lag of the firm-level average of the outcome variable. Column (5) includes the contemporaneous and the second lag of inflows at ties. Column (6) controls for the inverse hyperbolic sine transformation of lagged outflows. All models are estimated with year fixed effects and lender controls. Standard errors are clustered at the lender level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.23. Outside Options and Customer Complaints

	Number of Customer Complaints						Binarized
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IAT_{t-1}	0.005***	0.006***	0.006***	0.005***	0.003***	0.006***	0.004***
	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)
IAT_{t}					0.005*		
					(0.003)		
IAT_{t-2}					0.0002		
					(0.001)		
$Outflows_{t-1}$						0.003	
						(0.004)	
Observations	582,568	561,018	551,079	551,079	551,079	551,079	551,079
Outcome Mean	0.018	0.017	0.017	0.017	0.017	0.017	0.017
Lender Controls	✓	✓	✓	✓	✓	✓	✓
Lender fixed effects	✓	✓					
Year fixed effects	✓	✓	✓		✓	✓	✓
Advisor fixed effects		✓					
Lender × Advisor fixed effects			✓	✓	✓	✓	✓
State × Year fixed effects				✓			

Notes: The table reports results from estimating Equation (2.12) at annual frequency at the advisor level. The dependent variable is the number of complaints filed against the advisor in a given year, which is binarized in Column (7). The main independent variable is $InflowsAtTies_{it-1}$. Lender controls are the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, equity, and the first lag of the firm-level average of the outcome variable. Column (5) includes the contemporaneous and the second lag of inflows at ties. Column (6) controls for the inverse hyperbolic sine transformation of lagged outflows. All models are estimated with year fixed effects and lender controls. Standard errors are clustered at the lender level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2.B.24. Outside Options and Settled Customer Complaints

	Number of Settled Customer Complaints						Binarized
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
IAT_{t-1}	0.004***	0.004***	0.004***	0.004***	0.003***	0.004***	0.003***
	(0.001)	(0.001)	(0.0010)	(0.0009)	(0.0007)	(0.001)	(0.0007)
IAT_{t}					0.004*		
					(0.002)		
IAT_{t-2}					-0.0005		
					(0.0008)		
$Outflows_{t-1}$						0.003	
						(0.003)	
Observations	582,568	561,018	551,079	551,079	551,079	551,079	551,079
Outcome Mean	0.008	0.008	0.008	0.008	0.008	0.008	0.008
Lender Controls	✓	✓	✓	✓	✓	✓	✓
Lender fixed effects	✓	✓					
Year fixed effects	✓	✓	✓		✓	✓	✓
Advisor fixed effects		✓					
Lender × Advisor fixed effects			✓	✓	✓	✓	✓
State × Year fixed effects				✓			

Notes: The table reports results from estimating Equation (2.12) at annual frequency at the advisor level. The dependent variable is the number of settled customer complaints an advisor had in a given year, which is binarized in Column (7). The main independent variable is $InflowsAtTies_{it-1}$. Lender controls are the inverse hyperbolic sine transformations of total assets, net income, number of employees, EBIT, capital expenditure, total liabilities, short- and long-term debt, equity, and the first lag of the firm-level average of the outcome variable. Column (5) includes the contemporaneous and the second lag of inflows at ties. Column (6) controls for the inverse hyperbolic sine transformation of lagged outflows. All models are estimated with year fixed effects and lender controls. Standard errors are clustered at the lender level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

References

Adrian, Tobias, and Markus K. Brunnermeier. 2016. “CoVaR.” *American Economic Review* 106 (7): 1705–41. [129, 135, 157]

Borusyak, Kirill, Peter Hull, and Xavier Jaravel. 2025. “A Practical Guide to Shift-Share Instruments.” *Journal of Economic Perspectives* 39 (1): 181–204. [133, 158]

Wooldridge, Jeffrey M. 1999. “Distribution-Free Estimation of Some Nonlinear Panel Data Models.” *Journal of Econometrics* 90 (1): 77–97. [132, 152, 153]

Chapter 3

Gig Work as Income Insurance during Unemployment: Evidence from Germany

Joint with Giacomo Marcolin

3.1 Introduction

Since the 2008 financial crisis, digital labor platforms have emerged as a new form of employment, engaging an estimated 5.8 million people in the United States by 2023 (Garin et al. 2025). Compared to traditional freelancing arrangements, these platforms feature lower barriers to entry, allowing individuals to supply labor almost instantaneously (Garin et al. 2020; Friedrich et al. 2024). Access to gig work may therefore affect individuals' ability to absorb income shocks and smooth consumption (Koustas 2019), with downstream effects on job-search behavior and the take-up of unemployment benefits (Jackson 2022; Fos et al. 2025).

This chapter studies how the local availability of gig labor shapes workers' responses to negative income shocks following job separation, that is, either layoffs or quits. We exploit the staggered entry of gig delivery platforms between 2014 and 2017 across German districts to test whether platform presence increases participation in gig work and whether it substitutes for unemployment insurance. We also examine whether platform availability affects job-search activity and participation in active labor market policies (ALMPs) among recently unemployed workers.

Our empirical analysis draws on the Sample of Integrated Labour Market Biographies (SIAB), a large administrative dataset provided by the German Federal Employment Agency, merged with information on gig delivery platform availability across districts. Our central finding is that, following a job separation, workers are 5–10% more likely to transition into gig work when a platform is present in their district. Take-up is concentrated among workers who match the profile of typical gig workers: young, male, and highly skilled. We also find some evidence that gig work crowds out other forms of employment in the medium run. However, despite the increase in gig work, platform entry does not reduce reliance on unemployment benefits. Hence, rather than substituting for traditional unemployment insurance, gig work appears to provide a supplementary source of income for some unemployed workers following job separation.

We extend the analysis to examine longer-run effects on local labor markets. First, we document an increase in participation in ALMPs, suggesting that gig work may serve a signaling role toward potential employers and government agencies. Second, we find that the employment effects of platform availability persist and grow over time, which is inconsistent with a simple vacancy-filling story in which effects would be temporary and one-off. Instead, the evidence points toward workers using gig work as an income buffer while searching for other employment. In line with this interpretation, we observe an increase in the number of jobseekers who are not registered as unemployed following platform entry. Taken together, these findings complement prior evidence documenting the structural impact of gig platforms on local labor markets (Koustas 2019; Fos et al. 2025). Third, we examine whether platform availability encourages business formation by providing implicit insurance to potential entrepreneurs. We find no evidence of such effects, suggesting that the insurance value of gig work may be limited in our context.

Identifying these causal effects is challenging because platform availability is not randomly distributed across districts, but instead shaped by local demographic characteristics. A naive comparison between districts with and without gig platforms would therefore yield biased estimates. To address this endogeneity concern, we adopt two complementary difference-in-differences approaches that compare labor market outcomes across districts before and after platform entry. Our specifications include different sets of fixed effects to account for time-invariant heterogeneity and district-level employment trends, allowing us to isolate the impact of gig work opportunities on recently unemployed individuals. The causal interpretation of our estimates relies on a parallel trends assumption, which we support through a range of robustness checks.

The chapter makes three contributions. First, to our knowledge, no prior study examines the income insurance role of gig work in Europe. The distinction from the US context matters because more generous welfare systems and lower labor market flexibility may alter workers' incentives to take up gig work. Second, we are the first to study how gig work affects job-search behavior and participation in ALMPs. These dimensions are central to understanding the longer-run effects of gig platforms and the ways in which gig jobs differ from traditional employment. Third, we focus on food delivery platforms. Compared with ride-sharing platforms, which have been the focus of most existing empirical work, food delivery differs in job content and earnings. Importantly, entry barriers are substantially lower, since ride-sharing work typically requires car ownership. These lower barriers make our setting particularly informative for understanding the broader insurance potential of the gig economy.

Literature: A growing body of work examines the insurance value of gig work and its implications for labor market outcomes. Focusing on ride-sharing platforms as flexible second jobs, Koustas (2019) shows that gig workers are able to smooth consumption when they experience income losses in their primary jobs. Fos et al. (2025) find that the introduction of Uber had sizeable effects on local labor markets across the US, reducing the probability of recently laid-off workers claiming unemployment insurance or defaulting on credit. Furthermore, platform presence has been linked to lower default rates on student loans (Moser 2021) and to higher rates of new business formation, as gig labor can limit the income volatility related to entrepreneurial activity (Barrios, Hochberg, and Yi 2022). Given that this insurance mechanism is strongest in downturns (Fos et al. 2025), gig platforms may even attenuate the labor market consequences of aggregate shocks.

These benefits, however, are likely to come at a cost for parts of the workforce. Gig platforms can weaken the effectiveness of minimum wages (Glasner 2023), raising concerns about monopsony power (Adams-Prassl, Adams-Prassl, and Coyle 2021). Jackson (2022) documents the emergence of a new gig precariat, characterized by low attachment to the traditional labor market and lower long-run earnings. One explanation is that gig work crowds out training and traditional private sector employment, while hampering human capital accumulation and reducing job-search intensity. We contribute to this literature by studying how platform presence affects participation in ALMPs[1] and job-search behavior.

We also complement a descriptive literature on nonstandard work arrangements. Several studies document the size and growth of these forms of employment using surveys and administrative data (Collins et al. 2019; Farrell, Greig, and Hamoudi 2019; Katz and Krueger 2019; Garin et al. 2020; Garin et al. 2025) and characterize the workers involved (Abraham and Houseman 2019; Liu and Nazareno 2019; Boeri et al. 2020).[2] A common finding is that the gig economy represents an important and growing segment of nonstandard work, but often serves as a supplementary source of income. For instance, in Italy, 80% of gig workers use it to supplement their income (Boeri et al. 2020).

More broadly, our work connects to a literature on unemployment insurance, its consumption smoothing benefits, and its potential distortions (e.g., Baily 1978; Gruber 1997; Chetty 2008; Kolsrud et al. 2018). These papers show that unemployed workers are often unable to smooth consumption even when receiving benefits. We add to this literature by documenting that some unemployed individuals take up gig work and use it to buffer income losses incurred from unemployment shocks.

1. For a review of the literature on ALMPs, see Card, Kluve, and Weber (2017).
2. For a review, see Mas and Pallais (2020).

The remainder of the chapter is organized as follows. Section 3.2 describes the data and sample construction. Section 3.3 discusses the empirical strategy. Section 3.4 presents the main results, robustness checks, and analyses of the underlying mechanisms. Section 3.5 concludes.

3.2 Data

This section describes the raw data and data preparation. The main dataset consists of two components: a comprehensive district-level[3] record of entry and exit dates of gig delivery platforms in Germany, and a high-frequency individual-level panel of German workers covering the period 2012 to 2017.

3.2.1 Data on Platform Availability

We collect entry and exit dates of gig delivery platforms by contacting the companies directly and by reviewing local newspapers and company press releases. We identify all platform launches through the end of 2019, covering 37 cities served by at least one delivery platform. Since we use employment data only through the end of 2017, the analysis is restricted to platform entries occurring before that year.

During the period of interest, three delivery platforms (*Deliveroo*, *Foodora*, and *Lieferando*) operated in the German market, offering their own delivery services through so-called *riders*.[4] These firms constituted the main source of on-site gig work opportunities at the time, as the regulatory framework prevented many other large gig platforms, such as Uber and Lyft, from operating legally in Germany.[5]

Compared with most traditional jobs, work for these platforms entails minimal entry requirements, likely satisfied by a large share of the population we study: recently unemployed workers. The main prerequisites are: (*i*) an official ID confirming the worker is at least 18 years old, (*ii*) a work permit (for foreigners), (*iii*) a smartphone, and (*iv*) a means of transport to carry out deliveries (although some platforms rent bikes to *riders*). Applications are typically submitted online and processed within a few days, sometimes followed by a phone interview before hiring.

3. We use NUTS-3 districts, the finest geographic level in the administrative employment data. In most cases, these units correspond to district-free cities. The remaining units are municipal associations, which additionally include adjacent rural districts. We assign platform presence based on entry into the district-free city, as these cities account for the vast majority of the population.

4. Two other delivery platforms, *Lieferheld* and *Pizza.de*, operated in the same sector under a different business model. Rather than hiring and paying workers directly, these platforms matched restaurants with clients and therefore do not meet our definition of gig platforms. As these companies expanded much earlier and were already present in all cities in our sample before 2013 (see bit.ly/48x6yxp and bit.ly/3teo09Z), there is limited concern that their roll-out may bias our results.

5. Consistent with the typical *winner-takes-all dynamics* of platforms (Furman 2019), concentration increased in the food delivery market. In 2019, *Deliveroo* exited (bit.ly/3PRbDJL) and *Foodora*'s German operations were acquired by *Takeaway*, which also owns *Lieferando* (bit.ly/46qpo7I).

Figure 3.1. Number of Districts With Gig Platforms Over Time

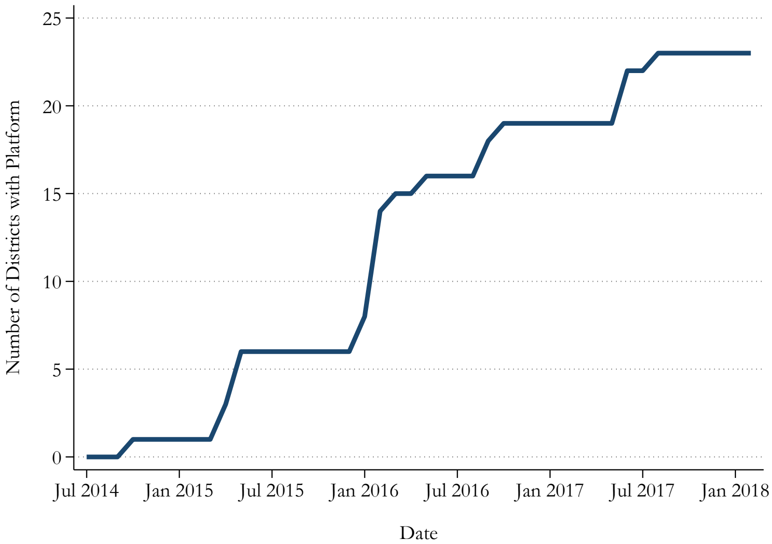

Notes: The figure plots the number of German districts with at least one gig delivery platform over time.

Figure 3.1 displays the number of districts with at least one gig platform over time. The sample includes 23 districts experiencing a platform entry between October 2014 and August 2017. Platforms often opened in multiple locations simultaneously by grouping launch events, and none exited during the sample period. Figure 3.A.1 shows a map of German districts, indicating the number of months a platform has been present by the end of 2017, with control districts shown in gray. Entry first targeted urban centers and by the end of the observation period had expanded to most major cities.

We explore the platform roll-out more systematically by estimating linear prediction models for whether and when a platform enters a district (similar to Deshpande and Li 2019). Specifically, we regress a dummy for platform entry and the time since first entry on a set of control variables chosen to proxy for factors potentially relevant to firms' entry decisions, including market size, local economic conditions, demand for food delivery products, labor supply, and the degree of urbanization. The results are shown in Figure 3.2. Population size and density are significant predictors on both the extensive margin and the timing of entry. That is, platforms tend to enter larger and more urbanized districts, and do so earlier in cities where these characteristics are more pronounced. By contrast, conditional on these variables, local economic and labor market conditions play no substantial role, although the estimated coefficients retain the expected sign. Similar results obtain when the explanatory variables are specified in growth rates rather than levels.

Figure 3.2. Predictors of Platform Entry

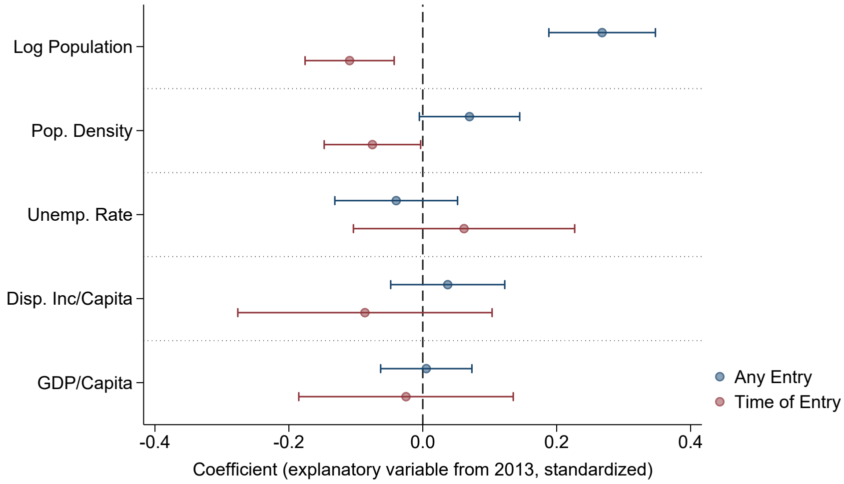

Notes: The figure shows results from two separate prediction exercises for platform entry (dummy) and timing of platform entry conditional on entry (measured in months since first platform entry in Germany, standardized between 0 and 1), each regressed on a constant and values of covariates in 2013 (standardized to mean zero and standard deviation one). The sample for the timing of platform entry includes all 23 districts in which platforms enter eventually, while the sample for any entry additionally contains the 50 control units used in the district-level analysis (city-districts with more than 100,000 inhabitants). The data on the explanatory variables (log population, population density, unemployment rate, disposable income per capita, and GDP per capita) have been obtained from the Federal Statistical Office of Germany. Point estimates are shown alongside 95% confidence intervals obtained from robust standard errors.

3.2.2 Employment Data

We use individual-level employment data from the Sample of Integrated Labour Market Biographies (SUF SIAB-Regional File 7517), a 2% random sample of the Integrated Employment Biographies (IEB) maintained by the Institute for Employment Research (IAB). The IEB covers all individuals with notifications to the social security system, including employment that is subject to social security contributions, benefit receipt, job-search status, and participation in active labor market programs (ALMPs). The data contain detailed employment histories with daily precision from 1975 to 2017, together with a set of individual characteristics.

Data Preparation: We process the data in three steps. First, we convert the spell-level information into an individual-month panel following Antoni, Ganzer, and Berge (2019) and Dauth and Eppelsheimer (2020).[6] We then identify the relevant job separation events, namely layoffs and quits, and retain the 12 months before

6. Details are provided in Appendix 3.C.

and after each event, following Fos et al. (2025).[7] To ensure that separations represent meaningful income shocks, we restrict the sample to individuals with at least 6 months of prior employment.

Second, we turn to the district-level component of the data construction, where the key challenge is that platform roll-out is staggered across districts. To avoid the well-documented problems of two-way fixed effects (TWFE) estimators under treatment effect heterogeneity (e.g., Goodman-Bacon 2021; Borusyak, Jaravel, and Spiess 2024), we *stack* the district-level data following Cengiz et al. (2019) and Deshpande and Li (2019). After merging the information on platform entry, we identify cohorts of districts treated in the same period. For each cohort, we define a separate sub-experiment that includes the treated districts and all not-yet-treated districts within a symmetric observation window around the entry event. The control group in each sub-experiment therefore consists of both not-yet-treated districts, which experience entry only after the relevant time window ends, and never-treated districts. For the latter, we restrict attention to the 50 never-treated district-free cities with more than 100,000 inhabitants.[8] We use a window of 19 months before and after entry for each sub-experiment, which allows us to retain 16 of the 23 entry events. Overall, this procedure yields seven sub-experiments in which districts experiencing platform entry are compared to districts where platforms have not yet entered. The construction of the individual-level dataset follows the same logic. However, the greater statistical power afforded by individual-level variation allows us to restrict the sample to ever-treated districts only, reducing concerns about selective platform entry. Further details on the differences in the preparation process are provided in Appendix 3.C.

As the third and final step, we apply Coarsened Exact Matching (CEM, Iacus, King, and Porro 2011) at the job separation episode level. Within each sub-experiment, we create cells using the exact values of the following coarsened variables at the time of job separation: birth cohort (5-year intervals), gender, educational level (3 tiers), wage (quintiles, specific to the 5-year birth cohorts), and employment status (6 groups). Subsequently, we drop all groups containing job separation episodes in treatment districts but no individuals in assigned control districts and vice versa. Note that our empirical design does not impose that the groups are identical in terms of observables, nor do we force a one-to-one matching of individuals in treated and control districts. Instead, this step guarantees overlap and reasonable balance along key socio-economic dimensions between the two groups.

7. The 12-month window also corresponds to the maximum duration of unemployment insurance benefits for most workers.

8. This matches the official definition of a large city in Germany and makes never-treated districts more comparable to treated ones. Indeed, 22 of the 23 entry events observed before the end of 2017 occur in such districts.

For the individual-level analysis, this procedure yields a sample with more than 94,000 individuals and a total of 103,633 distinct job separation events. Table 3.1 summarizes the means of the variables used for the CEM, separately for the treatment and the control group. The table also displays the difference between the groups, standardized by the pooled standard deviation. Treatment and control group are similar along most dimensions, and differences are relatively small in terms of pooled standard deviations. The only exceptions consist of wages and education, with individuals in the treatment group earning on average higher wages and being more educated. Beyond these summary statistics, CEM also ensures comparability between the groups over the entire distribution of the continuous variables. To exemplify this, Figure 3.A.2 plots kernel densities for age and wage at job separation. Overall, the distributions in the control group closely track the ones in the treatment group. What is more, due to the exact matching on individual cells, CEM balances not only the marginal distributions of individual covariates but also their joint distribution. For the district-level analysis, we conclude the data preparation by aggregating the data at the district-month level, separately by sub-experiment.

Outcome Variables: We analyze three groups of outcome variables: (*i*) employment and unemployment benefit receipt, (*ii*) job-search status, and (*iii*) participation in ALMPs.

The key outcome in the first group is an indicator for employment in a mini-job, the main contractual type used by two of the three platforms to hire riders (*Lieferando* and *Foodora*). At the time, this type of marginal employment contract allowed for monthly incomes up to 450 € or for short-term employment relationships (less than 70 days or 3 months), exempting them from income tax. Nevertheless, workers in mini-job contracts are subject to compulsory insurance in the statutory pension scheme (since 2013) and are thus observed in our administrative dataset. We also construct an indicator for employment that is subject to social security contributions, allowing us to assess whether gig work substitutes for other employment or is taken up in addition to it.

Regarding unemployment benefits (UB), Germany distinguishes between two types of support. First, *UB I* (German Social Code, chapter III) is a wage-replacement scheme for unemployed workers (or those working fewer than 15 hours per week) who contributed for at least 12 months during the previous two years. Benefits replace 67% of prior earnings for individuals with dependent children and 60% otherwise,[9] based on average remuneration during the 12 months preceding job separation. Benefit duration depends on prior contributions and is capped at 12 (24) months for individuals below (above) 50. Second, *UB II* (German Social Code, chapter II) is a means-tested welfare program for individuals who do

9. Information on workers' family status is not included in the SIAB-R-7517.

Table 3.1. Matching Variables: Treatment vs. Control Districts

	Treatment Mean	Control Mean	Standardized Difference
Female	0.482	0.477	0.009
Age	37.194	37.242	-0.004
Wage	71.943	64.991	0.127
Education:			
Vocational Training	0.521	0.564	-0.087
University	0.265	0.216	0.114
Employment Status:			
Employed (social security)	0.722	0.686	0.078
Trainee	0.049	0.056	-0.030
Marginal Employment	0.179	0.211	-0.081
Employed (partial retirement)	0.007	0.009	-0.019
Interns/Student Trainees	0.033	0.026	0.044
Other Employment	0.010	0.012	-0.021
Observations	1,768,170	3,478,073	

Notes: The table shows average values for the variables used for the CEM procedure, comparing treatment against control districts. Wages are at daily frequency (in €) and deflated with respect to 2015. The standardized difference is calculated as the difference in means divided by the pooled standard deviation, which is defined as the square root of the sample-size weighted average of the two standard deviations.

not qualify for or have exhausted *UB I*. Benefits are flat-rate and depend on household characteristics. Since the program is primarily administered by local authorities, our dataset does not contain information on *UB II* amounts. As our focus is on income insurance following job separation, we use *UB I* receipt and benefit amounts as our main outcomes. Importantly, benefit receipt does not preclude paid work. Under *UB I*, individuals may earn up to 165 € per month in addition to benefits, with a 100% marginal tax thereafter.

As the second group of outcomes, we study a set of dummies for individuals' registered job-search status. These include indicators for unemployed jobseekers, unemployed individuals unable to work, not unemployed jobseekers[10], and individuals who are not currently searching for a job.

Finally, we construct a set of indicators for ALMP participation.[11] We distinguish between (*i*) vocational integration measures, such as job application training or language courses aimed at facilitating labor market integration; (*ii*) employment-entry subsidies, which cover up to 50% of earnings upon re-employment; (*iii*) vocational training programs, typically targeted at younger workers and including

10. This group consists primarily of individuals anticipating unemployment.

11. Our analysis focuses on measures under German Social Code, chapter III, which accounted for roughly 99.8% (own calculations based on data provided by the Federal Employment Agency) of all ALMP expenditures in 2015.

Table 3.2. Average Outcomes Around Job Separation

	All Obs.		Pre-JS		Post-JS	
	N	Mean	N	Mean	N	Mean
Main Outcomes:						
Employed in mini-job	5,246,243	0.162	2,548,922	0.215	2,697,321	0.112
Employed subject to social security	5,246,243	0.599	2,548,922	0.658	2,697,321	0.542
Unemployment Benefit I received	5,246,243	0.053	2,548,922	0.012	2,697,321	0.092
Amount of UB I received	278,049	867.716	30,544	754.248	247,505	881.718
Job-Seeking:						
Unemployed jobseeker	5,246,243	0.077	2,548,922	0.039	2,697,321	0.113
Unemployed unable to work	5,246,243	0.002	2,548,922	0.001	2,697,321	0.002
Not unemployed jobseeker	5,246,243	0.074	2,548,922	0.079	2,697,321	0.069
Not-seeking work	5,246,243	0.017	2,548,922	0.017	2,697,321	0.017
ALMP:						
Vocational integration	5,246,243	0.004	2,548,922	0.002	2,697,321	0.005
Starting an employment	5,246,243	0.005	2,548,922	0.003	2,697,321	0.006
Vocational training	5,246,243	0.003	2,548,922	0.005	2,697,321	0.001
Vocational retraining	5,246,243	0.005	2,548,922	0.003	2,697,321	0.007

Notes: The monthly amount of *UB I* is computed by multiplying the daily benefit (in €, deflated with respect to 2015) by the number of days of *UB I* receipt in a given month.

career counseling or other support measures outside the standard apprenticeship system (Roesler et al. 2021); and (*iv*) vocational retraining programs, consisting of curriculum-based courses designed to provide specific skills (Osikominu 2013). The specific support provided under each measure varies substantially and is determined at the discretion of local job centers and caseworkers (Dauth 2020; Fitzenberger, Osikominu, and Paul 2023).

3.2.3 Summary Statistics

Table 3.2 reports sample sizes and averages for outcome variables in the individual-level dataset, both overall and relative to the job separation event. The unconditional monthly probability of mini-job employment is 16.2%. In the 12 months prior to separation, this probability is 21.5%, falling to 11.2% thereafter. We observe a similar decline in employment subject to social security. Unemployment benefit receipt is relatively rare prior to separation, at 1.2%, but increases to 9.2% afterwards. Conditional on receipt the increase in the amount of *UB I* benefits is only moderate.

As for job-seeking outcomes, averages move in the expected direction around job separation, with an increase in registered jobseekers and a decline in individuals not registered as unemployed jobseekers. Participation in ALMPs also rises after separation, with the exception of vocational training. This is not surprising, as layoffs and quits from vocational training contracts are themselves classified as job separations.

We next stratify the sample along the three dimensions used in the econometric specification. The individual-level design exploits within-unit variation around job

Table 3.3. Average Outcomes by Platform Availability, Treatment vs. Control

	Post-Entry		Pre-Entry		
	(1)	(2)	(3)	(4)	(5)
	Treatment	Control	Treatment	Control	DiD
Main Outcomes:					
Employed in mini-job	-0.081	-0.100	-0.102	-0.113	0.008
Employed subject to social security	-0.137	-0.105	-0.128	-0.109	-0.013
Unemployment Benefit I received	0.076	0.075	0.083	0.082	-0.001
Amount of UB I received	132.599	113.292	142.558	123.720	0.469
Job-Seeking:					
Unemployed jobseeker	0.070	0.072	0.080	0.074	-0.007
Unemployed unable to work	0.002	0.002	0.002	0.002	0.000
Not unemployed jobseeker	-0.007	-0.009	-0.010	-0.011	0.001
Not-seeking work	-0.001	-0.001	-0.001	0.000	0.000
ALMP:					
Vocational integration	0.003	0.003	0.002	0.002	0.000
Starting an employment	0.004	0.002	0.003	0.003	0.001
Vocational training	-0.002	-0.003	-0.003	-0.004	0.000
Vocational retraining	0.004	0.004	0.004	0.004	0.000

Notes: All values refer to post-job separation periods, from which we subtracted the equivalent pre-job separation values. The monthly amount of *UB I* is computed by multiplying the daily benefit (in €, deflated with respect to 2015) by the number of days of *UB I* receipt in a given month. All other variables are indicators equal to one if an individual can be assigned to the respective employment status. The values in Column (5) are obtained by taking differences-in-differences, i.e., as (5) = (1) – (2) – ((3) – (4)).

separations, comparing individuals in treated and control districts before and after platform entry. Table 3.3 illustrates this by reporting averages separately for treatment and control groups, with and without platform presence. To simplify the exposition, we collapse one dimension by focusing on differences between post- and pre-separation outcomes. Taking differences-in-differences across Columns (1) to (4) yields a first estimate of the effect of platform presence, shown in Column (5).

The results in the final column are broadly consistent with the expected effects of platform availability. Most importantly, the excess probability of mini-job employment is 0.8 percentage points, suggesting that recently separated individuals take up gig work. Given the decline in standard employment, this is unlikely to reflect general labor market trends. At the same time, changes in unemployment benefit receipt are negative but small. We also find a decline in registered unemployment, alongside a modest increase in non-unemployed jobseekers and new employment entries. This pattern is consistent with gig work serving as a temporary buffer while individuals search for more stable employment (Heiland 2019). Apart from a small increase in the probability of starting an employment with a hiring subsidy, we do not detect meaningful differences in ALMP participation.

3.3 Empirical Strategy

The aim of this chapter is to document how gig platform availability affects responses to unemployment shocks. We hypothesize that flexible platform work reduces labor market frictions by enabling rapid labor supply and providing a tool to buffer income losses. In the German context, we therefore expect an increase in mini-jobs, the employment form commonly used for platform-based work. Identifying these effects is not straightforward. Even with valid identification, gig work may not translate into higher observed mini-job incidence if it is infra-marginal, for example by displacing other mini-jobs rather than expanding overall take-up. In addition, a share of gig workers operates as freelancers, implying that our estimates provide a lower bound on the true effect. Beyond gig-work take-up, we study the effects of platform entry on other labor market outcomes.

3.3.1 District-Level Specification

The ideal experiment would consist of randomly offering recently unemployed workers the opportunity to engage in gig work while preventing others from doing so. Comparing outcomes across these groups would then identify the causal effect of gig-work availability on behavior during unemployment. However, as discussed in Section 3.2.1, platform entry is not random, as firms strategically choose when and where to operate. We address this challenge by exploiting the quasi-experimental variation generated by staggered gig-platform entry across German districts.

Using only *observations following a job separation*, we collapse the outcome variables to the district level. Considering an observation in district d, month t, and sub-experiment s, we estimate the following stacked difference-in-differences specification:

$$y_{dt} = \alpha \, TreatDistrict_{ds} \times PostEntry_{ts} + \phi_{ds} + \psi_{ts} + \Gamma X_{dt} + \varepsilon_{dts} \qquad (3.1)$$

where y_{dt} denotes the average district-level outcome, $TreatDistrict_{ds}$ is an indicator for treated districts, and $PostEntry_{ts}$ equals one if at least one gig delivery platform is present throughout month t in the treated districts.[12] The district (ϕ_{ds}) and month (ψ_{ts}) fixed effects are included separately for each sub-experiment, reflecting that time windows and the composition of treatment and control groups differ by design. This ensures identification comes from within- rather than across-sub-experiment variation (Cengiz et al. 2019).[13] X_{dt} is a vector of time-varying district-level controls

12. As platform launches occur throughout the month, we define platform availability based on presence throughout the entire month. Our results are robust to alternative definitions.

13. For the same reason, we do not include district-month fixed effects, even though the data structure would in principle allow for them. These would exploit cross-sub-experiment variation within the same district-month, which is the comparison the stacked design is intended to avoid.

capturing demographic and economic factors not absorbed by the fixed effects.[14] Regressions are weighted by the number of post-separation observations in each district-month cell, and standard errors are clustered at the district level, following Deshpande and Li (2019).

The coefficient α, the main parameter of interest, is the DiD estimator. It measures the conditional difference in the post-separation outcome y_{dt} between treated and control districts, before and after platform availability. Identification rests on a parallel trends assumption: absent platform entry, the outcome in the treatment group would have evolved as in the control group, conditional on fixed effects and the time-varying controls X_{dt}. Our stacked approach offers two advantages over standard two-way fixed effects (TWFE) specifications. First, TWFE regressions estimate weighted averages of treatment effects that can include "forbidden comparisons" between treated and already-treated units (De Chaisemartin and d'Haultfoeuille 2020). The stacked design avoids this by constructing a clean control group of not-yet-treated districts for each treated cohort. Second, in the individual-level regression, the stacked dataset allows us to fully saturate the specification with respect to all lower-level interactions of the treatment, reducing concerns about unobserved heterogeneity.

We highlight that α in Equation (3.1) measures a lower bound on actual mini-job take-up, while the coefficients for unemployment benefits, job search, and ALMP participation capture the effect of platform *availability* rather than of *working* on a platform. This focus on Intention-to-Treat (ITT) estimates is dictated by the setting but also desirable in its own right. From a policy perspective, the relevant question is whether platforms should be allowed to operate and, if so, under which limitations, which depends on the effect of their availability.

3.3.2 Individual-Level Specification

We extend the district-level approach to the individual level, which offers three advantages. First, the disaggregation allows us to include worker and district-month fixed effects, absorbing individual heterogeneity and aggregate time-varying shocks that the district-level specification cannot capture. Second, exploiting pre- vs. post-separation variation yields a triple-DiD design that relies on a less stringent parallel trends assumption. Third, the increased power lets us restrict the sample to ever-treated districts only, reducing selection concerns related to platform entry.

14. Controls constructed from the SIAB (measured at the time of job separation) are the share of German individuals in the dataset, the share of recently separated workers who are German, who hold a tertiary degree or vocational training, or who work in low-skill occupations, as well as their average wage (in logs) and age. We additionally include overall and male population size, the number of asylum seekers (all three in logs), GDP per capita and per worker, and disposable income per capita from the German Federal Statistical Office (*Regionalstatistik*).

To implement this, we retain all observations, before and after a job separation episode i, and estimate the following specification:

$$\begin{aligned} y_{it} = {} & \alpha\, PostJS_{it} \times TreatDistrict_{ds} \times PostEntry_{ts} \\ & + \beta\, PostJS_{it} \times TreatDistrict_{ds} + \gamma\, PostJS_{it} \times PostEntry_{ts} \\ & + \delta\, PostJS_{it} + \phi_{is} + \psi_{dts} + \varepsilon_{its} \end{aligned} \tag{3.2}$$

where $PostJS_{it}$ is a dummy equal to one in the month of separation and the 12 months following it, and zero in the 12 months preceding it. ϕ_{is} are episode-sub-experiment fixed effects, absorbing time-invariant episode characteristics that may correlate with the outcomes, such as *UB I* entitlement or age at separation. ψ_{dts} are district-month-sub-experiment fixed effects, conditioning on time-varying local labor market shocks that could correlate with both outcomes and the timing of job separation. The other indices and variables are defined as before. Standard errors are double-clustered at the episode and district-month level.

As before, the parameter of interest is α. Unlike in the simple DiD, however, it does not measure changes in post-separation outcomes relative to a no-entry counterfactual. Instead, the triple DiD exploits within-individual variation between pre- and post-separation outcomes, differencing out factors that vary over time across treated and control districts but affect both periods equally. Identification therefore requires that unobserved factors neither differ systematically between treatment and control districts nor differentially affect outcomes before and after separation. Equivalently, the triple DiD relaxes the standard parallel trends assumption: outcomes in treated and control districts need not evolve in parallel, provided the difference in trends is the same before and after separation (see Olden and Møen 2022).

We retain the stacked setup for two reasons. First, $TreatDistrict_{ds}$ and $PostEntry_{ts}$ vary by sub-experiment, allowing the second and third terms in Equation (3.2) to be defined cohort-by-cohort rather than pooled across separation timings, which would impose a more restrictive homogeneity assumption. Second, the stacked structure preserves a clean control group of not-yet-treated districts for each entry event. Estimating a TWFE regression on the panel of recently unemployed individuals with a single time-varying indicator for platform availability in their district would forgo both the triple DiD structure and the cohort-specific controls, reintroducing the very comparisons of later-treated against already-treated units that the stacked design is meant to avoid (De Chaisemartin and d'Haultfoeuille 2020).

Event Study: To assess the plausibility of parallel trends and trace the dynamics of platform entry, we also estimate an event-study version of Equation (3.2), replacing the binary $PostEntry_{ts}$ indicator with a set of leads and lags relative to entry, indexed by $m \in \{-19, \dots, 19\}$. We bin the endpoints to capture average treatment effects beyond the 19-month window and omit the first lead ($m = -1$) as the reference period.

3.4 Results

We begin with a descriptive analysis of the data, then present district-level baseline estimates and tests for heterogeneity in take-up, followed by the individual-level results. Sections 3.4.4 and 3.4.5 assess the plausibility of our identifying assumptions and examine the underlying economic mechanism.

3.4.1 Graphical Evidence

We first inspect the data by plotting event studies around job separations (Freyaldenhoven, Hansen, and Shapiro 2019; Clarke and Tapia-Schythe 2021), comparing pre- and post-separation outcomes between treatment and control groups during periods after platform entry.[15] For each group, we estimate:

$$y_{it} = \sum_{m \neq -1} \beta_m \, PostJS_{it}^{m} + \phi_s + \varepsilon_{its} \tag{3.3}$$

for $m \in \{-12, \ldots, 12\}$, where $PostJS_{it}^{m}$ indicates the job separation event happening in month $t+m$.[16] Sub-experiment fixed effects absorb latent differences across stacking layers, and the month before separation is omitted as the reference period. We include no further controls to provide a raw view of the data. Figure 3.3 plots the series $\{\beta_m\}_{m=-12}^{12}$ for each main outcome alongside 95% confidence intervals, with red (blue) lines referring to workers in districts with (without) platform availability.

Panel (a) of Figure 3.3 plots the probability of mini-job employment in event-time around job separation. As expected, both groups experience a drop in this outcome following a separation, but the reduction is smaller for workers in districts with platform availability. Under the assumption that mini-job employment proxies for gig work, this provides evidence that some individuals actually take up gig work following an employment termination, if they have the opportunity to do so.

Panel (b) shows the corresponding results for employment subject to social security. Pre-separation trajectories are very similar across groups. After separation, both groups experience a substantial drop, followed by a stronger recovery in districts without platform availability, hinting at substitution between gig work and non-marginal dependent employment. Comparing Panels (a) and (b), the difference appears driven by treated individuals remaining in mini-jobs while control individuals gradually transition into traditional employment. We scrutinize this further in Sections 3.4.3 and 3.4.5, where we find some evidence of such a substitutive relationship in the medium run.

15. This restriction ensures that differences between pre- and post-entry periods cannot drive our results. Figure 3.A.3 shows an alternative specification comparing pre- against post-entry periods within the treatment group, with very similar results.

16. Endpoints are not binned as the number of parameters is limited by the sample construction.

Figure 3.3. Main Outcomes Relative to Job Separation by Platform Availability

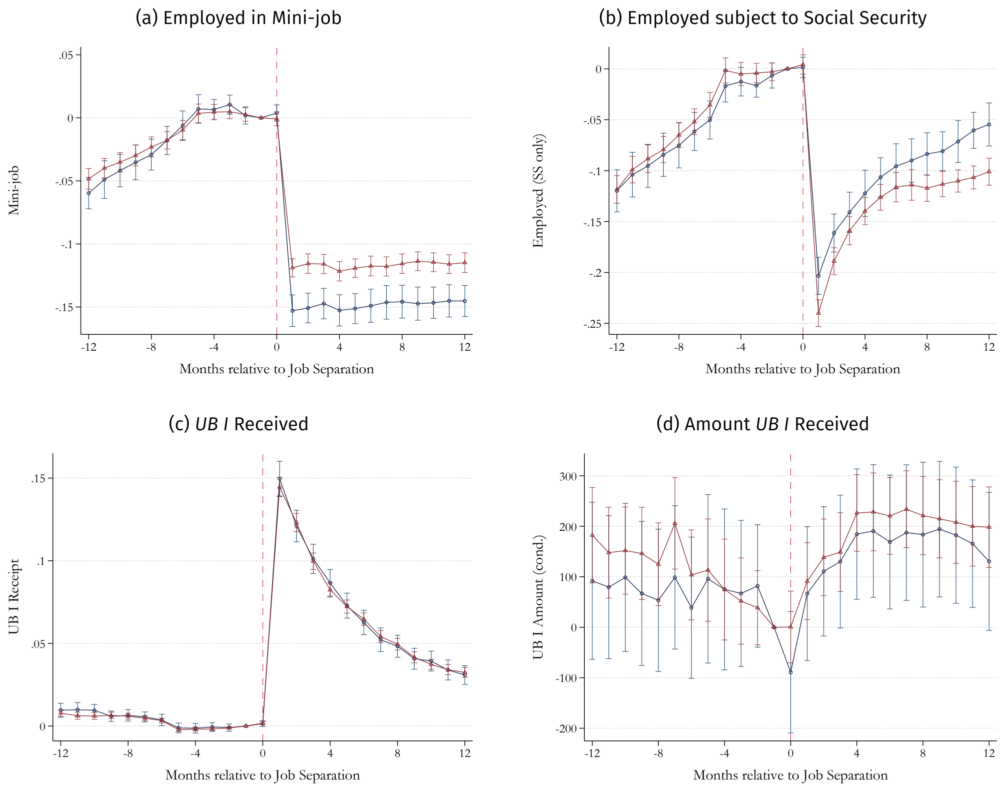

Notes: The figures plot averages of the outcome variables relative to a job separation by platform availability. Red (blue) lines refer to workers in treatment (control) districts after platform entry. Coefficients are obtained by estimating Equation (3.3) separately for each group. The month before separation is omitted. 95% confidence intervals based on standard errors double-clustered by episode and district-month.

The remaining panels of Figure 3.3 demonstrate that the relative increase in marginal employment is not associated with lower reliance on unemployment benefits, neither at the extensive (Panel c) nor at the intensive margin (Panel d).[17] Panel (c) exhibits a spike in benefit receipt right after separation, with a declining pattern likely reflecting gradual expiration or take-up of other employment. The amount of *UB I* received, conditional on receipt, is relatively flat across the 25-month window (Panel d).

Overall, Figure 3.3 illustrates three important points. First, pre-separation behavior nearly coincides across groups, supporting the parallel trends assumption. Second, the absence of significant post-separation differences in Panels (c) and (d) suggests that platform availability leaves *UB I* take-up unchanged and does not

17. Positive values prior to separation are possible and observed for both variables, since workers may receive benefits from previous unemployment spells.

front- or backload its receipt. Third, the results reduce concerns that platform entry correlates with voluntary separations rather than layoffs: since voluntary unemployment triggers a three-month *UB I* suspension, such selection would generate a visible difference around the third month post-separation, which is not apparent in the graphs.

3.4.2 District-Level Results

We now estimate Equation (3.1) at the district level, showing that gig-work take-up is discernible in the aggregate. Table 3.4 reports the main outcomes. Mini-job employment increases by 0.5 percentage points (5% of the control mean) in districts with platform availability, suggesting that recently unemployed workers rely on gig work to buffer income losses following separation. The estimate is smaller than the raw difference-in-differences in Table 3.3 (0.8 percentage points), reflecting the role of the controls and fixed effects. The coefficients on other employment and unemployment benefits are insignificant.

We also present the corresponding estimates for the job-search and ALMP outcomes in Tables 3.B.1 and 3.B.2. The coefficients are generally small and insignificant, with the only exception being a positive coefficient on the probability of starting an employment, that is, taking up a job with a hiring subsidy granted to employers. Given the low variation in the remaining outcomes, we conjecture that more power may be needed to detect other effects of platforms' presence.

Since gig work differs from traditional employment (Abraham and Houseman 2019; Collins et al. 2019), we expect both engagement in gig labor and substitution away from unemployment benefits to vary across worker types. We test for heterogeneity in gig-work and UB take-up by gender, age, education, and combinations thereof, estimating Equation (3.1) separately for the sub-samples obtained by stratifying according to these variables. This exercise also serves to validate our empirical strategy: if our approach captures variation driven by gig workers, effects on mini-job take-up should be stronger in groups with higher representation among platform workers. To benchmark these groups, we use a cross-sectional dataset received from *Lieferando* in September 2020 containing basic demographics on all 5,255 platform workers active at the time. Of these, 44% were German-born and 90% were male. The age distribution (Figure 3.A.4) is right-skewed with most mass between 20 and 30, and a mean (median) of 28.4 (26.5).[18] These figures broadly match Heiland (2019) and Urzì Brancati, Pesole, and Fernández-Macías (2020), who also document relatively high education levels among platform workers. Verifying that mini-job uptake is stronger among male, young, and educated workers hence speaks in favor of our identification strategy.

18. This is consistent with Collins et al. (2019), who use US tax data and find that participation in online platform work peaks around age 30 and declines beyond 35.

Table 3.4. District-Level Effects of Gig Platform Availability: Main Outcomes

	Mini-job	Employed (SS only)	UB I Receipt	UB I Amount (cond.)
	(1)	(2)	(3)	(4)
Platform	0.005**	-0.005	0.003	-0.745
	(0.002)	(0.005)	(0.003)	(12.662)
R^2	0.568	0.732	0.500	0.520
Observations	17,121	17,121	17,121	17,121
Control Mean	0.108	0.523	0.093	851.120
Controls	✓	✓	✓	✓
District-Exp. fixed effects	✓	✓	✓	✓
Month-Exp. fixed effects	✓	✓	✓	✓

Notes: OLS estimates. All regressions include district-sub-experiment and month-sub-experiment fixed effects and the control variables described in Section 3.3. Regressions are weighted by the number of post-separation observations in each month-district cell. Standard errors are clustered at the district level. Mini-job is an indicator for marginal employment, Employed (SS only) for employment subject to social security excluding mini-jobs, and UB I Receipt for receipt of unemployment benefits. *Platform* equals one if at least one platform has been present in the district for the entire month. The monthly amount of *UB I* is computed by multiplying the daily benefit (in €, deflated with respect to 2015) by the number of days of *UB I* receipt in a given month. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 3.4 reports the heterogeneity results for mini-job and *UB I* take-up across subgroups, with the estimated coefficient for the full sample shown for reference. While most coefficients are not statistically different from one another across subgroups, the sign and magnitude of the point estimates are informative and align with our priors. For workers below 40, both the mini-job and *UB I* estimates are positive but insignificant. Further restricting the sample to males, the mini-job coefficient slightly increases while the coefficient for *UB I* approaches zero. Among workers without tertiary education, we find no substantial mini-job effect. By contrast, the coefficient is positive and significant for high-skilled males below 40, at 0.01, amounting to roughly 13% of the variable's mean in the control group. As a placebo, we also present the corresponding estimates for females below 40 and males above age 40. As expected, the estimated coefficients for these latter groups are statistically indistinguishable from zero.[19]

As a further robustness check, in Figure 3.A.6, we repeat the baseline regression for each subgroup, interacting the controls with dummies for each sub-experiment (Panel a) and the post-platform periods within each sub-experiment (Panel b). In this way, we relax the assumption that the relationship between controls and out-

19. Figure 3.A.5 also reports estimates stratified by citizenship at the time of job separation. The baseline results hold for both Germans and foreigners, even though the limited number of foreigners precludes precise comparisons. Note that our sub-population of interest also differs from the broader rider population: 56% of *Lieferando*'s active platform workers in September 2020 were not born in Germany, compared to 12% in our sample.

Figure 3.4. District-Level Effects of Gig Platform Availability by Subgroup

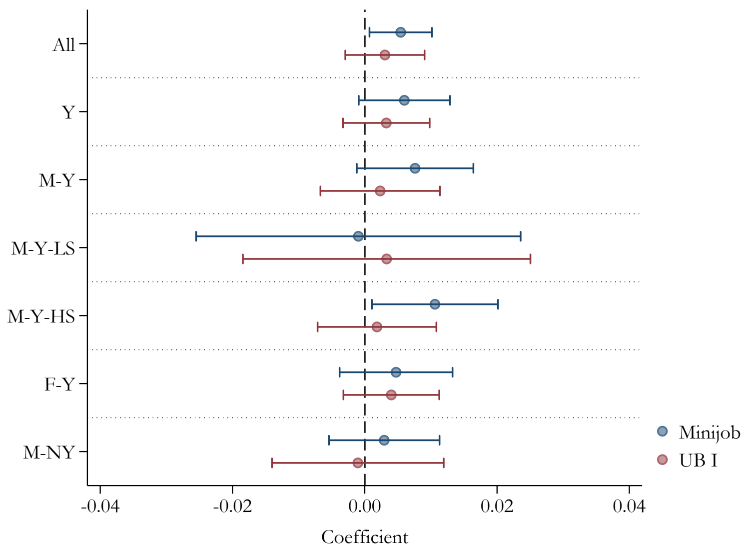

Notes: The figure plots the district-level OLS estimates for the coefficient of interest obtained from separate regressions for each subgroup. *Y/NY*: below/above age 40 at job separation, *M/F*: male/female, *HS/LS*: tertiary/no tertiary education. Confidence intervals based on robust standard errors clustered at the district level. All regressions include fixed effects at the district-sub-experiment and month-sub-experiment level as well as the control variables described in Section 3.3. The regressions are also weighted by the number of post-job separation observations in each month-district cell. Mini-job is an indicator for working under a contract classified as marginal employment. *UB I* indicates any receipt of unemployment benefits.

comes is constant across treatment cohorts (Panel a) or before and after platform entry (Panel b). Even though precision decreases, the point estimates remain similar to the baseline, supporting the robustness of our approach.

3.4.3 Individual-Level Results

We now turn to the individual-level specification. As discussed in Section 3.3, the triple DiD exploits both the staggered entry of platforms across districts and within-individual variation between pre- and post-separation outcomes. All specifications include episode and district-month fixed effects within sub-experiments. The coefficient of interest is on the triple interaction, capturing the effect of gig labor availability.[20]

The results obtained by estimating Equation (3.2) for the main outcomes are shown in Table 3.5, confirming the patterns from the district-level analysis. Column (1) shows that workers are significantly more likely to be in mini-job employ-

20. All lower-level interactions not subsumed by the fixed effects are included in the regressions throughout the paper, but not reported for brevity.

Table 3.5. Individual-Level Effects of Gig Platform Availability: Main Outcomes

	Mini-job	Employed (SS only)	UB I Receipt	UB I Amount (cond.)
	(1)	(2)	(3)	(4)
PostJS x Platform	0.017***	-0.003	0.000	-5.133
	(0.005)	(0.006)	(0.003)	(37.395)
R^2	0.558	0.654	0.378	0.859
Observations	3,265,524	3,265,524	3,265,524	138,506
Control Mean	0.170	0.589	0.054	849.980
Interactions	✓	✓	✓	✓
District-Month-Exp. fixed effects	✓	✓	✓	✓
Episode-Exp. fixed effects	✓	✓	✓	✓

Notes: OLS estimates, with standard errors double-clustered at the individual and district-month level in parentheses. All regressions include fixed effects at the district-month-sub-experiment and episode-sub-experiment level. *Interactions* refers to the non-absorbed components of the treatment variable, i.e., *PostJS* × *TreatDistrict*, *PostJS* × *PostEntry*, and *PostJS*. Mini-job is an indicator for working under a contract classified as marginal employment. *PostJS* is an indicator for all post-job separation periods. *Platform* is a dummy variable equal to one if at least one platform has been present in the district for the entire month. The monthly amount of *UB I* is computed by multiplying the daily benefit (in €, deflated with respect to 2015) by the number of days of *UB I* receipt in a given month. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

ment when a platform is available. The estimate is 1.7 percentage points, significant at the 1% level, and roughly three times larger than the district-level coefficient. In terms of the control mean, the estimate increases somewhat to around 10%. For the other outcomes, the signs of the coefficients match those obtained before across the board, but they remain insignificant.

We provide results from a further validation exercise in Figure 3.A.7, where we study gig work take-up by 5-year age brackets. Mini-job take-up increases significantly in response to platform entry for workers aged 20 to 30, with no change for older groups compared to the counterfactual. This pattern closely tracks the age distribution of *Lieferando* riders in Figure 3.A.4.[21]

In Table 3.6, we provide the corresponding estimates for the job-seeking outcomes. Column (1) shows the results for unemployed jobseekers, which appear not to be affected by platform presence. In Column (2), we report the estimates for an indicator for unemployed individuals who are unable to work. In line with expectations, the coefficient is precisely estimated at zero. Column (3) reveals a significant increase in non-unemployed individuals searching for a job. One explanation for this pattern is that individuals take up gig work while continuing to search for traditional and better-paying alternatives. Lastly, we find no substantial change in the number of non-seeking workers (Column 4).

To complement the portfolio of labor market outcomes, Table 3.7 shows the results for ALMP participation. For vocational integration, we do not observe any

21. Jackson (2022) finds stronger engagement in gig labor following unemployment shocks for older workers, likely reflecting differences in the type of gig labor and the welfare system considered.

Table 3.6. Individual-Level Effects of Gig Platform Availability: Job-Seeking Outcomes

	Unempl. Jobseeker (1)	Unempl. Unable (2)	Not Unempl. Jobseeker (3)	Not Seeking (4)
PostJS x Platform	0.002	0.000	0.008**	-0.002
	(0.003)	(0.000)	(0.003)	(0.001)
R^2	0.422	0.105	0.403	0.537
Observations	3,265,524	3,265,524	3,265,524	3,265,524
Control Mean	0.080	0.002	0.075	0.018
Interactions	✓	✓	✓	✓
District-Month-Exp. fixed effects	✓	✓	✓	✓
Episode-Exp. fixed effects	✓	✓	✓	✓

Notes: OLS estimates, with standard errors double-clustered at the individual and district-month level in parentheses. All regressions include fixed effects at the district-month-sub-experiment and episode-sub-experiment level. *Interactions* refers to the non-absorbed components of the treatment variable, i.e., *PostJS* × *TreatDistrict*, *PostJS* × *PostEntry*, and *PostJS*. All outcome variables are indicators equal to one if an individual is assigned to the respective job-search status. *PostJS* is an indicator for all post-job separation periods. *Platform* is a dummy variable equal to one if at least one platform has been present in the district for the entire month. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

significant effect. As to the other outcomes, we detect significant increases in the number of individuals starting an employment[22] and vocational training, whereas we observe a quantitatively small decline in the probability of participating in vocational retraining measures. Summing over the coefficients, these results suggest an overall increase in ALMP participation in response to platform availability. This pattern stands in contrast to the negative consequences of gig work highlighted in Jackson (2022). One potential explanation for our result is that gig work acts as a signaling device (for experimental evidence, see Adermon and Hensvik 2022) for motivated individuals, who consequently are more likely to be assigned to training measures by their caseworkers. The decline in vocational retraining is consistent with this interpretation. Among the ALMP measures we consider, retraining is the only one requiring workers to leave the labor market for an extended period to acquire new skills. Workers using gig work as a temporary income buffer thus face a higher opportunity cost of committing to such measures and may prefer to remain in the labor market.

3.4.4 Identification and Heterogeneity Across Platforms

Identification: In this section, we present results from event studies that allow us to discuss both the plausibility of the parallel trends assumption and potential economic mechanisms. Figure 3.5 plots the estimated coefficients on the triple in-

22. Note that the subsidies specifically exclude mini-job employment.

Table 3.7. Individual-Level Effects of Gig Platform Availability: Participation in ALMPs

	Voc. Integration (1)	Start Employment (2)	Voc. Training (3)	Voc. Retraining (4)
PostJS x Platform	-0.001	0.002**	0.002**	-0.002**
	(0.001)	(0.001)	(0.001)	(0.001)
R^2	0.218	0.295	0.556	0.416
Observations	3,265,524	3,265,524	3,265,524	3,265,524
Control Mean	0.004	0.005	0.003	0.005
Interactions	✓	✓	✓	✓
District-Month-Exp. fixed effects	✓	✓	✓	✓
Episode-Exp. fixed effects	✓	✓	✓	✓

Notes: OLS estimates, with standard errors double-clustered at the individual and district-month level in parentheses. All regressions include fixed effects at the district-month-sub-experiment and episode-sub-experiment level. *Interactions* refers to the non-absorbed components of the treatment variable, i.e., *PostJS* × *TreatDistrict*, *PostJS* × *PostEntry*, and *PostJS*. All outcome variables are indicators equal to one if an individual is recorded in the respective ALMP. *PostJS* is an indicator for all post-job separation periods. *Platform* is a dummy variable equal to one if at least one platform has been present in the district for the entire month. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

teraction $PostJS_{it} \times TreatDistrict_{ds} \times PostEntry^{m}_{ts}$ for the main outcomes, tracing the dynamics of treatment effects in the months before and after platform entry.[23]

We first use the event studies to check for anticipation effects: if the effects are driven by platform availability, recently unemployed workers should not respond to their future entry. Across all panels, we find no strong evidence of pre-trends. Panel (a) shows a slight pre-entry uptick in mini-job employment, but the coefficients are insignificant and substantially smaller than the post-entry increase. Panels (b) and (d) display minor deviations more than a year before platform entry, but the coefficients close to treatment are insignificant and near zero. The discontinuous changes following platform entry are thus unlikely to reflect differential trends in omitted variables.

The estimated post-treatment coefficients also let us trace the dynamic treatment effect following platform entry. For mini-job employment, we observe no substantial effect in the first four months but a strong and persistent increase thereafter, consistent with platforms gradually expanding their operations within the cities they enter. Unemployment benefit take-up rises on impact and remains slightly elevated, though the coefficients are insignificant. The patterns in Panels (b) and (d) are less clear, with most coefficients fluctuating around zero.

Endogeneity in the timing of platform entry is unlikely to drive the results. As documented in Figure 3.2, entry timing is strongly predicted by population size and density. Both variables are relatively stable over time, alleviating concerns about

23. Analogous results for the other outcomes are in Figures 3.A.8 and 3.A.9.

Figure 3.5. Dynamic Effects Relative to Platform Entry: Main Outcomes

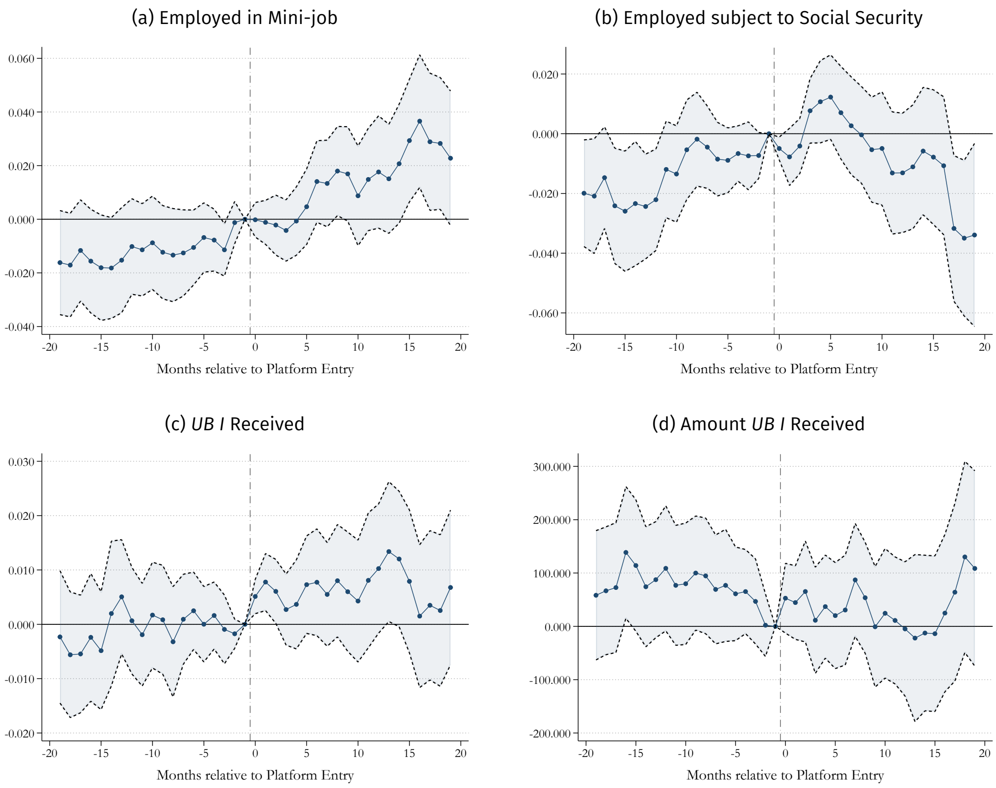

Notes: These figures report estimated coefficients on the interaction terms $PostJS_{it} \times TreatDistrict_{ds} \times PostEntry^{m}_{ts}$ of the event-study specification described in Section 3.3, for $m \in \{-19, ..., 19\}$ for each outcome of interest. The coefficient on the interaction term $PostJS_{it} \times TreatDistrict_{ds} \times PostEntry^{-1}_{ts}$ is standardized to zero as reference period. Both endpoints are binned. Shaded areas display 95% confidence intervals based on standard errors double-clustered by individual and district-month.

endogenous entry driven by other time-varying characteristics, and are absorbed by the district-month fixed effects in the individual-level specifications.

Finally, we check whether the results could be driven by general equilibrium effects. The key issue is that platform entry could crowd out jobs in related sectors such as traditional food delivery, pushing displaced workers into gig work and mechanically generating a positive coefficient even if these jobs only substitute for others. To test this hypothesis, we estimate the district-level specification with two alternative dependent variables: the number of job separations overall and the number from mini-jobs, both normalized by the number of employed. Figure 3.A.10 plots the coefficients. For overall job separations, the estimates are close to zero across all subgroups. More importantly, the coefficients for the share of separations from mini-jobs are similarly small and insignificant, reducing concerns that platform entry causes job separations in low-wage sectors in the first place.

Heterogeneity Across Platforms: We corroborate our results by exploiting institutional differences across platforms. As noted in Section 3.2.1, only two of the three platforms hired gig workers under mini-job contracts, and market statistics suggest that *Lieferando* soon became the dominant firm.[24] We thus expect *Lieferando* to be the primary driver of our mini-job results. We test this by augmenting the district-level specification with an indicator for *Lieferando* presence. Collinearity prevents us from interacting the two treatment variables, since *Lieferando* presence implies any-platform presence by construction.

Figure 3.A.11 plots the coefficients on the *Lieferando* indicator by subgroup. For mini-job take-up, the estimates are positive and significant for the full sample and grow when successively restricting to young, male, and high-skilled individuals. The magnitudes exceed those in the baseline regressions, indicating that this platform indeed drives our main results. By contrast, the coefficients on the other outcomes are small and insignificant.[25]

3.4.5 Economic Mechanism

Permanent vs. Transitory Effects of Platform Entry: Following Fos et al. (2025), we test two competing hypotheses about the persistence of the changes induced by gig platforms. Under the first, they act like ordinary employers: their entry generates a one-time increase in labor demand, implying that the effect on mini-job take-up should fade once the initial vacancies are filled. Under the second, platforms produce a structural change by reducing search frictions, shifting the Beveridge curve inward for parts of the labor market. In this case, effects should persist as workers continue to rely on platforms as a temporary buffer following separations.

The event studies in Figure 3.5 provide initial support for the second hypothesis: despite some fluctuation, the effects persist and even grow over time. To test this more systematically, we re-estimate Equation (3.2) after removing all observations from the first year and, separately, the first two years following platform entry. If effects were transitory, the coefficients should become small and insignificant. Tables 3.B.3 to 3.B.5 show that the estimates remain remarkably consistent with the main analysis, despite these cuts eliminating roughly 64% and 93% of post-separation observations. Mini-job take-up remains at 1.6 percentage points when removing the first year and rises to 3.5 percentage points when also removing the second. The growing magnitude suggests that workers only gradually become aware of these work opportunities and begin to exploit them as an alternative source of income. Taken together, we find support for the second mechanism. For recently un-

24. For data on 2017, see bit.ly/3ZBnE9l.

25. The equivalent individual-level specification yields very similar estimates: for individuals below 40, the coefficient is 0.0087 ($t = 1.80$), compared to 0.0083 ($t = 2.05$) at the district level.

employed workers, gig platforms play a different role compared to traditional firms, and the benefits of their availability persist for an extended period.

The one notable exception is a highly significant negative coefficient on employment subject to social security in Panel A of Table 3.B.3. Consistent with Figure 3.5, this points to substitution between gig work and traditional employment in the second year after platform entry, plausibly as workers become more familiar with platform work and consider it a viable alternative to traditional employment. Beyond that window, the coefficient becomes insignificant but remains negative.

Gig Platforms and Entrepreneurship: So far, we have addressed the insurance role of gig work through the lens of employment and labor market institutions. We round off our analysis by studying the impact of gig platforms on business formation. Barrios, Hochberg, and Yi (2022) and Denes, Lagaras, and Tsoutsoura (2025) show that the launch of gig platforms in the US spurred entrepreneurship through two channels: the gig economy's capacity to serve as an income supplement, and the resulting option value, which reduces downside risk for entrepreneurs.

To examine the impact of gig platforms on entrepreneurship, we use annual data on business formation from two sources[26] as well as scraped information on startups.[27] As in our district-level regressions, we construct a stacked dataset between 2013 and 2017 covering 23 treatment and 50 control districts. We replicate the specification from Barrios, Hochberg, and Yi (2022), which includes district and year fixed effects, district-specific linear trends, and a set of control variables.[28]

Table 3.B.6 reports the results separately for each outcome (in logs) with and without controls. For the business formation statistics, we find no effect of platform availability, with insignificant and slightly negative coefficients. For the startup data, we obtain a highly significant positive estimate that vanishes once controls are added. Under the mild assumption that the two channels do not offset each other, we thus find neither the income supplement channel nor the option value channel to be relevant in our setting.

Several factors may account for the divergence from Barrios, Hochberg, and Yi (2022). First, our data are at lower frequency and have fewer cross-sectional units, possibly attenuating estimates and reducing power. Second, food delivery platforms differ from Uber, and the complier population that starts businesses in response to ride-sharing availability need not extend to our setting. Third, institutional differences may modulate the link between the gig economy and entrepreneurship: Ger-

26. The first is obtained from the regional statistics of the Federal Statistical Office of Germany (Destatis), while the second is complemented by survey information from the *Institut für Mittelstandsforschung* (IfM). The correlation between the two measures is high, at 0.97.

27. We scraped all available data for Germany from www.eu-startups.com. The correlation with the other measures is around 0.75.

28. The control variables are log population, the first lag of log GDP per capita, and the first lag of the unemployment rate.

many's relatively generous welfare state (especially for those with high prior contributions) already buffers entrepreneurial risk, consistent with our null effect on *UB I* take-up. Our results imply that recently unemployed individuals start gig jobs, but to top up their income from other sources rather than to reduce reliance on benefits.

3.5 Conclusion

This chapter examines the role of gig work as an alternative income insurance mechanism. We exploit the staggered entry of gig delivery platforms across German districts between 2014 and 2017 to explore how platform availability affects workers' behavior around a job separation event. Using two identification strategies, we show that recently unemployed workers are significantly more likely to work under contractual forms commonly used by gig delivery platforms once these platforms are available in their district. Take-up is particularly pronounced for male, young, and highly educated workers.

We find no significant effects of platform entry on reliance on unemployment benefits or on entrepreneurial activity, but some evidence of substitution between gig work and other forms of employment. These findings contrast in part with studies on ride-sharing platforms in the US (Koustas 2019; Barrios, Hochberg, and Yi 2022; Fos et al. 2025), highlighting the role of institutional context in determining the insurance value of gig work.

At the same time, platforms appear to induce lasting changes in the labor force. We observe a significant increase in training participation, suggesting a signaling role of gig work (in line with Adermon and Hensvik 2022), and find that recently unemployed workers rely on gig work as a temporary income source while searching for more stable or better-paid employment. In sum, these results point tentatively toward lower attachment to the traditional labor market, but not to a degradation of human capital in response to platform entry, contrasting with Jackson (2022).

Our analysis is subject to several limitations. We rely on extensive-margin measures, as we observe neither who actually takes up gig work nor who is eligible for and applies for unemployment insurance. More detailed data would allow us to study intensive-margin effects and to characterize the channels at play more sharply. A longer horizon would also be necessary to assess the impact on workers' income trajectories, by tracking wages and employment of former gig workers and exploring general equilibrium effects of the platform economy (Berger, Chen, and Frey 2018). Overall, our results suggest that gig platforms can have important indirect effects on the welfare of the labor force, which should be taken into account when designing labor market policies.

References

Abraham, Katharine, and Susan Houseman. 2019. "Making Ends Meet: The Role of Informal Work in Supplementing Americans' Income." *RSF: The Russell Sage Foundation Journal of the Social Sciences* 5 (5): 110–31. [177, 191]

Adams-Prassl, Abi, Jeremias Adams-Prassl, and Diane Coyle. 2021. "Uber and Beyond: Policy Implications for the UK." Working Paper, 1, The Productivity Institute. [177]

Adermon, Adrian, and Lena Hensvik. 2022. "Gig-Jobs: Stepping Stones or Dead Ends?" *Labour Economics* 76: 102171. [195, 200]

Antoni, Manfred, Andreas Ganzer, and Philipp vom Berge. 2019. "Sample of Integrated Labour Market Biographies Regional File (SIAB-R) 1975-2017." *FDZ-Datenreport, Nuremberg.* [180, 218]

Baily, Martin Neil. 1978. "Some Aspects of Optimal Unemployment Insurance." *Journal of Public Economics* 10 (3): 379–402. [177]

Barrios, John, Yael Hochberg, and Hanyi Yi. 2022. "Launching with a Parachute: The Gig Economy and New Business Formation." *Journal of Financial Economics* 144 (1): 22–43. [177, 199, 200]

Berger, Thor, Chinchih Chen, and Carl Benedikt Frey. 2018. "Drivers of Disruption? Estimating the Uber Effect." *European Economic Review* 110: 197–210. [200]

Boeri, Tito, Giulia Giupponi, Alan Krueger, and Stephen Machin. 2020. "Solo Self-Employment and Alternative Work Arrangements: A Cross-Country Perspective on the Changing Composition of Jobs." *Journal of Economic Perspectives* 34 (1): 170–95. [177]

Borusyak, Kirill, Xavier Jaravel, and Jann Spiess. 2024. "Revisiting Event-Study Designs: Robust and Efficient Estimation." *Review of Economic Studies* 91 (6): 3253–85. [181]

Card, David, Jochen Kluve, and Andrea Weber. 2017. "What Works? A Meta Analysis of Recent Active Labor Market Program Evaluations." *Journal of the European Economic Association* 16 (3): 894–931. [177]

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs." *Quarterly Journal of Economics* 134 (3): 1405–54. [181, 186]

Chetty, Raj. 2008. "Moral Hazard Versus Liquidity and Optimal Unemployment Insurance." *Journal of Political Economy* 116 (2): 173–234. [177]

Clarke, Damian, and Kathya Tapia-Schythe. 2021. "Implementing the Panel Event Study." *Stata Journal* 21 (4): 853–84. [189]

Collins, Brett, Andrew Garin, Emilie Jackson, Dmitri Koustas, and Mark Payne. 2019. "Is Gig Work Replacing Traditional Employment? Evidence from Two Decades of Tax Returns." Working Paper, Internal Revenue Service. [177, 191]

Dauth, Christine. 2020. "Regional Discontinuities and the Effectiveness of Further Training Subsidies for Low-Skilled Employees." *ILR Review* 73 (5): 1147–84. [184]

Dauth, Wolfgang, and Johann Eppelsheimer. 2020. "Preparing the Sample of Integrated Labour Market Biographies (SIAB) for Scientific Analysis: A Guide." *Journal for Labour Market Research* 54 (1). [180, 218]

De Chaisemartin, Clément, and Xavier d'Haultfoeuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review* 110 (9): 2964–96. [187, 188]

Denes, Matthew, Spyridon Lagaras, and Margarita Tsoutsoura. 2025. "Entrepreneurship and the Gig Economy: Evidence from US Tax Returns." *Journal of Financial Economics* 173: 104156. [199]

Deshpande, Manasi, and Yue Li. 2019. "Who Is Screened Out? Application Costs and the Targeting of Disability Programs." *American Economic Journal: Economic Policy* 11 (4): 213–48. [179, 181, 187]

Farrell, Diana, Fiona Greig, and Amar Hamoudi. 2019. "The Evolution of the Online Platform Economy: Evidence from Five Years of Banking Data." *AEA Papers and Proceedings* 109: 362–66. [177]

Fitzenberger, Bernd, Aderonke Osikominu, and Marie Paul. 2023. "The Effects of Training Incidence and Duration on Labor Market Transitions." *Journal of Econometrics* 235 (1): 256–79. [184]

Fos, Vyacheslav, Naser Hamdi, Ankit Kalda, and Jordan Nickerson. 2025. "Gig Labor: Trading Safety Nets for Steering Wheels." *Journal of Financial Economics* 163: 103956. [175–177, 181, 198, 200]

Freyaldenhoven, Simon, Christian Hansen, and Jesse Shapiro. 2019. "Pre-Event Trends in the Panel Event-Study Design." *American Economic Review* 109 (9): 3307–38. [189]

Friedrich, Martin, Ramona Jost, Julia Lang, and Christoph Müller. 2024. "Gig-Work in the German Delivery-Services Sector: Employment Has Increased Significantly in Recent Years." In *IAB-Forum,* 1–9. 07. Institut für Arbeitsmarkt- und Berufsforschung (IAB). [175]

Furman, Jason. 2019. "Prepared Testimony for the Hearing: Online Platforms and Market Power, Part 3: The Role of Data and Privacy in Competition." *US House of Representatives.* [178]

Garin, Andrew, Emilie Jackson, Dmitri Koustas, and Carl McPherson. 2020. "Is New Platform Work Different from Other Freelancing?" *AEA Papers and Proceedings* 110: 157–61. [175, 177]

Garin, Andrew, Emilie Jackson, Dmitri Koustas, and Alicia Miller. 2025. "The Evolution of Platform Gig Work, 2012-2023." *NBER Chapters.* [175, 177]

Glasner, Benjamin. 2023. "The Minimum Wage, Self-Employment, and the Online Gig Economy." *Journal of Labor Economics* 41 (1): 103–27. [177]

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics* 225 (2): 254–77. [181]

Gruber, Jonathan. 1997. "The Consumption Smoothing Benefits of Unemployment Insurance." *American Economic Review* 87 (1): 192–205. [177]

Heiland, Heiner. 2019. "Plattformarbeit im Fokus. Ergebnisse einer Explorativen Online-Umfrage zu Plattformvermittelter Kurierarbeit." *WSI-Mitteilungen* 72 (4): 298–304. [185, 191]

Iacus, Stefano, Gary King, and Giuseppe Porro. 2011. "Multivariate Matching Methods That Are Monotonic Imbalance Bounding." *Journal of the American Statistical Association* 106 (493): 345–61. [181]

Jackson, Emilie. 2022. "Availability of the Gig Economy and Long Run Labor Supply Effects for the Unemployed." Working Paper, Michigan State University. [175, 177, 194, 195, 200]

Katz, Lawrence, and Alan Krueger. 2019. "Understanding Trends in Alternative Work Arrangements in the United States." *RSF: The Russell Sage Foundation Journal of the Social Sciences* 5 (5): 132. [177]

Kolsrud, Jonas, Camille Landais, Peter Nilsson, and Johannes Spinnewijn. 2018. "The Optimal Timing of Unemployment Benefits: Theory and Evidence from Sweden." *American Economic Review* 108 (4-5): 985–1033. [177]

Koustas, Dmitri. 2019. "Consumption Insurance and Multiple Jobs: Evidence from Rideshare Drivers." Working Paper, University of Chicago. [175–177, 200]

Liu, Cathy Yang, and Luísa Nazareno. 2019. "The Changing Quality of Nonstandard Work Arrangements: Does Skill Matter?" *RSF: The Russell Sage Foundation Journal of the Social Sciences* 5 (4): 104–28. [177]

Mas, Alexandre, and Amanda Pallais. 2020. "Alternative Work Arrangements." *Annual Review of Economics* 12: 631–58. [177]

Moser, Rodrigo. 2021. "Driving Up Repayment: The Impact of the Gig-Economy on Student Debt." Working Paper, Washington University. [177]

Olden, Andreas, and Jarle Møen. 2022. "The Triple Difference Estimator." *Econometrics Journal* 25 (3): 531–53. [188]

Osikominu, Aderonke. 2013. "Quick Job Entry or Long-Term Human Capital Development? The Dynamic Effects of Alternative Training Schemes." *Review of Economic Studies* 80 (1): 313–42. [184]

Roesler, Konrad, Christiane Papenroth, Antje Hess, Sylke Gollin, Joachim Fritz, Angela Kahler, Sebastian Lorenz, and Sylvie Böhme. 2021. "Statistik zu Maßnahmen und Teilnehmenden an Maßnahmen der Arbeitsförderung." *Bundesagentur für Arbeit.* [184]

Urzì Brancati, Maria Cesira, Annarosa Pesole, and Enrique Fernández-Macías. 2020. "New Evidence on Platform Workers in Europe: Results from the Second COLLEEM Survey." Working Paper, Joint Research Centre (JRC). [191]

Appendix 3.A Additional Figures

Figure 3.A.1. Months Since Platforms' Entry at the End of 2017 by District

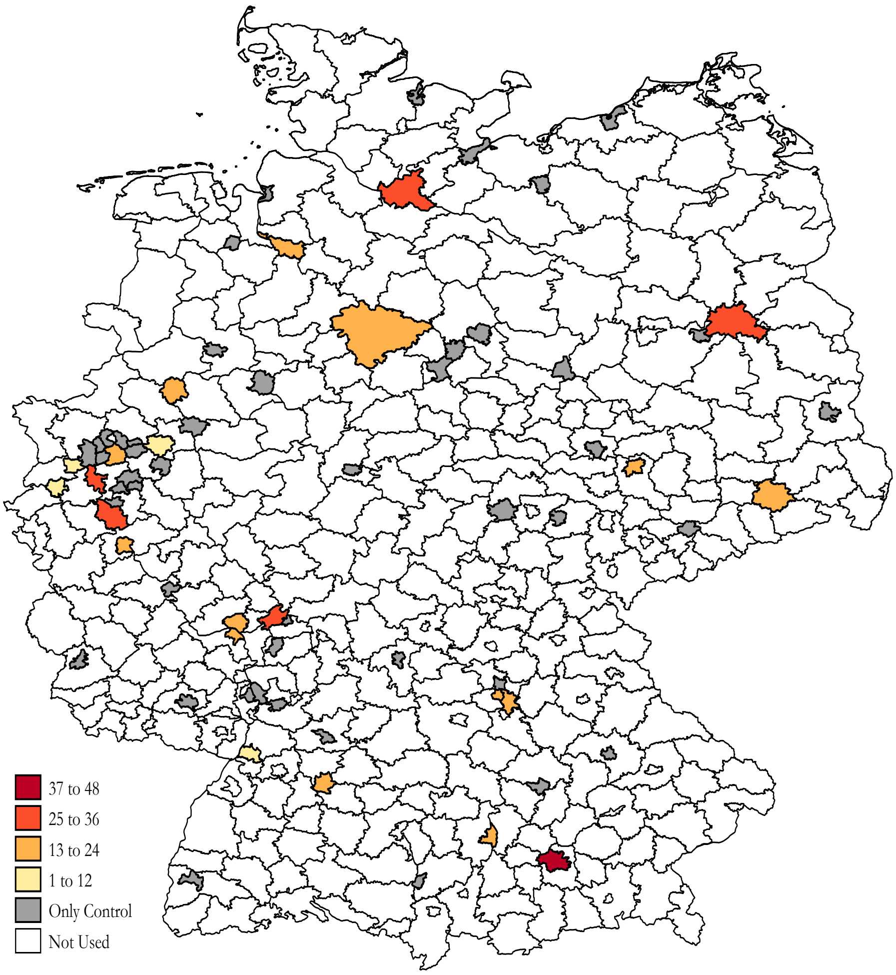

Notes: The figure shows a map of German districts, specifying the number of months at least one platform has been present in each district by the end of 2017. Gray color indicates control districts in the district-level specification.

Figure 3.A.2. Distribution of Matching Variables: Treatment vs. Control Districts

(a) Age at Job Separation (b) Wage at Job Separation

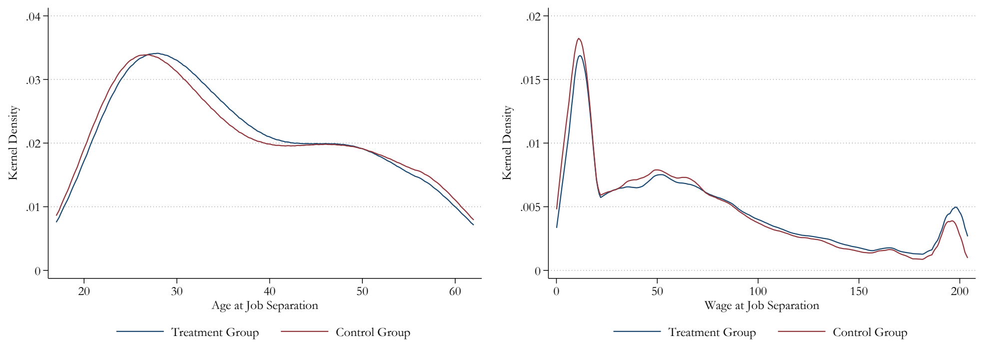

Notes: The plots show kernel densities of age (Panel a) and daily wage (Panel b) at job separation, separately for the treatment and control groups, generated following the description outlined in Section 3.2.2.

Figure 3.A.3. Main Outcomes Relative to Job Separation by Platform Availability (Alternative Comparison Group)

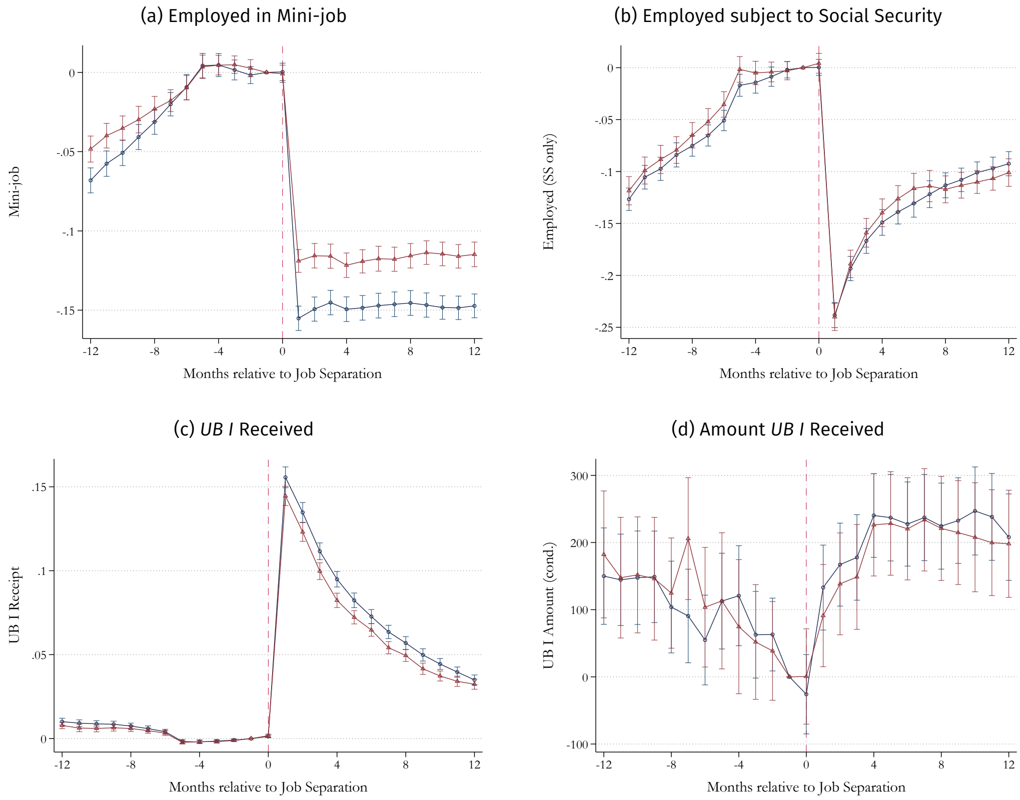

Notes: The figures plot averages of the outcome variables relative to a job separation by platform availability. Red (blue) lines refer to periods after (before) platform entry, both restricted to workers in treatment districts. Coefficients are obtained by estimating Equation (3.3) separately for each group. The month before separation is omitted. 95% confidence intervals based on standard errors double-clustered by episode and district-month.

Figure 3.A.4. Age Distribution of *Lieferando* Platform Workers in 2020

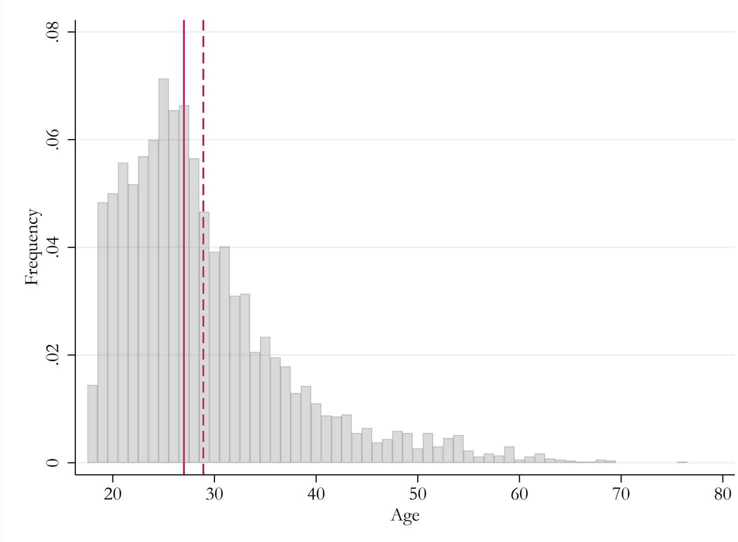

Notes: The figure shows the age distribution of platform workers working for *Lieferando* in September 2020. The dashed (solid) vertical line depicts the mean (median) of the distribution, which is at 28.4 (26.5).

Figure 3.A.5. District-Level Effects of Gig Platform Availability by Subgroup and Citizenship

(a) German Citizens (b) Non-German Citizens

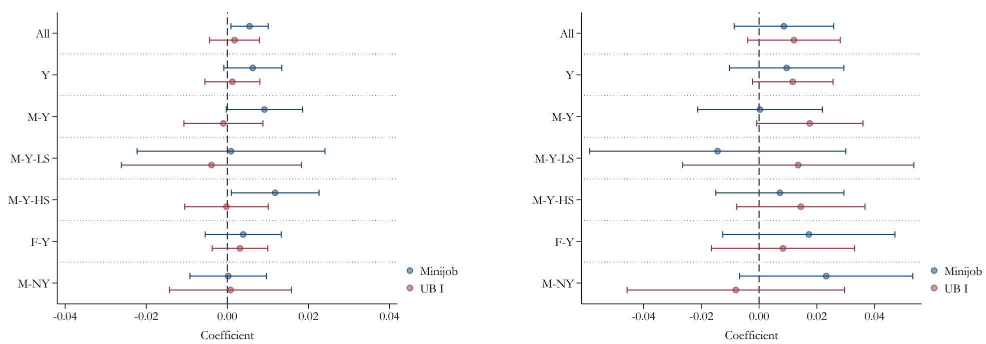

Notes: The figures plot the district-level OLS estimates for the coefficient of interest obtained from separate regressions for each subgroup. *Y*/*NY*: below/above age 40 at job separation, *M*/*F*: male/female, *HS*/*LS*: tertiary/no tertiary education. Confidence intervals based on robust standard errors clustered at the district level. All regressions include fixed effects at the district-sub-experiment and month-sub-experiment level as well as the control variables described in Section 3.3. The regressions are also weighted by the number of post-job separation observations in each district-month cell. Mini-job is an indicator for working under a contract classified as marginal employment. *UB I* indicates any receipt of unemployment benefits. In Panel (a), we restrict the sample to German citizens, while in Panel (b), we only use observations from non-German citizens (at time of job separation).

Figure 3.A.6. District-Level Effects of Gig Platform Availability by Subgroup (Interacted Controls)

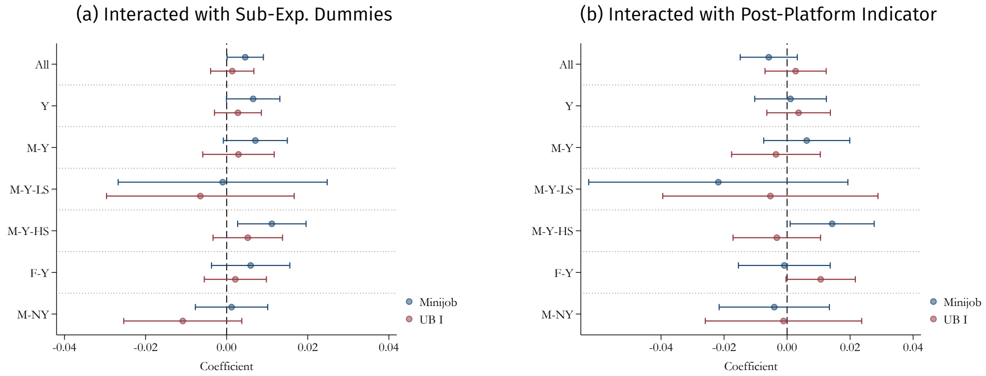

Notes: The figures plot the district-level OLS estimates for the coefficient of interest obtained from separate regressions for each subgroup. *Y*/*NY*: below/above age 40 at job separation, *M*/*F*: male/female, *HS*/*LS*: tertiary/no tertiary education. Confidence intervals based on robust standard errors clustered at the district level. All regressions include fixed effects at the district-sub-experiment and month-sub-experiment level as well as the control variables described in Section 3.3. The regressions are also weighted by the number of post-job separation observations in each district-month cell. Mini-job is an indicator for working under a contract classified as marginal employment. *UB I* indicates any receipt of unemployment benefits. In Panel (a), the control variables have been interacted with dummies for each sub-experiment, while in Panel (b), the controls have been interacted with an indicator for post-platform periods.

Figure 3.A.7. Individual-Level Effects of Gig Platform Availability on Mini-Job Employment by Age

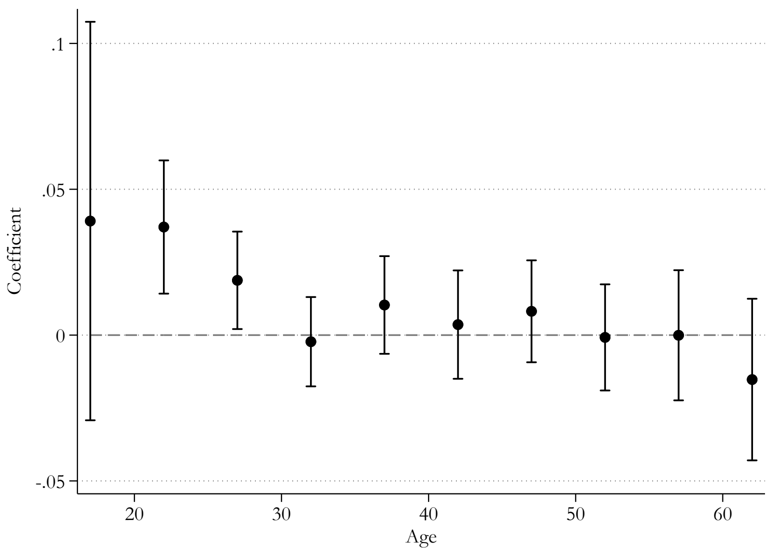

Notes: The figures plot the individual-level OLS estimates for mini-job take-up obtained by estimating Equation (3.2) separately for each 5-year age bracket. Confidence intervals based on robust standard errors double-clustered at the individual and district-month level. All regressions include fixed effects at the district-month-sub-experiment and episode-sub-experiment level.

Figure 3.A.8. Dynamic Effects Relative to Platform Entry: Job-Seeking Outcomes

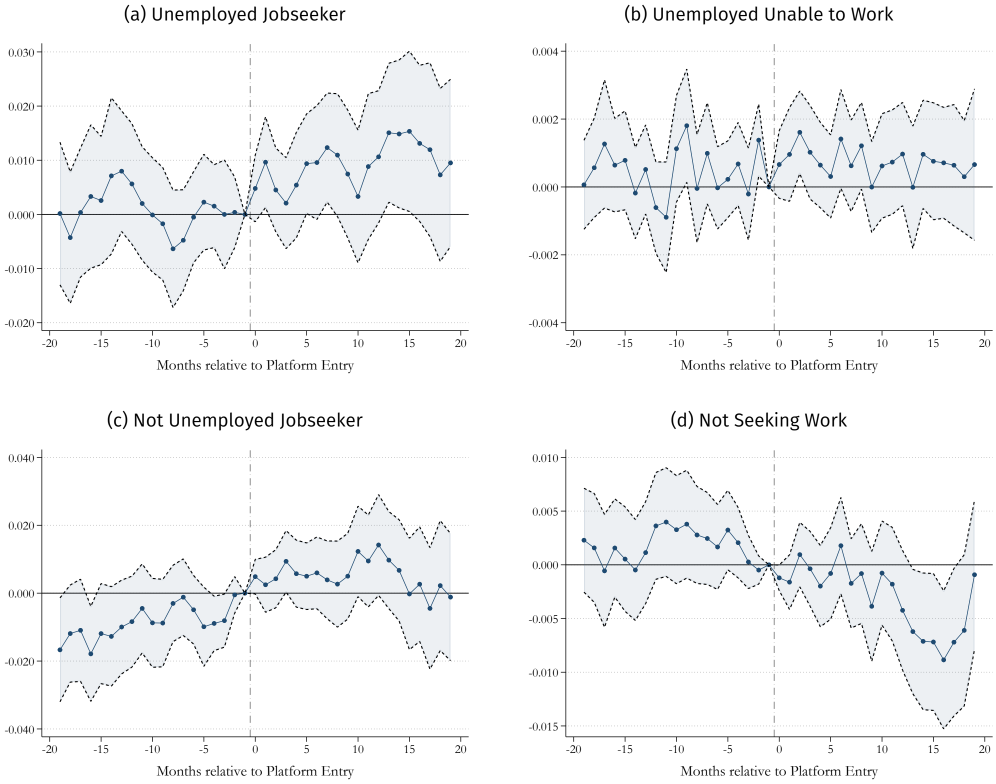

Notes: The figures report estimated coefficients on the interaction terms $PostJS_{it} \times TreatDistrict_{ds} \times PostEntry^{m}_{ts}$ of the event-study specification described in Section 3.3, for $m \in \{-19, ..., 19\}$ for each outcome of interest. The coefficient on the interaction term $PostJS_{it} \times TreatDistrict_{ds} \times PostEntry^{-1}_{ts}$ is standardized to zero as reference period. Both endpoints are binned. Shaded areas display 95% confidence intervals based on standard errors double-clustered by individual and district-month.

Figure 3.A.9. Dynamic Effects Relative to Platform Entry: Participation in ALMPs

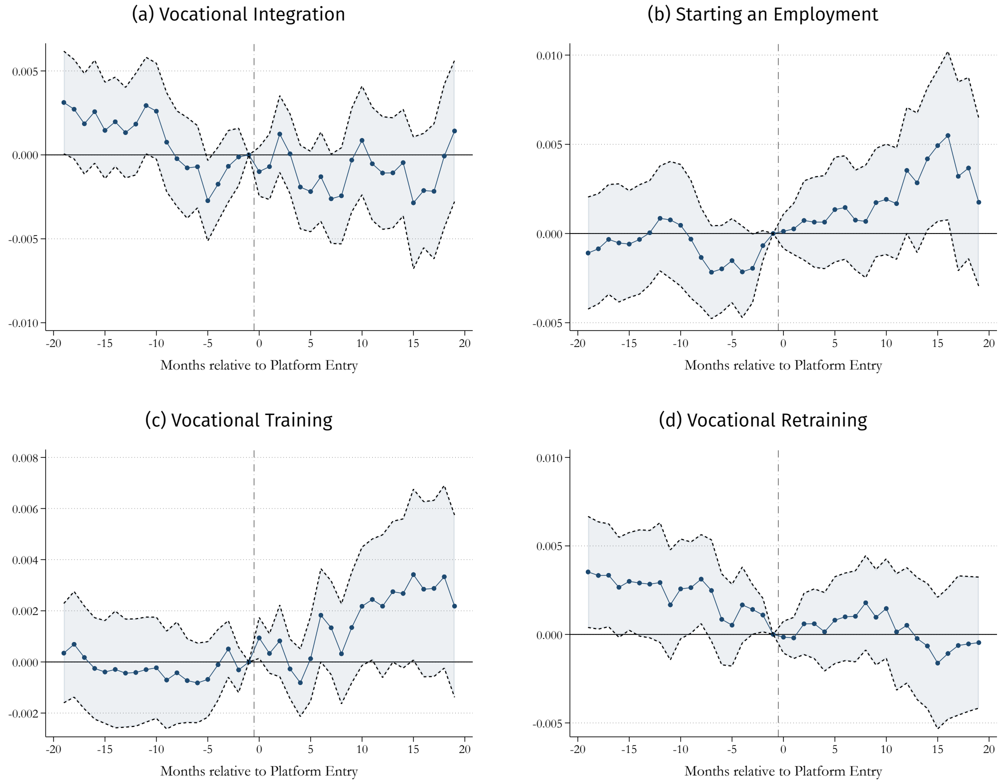

Notes: The figures report estimated coefficients on the interaction terms $PostJS_{it} \times TreatDistrict_{ds} \times PostEntry^{m}_{ts}$ of the event-study specification described in Section 3.3, for $m \in \{-19, ..., 19\}$ for each outcome of interest. The coefficient on the interaction term $PostJS_{it} \times TreatDistrict_{ds} \times PostEntry^{-1}_{ts}$ is standardized to zero as reference period. Both endpoints are binned. Shaded areas display 95% confidence intervals based on standard errors double-clustered by individual and district-month.

Figure 3.A.10. Gig Platform Availability and Number of Job Separations: District-Level by Subgroup

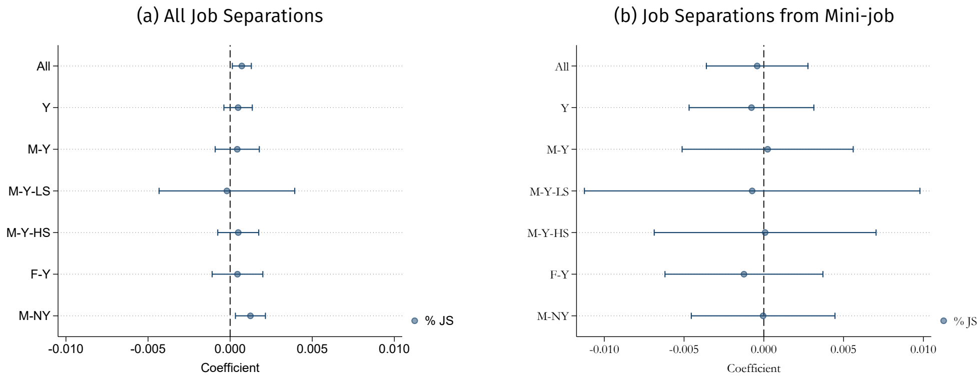

Notes: The figures plot the district-level OLS estimates for the coefficient of interest obtained from separate regressions for each subgroup. The setup is equivalent to the district-level DiD regressions reported in Section 3.4.2. *Y*/*NY*: below/above age 40 at job separation, *M*/*F*: male/female, *HS*/*LS*: tertiary/no tertiary education. Confidence intervals based on robust standard errors clustered at the district level. All regressions include fixed effects at the district-sub-experiment and month-sub-experiment level as well as the control variables described in Section 3.3. The regressions are also weighted by the number of post-job separation observations in each district-month cell. The dependent variables are the ratio of job separations overall (Panel a) and from mini-jobs (Panel b) normalized by the number of employed (both restricted to being employed for at least 6 months prior to job separation).

Figure 3.A.11. District-Level Effects of Gig Platform Availability by Subgroup (Robustness, Coefficients for *Lieferando*)

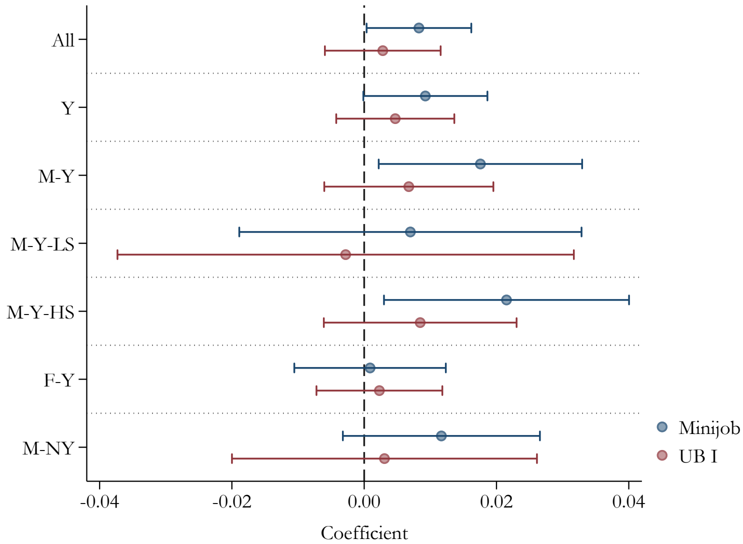

Notes: The figure plots the district-level OLS estimates of α obtained from the following regression by subgroup: $y_{dt} = \alpha Lieferando_{dt} + \beta\, TreatDistrict_{ds} \times PostEntry_{ts} + \phi_{ds} + \psi_{ts} + \Gamma X_{dt} + \varepsilon_{dts}$, where $Lieferando_{dt}$ is an indicator equal to one after entry of this platform in the district and sub-experiment. The coefficient of interest is α. Given the construction of the indicator for any platform being present, we cannot interact the two treatment variables due to collinearity. The remaining variables and indices follow the notation introduced in Section 3.3. *Y*/*NY*: below/above age 40 at job separation, *M*/*F*: male/female, *HS*/*LS*: tertiary/no tertiary education. Confidence intervals based on robust standard errors clustered at the district level. All regressions include fixed effects at the district-sub-experiment and month-sub-experiment level as well as the control variables described in Section 3.3. The regressions are also weighted by the number of post-job separation observations in each district-month cell. Mini-job is an indicator for working under a contract classified as marginal employment. *UB I* indicates any receipt of unemployment benefits.

Appendix 3.B Additional Tables

Table 3.B.1. District-Level Effects of Gig Platform Availability: Job-Seeking Outcomes

	Unempl. Jobseeker (1)	Unempl. Unable (2)	Not Unempl. Jobseeker (3)	Not Seeking (4)
Platform	0.000	0.000	0.000	-0.001
	(0.003)	(0.000)	(0.002)	(0.001)
R^2	0.668	0.105	0.613	0.377
Observations	17,121	17,121	17,121	17,121
Control Mean	0.113	0.002	0.056	0.017
Controls	✓	✓	✓	✓
District-Exp. fixed effects	✓	✓	✓	✓
Month-Exp. fixed effects	✓	✓	✓	✓

Notes: OLS estimates. All regressions include district-sub-experiment and month-sub-experiment fixed effects and the control variables described in Section 3.3. Regressions are weighted by the number of post-separation observations in each district-month cell. Standard errors are clustered at the district level. All outcome variables are indicators equal to one if an individual is assigned to the respective job-search status. *Platform* is a dummy variable equal to one if at least one platform has been present in the district for the entire month. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.2. District-Level Effects of Gig Platform Availability: Participation in ALMPs

	Voc. Integration (1)	Start Employment (2)	Voc. Training (3)	Voc. Retraining (4)
Platform	0.000	0.001**	0.000	0.001
	(0.000)	(0.001)	(0.000)	(0.001)
R^2	0.227	0.326	0.261	0.248
Observations	17,121	17,121	17,121	17,121
Control Mean	0.005	0.006	0.002	0.007
Controls	✓	✓	✓	✓
District-Exp. fixed effects	✓	✓	✓	✓
Month-Exp. fixed effects	✓	✓	✓	✓

Notes: OLS estimates. All regressions include district-sub-experiment and month-sub-experiment fixed effects and the control variables described in Section 3.3. Regressions are also weighted by the number of post-job separation observations in each district-month cell. Standard errors are clustered at the district level. All outcome variables are indicators equal to one if an individual is recorded in the respective ALMP. *Platform* is a dummy variable equal to one if at least one platform has been present in the district for the entire month. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.3. Main Specification Without First Years After Platform Entry, Main Outcomes

	Mini-job	Employed (SS only)	UB I Receipt	UB I Amount (cond.)
	(1)	(2)	(3)	(4)
Panel A: 1st year removed				
PostJS x Platform	0.016	-0.037***	0.004	-45.870
	(0.010)	(0.012)	(0.007)	(86.993)
R^2	0.593	0.686	0.417	0.867
Observations	2,502,928	2,502,928	2,502,928	107,268
Control Mean	0.170	0.588	0.054	849.980
Interactions	✓	✓	✓	✓
District-Month-Exp. fixed effects	✓	✓	✓	✓
Episode-Exp. fixed effects	✓	✓	✓	✓
Panel B: 1st and 2nd years removed				
PostJS x Platform	0.035**	-0.015	-0.019	200.000
	(0.015)	(0.026)	(0.014)	(152.655)
R^2	0.602	0.697	0.426	0.869
Observations	2,151,021	2,151,021	2,151,021	94,080
Control Mean	0.170	0.588	0.054	849.980
Interactions	✓	✓	✓	✓
District-Month-Exp. fixed effects	✓	✓	✓	✓
Episode-Exp. fixed effects	✓	✓	✓	✓

Notes: OLS estimates, robust standard errors double-clustered at the individual and district-month level in parentheses. All regressions include fixed effects at the district-month-sub-experiment and episode-sub-experiment level. *Interactions* refers to the non-absorbed components of the treatment variable, i.e. *PostJS* × *TreatDistrict*, *PostJS* × *PostEntry*, and *PostJS*. Mini-job is an indicator for marginal employment, Employed (SS only) for employment subject to social security excluding mini-jobs, and UB I Receipt for receipt of unemployment benefits. *PostJS* is an indicator for all post-job separation periods. *Platform* is a dummy variable equal to one if at least one platform has been present in the district for the entire month. The monthly amount of *UB I* is computed by multiplying the daily benefit (in €, deflated with respect to 2015) by the number of days of *UB I* receipt in a given month. In Panel A (B), we remove all observations in the first (and second) year following platform entry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.4. Main Specification Without First Years After Platform Entry, Job-Seeking Outcomes

	Unempl. Jobseeker (1)	Unempl. Unable (2)	Not Unempl. Jobseeker (3)	Not Seeking (4)
Panel A: 1^{st} year removed				
PostJS x Platform	0.012*	0.000	-0.008	-0.004
	(0.007)	(0.001)	(0.006)	(0.003)
R^2	0.457	0.127	0.430	0.567
Observations	2,502,928	2,502,928	2,502,928	2,502,928
Control Mean	0.080	0.002	0.075	0.018
Interactions	✓	✓	✓	✓
District-Month-Exp. fixed effects	✓	✓	✓	✓
Episode-Exp. fixed effects	✓	✓	✓	✓
Panel B: 1^{st} and 2^{nd} years removed				
PostJS x Platform	-0.008	0.000	-0.017	0.000
	(0.014)	(0.001)	(0.013)	(0.005)
R^2	0.469	0.128	0.435	0.574
Observations	2,151,021	2,151,021	2,151,021	2,151,021
Control Mean	0.080	0.002	0.075	0.018
Interactions	✓	✓	✓	✓
District-Month-Exp. fixed effects	✓	✓	✓	✓
Episode-Exp. fixed effects	✓	✓	✓	✓

Notes: OLS estimates, robust standard errors double-clustered at the individual and district-month level in parentheses. All regressions include fixed effects at the district-month-sub-experiment and episode-sub-experiment level. *Interactions* refers to the non-absorbed components of the treatment variable, i.e., *PostJS* × *TreatDistrict*, *PostJS* × *PostEntry*, and *PostJS*. All outcome variables are indicators equal to one if an individual is assigned to the respective job-search status. *PostJS* is an indicator for all post-job separation periods. *Platform* is a dummy variable equal to one if at least one platform has been present in the district for the entire month. In Panel A (B), we remove all observations in the first (and second) year following platform entry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.5. Main Specification Without First Years After Platform Entry, Participation in ALMPs

	Voc. Integration (1)	Start Employment (2)	Voc. Training (3)	Voc. Retraining (4)
Panel A: 1st year removed				
PostJS x Platform	-0.003*	0.005***	0.003*	-0.003
	(0.001)	(0.002)	(0.002)	(0.002)
R^2	0.250	0.333	0.592	0.455
Observations	2,500,765	2,500,765	2,500,765	2,500,765
Control Mean	0.004	0.005	0.003	0.005
Interactions	✓	✓	✓	✓
District-Month-Exp. fixed effects	✓	✓	✓	✓
Episode-Exp. fixed effects	✓	✓	✓	✓
Panel B: 1st and 2nd years removed				
PostJS x Platform	-0.005**	0.003	0.004	-0.004
	(0.002)	(0.003)	(0.004)	(0.003)
R^2	0.264	0.340	0.600	0.469
Observations	2,151,021	2,151,021	2,151,021	2,151,021
Control Mean	0.004	0.005	0.003	0.005
Interactions	✓	✓	✓	✓
District-Month-Exp. fixed effects	✓	✓	✓	✓
Episode-Exp. fixed effects	✓	✓	✓	✓

Notes: OLS estimates, robust standard errors double-clustered at the individual and district-month level in parentheses. All regressions include fixed effects at the district-month-sub-experiment and episode-sub-experiment level. *Interactions* refers to the non-absorbed components of the treatment variable, i.e., *PostJS* × *TreatDistrict*, *PostJS* × *PostEntry*, and *PostJS*. All outcome variables are indicators equal to one if an individual is recorded in the respective ALMP. *PostJS* is an indicator for all post-job separation periods. *Platform* is a dummy variable equal to one if at least one platform has been present in the district for the entire month. In Panel A (B), we remove all observations in the first (and second) year following platform entry. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 3.B.6. Gig Platform Availability and Business Formation

	Business formation				Number of startups	
	Destatis		IfM			
	(1)	(2)	(3)	(4)	(5)	(6)
Platform	-0.028	-0.013	-0.002	-0.008	0.614***	0.153
	(0.049)	(0.039)	(0.029)	(0.031)	(0.097)	(0.178)
Log Population		-0.197		-0.133		6.086
		(1.891)		(1.530)		(6.888)
Log GDP/Capita (Lag)		-0.175		0.162		0.159
		(0.238)		(0.198)		(1.218)
Unemp. Rate (Lag)		9.574		2.456		-84.010**
		(12.453)		(5.949)		(41.112)
R^2	0.991	0.998	0.995	0.999	0.942	0.971
Observations	955	758	959	762	959	762
District-Exp. fixed effects	✓	✓	✓	✓	✓	✓
Year-Exp. fixed effects	✓	✓	✓	✓	✓	✓
District-Exp. Trend	✓	✓	✓	✓	✓	✓

Notes: OLS estimates, robust standard errors clustered at the district level in parentheses. All regressions include fixed effects at the district-sub-experiment and year-sub-experiment level. The specifications shown in Columns (2), (4), and (6) additionally include district-sub-experiment-specific trends. The dependent variables are defined in logs. Data on business formation in Columns (1) and (2) are obtained from the Federal Statistical Office of Germany (Destatis), while the data on business formation in Columns (3) and (4) come from the *Institut für Mittelstandsforschung* (IfM). The information on the number of startups used in Columns (5) and (6) has been scraped from www.eu-startups.com. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix 3.C Details on Employment Data

Background: The Sample of Integrated Labour Market Biographies (SUF SIAB-Regional File 7517) is a 2% random sample of the Integrated Employment Biographies (IEB) of the Institute for Employment Research (IAB).[29] The data include detailed information on the employment status of individuals with daily precision from 1975 to 2017, alongside a set of individual characteristics. The entire population of the IEB comprises all those individuals who have been in employment subject to social security or classified as marginal part-time employment (since 1999); or who have received social security benefits based on Social Code Book III (SGB III), such as the *UB I*, or on Social Code Book II (SGB II), such as the *UB II* (since 2005); or who have been officially registered as jobseekers (since 1997); or who have participated in ALMPs (since 2000). The SIAB-R-7517 dataset counts a total of 1,827,903 individuals, whose employment biographies are documented in 62,340,521 non-overlapping observations.[30]

Data Preparation: A number of steps are necessary to construct the final dataset. To begin, labor market biographies are recorded as spells of varying length that may overlap. Following Antoni, Ganzer, and Berge (2019) and Dauth and Eppelsheimer (2020), we clean the data, impute censored wages, and convert them into a monthly panel. To make the dataset tractable while preserving its informational content, we aggregate parallel episodes by keeping only the main one[31] and retain information on the remaining ones by recording aggregate earnings and unemployment insurance benefits in each period. This is necessary because individuals may hold multiple jobs or receive benefits simultaneously.

We then identify job separation events for each worker, i.e., layoffs and quits. A limitation of the data is that anonymization does not allow us to distinguish between these types of events, recording both as the end of employment. While more detailed information would be desirable, German unemployment insurance does not exclude voluntary separations, although it entails less generous benefits in such cases. Compared to countries such as Italy or the United States, sanctions for voluntary unemployment are relatively mild, including a standard waiting period of 12 weeks (Venn 2012). Thus, although the unobserved reason for separation affects benefit generosity, it is not a decisive determinant of eligibility.

29. This study uses the factually anonymous Sample of Integrated Labour Market Biographies (version 1975–2017). Data access was provided via a Scientific Use File supplied by the Research Data Centre (FDZ) of the German Federal Employment Agency (BA) at the Institute for Employment Research (IAB). DOI: 10.5164/IAB.FDZD.1904.en.v1

30. The original sample included 45,859,427 observations, also called episodes, that we split to create the non-overlapping observations.

31. As in Dauth and Eppelsheimer (2020), the main episode is defined as the one with the longest tenure.

For each job separation, we retain the 12 months before and after the month of the event. We thus obtain a set of 25-month windows centered on each separation event. For a small subset of workers, these windows overlap due to multiple separations occurring in close succession. In these cases, we randomly select one separation event and its corresponding 25-month window from the overlapping set.

Given that individuals in the SIAB-R-7517 are observed only when they engage in activities reported to social security agencies, such as employment or receipt of unemployment benefits, they are unobserved when not engaged in any such activity. We exploit this feature to impute missing individual-month observations, assigning zeros to selected outcomes of interest.[32] Accordingly, workers are recorded as not in employment subject to social security (including mini-jobs), not registered in ALMPs, not registered as jobseekers, and not receiving unemployment benefits. This procedure rectangularizes the data into a balanced panel of job separation events, each observed over a 25-month window.[33]

District vs. Individual-Level Data Preparation: The sample used for the individual-level analysis differs slightly from that used at the district level. Both datasets are based on spell-level information, from which we construct non-overlapping job separation episodes centered on the timing of separation. We then stack the data in two slightly different ways.

A first difference is that the individual-level analysis uses only the 23 ever-treated districts. This restriction further reduces selection concerns, as all of these districts eventually experience platform entry, and does not materially affect power given the large number of observations at the individual level. Since gig platforms initially targeted large urban centers, this sample still covers a substantial share (26.3%) of all job separation events.

Second, in the individual-level analysis we do not require a common observation window across districts or equal-length sub-experiments. Instead, each sub-experiment starts in January 2013 and ends when no valid control districts remain. The control group consists of individuals in not-yet-treated districts, where entry occurs at least four months after the platform entry in the corresponding treated districts. This setup retains 21 of the 23 entry events across nine sub-experiments. It is motivated by the fact that, in addition to staggered platform roll-out, we exploit variation in the timing of job separations across individuals, ensuring treatment variation throughout the sample period, including its final months. We further ensure comparability using CEM as described in Section 3.2.2.

32. We impute other variables, such as plausibly time-invariant worker characteristics, using their last observed value.

33. The only exceptions arise when separations occur within 12 months of the start or end of the observation window. This affects only 2013 and 2017, as the sample runs from January 2013 to December 2017.

References

Antoni, Manfred, Andreas Ganzer, and Philipp vom Berge. 2019. "Sample of Integrated Labour Market Biographies Regional File (SIAB-R) 1975-2017." *FDZ-Datenreport, Nuremberg*. [180, 218]

Dauth, Wolfgang, and Johann Eppelsheimer. 2020. "Preparing the Sample of Integrated Labour Market Biographies (SIAB) for Scientific Analysis: A Guide." *Journal for Labour Market Research* 54 (1). [180, 218]

Venn, Danielle. 2012. "Eligibility Criteria for Unemployment Benefits: Quantitative Indicators for OECD and EU Countries." Working Paper, 131, OECD (Social, Employment and Migration Working Papers). [218]