Talent Poaching and Job Rotation

Diego Battiston, Miguel Espinosa and Shuo Liu*

May 2024

Abstract

The value of a firm's service lies both in its workers and its relationship with clients. In this paper, we study the interaction between client-specific experience accumulated by workers, poaching behaviour from clients and strategic rotation of workers by firms. Using detailed personnel data from a security-service firm, we show that an increase in client-specific experience increases both the productivity of workers and their probability of being poached. The firm reacts to this risk by rotating workers across multiple clients, and more frequently so to those workers more likely to be poached. Furthermore, we find that after a policy change that prohibited poaching, the firm sharply decreased the frequency of rotation which in turn increased workers' productivity. We propose a theoretical model that guides the empirical patterns and allows us to argue their external validity beyond our specific empirical setting. Finally, we provide survey evidence from the security service sector, demonstrating the consistency between our findings and industry observations.

1 Introduction

A well-documented and widespread feature of labor markets is that firms take actions to avoid their workers leave and work for competitors (Lipsitz and Starr, 2022). This concern has become less important over time because, across industries and countries, firms increasingly rely on service providers to undertake jobs that were previously carried out by their workers (e.g., Goldschmidt and Schmieder, 2017; Dorn et al., 2018). However, this significant labor market change increases the prominence of a concern that has received less attention but is also important for service firms: their workers can leave and work for clients.

On the job, outsourced workers accumulate experience that makes them more productive to clients. However, after a worker has acquired sufficient skills specific to a client, that client may want to hire the worker in-house. Anticipating this potential loss of both employees and clients,

*Battiston: School of Economics, University of Edinburgh. Espinosa (*corresponding author*): Department of Management & Technology, Bocconi University. Liu: Guanghua School of Management, Peking University. We are grateful for useful comments and suggestions to Jérôme Adda, Anik Ashraf, Jordi Blanes-i-Vidal, Francesco Decarolis, Alexia Delfino, Florian Englmaier, Josh Feng, Guido Friebel, Robert Gibbons, Albrecht Glitz, Gianmarco Leon, Igor Letina, Fei Li, Rocco Macchiavello, Margaret Meyer, Dimitri Migrow, Massimo Motta, Claudio Panico, Eric Posner, Yona Rubinstein, Armin Schmutzler, Sandro Shelegia, Carlos Serrano, Christopher Stanton, Claudia Steinwender, Evan Starr, Xi Weng as well as participants at several seminars and workshops for useful feedback. Juan Lugo, Angela Rodriguez, and Jairo Galvis provided excellent research assistance. Finally, we acknowledge our partner organization for sharing the data and helping us understand the institutional setting. Shuo Liu acknowledges financial support from the National Natural Science of China (Grants No. 72192844 and 72322006).

service firms may take costly actions to prevent poaching.[1] We argue that among the set of tools available to deter poaching, one of them consists of rotating workers from one client to another. By doing so, service firms hinder workers' acquisition of client-specific skills (henceforth CSS), so that workers remain sufficiently unattractive to the clients.

We are not aware of any existing study that quantifies how severe the phenomenon of talent poaching from clients is. Nevertheless, media coverage and public discussions suggest that many and various types of firms and clients do care about this type of poaching. For instance, there is registered involvement in poaching suppliers' employees for leading companies such as Apple (Bradshaw, 2015, 2017) and less eye-catching multi-million dollar firms like Guardsmark.[2] More generally, the phenomenon has been documented for a diverse set of occupations and industries, including nursing (DLA Labor Dish Editorial Board, 2014), cleaning (Shubber, 2018), engineering (Chaput, 2018), marketing (Liffreing, 2018), managerial services (StevensVuaran Lawyers, 2019), travel advising (Pestronk, 2019), and game publishing (Schreier, 2020) among many others. It is therefore unsurprising that the issue has drawn public attention in various countries, such as Australia (StevensVuaran Lawyers, 2019), Canada (Chaput, 2018) and the US (Bennet, 2018).

Despite the prevalence and importance of this poaching problem, research on this topic has been limited, probably due to the lack of a comprehensive database that collects information on the transition and performance of service workers across multiple clients and their poaching behaviour. To overcome this challenge, we concentrate on the security-service industry, where we have access to detailed data from a single firm that allows for an in-depth examination of a phenomenon. Focusing on this particular case study provides an appropriate framework to investigate the issue of poaching for two reasons. First, in the middle of our sample period, a non-poaching policy was implemented by the government in the country where our partner firm is located, giving exogenous variation to the extent that poaching behaviour is allowed. Second, we have access to a very extensive dataset. During 74 months, the firm allocated 589 guards to 116 residential buildings daily. For each guard, we know her socio-demographic information as well as when and where she worked. For each building, we have information about its size and location. Additionally, the data contains two measures of poaching intensity: whether a guard received a formal solicitation from a building, and whether a guard was hired in-house by a building. Finally, we also have information on one of the

[1] This type of strategic response is a familiar problem in antitrust law. When firms are restricted from engaging in anti-competitive behavior, they might turn to alternative inefficient practices to achieve market power (e.g. excessive product differentiation; Makadok and Ross, 2013), which can in turn lead to adverse welfare consequences.

[2] See the United States District Court (E.D. Kentucky, Covington Division) case *Borg-Warner Protective Services v. Guardsmark, Inc.* 946 F. Supp. 495, 27 Nov. 1996.

most important measures of guards' productivity: crimes committed in buildings during guards' working shifts. A natural limitation with any type of organizational study like ours is that may have specific dynamics of workers' outside options and firms' cost-benefit ratios. We are explicit about them throughout the manuscript.

We present three main empirical results. First, guards with more client-specific experience are more effective at reducing crime but are also more likely to be poached. Second, the security firm responds to this poaching concern by rotating guards across buildings, especially those with a higher poaching risk (e.g. men living in large households). Third, the anti-poaching legislation reduced both rotation and crime.

The first result studies the relationship between the CSS of a worker and the poaching decision of the client. We find that an increase in the duration that a guard has worked for a specific building increases her probability of being poached by that building, even after controlling for her total working experience. We argue that this is because the skill that a guard acquires by working with the same client is important for her productivity: As a guard accumulates more working shifts in a building, both the probability that a crime occurs in that building and the expected value of stolen properties decrease.[3] These results are robust to different alternative exercises, including an instrumental variable approach based on the system that the firm used to allocate guards to shifts and an event study. In particular, the latter exercise allows us to understand better how crime rates vary before and after rotation events. Notably, we find that once a guard is rotated to a new building (thereby resulting in a reset on the accumulation of CSS), there is a subsequent and large increase in crime incidence and the value of property lost in that building.

The second empirical result shows that the firm tends to rotate more often those guards at a higher risk of being poached. To estimate the poaching risk, we exploit the fact that buildings prefer to hire directly guards with certain baseline characteristics. Based on these features, we construct a cross-section worker-specific index of poaching risk using a machine learning approach. Our analysis demonstrates a strong correlation between the rotation of guards and this poaching risk index. Specifically, a one-standard-deviation increase in the estimated risk of poaching is associated with 1.5 additional percentage points in the probability of rotation. This estimate is sizable as it corresponds to 40% of the mean of the dependent variable.

The third and last main empirical result exploits a policy change that *de facto* limited buildings

[3] Huckman and Pisano (2006) find a similar relationship between the quality of a cardiac surgeon's performance at a given hospital and her recent volume of surgeries at that hospital.

from directly hiring guards in-house. If the security company rotates workers to curtail their CSS acquisition, and therefore to decrease the probability of being captured by the clients, this rotation should decrease following the policy change. Consistent with this intuition, we show that guards who were more likely to be poached before the policy introduction experienced less intensive rotation after the policy took effect. More precisely, a one-standard-deviation increase in the poaching risk is associated with a reduction of 2 percentage points in the probability of rotation after the policy change. The magnitude of this effect is large (58%) compared to the average monthly rotation before the policy took effect. We complement this result by showing that guards who were rotated less frequently before the policy change exhibited greater productivity increases, as evidenced by larger decreases in crime.

Taken together, our empirical findings suggest that the firm strategically rotated its workers excessively to prevent them from being poached. Then, when a non-poaching policy was implemented, the firm reduced rotation, allowing workers to acquire larger CSS and as a consequence, crime rates decreased. An important lesson from our results is that in environments where service companies take costly actions to avoid poaching, a policy that prohibits poaching can increase the productivity of workers.

One potential concern regarding our results is that they may be driven by the specific empirical setting that we examined. To advance in the broad applicability of the mechanism proposed by our paper, we pursue two supplementary approaches.

Firstly, we present extensive survey evidence from firms in the security sector, as well as anecdotal evidence from firms in other industries such as legal, software development, and cleaning services. The qualitative evidence gathered from these diverse sources not only bolsters the validity of assumptions made in our study but also strengthens the generalizability of our core findings. For instance, we find that vertical poaching is a salient issue in the aforementioned industries and that in many cases rotation emerges as a strategic managerial response to the poaching problem.[4]

Secondly, we propose a theoretical model that captures the strategic tension arising from empirical settings that as ours are prone to both poaching and rotation. In our model, a firm employs a pool of workers and transacts with a client. At the outset, the client, who lacks the necessary in-house labor, outsources a production activity to the firm. As a worker gains productivity-enhancing experience by performing the client's activity, the client may find it cost-efficient to hire that worker directly. We show that, even with other retention tools available (e.g., pecuniary incentives, ameni-

[4]For more information on the survey, see Section A.1 in the Online Appendix.

ties -Bidwell et al. (2015)-, or non-poaching contracts -Starr et al. (2021)-), the firm may prefer to preempt poaching by inefficiently rotating workers before they reach a client-experience threshold. In equilibrium, workers with more desirable characteristics (e.g., larger baseline productivity) face higher poaching risks and are rotated more frequently. Accordingly, our model corroborates that a non-poaching policy can enhance worker productivity by eliminating strategic over-rotation.

Related literature. Researchers have long been aware that job rotation can impede skill accumulation and decrease job-specific productivity (Ickes and Samuelson, 1987; Groysberg et al., 2008). However, to justify the widespread use of rotation in organizations, some argue that the learning benefits of rotation can outweigh the potential productivity loss. This applies to both employee learning, which emphasizes that rotation can increase the general human capital of workers by allowing them to be exposed to a wide range of experiences, as well as employer learning, which stresses that rotation can be an effective tool for firms to learn about relevant characteristics (e.g. productivity) of different workers and/or tasks (Meyer, 1994; Ortega, 2001). Another strand of research focuses on the incentive aspect of rotation. The general insight is that many agency problems between firms and workers can be alleviated by including job rotation as part of the organizational design (e.g. Hertzberg et al., 2010). As we will show, these familiar hypotheses do not seem to be consistent with our empirical setting. Instead, our paper proposes and demonstrates a totally different rationale for job rotation — it can be used as an organizational remedy to mitigate poaching risk.

There is also literature studying how poaching affects on-the-job training (e.g., Acemoglu, 1997; Moen and Rosén, 2004). In this literature, firms provide both general and job-specific skill training to their workers. It has been well understood that if firms cannot avoid poaching from their competitors (because *non-poaching agreements* between employers operating in the same product market are illegal), the provision of general skill training will be insufficient. We contribute to this literature by showing that in the complementary case where firms cannot avoid poaching from their clients, the acquisition of job-specific skills may also be distorted.

Finally, some literature has also identified factors that can constrain worker mobility even in the absence of employer intervention, such as organizational status (Bidwell et al., 2015), search and switching costs (Wright et al., 1994), or limited information about outside options (Campbell et al., 2012). Important for our work is that the problem of firm-sponsored general-skill provision can be alleviated by *non-compete clauses* (e.g., Marx et al., 2009; Starr et al., 2020, 2021; Lipsitz and Starr, 2022). This type of clause limits workers from leaving their current employers and working for other firms in the same industry, sometimes within a pre-specified geographic area and period. Similarly, the employers in our setting also take actions (job rotation) to hinder workers from quitting their job and working for another employer (who in this case is a client). However, while policymakers tend to be against non-compete clauses (e.g., Dougherty, 2017), we offer a new theoretical rationale and empirical evidence to make the case for a non-poaching policy: it can enhance productivity (e.g., improve crime prevention in our setting).

2 Institutional Setting

We partnered with a private firm in Colombia that provides security services to residential buildings. We have detailed 12-hour shift data of the firm's transactions from February 1992 to April 1998. Our sample consists of 589 security guards allocated to 116 buildings. For each guard, we have information on when and where she worked, previous professional experience, age, gender, and residential address. For each building, we know who worked there and when, where it is located, the number of flats, the required number of guards, and the type of crime that occurred (if any).

2.1 Relationship Between the Security Firm and Buildings

The allocation of guards to buildings works as follows: A guard works successively for 12 days in shifts of 12 hours each: six consecutive days during the *day shift* (6 am - 6 pm) and the following six days during the *night shift* (6 pm - 6 am).[5] After 12 working days, the guard rests for two days. Most guards are allocated to work in a unique building for several months. However, about 15% of guards work exclusively covering the resting days of their colleagues. As a result, they work across multiple buildings during the 12-day period (see Figure A.1 for an example of guards's shift schedule). We refer to the above two types of guards as type-I and type-II, respectively.

The private security firm transacts with multiple residential buildings. During the whole sample period Colombian legislation prohibited any type of firm from using any formal contracts (e.g., non-compete clauses) to restrict the possibility of workers being poached by other firms in the same product market. However, before 1994, it was legally possible that residential buildings poached security guards. We argue in this paper that the security firm rotated workers from one building to

[5] There are very few occasions when guards slightly depart from this schedule. For instance, illness episodes of one guard can result in other guards working overtime.

another to avoid poaching. When these rotations occurred, they were typically communicated to both the building and the guard about one week before the rotation date.

Workers do not necessarily have the same preferences between working directly for clients and being employed by security companies. There are trade-offs to consider. Average wages are determined by the market and do not significantly differ between internal and external hiring. Working directly for clients provides guards with a more amicable environment and assures that in expectation they will be working in the same place for a long period. The latter factor is appreciated by the guards because security companies often fail to consider workers' home locations and transportation expenses in rotation planning. In contrast, employment through security companies provides advantages like better training opportunities, as well as job security independent of particular clients. Our conversations with several firms in the sector indicate that the contracts that workers sign do not substantially differ whether they are with security companies or with clients directly.

Conversations with buildings that initiated poaching show that usually, they have other potential guards lined up to cover the remaining shifts before poaching occurs. These potential substitutes (of the poached guard) often include former in-house guards, or referrals from guards or residents. When a building poaches a guard, the firm terminates the contract with the building. The non-poached guards working in the poaching building are typically transferred to other buildings. For every poaching case, we observe the identity of the hired guard, the building that initiated the poaching, and the exact date that the guard left the firm. In our data, all poaching episodes happened while the guard was working in the building.

Although the buildings could post a vacancy and hire guards directly before the policy change, our survey evidence indicates that most of the buildings preferred to outsource these positions because (i) the security company has a comparative advantage in performing the job due to the economies of scale (e.g., it may bind the needs of different clients through the training and management of a large set of employees), and (ii) the company (acting as an insurance provider) pays a fraction or the totality of the stolen items to the building if a crime occurs. The magnitude of this fraction depends on the proven responsibility that the guard had in the crime. We do not have information on the amount of money paid for each crime episode. Insurance companies also offer services to buildings with in-house guards.

Finally, buildings always provide all necessary materials and amenities (like a staircase, heater, etc.) that can increase guards' productivity from the start of the contract. Failure to meet this condition will leave the building uninsured in case of crime.

Rationales for Rotation. There are multiple reasons why firms in our empirical context decided to rotate workers. We classify them into two categories: strategic and non-strategic. The former relates to the firms' motive to deter vertical poaching (choices that firms can use to retain human capital in the face of external pressures from clients), while the latter encompasses motives that are unrelated to that goal. In our survey, firms reported numerous non-strategic reasons for rotation: guards requesting reassignment to buildings closer to their homes; guards needing sick leave; and guards asking for time off to attend to personal matters, etc. We consider these latter reasons for rotation to be mainly idiosyncratic. As such, we expect them to occur throughout the sample period and therefore should not be affected by the policy change. This implies that there may be some level of rotation absent strategic motives. Our survey evidence also indicates that rotation did not damage the firm reputation, since it was a common practice in the sector in the 1990s.

2.2 Client-Specific Skills in Our Context

One of the most important tasks of a guard is to control entry into buildings and invigilate for the presence of potential criminals. When a visitor arrives, the guard contacts the resident that the visitor wishes to see and verifies if the guest is welcome. If the reply is positive, the guard registers some basic information about the visitor (name, national ID number, time of arrival) and lets him/her in. This process takes about 5-7 minutes, and both guards and frequent visitors prefer skipping it due to transaction costs.

The best guards reduce transaction costs by distinguishing residents and frequent visitors from the rest. Recognizing those residents and visitors is a CSS. Naturally, this skill increases over time as guards become more familiar with the identities of residents or those who visit the building frequently. However, without sufficient experience in the building, a guard is not able to screen unwanted visitors (e.g., thieves) from others. Hence, an inexperienced guard either makes everyone pay the transaction costs or overlooks the entry of unwanted visitors.

According to our partner firm and survey, the CSS of a guard also include the understanding of the inner workings of the building. Guards accumulate this knowledge over time, allowing them to prevent crime more efficiently. This is the case because criminals not only try to enter the building by registering with the guard at the entry but also by other means. The longer a guard works in a building, the more likely she will be able to detect when and how criminals try to enter it. In this sense, the performance of a new guard may be different from a more experienced peer, as the latter would be more likely to anticipate how criminals might try to steal residents' property.

3 Theory

Before proceeding to analyze the data, we present a dynamic agency model that accentuates the key tension arising from our empirical setting: the accumulation of client-specific skills increases not only productivity but also poaching risk. Our goal with this model is twofold: First, we aim to develop formal and testable predictions to guide the subsequent empirical analysis. Second, we seek to shed light on the generalizability of our proposed mechanism – specifically, the circumstances under which we would expect service-providing firms to utilize rotation over alternative strategies to counter client-poaching.

Our model is underpinned by two key assumptions, which we will validate empirically. The first assumption is that the firm-client relationship ends when an attempt to poach occurs, irrespective of its success. The second assumption is that the departure of employees to a client incurs substantial costs for the firm. Should poaching not result in significant costs – for instance, if having a former employee working on the client side can facilitate future business opportunities (e.g. Somaya et al., 2008) – the firm might actually be inclined to encourage rather than deter poaching.

Multiple sources of empirical evidence support the validity of the first assumption in our setting. In our data, every client that poached a worker disappeared from the sample. Indeed, our partner firm confirmed that clients typically cease requiring services after poaching attempts, and conducting future business with those clients is unlikely due to the loss of trust. Our survey evidence corroborates that this reaction to clients' poaching behavior is common in the security services industry. As for the second assumption, we justify it by providing a back-of-the-envelope calculation of the cost imposed on the firm due to client poaching (see Appendix Section A.3 for details). Our calculation indicates that a poaching episode costs around $\$5,069$ to the firm in total. This cost is substantial, amounting to about 20 times the monthly minimum wage.[6]

3.1 Model

We consider a client (he) that repeatedly engages in a production activity at period $t = 0, 1, 2, ... +\infty$. Performing the activity requires a unit input of labor (from a worker, she) at every period. At the beginning of the game, the client does not have an in-house worker, so he outsources the activity to a service firm that specializes in providing such a workforce.

[6] In the Online Appendix, we also provide an estimate for the cost of rotation. Our back-of-the-envelope calculation indicates the cost of rotating a guard is approximately \$205. This amount is substantially smaller than the estimated $\$5,069$ loss per poaching episode.

The productivity of a worker depends on the number of periods $e \in \mathbb{N}$ that she has accumulated in serving the client. Specifically, a worker generates a surplus $z(e)$ that strictly increases with her client-specific experience. We interpret this surplus gain as a product of the worker's adaptation to the working environment and proficiency in required tasks, which fosters welfare improvements such as enhanced protection for the client, smoother interaction between worker and client, and reduced effort costs for the worker. All players are risk-neutral and have the same discount factor $\delta \in (0, 1)$.

In every period, the players interact with each other according to the following timeline (see Figure A.2 for a graphical illustration). Initially, the service firm selects a worker to assign to the client. The firm can either send the same worker to the client as in the previous period or appoint a new worker to replace the previous one. The client then decides whether to accept the firm's service or not. If accepted, the client receives a fixed payoff $\underline{v}$ and pays a fee p to the firm, where $0 < p < \underline{v}$, and the stage game ends. In this case, the flow payoffs accrued to the firm and the worker are given by $\pi_t = p - \underline{w} + \theta + \alpha z(e_t)$ and $u_t = \underline{w} + (1 - \alpha)z(e_t)$, respectively. Here, $\underline{w} > 0$ is the default wage that the firm pays to its employees (e.g., minimum wage as our qualitative evidence shows), $\theta \in \mathbb{R}_+$ captures the baseline productivity of a worker, and $\alpha \in (0, 1]$ measures how surplus generated from client-specific experience is divided between the firm and the worker.[7]

Should the client opt not to purchase the service from the firm, he may attempt to poach the worker by proposing a wage offer w. The firm can then respond with a counteroffer w' (which encompasses both wages and amenity changes), but the worker is free to decide whether to stay or to leave. Additionally, the client incurs a fixed cost $c_t = \underline{c} + \varepsilon_t \in \mathbb{R}_+$ for initiating poaching. This cost incorporates expenses for acquiring new equipment or recruiting complementary co-workers, administrative hassles, and potential loss of reputation. It may also reflect the binding nature and effectiveness of any non-poaching agreement between the involved parties: higher costs indicate situations where violating such agreements is more difficult or punitive for clients (e.g., higher legal costs in expectation). Specifically, the constant $\underline{c}$ represents baseline poaching cost, while the term ε_t captures stochastic fluctuations. We assume that each ε_t is an independent draw from a commonly known distribution $\Pr(\varepsilon_t = \varepsilon_L) = 1 - \Pr(\varepsilon_t = \varepsilon_H) = \lambda \in (0, 1)$, where $\varepsilon_L < \varepsilon_H$, and the client privately observes the realization of ε_t before making the poaching decision.[8]

[7]The assumption of no client share under outsourcing aligns with our empirical context, because the buildings were fully insured (against losses from theft) by the security firm. Although the buildings might have the option to secure insurance elsewhere even if they hire guards directly, it is unlikely that the associated cost would be significantly lower than that charged by the firm. In any case, our results will persist provided that the client derives substantially less benefit from workers' productivity gains under outsourcing compared to in-house production (e.g., we could allow the firm to transfer some of the increasing surplus to the client in the form of a decreasing service fee).

[8]Throughout our analysis, we will set the values of ε_H and ε_L sufficiently apart whenever needed, ensuring that

Per the first key assumption discussed at the start of this section, the contractual relationship between the service firm and the client will end irreversibly, no matter the poaching result. Specifically, if the client's recruitment attempt fails, he will thereafter receive zero payoff, while the firm will receive a constant flow payoff $\underline{\pi} - w'$, where $\underline{\pi}$ captures the firm's anticipated profit from finding a new client later, and w' is the wage promised to the worker in the counter-offer. Alternatively, upon successful recruitment with a wage w, the client gains the ability to produce in-house going forward. In this case, the flow payoffs for the client and the worker will be $v_t = -w + \theta + \beta z(e_t)$ and $u_t = w + (1 - \beta)z(e_t) + \gamma$, respectively. Here, the parameter $\beta \in (0, 1]$ determines the surplus split between the two parties, while $\gamma \in \mathbb{R}$ represents the worker's relative preference for working for the firm versus the client.[9] Meanwhile, the firm, now needing to find both a new client and a new employee, will receive a constant flow payoff $\underline{\pi} - \kappa$ going forward, where $\kappa > 0$ incorporates factors such as hiring and training costs. Aligned with the second key assumption discussed earlier (that losing workers to a client is very costly for the firm), we take the value of $\underline{\pi} - \kappa$ to be sufficiently small, thereby avoiding any complications from the firm preferring its workers to be poached.

3.2 Equilibrium analysis

The equilibrium analysis of our model is non-trivial due to the mutual influence between the client's poaching decision and the firm's rotation scheme. However, the client's preference is to first utilize the firm's service, and then switch to in-house production later on once the assigned worker has gained enough experience to become highly skilled at the task. Leveraging this monotonicity of the client's poaching incentive, our main theoretical result below establishes the existence of an equilibrium in which the firm strategically implements a stationary rotation scheme.

Proposition 1. *Assuming λ is small enough, there exists a Perfect Bayesian equilibrium in which:*

(i) every worker is rotated by the firm after having served the client for T_H periods, and

(ii) poaching of a worker occurs before rotation by the firm if the client draws a low poaching cost and the worker has served for longer than $T_L < T_H$ periods,

where the values of T_H and T_L are uniquely determined by the model's parameters.

---

the timing of the client's poaching decision will depend in a meaningful way on which poaching cost is drawn (see Appendix Section A.5 for the exact parametric assumption made).

[9] In general, γ can incorporate both intrinsic preference and any direct cost that the worker pays by accepting the client's poaching offer, such as repayment to the firm for breaching a non-compete clause.

Proposition 1 highlights that even if the firm has other tools at its disposal to counter the risk of poaching, such as offering higher wages and/or better amenities to its employees, rotation can still be the preferred strategy. Intuitively, since the outsourcing relationship will end whenever poaching occurs, the maximum wage that the firm is willing to pay to retain a worker is capped at the replacement cost κ. Thus, if the productivity gain from the worker's CSS accumulation eventually outweighs even a high poaching cost, the client will outbid the firm to hire the skilled worker directly. In this scenario, the compensation package that the firm can offer will be insufficient to deter poaching. However, rotation remains an effective pre-emptive tactic against poaching: By optimally setting T_H, the firm imposes a ceiling on the attractiveness of any individual worker to the client. This allows the firm to garner some productivity gains from growing experience while keeping the poaching risk low.

Next, we summarize the testable predictions that emerge from our model, each based on a comparative statics result relating to the equilibrium described in Proposition 1. The first prediction, detailed in the proposition below, is that clients are more likely to poach workers who have accumulated more client-specific experience (and who are also more productive).

Proposition 2. *As a worker accumulates more experience specific to a client, her likelihood of being poached increases upon being assigned to that same client again.*

The second prediction is that workers at higher poaching risk will be rotated more frequently by the firm. Specifically, our next proposition formalizes the driving force behind this correlation: both poaching risk and rotation frequency are positively associated with a worker's baseline productivity.

Proposition 3. *Consider two groups of workers, where workers in the first group have higher baseline productivity than those in the second group. Then, workers from the first group will also: (i) face higher risks of poaching whenever assigned to a client, and (ii) be rotated more frequently by the firm.*

Last, our model predicts a negative relationship between rotation frequency and poaching costs.

Proposition 4. *The frequency at which the firm rotates its workers decreases as the baseline poaching cost increases. Specifically, if the poaching cost is sufficiently high, the firm will cease using worker rotation as a strategy to counter client-poaching.*

To sum up, our theoretical analysis demonstrates that the threat of poaching can lead to excessive job rotation, destroying valuable human capital. Implementing a non-poaching policy would halt this vicious dynamic — if poaching were prohibited, the rotation should be merely driven by factors exogenous to our model (e.g. sick leave of workers). This would enable a greater accumulation of CSS, thereby increasing the productivity of workers. However, the policy may not improve welfare equally for all agents in the economy. The firm will unambiguously benefit from the policy because its business with the client will be protected and it can capture more surplus from the transaction due to the larger CSS of the workers. In contrast, the workers could be worse off as the policy change cuts off their access to valuable outside options. Clients would also be affected, as they would no longer be able to poach workers that they like. Overall, we caution that the net welfare impact of the policy can be ambiguous because it may depend on the relative magnitudes of these countervailing effects.

4 Data and Empirical Analysis

4.1 Descriptive Statistics

Table 1 provides descriptive statistics of our database. The table summarizes key predetermined characteristics of the guards, such as previous experience working as a security guard, military training, and various socio-economic variables. Most guards are male, have military training, and about half of them have experience working as security guards before joining the firm. There is a large variation in terms of age and migration status among the guards. On average, guards reside with 5 additional family members, with only 7% of them living alone. About 80% of the guards joined the firm before our sample period started. We do not have wage information for every guard, but we know that the majority of guards earn the minimum wage during the entire sample period and their earnings do not depend on building-specific experience.[10] The monthly service fee that the firm charges for providing a guard position in a building (which requires of three guards) is about 5 times the monthly minimum wage.

[10]We have wage information for a small subset of guards. Despite inherent limitations due to measurement errors, our analysis reveals that year-fixed effects and the years since the guard joined the firm account for more than 90% of the variation in real wages (indeed, a single-year linear trend explains 73% of the variation). These findings are consistent with the company's narrative that wages change in a very mechanical and predictable way based on minimum wage and tenure.

Table 1: **Characteristics of Guards and Buildings**

	(1) Mean	(2) Sd	(3) Min	(4) Max
Guard Characteristics				
Number of Guards	589			
Type-I Guard	0.88	0.32	0	1
Male	0.78	0.41	0	1
Military experience	0.64	0.48	0	1
Neighborhood strata	1.89	0.57	1	5
Household size	5.50	3.43	0	12
Lives alone	0.07	0.25	0	1
Age	35.93	9.15	20	71
Past experience as guard (months)	31.50	51.24	0	285
Has experience as guard	0.49	0.50	0	1
Tenure (months)	23.47	17.27	0	89
Immigrant	0.42	0.49	0	1
Recent immigrant	0.19	0.39	0	1
Started job on/before January 1992	0.81	0.40	0	1
N of shifts worked in the month	24.32	5.56	1	54
Max tenure in the building (in months)	16.57	17.38	0	65
N of buildings per month (Type-I)	1.03	0.18	1	4
N of buildings per month (Type-II)	2.09	0.77	1	5
Rotated to a new building during the month (Type-I)	0.02	0.15	0	1
Rotated to a new building during the month (Type-II)	0.05	0.21	0	1
Avg. shifts worked per building (Type-I)	26.15	2.56	1	29
Avg. shifts worked per building (Type-II)	9.34	4.12	1	27
Building Characteristics				
N of buildings	116.00			
N of guards	4.37	2.49	2	13
N of flats	98.05	57.15	20	299
Neighborhood strata	2.78	1.28	1	6
N of crimes per month in the building	1.50	3.35	0	34
Value of property lost (USD)	46.37	116.44	0	1,421
Value of property lost (USD) if crime occurs	143.04	167.32	0	1,421

Table 1 also reports variables related to the rotation of guards across buildings. On average, for every building that they are assigned to, a type-I guard accumulates 26 shifts per month, while a type-II accumulates 9 shifts a month. Further, type-I guards work on average in 1.03 buildings per month, and only 2% of them rotate each month. This contrasts with type-II guards who work on average in 2 different buildings each month and rotate to a new building with a monthly probability of 5%.[11]

Finally, the bottom part of Table 1 presents summary statistics for buildings. Buildings are relatively large, with an average of 98 flats, and require 4.4 different guards to cover all the shifts during a month. The average strata of the neighborhoods where the buildings are located is 2.8. The

[11]Figure A.3 shows that the typical rotation happens before the peak of workers' performance. The figure also shows that productivity, as measured by crime incidence, decreases monotonically over several months with a change in slope only around the 20th month.

strata value captures several measures of the quality of housing on a scale from 1 to 6. Neighborhoods with larger strata tend to be safer. The average building experiences about 1.5 crimes in a month. The most common crime is burglary. Stolen properties frequently include items from the common space of the building (ladders, fridges, automobiles, bicycles, motorcycles) as well as electronic appliances and jewelry from flats. The average value of property stolen is about 46 USD. This corresponds to approximately 21% of the 1993 Colombian monthly minimum wage.

4.2 Client-Specific Experience, Worker Productivity and Poaching

Building-specific experience and guard's productivity. As characterized by our theoretical model, client-specific experience improves workers' productivity over time. Although we do not observe all the possible dimensions of each guard's performance (e.g., efficiency of visitor entry registration, trust between residents and the guard, etc.), we do have information about the incidence of crime. According to the security firm and buildings, crime is the single most important measure of productivity in this setting.

The importance of building-specific experience for crime prevention has been emphasized both by our partner firm and by other surveyed security companies. For instance, one firm stated: "*[T]he best guards are those that spend a significant amount of time in the building. Spending time with a client helps them to understand the specific location that criminals can use to enter the building.*" Likewise, interviewed companies from other sectors that assign workers to clients also noted that client-specific experience is an important determinant of worker productivity. For instance, one cleaning company stated: "*[For our staff] to work efficiently, they need to work in the same environment consistently.*"[12]

To provide more robust evidence on the role of building-specific experience, we estimate the following equation at the guard-building-week level:

$$Crime_{ibt} = \beta ExpInBuilding_{ibt} + \eta TotalExp_{it} + \pi Wd_{ibt} + \delta_{ib} + \gamma_t + \epsilon_{ibt}, \tag{1}$$

where $Crime_{ibt}$ is an indicator for the occurrence of crime in a shift when the guard i was working at building b during week t. We also consider an alternative dependent variable: the inverse-hyperbolic-sine transformed (IHST) value of property stolen if a crime occurs.[13] Our main explanatory variable

[12] For more survey evidence, refer to section A.9 in the Online Appendix.

[13] As a robustness exercise, we have estimated equation (1) with the value of property stolen in levels, and the effect relative to the mean is roughly similar.

$ExpInBuilding_{ibt}$ is the number of shifts that the guard worked in the building (expressed in months). We include pair-specific fixed effects δ_{ib} and exploit the variation in building-specific experience *within* each guard-building pair over time. We also include week fixed effects γ_t to avoid confounding the effect of building-specific experience with systematic changes in crime over time. To isolate the direct effect of CSS on performance, we control for the overall experience of the guard $TotalExp_{it}$. Finally, we control for Wd_{ibt}, the number of days the guard worked in the building during the week, as the likelihood of encountering a crime is higher for guards who worked more days that week.

The first column in Panels A and B of Table 2 shows the estimates of equation (1). The estimated coefficients of building-specific experience are negative and significant. Column (2) shows that results remain similar if we control for the schedule characteristics of the guard. These estimates are large relative to the mean of the dependent variables. For instance, four additional months of experience in a building is associated with an approximately 1 percentage point reduction in the probability of crime (25% of the mean) and decreases the monetary cost of crime by more than 12%. Equivalently, an increase of one standard deviation in building-specific experience is associated with a reduction of the probability of crime equal to 84% of its mean and 40% of the monetary cost of crime.

While equation (1) includes an extensive set of controls to account for a broad spectrum of potential confounders, it is important to interpret the results cautiously due to the absence of experimental variation. This limitation restricts the ability to establish a definitive causal relationship. To reduce this concern, in Column (3) of Table 2 we exclude from the estimation the last quarter of the guard in the building. Intuitively, crimes that prompt rotation would occur more frequently in a guard's final months in a building, as rotation may happen soon after incidents. However, estimates from this column do not change significantly with respect to the previous columns, which suggests a limited role of this type of reverse causation in explaining our results.[14]

In Appendix Table A.1, we show that our estimates are broadly similar under several additional robustness checks. In Column (1), we use an instrumental variable approach that leverages the haphazard assignment of guards into types and the differential rate at which Type-I and Type-II guards accumulate building-experience over time. This IV approach, which we discuss in detail in Appendix Section A.6, aims to reduce concerns regarding some time-varying confounding factors

[14] This finding is also consistent with the fact that the occurrence of crime is not more likely in the days before rotation, as we show in Appendix Figure A.4.

that the fixed effects may not be able to absorb.

Table 2: Productivity and Client-Specific Experience

	(1)	(2)	(3)	(4)
	Panel A: Crime occurred During Guard's Shift			
Experience in Building (months)	-.0026***	-.0026***	-.0029***	-.0024***
	(.0003)	(.0003)	(.00032)	(.00027)
Total Experience (months)	.00076*	.00076*	.00074*	.00037
	(.00043)	(.00043)	(.00043)	(.00037)
N	121,132	121,132	107,921	107,921
R2	.089	.1	.1	.16
Mean Depvar	.042	.042	.044	.044
	Panel B: IHST Value of Property Lost in Crime			
Experience in Building (months)	-.032***	-.032***	-.035***	-.03***
	(.0035)	(.0035)	(.0037)	(.0032)
Total Experience (months)	.0094*	.0094*	.0091*	.0059
	(.0051)	(.005)	(.005)	(.0043)
N	121,132	121,132	107,921	107,921
R2	.087	.1	.1	.16
Mean Depvar	.51	.51	.53	.53
Guard X Building FE:	YES	YES	YES	YES
Week FE:	YES	YES	YES	YES
Days Worked Week:	YES	YES	YES	YES
Shift and Weekend controls:	NO	YES	YES	YES
Excl Last Guard-Build Quarter:	NO	NO	YES	YES
Neighb X Month FE:	NO	NO	NO	YES

N guards = 567; N buildings = 116. All regressions are at guard x week x building level. The independent variable is the accumulated experience of the guard in the building (measured in months). In Panel A, the dependent variable is an indicator of a crime occurring during a shift when the guard was working in the building during the week. In Panel B, the dependent variable is the (inverse hyperbolic sine transformation of the) estimated value of the property stolen or destroyed during the crime. All regressions control for the total experience of the guard and number of shifts that the guard worked during the week. Columns (2) to (4) include additional controls for the share of days that the guard worked on night shifts during the week and an indicator for whether the guard worked at least one weekend shift during the week. Columns (3) and (4) exclude the last quarter the guard worked in the building. Column (4) controls for the interaction between the area of the building and the month. Robust standard errors clustered at the guard level are reported in parentheses.

Given that equation (1) relies on linearity assumptions of the independent variables, in Column (2) of Table A.1 we also show that effects are robust to controlling non-parametrically for the total experience of the guard. Additionally, in Column (3) of Table A.1 we exclude from estimations the first month of each guard in the firm to address the possibility that results are driven by a period when crime could be disproportionately high due to the lack of overall experience in the job. Finally, in Column (4) of Table A.1, we estimate equation (1) dropping all the observations corresponding to the first building in which we observe the guard.

An event study of guards' rotation. To provide further empirical evidence on the importance of guards' building-specific experience, we conduct an event study examining the evolution of crime around the time a guard is rotated to a new building relative to the non-rotating guards at the old or the new building. Details can be found in Appendix Section A.7. We find that guards experience a relative increase in crime outputs immediately after they are rotated. Notably, this increase is more pronounced for guards who have accumulated longer experience at their previous building.

Overall, the event study results align with the findings of Table 2, suggesting a negative impact on guard performance due to the loss of building-specific experience after rotation. The findings of Tables 2 and A.2 are important for two reasons. First, they challenge the hypothesis that rotation is implemented to avoid collusion with criminals (Choi and Thum, 2003; Rose-Ackerman, 2010; Bhuller et al., 2020). Under this hypothesis, the longer a guard works in a building, the more likely she may cooperate with criminals and therefore the more likely crime will happen. Contrary to this idea, our results show that crime decreases as guards spend more time in the building. Hence, in the current empirical setting, the main purpose for rotation does not seem to be deterring guards from colluding with criminals.

Second, the results are consistent with the notion that rotation can be inefficient as it destroys skills that positively affect productivity. Therefore, a natural question is why service firms do it. Our theory suggests that rotation can benefit the firm if the accumulation of building-specific experience, absent rotation, increases the risk of guards being poached. In the remainder of this section, we provide empirical evidence consistent with this rationale by showing that buildings prefer to poach guards with greater building-specific experience.

Building-specific experience and observed poaching. Proposition 2 of our theoretical model predicts that a higher building-specific experience increases the probability of poaching. This prediction aligns closely with the narrative presented by our partner firm and is frequently echoed in our survey responses from other security companies. For instance, one firm stated: "*We realized that buildings were poaching guards that spend significant time with them. We did not worry about the newly allocated guards.*".

Table 3: Poaching and Client-Specific Experience

	(1)	(2)	(3)	(4)
Panel A:	**Client-Specific Experience at Poaching (months)**			
		Non-Poached Guards		
	Poached Guard	**All**	**Type-I**	**Same Building**
Mean	13.27	8.79	9.34	7.33
Median	13.18	8.07	8.26	4.46
75th pctile	16.89	13.97	14.92	15.44
Panel B:	**Duration Model (Hazard Ratios)**			
Experience in Building (months)	1.3**		1.5***	
	(.16)		(.2)	
N Past Rotations		.24***		.13***
		(.13)		(.089)
p-val prop hazard	.82	.81	.11	.14
Building RE	YES	YES	YES	YES
Total Experience	YES	YES	YES	YES
Guard Chars	NO	NO	YES	YES
Build Chars	NO	NO	YES	YES

N guards = 454; N buildings = 116. This table investigates the relationship between client-specific experience and poaching. The sample is for the period before the introduction of the Law. **Panel A** reports the mean, median and 75th percentile of the client specific experience (in months) at the week when poaching took place. Column (1) refers to the poached guard. Column (2) includes all the non-poached guards during weeks when a poaching episode took place. Column (3) accounts for all non-poached guards that are type-I during the week when poaching took place. Column (4) refers to all non-poached guards working in the same building of the poached guard during the week when poaching took place. **Panel B** reports the hazard ratios estimated from a Cox proportional hazards model for the time (specified in weeks) the guard spends in the firm before being poached. The model is right-censored for the date when the law was introduced and includes 454 guards observed for a maximum of 109 weeks (39805 total observations). Type-II guards are assigned to the building where they have worked the most days during the week. The model accounts for unobserved heterogeneity across buildings by incorporating a building-specific random effect. In Columns (1) and (3), the main independent variable is the building-specific experience of the guard (measured in months). In Columns (2) and (4), the main independent variable is the cumulative number of rotations prior to the observation's week. All duration models control for the total experience of the guard (measured in months). Columns (3) and (4) also control for guard characteristics (gender, previous experience, household structure, migration status and type) and building characteristics (size, tenure and socioeconomic strata of the area). The table reports the p-value of a global Chi2 test of the proportional hazard assumption of the Cox model based on Schoenfeld residuals.

To empirically substantiate this relationship within our data, we leverage information from the poaching episodes where guards were hired in-house by buildings that were contractually engaged with the firm at the time of poaching.[15] A descriptive analysis (Panel A in Table 3) indicates that poached guards typically have significantly more building-specific experience than their peers, often ranking in the top 30% for tenure. Additionally, guard rotation is strongly negatively correlated with

[15] In 70% of these cases, buildings poached only one guard. We lack shift data for 4 poached guards.

poaching, while building-specific experience is associated with higher poaching likelihood (Appendix Table A.4).

We further explore this association using a duration model, a method well suited for analyzing the temporal relationship between poaching incidents and guards' rotation patterns (or their accumulated experience in a specific building). This model effectively characterizes the time-to-event nature of the poaching data and incorporates truncation and censoring issues more naturally (Bazen, 2011). In Panel B of Table 3, we display the estimated hazard ratios of a Cox proportional hazard model, analyzed at the week level.[16] The baseline model only controls for the total experience of the guard, but this relationship is robust even after controlling for guard and building characteristics.[17] We find that the hazard ratios are substantially larger than one for the building-specific experience and significantly lower than one for the number of times a guard was rotated. The results suggest that each rotation of a guard is associated with a 70% decrease in the baseline hazard of poaching, while each additional month of building-specific experience increases the baseline hazard by 30%.[18]

In sum, our qualitative evidence, robust cross-sectional analysis, and duration analysis collectively and consistently point to a notable increase in the probability of poaching as building-specific experience grows.

5 A Non-Poaching Policy Change

At the beginning of the 1990s, Colombian guerrilla groups heavily victimized the country's civil population. In response to this crisis, a civil-led initiative emerged, advocating for private security forces to provide safety services from these terrorist groups. To ensure activities were conducted under the right and legal institutional framework, and to facilitate and regulate the citizens' initiative, the Colombian government approved the *Decree 356 of 1994*, which mandates clients interested in acquiring any type of security services to access those services only through a company. The decree defines a security company as one with a significant amount of financial assets, which *de facto* limits the possibility that one guard establishes a security company to work as an in-house provider. As a consequence, the introduction of the new law inhibited buildings from hiring guards directly. The

[16] See Van Den Berg (2001) and Fisher and Lin (1999) for discussions of duration models with time-varying controls.

[17] The model is censored at the moment the law was introduced, as poaching was no longer possible after that point. We also account for the heterogeneity at the building level by introducing building-level random effects. We test the proportional hazard assumption using Schoenfeld residuals, and in all cases, we cannot reject the null hypothesis of proportionality (all p-values are above 10% significance level).

[18] As the baseline hazard of poaching is relatively small, these changes do not lead to dramatic shifts in the absolute likelihood of poaching.

partner firm and other interviewed firms mentioned that there were no changes in guards' earnings or service fees charged to buildings around the policy change.

We use this policy change to provide evidence for the central mechanism highlighted by our theoretical model: if the security company rotates guards to trade off client-specific productivity and poaching risk, the rotation of guards should decrease once the law takes effect. Indeed, after the decree was introduced, the unconditional probability that a guard rotates in a given month dropped from 4% to 2%. Figure 1 plots the time series of the average rotation across guards and provides evidence of this pattern.[19]

Figure 1: Evolution of Total Average Rotation

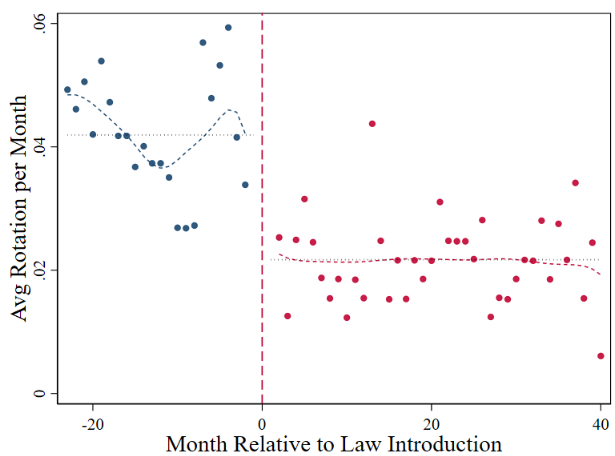

This figure shows the monthly average rotation for type-I guards, with each dot representing the average for all guards in that month. Dashed curves illustrate the local polynomial estimation of average rotation trends before and after the policy change. Dotted lines indicate the average rotation for each period. On average, 295 guards worked each month.

A main limitation we face is the absence of a natural exogenous control group. To overcome this challenge, we compare the change in rotation across guards that had different probabilities of being poached before the policy change. Intuitively, we exploit the fact that guards differ by their baseline characteristics, which make them more or less attractive to be poached by buildings. As implied by Proposition 3, the security firm should rotate more often those guards that are more attractive to buildings – but only before the policy change, when buildings could poach guards. Therefore,

[19]Survey evidence shows no changes around the policy change in either the violence of crime or types of skills that workers need. Some firms reported an increase in the demand for their services. A potential concern is that the poaching risk could have increased if the law was not strongly enforced since the higher demand would lead buildings to anticipate higher fees. In response, they may have tried to poach guards before the law was fully enforced. However, this is at odds with the fact that rotation dropped immediately after the policy change, particularly for high-risk guards. We also evaluated wage data for a subset of guards from the partner company. We find that most of the variation in real wages is explained by aggregated time trends and that individual characteristics do not significantly explain the wage differences. Importantly, we find that wages did not change differentially for guards with different poaching risks, neither before nor after the policy. However, we caveat these results because of the large measurement error of wage data.

we examine whether the frequency of rotation dropped *relatively* more, once the degree came into effect, for guards who were more likely to be poached before the policy change.

We start by estimating an index that reflects the probability that a guard is poached based on her observable characteristics. We focus our analysis on type-I guards to estimate the relationship between observed poaching and the predetermined characteristics of the guard. The use of these characteristics is aligned with anecdotal evidence given by our partner firm. The company stated that, for instance, the size of the household of the guard may predict whether or not a building is attracted to that specific guard (conditional on sufficient CSS). Buildings prefer guards living in large households because, in case of the absence of the guard, she can more easily find a trustworthy replacement for the working shift. Since the sample of poached guards is not particularly large (28 episodes) and given the potentially large number of characteristics (and interactions between them) that could predict poaching, we use a machine learning procedure (Random Forest) to construct an index of poaching risk for each guard. The index is standardized to a mean of zero and a standard deviation of one. The details of this procedure are described in Appendix Section A.8.

5.1 Rotation of Guards due to Poaching Risk

In this subsection, we present some descriptive evidence consistent with the fact that, before the policy change, the firm rotated more often those guards with higher poaching risk. We measure rotation with a dummy that takes the value 1 if the guard is reallocated to work into a new building during the month and 0 otherwise. As an alternative, we also use an intensive margin measure that counts the number of buildings in which the guard worked each month.

Figure 2 shows the cumulative share of guards rotated overtime before the law was introduced (Panel A). Rotation patterns diverge significantly between high-risk (above median) guards and low-risk (below median) guards. As anticipated, guards with a higher likelihood of being poached are rotated more intensively. Panel B of the figure presents a similar comparison for the period after the law took effect. While there is an overall decline in rotation, the reduction is particularly pronounced for the high-risk group after the policy change.

Figure 2: Rotation of High vs. Low Poaching Risk Guards. Cumulative Rates

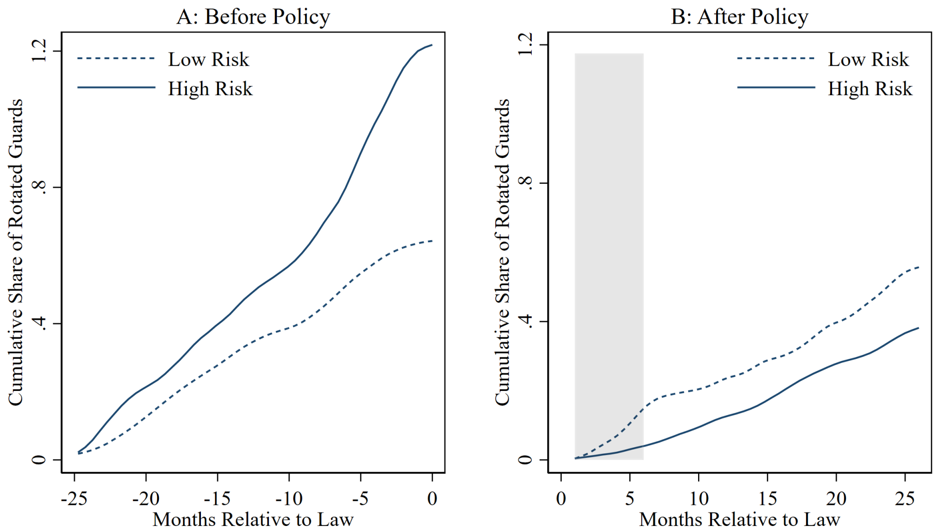

This figure shows the estimated cumulative share of guards who have rotated over time, calculated as the cumulative sum of rotations divided by the total number of guards, measured weekly for 22-month periods. It differentiates between high-risk (above median poaching risk) and low-risk guards (below median). The lines represent kernel-weighted local polynomial smoothing (Epanechnikov Kernel, ROT bandwidth) applied to daily data. Panel A covers the pre-policy period, and Panel B the post-policy period, highlighting the transition period's increased rotation of low-risk guards. Each panel begins with a zero cumulative share on the period's first day.

We also regress the measures of rotation on the estimated risk of poaching, controlling for time-varying characteristics of the guard and monthly fixed effects. As predicted by Proposition 3, the first two columns of Table 4 show that prior to the policy change, the firm rotated more often guards with a higher risk of being poached. Specifically, a one standard deviation increase in the estimated risk of poaching is associated with 1.5 additional percentage points in the probability of rotation. This is equivalent to 44% of the monthly average rotation rate in the year before the policy change. The correlation between poaching risk and the number of buildings worked per month is also positive and highly significant. Indeed, the coefficients of Columns (1) and (2) are also similar in magnitude because few guards rotated more than once in a month.

5.2 The Effect of the Policy on Rotation

The threat of buildings poaching guards dropped substantially after the introduction of the 1994 Decree. In fact, our data contains no poaching episodes after the policy took effect. The descriptive evidence shown in Figure 2 indicates a potentially disproportionate reduction in rotation for guards with ex-ante high poaching risk after the policy introduction. For example, before the policy change, the monthly average rotation probabilities were 4.5% for high-risk (above median) guards and 2.2%

for low-risk guards (below median). After the policy change, the rotation probability of high-risk guards dropped to 1.3%, but for low-risk guards, it remained at 2.1% (see last rows of Table 4).

Table 4: Correlation between Rotation and Risk of being Poached

	(1)	(2)	(3)	(4)
	Year Before Policy		**Year After Policy**	
Dependent Variable	**Rotated**	**N Builds Worked**	**Rotated**	**N Builds Worked**
Poaching Risk	.015***	.014***	-.0031	-.0041
	(.0027)	(.0044)	(.0026)	(.0025)
	[.0041]	[.0069]	[.0036]	[.0035]
N	3,464	3,464	3,293	3,293
R2	.012	.01	.0056	.0069
Mean Depvar	.034	1.041	.017	1.015
Average Rotation by Risk:				
Low (below median):	.022		.021	
High (above median):	.045		.013	

N guards = 328; N buildings = 109. This table investigates the correlation between the estimated risk of being hired by a building and the rotation of guards. The poaching risk index is standardized to a mean of zero and a SD of one. Columns (1) and (2) use the sample period corresponding to one year before the policy introduction. Columns (3) and (4) repeat the estimation for the sample period corresponding to the year following the policy introduction. The sample only includes guards who joined the firm at least one year before the policy. In Columns (1) and (3), the dependent variable is an indicator of whether the guard was rotated to a new building during the month. In Columns (2) and (4), the dependent variable is the number of buildings in which the guard worked during the month. Each regression controls for the (log) tenure of the guard in the firm and month fixed effects. We exclude guards hired one month before or after the policy change. The poaching risk index is standardized to a mean of zero and a SD of one. Robust standard errors are clustered at the guard level and are indicated in parentheses (with asterisks denoting significance for these s.e.). The square brackets report the standard error of the coefficient obtained by 200 bootstrap repetitions of the whole two-step procedure, where for each bootstrap sample, in the first step we estimate the risk of poaching and in the second step the main regression. The last two rows of the table display the (raw) average rotation of guards in year before/after the policy change grouped by low risk of poaching (guards below the median of the risk distribution) and high risk of poaching (guards above the median of the risk distribution).

In Columns (3) and (4) of Table 4, we repeat the estimation from Columns (1) and (2) for post-policy data. Results indicate that the relationship between rotation and poaching risk, which was large and significant in the year before the law, became small and insignificant in the year after the policy took effect (which can also be interpreted as a placebo test). This result is consistent with the patterns shown in Figure 2 – low-risk guards were rotated more often than high-risk guards in the first months after the policy was introduced.[20] We interpret this sharp change as suggestive evidence that poaching risk stopped being a determinant of rotation after the policy change. This

[20]Excluding these months makes the coefficients of Columns (3) and (4) of Table 4 get closer to zero.

observation also contests the hypothesis that rotation was solely determined by the firm favoring guards with specific attributes (which could make rotation less costly), such as adaptability or client orientation. Even if these traits were confounded with our measure of poaching risk, we would expect their link to rotation to persist after the policy change. The lack of post-policy association reinforces our argument that poaching concerns substantially influenced guard rotation, a dynamic that the 1994 Decree effectively mitigated.[21] Importantly, we do not suggest that rotation should or will be entirely eliminated, as various reasons unrelated to poaching concerns could still justify rotation and remain relevant throughout our sample period.

To account for potential confounders and common shocks, as well as to isolate the causal impact of the policy change, we run the following specification at the guard-month level:

$$Rotation_{it} = \beta RiskPoaching_i \times After_t + \phi X_{it} + \eta_i + \gamma_t + \delta_{bit} + \varepsilon_{it}, \tag{2}$$

where the dependent variable measures the rotation of guard i during month t. The policy effect (β) is identified from the interaction between the estimated risk of poaching and a dummy taking one for post-policy periods.[22] Our estimation includes time-varying characteristics of the guards (X_{it}) like the number of days worked during the month and the tenure within the firm. We absorb any permanent differences in rotation levels across guards by including guard-fixed effects (η_i), and we account for time aggregated variation by including month-fixed effects (γ_t). We also include fixed effects for the building where the guard completed most shifts during the month (δ_{bit}) to control for changes in the rotation due to systematic differences between buildings where the guard works.[23]

[21] The effect was also unlikely to be explained by an overall change in the demand of guards, given its differential impacts. The policy change may have increased competition across service firms. This may increase the outside options for guards, making them more inclined to leave a firm that rotates them frequently and makes them worse off. If this is the primary effect of increased competition, firms will be less likely to use anti-poaching instruments that do not improve the situation of the workers. However, heightened competition among security firms could also raise the cost of replacing a guard after poaching (due to potentially higher search frictions), thereby strengthening the firms' incentives to engage in some anti-poaching behavior. Hence, the overall effect may not be entirely clear.

[22] This specification, resembling those in Bleakley (2010), Duflo (2001), and Card and Krueger (1994), should be interpreted cautiously as discussed in Callaway et al. (2021), particularly regarding the magnitude of coefficients in continuous treatment DiDs. Assuming guards with very low risk (below 25th percentile) as never-treated and those with high risk (above 75th percentile) as fully treated, a binary treatment DiD approach yields comparable, significant policy effects. Similarly, a basic 2×2 DiD with pre- and post-policy periods and extreme risk groups gives similar results as detailed in Appendix Table A.7.

[23] Including dummies for every building where the guard worked during the month (instead of just the one where the guard spent the most time) results in perfect collinearity with our rotation measure. As a robustness check, we re-estimate the main specification at the guard-date level (a guard can work in at most one building per day). Results are very similar to the main specification if we scale up the coefficients to the monthly level (see Appendix Table A.5).

Table 5: Effect of the Policy on Guards's Rotation and Crime

	(1)	(2)	(3)	(4)	(5)	(6)
			Panel A: Rotation			
	Rotated	**N Builds Worked**	**Rotated**	**N Builds Worked**	**Rotated**	**N Builds Worked**
Post × Poaching Risk	-.019***	-.017***	-.023***	-.022***	-.017***	-.016**
	(.0029)	(.0039)	(.0051)	(.0061)	(.0055)	(.0065)
	[.005]	[.009]	[.007]	[.011]	[.007]	[.01]
N	17,119	17,119	17,119	17,119	17,116	17,116
Mean Depvar	.027	1.028	.027	1.028	.027	1.028
			Panel B: Crime			
	N of Crimes	**IHST Value Prop Lost**	**N of Crimes**	**IHST Value Prop Lost**	**N of Crimes**	**IHST Value Prop Lost**
Post × Poaching Risk	-.032*	-.17*	-.031	-.19*	-.049**	-.32***
	(.017)	(.094)	(.019)	(.1)	(.019)	(.1)
	[.025]	[.127]	[.026]	[.146]	[.027]	[.145]
N	17,119	17,119	17,119	17,119	17,116	17,116
Mean Depvar	.25	1.8	.25	1.8	.25	1.8
Indiv Chars:	YES	YES	YES	YES	YES	YES
Month FE:	YES	YES	YES	YES	YES	YES
Guard FE:	YES	YES	YES	YES	YES	YES
Building FE:	YES	YES	YES	YES	YES	YES
Guard Trends:	NO	NO	YES	YES	YES	YES
Guard X Transition:	NO	NO	NO	NO	YES	YES

N guards = 356; N buildings = 116. This table investigates the effects of the introduction of the decree on guard's rotation (Panel A) and crime (Panel B). Each column reports the coefficient of the interaction between an indicator for the period after the law was introduced and the estimated poaching risk of the guard. The poaching risk index is standardized to a mean of zero and a SD of one. In Panel A, the dependent variable in Columns (1), (3) and (5) is an indicator of whether the guard was rotated to a new building during the month and in Columns (2), (4) and (6) the dependent variable is the number of buildings in which the guard worked during the month. In Panel B, the dependent variable in In Columns (1), (3) and (5) is the number of crimes that occurred in the building during the shifts when the guard was working and in Columns (2), (4) and (6), the dependent variable is the (IHST) value of the property lost in the month for the crimes occurred in the building during the shifts when the guard was working. All regressions use observations at the guard-month level, and include fixed effects of guard, month and the building where the guard worked most time during the month. All regressions also control for the log total number of days the guard worked during the month, the log-experience of the guard and an indicator for the first month of the guard in the firm. Columns (3) to (6) include guard-specific linear trends. Columns (5) and (6) control for the interaction between guard fixed effect and an indicator for the two quarters after the law was introduced. We exclude guards hired one month before or after the policy change. Robust standard errors are clustered two-ways at the guard-month level and are shown in parenthesis (with asterisks denoting significance for these s.e.). The square brackets report the standard error of the corresponding coefficient obtained by 200 bootstrap repetitions of the whole two-step procedure (i.e., the estimation of the poaching risk and the main regression).

Panel A of Table 5 reports the estimates of equation (2) (Column 1).[24] The results confirm that guards with a larger risk of poaching were rotated less often after the policy change. A one standard

[24]We report standard errors multi-way clustered at both guard and month levels, as well as two-step bootstrapped standard errors.

deviation increase in poaching risk is associated with a 2-percentage-point reduction in the rotation probability, a very large effect relative to the 2.7% monthly average.

Figure 3 depicts the leads and lags version of equation (2) and displays the estimated coefficients of the variables $RiskPoaching_i \times Q_t^j$ where Q_t^j is an indicator for the j^{th} quarter relative to the policy introduction ($j = 0$). The figure shows no evidence of pre-trends in rotation, but a sharp decrease for high-risk guards. The first months after the policy change display larger negative coefficients. This aligns with the descriptive evidence shown in Figure 2, where low-risk guards were rotated more intensively for a short period after the law was introduced.

In Appendix Section A.4, we discuss various exercises that demonstrate the robustness of the result that reducing the risk of poaching leads to a decrease in rotation. These exercises include accounting for guard-specific trends (Column (2) in Table 5) or excluding the transition period (Column (3) in Table 5). In the next subsection, we show that this lower rotation rate is also associated with a decrease in crime rates and the value of property stolen.

Figure 3: Effects of the Decree 356 on the Rotation of Guards

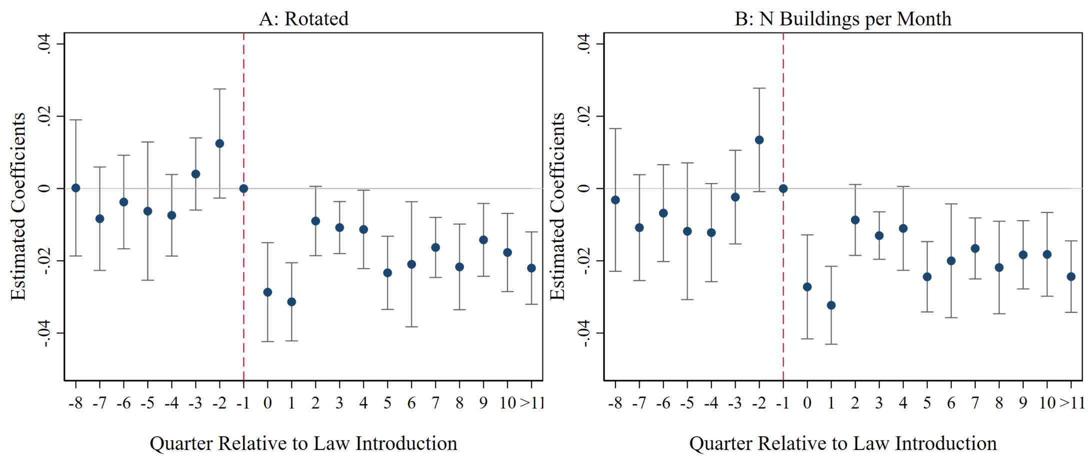

This figure displays the estimated coefficients and the 95% confidence intervals of interaction between a guard's rotation schedule and the estimated risk of being poached by a building, with leads and lags indicators relative to the quarter when the degree was introduced. The omitted category is the interaction with the quarter period previous to the introduction of the law. The dependent variable in Panel A is an indicator for whether the guard was rotated to a new building during the month. In Panel B, the dependent variable is the average number of shifts per building worked by the guard during a given month. All regressions control for guard and month fixed effects. Additional controls include the total number of days that the guard worked during the month, the (log) tenure in the firm, a fixed effect for the building where the guard worked most days in the month and an indicator for the first month the guard worked in the building. Observations are at the guard-month level. We exclude guards hired one month before or after the policy change. Standard errors are multi-way clustered at the guard-month level. $N = 17,119$.

5.3 The Effect of the Policy on Crime

The main insight of the theoretical model is that a firm may deliberately forgo potential productivity gains and *excessively* rotate workers in the presence of poaching risk, which can constrain the surplus generated from the firm-client relationship. In this sense, an important implication of non-poaching policies is that they may increase the productivity of workers by preventing the strategic destruction of client-specific human capital. To explore this implication, we estimate the reduced form effect of the law on crime. We exploit the same specification as in equation (2) but the dependent variables are the number of crimes that occurred while the guard was on duty during the month and the (IHST) value of property lost due to crime. The estimates capture the relative decrease in crime among guards with higher versus lower poaching risk. While one natural interpretation of this pattern is that decreased rotation mediates this effect, we acknowledge that our results might also capture other potential impacts of the policy beyond changes in rotation.

As reported in Panel B of Table 5, the estimated effect of rotation on crime, albeit less robust than the results for rotation,[25] is negative and large relative to the mean number of crimes: an additional standard deviation of the poaching risk is associated with a monthly reduction of the number of crimes in the range of 0.031 to 0.049. This effect is about 13% to 20% of the average number of crimes per month. Similarly, a one standard deviation increase in poaching risk is followed by a reduction in the cost of property lost in the order of 17% to 32%.[26] Appendix Figure A.6 reports the corresponding leads and lags estimates when crime is the dependent variable. The policy effect on crime appears to be stronger over time.[27]

Taken all together, the results of this section provide evidence consistent with (i) a sharp decline in rotation after the policy change due to the lower risk that buildings poach guards, and (ii) a consequent reduction in crime due to guards being rotated less frequently and accumulating more skills specific to the buildings that they serve.

6 Final Discussion

While poaching is recognized as an important issue in the service sector, to the best of our knowledge it has not been quantified precisely, likely due to limited data availability. One potential approach to address this lack of data is to examine indirect measures of the poaching problem, such as how commonly employers take actions to deter their employees from being hired away by clients. A common action in this regard is the implementation of non-solicitation agreements, which are

[25]Estimates are significant at 5% in Columns (5), (6) and significant at 10% in Columns (1), (2) and (4) but marginally non-significant when using two-step bootstrapped standard errors except for last specification.

[26]Reassuringly, the estimated effect relative to the mean are significant and roughly similar (18% to 30% of the mean) when the dependent variable is in levels.

[27]Results for the binary high-risk ($>$ 75th pctile) versus low-risk ($<$ 25th pctile) specification and the simple 2×2 specification (two groups and two periods) are reported in Appendix Table A.8.

contractual clauses that prohibit employees from contacting former clients about providing services. A survey by Balasubramanian et al. (2021) of a large employer sample found that 77% of the firms use non-solicitation agreements, suggesting that the issue of vertical poaching from clients is important and ubiquitous.

In this article, we have made a first step in understanding how service-providing firms respond to the threat of clients poaching their workers, focusing primarily on the firm's strategic decision. To the best of our knowledge, no existing dataset offers the granularity of data required for such a thorough investigation across firms and industries. Thus, following established examples (e.g. Bidwell and Keller, 2014), we focus on a single organization that provided us with granular data on a period where an exogenous shock occurred. Although this approach has inherent limits on generalizability, it provides the required detail for studying the specific mechanisms involved. Using this data, we show that the building-specific experience of a security guard decreases crime even after controlling for the guard's total experience. As the ability to prevent crime is desirable from the buildings' perspective, the risk of a guard being poached is also increasing in that guard's building-specific experience. Anticipating the association between building-specific experience and poaching, the security firm strategically rotates its workers, at a level exceeding the one that it would choose if poaching were prohibited. We also show that a policy change that forbids in-house contracting reduced crime rates, suggesting that prohibiting talent poaching can have a positive effect on welfare.

We have conducted a detailed analysis of our partner firm. However, one may ask about the broader relevance of our research question and the generalizability of the findings. In Appendix Section A.9 we delve deeper into these issues by advancing on three fronts: First, we use survey evidence to argue that our partner firm is representative of the industry. Second, we analyze when poaching is an organizational problem and how rotation interacts with other potential solutions. Third, we provide qualitative and empirical evidence from several different industries.

On the first point, we show that no single relevant attribute of our partner organization makes it unique. We have surveyed more than 20 security firms to confirm that not only our partner organization is representative of a large industry, but also that the mechanisms proposed and studied here are relevant to other organizations in this industry.

On the second point, we argue that poaching poses a major challenge for service firms when their workers have the ability to move to client organizations and when such turnover is very costly for the firm. In general, the impact tends to be more pronounced for the firm when the poached employees are more difficult to replace (e.g., due to their specialized skills and experiences). There are two types of forces that restrict workers from moving to other employers, including clients: demand-side and supply-side factors (Campbell et al., 2012). On the demand side, mobility is limited when the work cultivates a large level of firm-specific skills (as opposed to client-specific skills), when service firms and their clients are asymmetrically informed about the skills of the worker, when there is not enough volume of work to justify bringing the worker in-house (these last two imply that the worker has outside options besides the employing service firm), and when the client's poaching costs are low. On the supply side, mobility is constrained by switching costs or guards underestimating client demand. We believe that the prevalence of poaching in our context can be attributed to the low costs of mobility and poaching, and the ease with which guards learn about clients' demands. In Appendix Section A.9, we also argue that rotation may be a desirable strategy even when firms have access to legal and/or managerial practices to prevent their workers from being poached.

On the final point, and to demonstrate the broader relevance of our analysis beyond the security service context, we conducted interviews with managers from three additional industries. These interviews offer additional qualitative insights that deepened our understanding of the prevalence of vertical poaching across various sectors, as well as the use of rotation as a preferred deterrent strategy by service firms. Overall, the evidence indicates that vertical poaching is a common and important issue, and rotation is used with some regularity as an anti-poaching tactic, though its implementation varies based on specific market conditions and legal environments. Lastly, we provide empirical evidence supporting the relevance of vertical poaching in the lobbying industry.

With the preceding evidence, we have contended that the phenomenon of poaching is both relevant and widespread. However, there are scenarios in which service-providing firms might hold a more positive view about their employees being poached by clients, especially if these workers can ensure a future stream of transactions with their original employers (Somaya et al., 2008). Our setting is not appropriate to analyze such empirical situations, primarily because in our case the client obtains the necessary service either fully in-house or fully outsourced. We expect that the benefits of poaching are more significant in settings with different characteristics, for instance, those in which the client would require a fraction of the labor force in-house and acquire the remaining labor input through outsourcing. Exploring these alternate settings is outside the scope of this paper, but future work in this direction is warranted.

We conclude the paper by reiterating that, despite the comprehensiveness of our empirical analysis, it is not without limitations. First, we were not able to observe all dimensions of guard performance, such as the time taken to register guests or the ability to recognize frequent visitors. Second, our study focused on a single decision-maker – our partner firm, which may limit the generalizability of our findings. Although we have endeavored to broaden our scope through industry-wide surveys and cross-sector interviews, the lack of micro-data from multiple firms and sectors remains a constraint. Third, while our findings reveal a strong and robust link between building-specific experience and productivity, the relationship is bounded by the non-exogenous nature of rotations. Fourth, the analysis of the non-poaching policy's effects is hampered by the absence of a natural control group. To deal with this limitation, we constructed a poaching risk index for each guard using machine learning techniques and leveraged the variation in this measure to assess the effects of the policy. We hypothesized that the policy primarily influenced rotation through the reduction in poaching risk, and substantiated this with both indirect empirical evidence (such as the lack of pre-trends) and anecdotal or survey evidence from firms in our focused industry. Finally, this paper focused on the worker-level effects of poaching. A promising alternative, left for future work, is studying the broader organizational effects of poaching. The extensive literature on the work-unit consequences of staffing events (as evidenced by notable studies such as, Nyberg and Ployhart (2013); Reilly et al. (2014); De Stefano et al. (2019); Sajjadiani et al. (2023)) underscores the potential significant ripple effects these events have, impacting not only individual workers but also work-units. However, our setting presents certain constraints to study such effects. For example, poaching events cause the removal of entire units (buildings) as well as there is limited interaction among some guards due to non-overlapping shifts.

References

Acemoglu, D. (1997). Training and Innovation in An Imperfect Labour Market. *Review of Economic Studies*, 64(3):445–464.

Balasubramanian, N., Starr, E., and Yamaguchi, S. (2021). Bundling Employment Restrictions and Value Capture from Employees. Working paper.

Bazen, S. (2011). Duration Models. In *Econometric Methods for Labour Economics*, chapter 5, pages 97–121. Oxford University Press.

Bennet, J. (2018). 'Money and Greed': How Non-compete Clauses Force Workers to Fight for Rights. *The Guardian*. Available at https://www.theguardian.com/us-news/2018/oct/24/non-compete-clause-low-wage-workers-lawsuits-rights.

Bhuller, M., Dahl, G. B., Løken, K. V., and Mogstad, M. (2020). Incarceration, Recidivism, and Employment. *Journal of Political Economy*, 128(4):1269–1324.

Bidwell, M. and Keller, J. (2014). Within or Without? How Firms Combine Internal and External Labor Markets to Fill Jobs. *The Academy of Management Journal*, 57(4):1035–1055.

Bidwell, M., Won, S., Barbulescu, R., and Mollick, E. (2015). I Used to Work at Goldman Sachs! How Firms Benefit from Organizational Status in the Market for Human Capital. *Strategic Management Journal*, 36(8):1164–1173.

Bleakley, H. (2010). Malaria Eradication in the Americas: A Retrospective Analysis of Childhood Exposure. *American Economic Journal: Applied Economics*, 2(2):1–45.

Bradshaw, T. (2015). Batteries Suit Offers Clues to Apple's Car Ambitions. *Financial Times*. Available at https://www.ft.com/content/b061736c-b874-11e4-a2fb-00144feab7de.

Bradshaw, T. (2017). The Blessing and Curse of Being an Apple Supplier. *Financial Times*. Available at https://www.ft.com/content/3d49b76a-1b76-11e7-a266-12672483791a.

Callaway, B., Goodman-Bacon, A., and Sant'Anna, P. H. (2021). Difference-in-Differences With a Continuous Treatment. Working paper.

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

Card, D. and Krueger, A. (1994). Minimum wages and employment: A case study of the fast-food industry in new jersey and pennsylvania. *American Economic Review*, 84(4):772–93.

Chaput, M. (2018). Small Business and Employee Poaching. *The MileIQ Blog*. Available at https://www.mileiq.com/en-ca/blog/small-business-employee-poaching-canada/.

Choi, J. P. and Thum, M. (2003). The Dynamics of Corruption with the Ratchet Effect. *Journal of Public Economics*, 87(3-4):427–443.

De Stefano, F., Bonet, R., and Camuffo, A. (2019). Does losing temporary workers matter? the effects of planned turnover on replacements and unit performance. *Academy of Management Journal*, 62(4):979–1002.

DLA Labor Dish Editorial Board (2014). "No Hire" Agreements: Unlawful Collusion or Necessary Self-Protection? *The Labor Dish*. Available at https://www.labordish.com/2014/09/no-hire-agreements-unlawful-collusion-or-necessary-self-protection/.

Dorn, D., Schmieder, J. F., and Spletzer, J. R. (2018). Domestic Outsourcing in the United States. U.S. Department of Labor Technical Report.

Dougherty, C. (2017). How Noncompete Clauses Keep Workers Locked In. *The New York Times*. Available at https://www.nytimes.com/2017/05/13/business/noncompete-clauses.html.

Duflo, E. (2001). Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment. *The American Economic Review*, 91(4):795–813.

Fisher, L. D. and Lin, D. Y. (1999). Time-dependent covariates in the cox proportional-hazards regression model. *Annual Review of Public Health*, 20(1):145–157. PMID: 10352854.

Goldschmidt, D. and Schmieder, J. F. (2017). The Rise of Domestic Outsourcing and the Evolution of the German Wage Structure. *The Quarterly Journal of Economics*, 132(3):1165–1217.

Groysberg, B., Lee, L.-E., and Nanda, A. (2008). Can They Take It with Them? The Portability of Star Knowledge Workers' Performance. *Management Science*, 54(7):1213–1230.

Hertzberg, A., Liberti, J. M., and Paravisini, D. (2010). Information and incentives inside the firm: Evidence from loan officer rotation. *Journal of Finance*, 65(3):795–828.

Huckman, R. S. and Pisano, G. P. (2006). The Firm Specificity of Individual Performance: Evidence from Cardiac Surgery. *Management Science*, 52(4):473–488.

Ickes, B. W. and Samuelson, L. (1987). Job Transfers and Incentives in Complex Organizations: Thwarting the Ratchet Effect. *The RAND Journal of Economics*, 18(2):275–286.

Liffreing, I. (2018). 'Please don't Steal our People': Confessions of an Agency Exec on Clients Poaching Agency Talent. *Digiday*. Available at https://digiday.com/marketing/please-dont-steal-people-confessions-agency-exec-clients-poaching-agency-talent/.

Lipsitz, M. and Starr, E. P. (2022). Low-wage Workers and the Enforceability of Non-Compete Agreements. *Management Science*, 68(1):143–170.

Makadok, R. and Ross, D. G. (2013). Taking Industry Structuring Seriously: A Strategic Perspective on Product Differentiation. *Strategic Management Journal*, 34(5):509–532.

Marx, M., Strumsky, D., and Fleming, L. (2009). Mobility, Skills, and the Michigan Non-Compete Experiment. *Management Science*, 55(6):875–889.

Meyer, M. A. (1994). The Dynamics of Learning with Team Production: Implications for Task Assignment. *The Quarterly Journal of Economics*, 109(4):1157–1184.

Moen, E. R. and Rosén, Å. (2004). Does Poaching Distort Training? *Review of Economic Studies*, 71(4):1143–1162.

Nyberg, A. J. and Ployhart, R. E. (2013). Context-emergent turnover (cet) theory: A theory of collective turnover. *Academy of Management Review*, 38(1):109–131.

Ortega, J. (2001). Job Rotation as a Learning Mechanism. *Management Science*, 47(10):1361–1370.

Pestronk, M. (2019). Agency's Travel Advisors Poached by Client. *Travel Weekly*. Available at https://www.travelweekly.com/Mark-Pestronk/Agencys-travel-advisors-poached-by-client.

Reilly, G., Nyberg, A. J., Maltarich, M., and Weller, I. (2014). Human capital flows: Using context-emergent turnover (cet) theory to explore the process by which turnover, hiring, and job demands affect patient satisfaction. *Academy of Management Journal*, 57(3):766–790.

Rose-Ackerman, S. (2010). The Law and Economics of Bribery and Extortion. *Annual Review of Law and Social Science*, 6(1):217–238.

Sajjadiani, S., Kammeyer-Mueller, J., and Benson, A. (2023). Who is leaving and why? the dynamics of high-quality human capital outflows. *Academy of Management Journal*, 66(6):1929–1953.

Schreier, J. (2020). Game Publisher Cancels Contract With Developer, Then Tries to Poach Its Entire Team. *Bloomberg*. Available at https://www.bloomberg.com/news/articles/2020-06-03/kerbal-space-program-2-release-disrupted-by-corporate-strife.

Shubber, K. (2018). Cushman v the Cleaner: the Fight over Non-competes. *Financial Times*. Available at https://www.ft.com/content/bfb69d30-ce44-11e8-b276-b9069bde0956.

Somaya, D., Williamson, I. O., and Lorinkova, N. (2008). Gone but Not Lost: The Different Performance Impacts of Employee Mobility between Cooperators versus Competitors. *The Academy of Management Journal*, 51(5):936–953.

Starr, E., Prescott, J. J., and Bishara, N. (2020). The Behavioral Effects of (Unenforceable) Contracts. *Journal of Law, Economics, and Organization*, 36(3):633–687.

Starr, E., Prescott, J. J., and Bishara, N. (2021). Noncompete Agreements in the US Labor Force. *Journal of Law and Economics*, 64(1):53–84.

StevensVuaran Lawyers (2019). Can a Contract Prevent Your Customers From Poaching Your Employees? Blog Post. Available at http://stevensvuaranlawyers.com/customers-poaching-your-employees/.

Van Den Berg, G. J. (2001). Duration Models: Specification, Identification and Multiple Durations. In Heckman, J. J. and Leamer, E., editors, *Handbook of Econometrics*, volume 5, chapter 55, pages 3381–3460. Elsevier.

Wright, P. M., McMahan, G. C., and McWilliams, A. (1994). Human Resources and Sustained Competitive Advantage: A Resource-Based Perspective. *International Journal of Human Resource Management*, 5(2):301–326.