Bad apples in a barrel. Field evidence about corruption in organizations

Romain Ferrali*

September 28, 2017

Abstract

Corruption occurs in organizations: bureaucratic agencies, legislatures, and firms can all fall prey to corruption. Yet, we have little understanding of how the interpersonal relationships that take place within these organizations affect opportunities for corruption, for this requires measuring such relationships, and fraudulent behavior of individual members. This paper uses a fine-grained, original dataset on the daily operations of a large company that allows measuring, over time, dishonest behavior at the individual level and interactions among employees. I identify agents that are likely to be dishonest. I use a series of natural experiments, as well as a structural model, to show that the organizational structure of the firm affects the emergence, extent, and characteristics of corruption. In particular, increased monitoring by other agents deters corrupt individuals from engaging in dishonest behavior. The results suggest it would be sensible to redesign government agencies to reduce isolation among bureaucrats.

*Department of Politics, Princeton University, ferrali@princeton.edu. I thank Matias Iaryczower, Kosuke Imai, Jenn Larson, John Londregan, as well as seminar participants at the Princeton Political Economy seminar and the Imai Research Group for their helpful comments. I gratefully acknowledge support from the Mamdouha S. Bobst Center for Peace and Justice.

Corruption, "the abuse of entrusted power for private gain" (Transparency International) occurs within organizations: corrupt bureaucrats and politicians are all affiliated to government agencies, or other political organizations. Although we know that organizations define interpersonal relationships that affect outcomes (Chandler, 1977), we have little understanding of how such interpersonal dynamics affect corruption: which members are, by their position, more susceptible to corruption? Do honest members deter the corrupt behavior of these "bad apples," or do they succumb to their influence? We have little to say about these simple questions, for answering them requires measuring both individual-level corruption and social interactions within an organization. Because obtaining such data is difficult, available evidence on corrupt behavior in organizations often comes from the lab (Gino and Pierce, 2009; Gino, Ayal and Ariely, 2009), which comes with severe restrictions in terms of external and ecological validity. Available field studies typically focus on an few known corrupt agents but lacks systematic evidence on the rest of their peers (Aven, 2015; Gambetta, 1996; Vannucci and Della Porta, 2013), or only considers a fraction of the organization, such as top management (eg. Khanna, Kim and Lu, 2015).

Lacking solid micro-foundations for how corrupt and honest agents interact in organizations, we miss two major aspects of corruption. First, we cannot readily explain why corruption sometimes involves isolated individuals, and sometimes involves conspiracies of many accomplices. Second, we cannot explain how organizational structure affects corruption. As such, we cannot explain why corruption varies within a bureaucracy, where institutional features are held constant, but the organizational layout varies widely within and across agencies. In particular, we have little to say about policies that are frequently observed in practice, where bureaucracies combat corruption through major reforms of their organizational chart (eg. Bennet, 2012; Friedman, 2012; Hausman, 2011).

This paper answers the following question: do interactions between honest and dishonest colleagues spread, or deter corruption? To do so, it exploits a unique dataset on the daily operations of a call-center company in which clerks allocate incoming claims to a network of service providers. The data records the interactions of clerks with the internal company software used to allocate claims to service providers. This allows the detection of suspicious behavior during claim allocation, revealing suspicious deviations from company policy and potential preferential allocation of claims to specific service providers. The data also measures rich networks of interactions among clerks, enabling an investigation of how interpersonal relationships affect corrupt behavior.

The approach uses a large private organization as a proxy for a government agency. Analogies are a common approach in studies of the bureaucracy. The various organizational units that make up the bureaucracy are very heterogeneous, which makes comparing across them difficult. As a result, since the proverbial forest ranger (Kaufman, 1960), researchers often rely on single-agency case studies to extend conclusions to an entire bureaucracy (eg. Carpenter, 2001; Gordon, 2011; Haeder and Yackee, 2015). By considering a single organization, this paper places itself within this tradition. However, because obtaining such data for a bureaucracy proved unfeasible, I consider instead a private organization whose features readily map to bureaucracies. Broadly speaking, this organization is similar to a bureaucracy in that it is a large organization of relatively educated, white-collar workers. Equally as important, states also operate call-centers that can fall prey to corruption. Police dispatch operators have been caught dissimulating critical information, leaking it to accomplices, or dispatching fake radio calls to cover up for illegal operations.

Of course, it is difficult to simply observe an organization and find corruption. The setting combines features that allow detecting statistically likely instances of dishonest behavior among clerks. This is because clerks perform a standardized task where fraud is well-defined: instead of allocating claims to the best service provider, clerks may allocate them to their chosen provider, in return for a kickback. The company software governs the allocation of claims by randomly drawing the provider to which it should be allocated, according to a weighting defined by management. Clerks check whether the provider is available. If the provider is not available, they may *skip*; that is, ask for another draw. Clerks may also *force*; that is, they may opt out of random drawing, and manually select a provider. For management, verifying whether past skipping and forcing behavior was legitimate is costly, requiring them to manually match phone records to claims and listening to the recorded call. This gives an opportunity for dishonest clerks to use these tools wrongfully: they may pretend that the randomly drawn provider was unavailable and allocate the claim a favored provider instead. As such, dishonest clerks should have disproportionately high skipping and forcing rates.

I exploit these features to show that interpersonal interactions deter corrupt behavior. I report three core findings.

The first finding shows statistical evidence for dishonest behavior among clerks. Because the task is standardized, if all clerks were honest, they should behave similarly. This allows pinning down the distribution of skipping rates if all clerks were honest. I show that the claims of a small amount of clerks are awarded after significantly more draws that what would be expected under honest behavior. This suggests that these clerks behaved dishonestly, by deliberately skipping selected service providers.

The other findings establish that interpersonal interactions deter corrupt behavior. Using a series of natural experiments, I show that plausibly exogenous variation in the number of colleagues monitoring one's task caused some clerks to reduce their skipping and forcing rates. I exploit two plausibly exogenous events: first, new hires, whose type is yet unknown to their more senior colleagues, and may trigger additional caution; second, temporary allocations of clerks to markets they seldom operate in, due to unforeseen peaks in demand. I show that both events have heterogeneous effects: while they have no effect for most clerk-provider dyads, they lead to decreases in skipping and forcing rates for some dyads.

Finally, I use a structural model to unpack the patterns that natural experiments identified in a reduced form. The natural experiments label as dishonest the clerks who reduce their suspicious behavior when surrounded by other clerks. Doing so, they may conflate identifying dishonest types with the behavioral hypotheses attached to such types. This approach may be spurious if our assumptions on the behavior of honest and dishonest types are wrong. The structural model separates the two components of the problem, by estimating the latent probability that a provider is unavailable, and that a claim has been unduly skipped. This allows assigning honest and dishonest types to clerk-provider dyads, and identifying the behavior of dishonest types. Results confirm and sharpen the causal evidence: I find that dishonest clerks are more likely to skip unduly for high-value claims, and when subjected to less monitoring from their peers.

By providing field evidence that interpersonal interactions in organizations may deter corrupt behavior, this paper makes two contributions. First, it gives behavioral foundations to a literature in political science that has examined the impact of organizational structure on corruption from a macro perspective (Evans, 1995; Rauch and Evans, 2000; Charron et al., 2017; Carpenter and Moss, 2013). Although this literature points at features of the organization–such as meritocracy–that should reduce corruption, it is unable to make specific claims as to which organizational structures should reduce corruption because it lacks micro-foundations. Second, this paper contributes to a literature that studies the behavioral foundations of corruption in organizations (Gambetta, 1996; Vannucci and Della Porta, 2013; Gino and Pierce, 2009; Aven, 2015). It considers a setting that allows examining relationships between honest and dishonest agents in a setting that combines the ecological validity of a field study with the advantages in measurement and causal inference that are usually afforded by the lab.

The remainder of this paper proceeds as follows: I first review the literature on corruption in organizations (section 1), then provide additional details about the context and the data (section 2). I show descriptive evidence of fraudulent behavior (section 3), causal evidence that additional monitoring by peers deters corrupt behavior (section 4), as well as comparable structural evidence (section 5). I conclude by outlining avenues for future research (section 6).

1 Literature review

Tackling corruption is a major issue in developing and developed countries alike: bribery alone is estimated to cost about 2 percent of global GDP (International Monetary Fund 2016). Corruption hinders economic growth (Mauro, 1995) and undermines the legitimacy of state institutions.

The bulk of scholarly research on how to curtail corruption has focused on norms and institutions. As a result, we have a good understanding of how various policy instruments affect the extent of corruption. Two such instruments have generated most attention: wages, and various instruments of monitoring.

Corruption, however, occurs in organizations, and we have little understanding of how organizations affect corruption. Corrupt bureaucrats and politicians are *embedded* in organizations: bureaucratic agencies, and political bodies (Granovetter, 1985; Zukin and DiMaggio, 1990). Despite an early interest in how the organizational structure of the bureaucracy affects political outcomes (Weber, 1948; Crozier, 2009), we are unable to answer simple questions about the relationship between organizational structure and corruption: which bureaucrats are, by their position, most likely to engage in corruption? When do corrupt bureaucrats "infect" their honest colleagues, and when do they get deterred by them? Which organizational structures best mitigate corruption?

In political science, extant work on organizations and corruption has generated valuable insights as to the macro-implications of organizational structure on corruption. This line of work considers how some macro-structural features of an agency affect the extent of corruption.

Evans (1995) considered ties between the agency and the rest of society. Rauch and Evans (2000) and Charron et al. (2017) consider the impact of an array of organizational features that define the "Weberian-ness" of a bureaucracy, such as meritocratic recruitment procedures. Carpenter and Moss (2013) examine mechanisms that may lead to agency capture, highlighting the role of rent-seeking by politicians, industry threats of political or legal retaliation, and socio-cultural capture, stemming from close, repeated interaction.

Yet, this line of research is unable to make *specific* predictions as to which organizational structures best mitigate corruption, and how corruption spreads from agent to agent, because these macro-approaches lack strong behavioral foundations. In most of these accounts, it is unclear (1) how corrupt agents interact with each other, and (2) how they interact with honest agents. As a result, it is unclear which specific organizational structures underpin the macro-structural features highlighted by this literature.

Investigating the behavioral foundations of corrupt behavior in organization faces formidable measurement problems. The task faces two challenges. First, it requires data that are granular enough to allow separating honest and dishonest behavior, and keeping track of interactions among all agents. Second, as is the case with many network studies, the set of neighbors that may influence an individual's actions is hard to define (Pietryka and DeBats, 2017).

Existing evidence is partial or has relatively low ecological validity, because it addresses these challenges using either partial field studies, or lab experiments. Field studies typically consider past cases of fraud (Gambetta, 1996; Vannucci and Della Porta, 2013; Aven, 2015). This approach has detailed information on the "bad apples," but much less information on the good ones. It makes nuanced claims on how corrupt agents interact with each other, highlighting the importance of trust in an environment where contracts cannot be enforced by a neutral third-party. It has, however, much less to say about relationships between good and bad apples. Other field studies consider a large number of organizations, but have little information on relationships inside each of them, which severely restricts the range of inferences they can make. For instance, Khanna, Kim and Lu (2015) considers the top executives of about 2,700 US firms, and construct interactions based off whether an executive was appointed during the CEO's tenure. They show that better connected boards are more likely to commit financial fraud. Finally, lab experiments solve the measurement problems by trading off ecological validity for the control afforded by the lab (Gino and Pierce, 2009; Gino, Ayal and Ariely, 2009; Gino and Galinsky, 2012; Pitesa and Thau, 2013). Although the approach allows measuring precisely the concepts of interest and making solid causal claims, it is unclear how much the social relationships and dishonest behavior that are engineered in contrived lab settings reflect the durable, multi-layered relationships, and the variety of dishonest behaviors that may occur in real organizations.

I address these measurement problems by leveraging high-resolution data that measure interpersonal interactions very precisely, in a setting that combines work processes that are structured enough to allow for statistical detection of dishonest behavior, and contains enough as-if-random variation to make causal claims. This allows examining the behavioral foundations in a setting that combines the ecological validity of field studies with the internal validity of experimental work.

	Variable	Value
Dependent variable	Forcing rate	0.23
	Mean chain length	1.44
Claim characteristics	Number of claims	69,445
	Revenue	$2.90m
	Daily number of claims	534
	Daily revenue	$22,333
	Beginning of period	11/1/2016
	End of period	3/10/2017
	Number of markets	238
	Number of service providers	325
	Number of clerks	177
Clerk characteristics	Monthly wage	$424
	Percent females	0.57
	Age	28.29
	Turnover rate	0.70

Table 1: **Sample descriptive statistics**. Mean chain length reports the mean chain length of claims that were not forced. Salary represents about 1.7 times minimum wage. Turnover rate is computed for year 2016.

2 Context and data

The data describes the daily operations of a call-center company based in Casablanca, Morocco, between Novembers 1st, 2016 and March 10th, 2017. This call-center operates a network of 325 partnering service providers. When calling, customers file claims to the company, who dispatches a service provider to satisfy the customer's demand. These claims are typically urgent, and need to be serviced within the shortest amount of time–5 minutes on average for urgent services, and 30 minutes on average for non-urgent services. The company compensates providers for their service, with the subscriptions of its customers. In our sample, average monthly compensation is $2,337.

The data is a backup of the database of the company's internal software. The database contains all interactions of clerks with the company software, as well as characteristics of the 69,445 claims awarded during the period. Table 1 shows descriptive statistics.

This context has the advantage that corruption is well-defined. Clerks can behave dishonestly by either creating fictitious claims, or misallocating them. Creating fictitious claims is virtually impossible, for management monitors all claims that involve financial transactions. As such, the only kind of corrupt behavior is the preferential allocation of claims to service providers, in return for a kickback. Some of such schemes have been detected by company management in the past, with the last prominent incident dating back to 2012: a couple was colluding with several providers, and got caught after they approached an honest provider, who signaled the incident to management.[1]

The incident highlights a weakness in managerial processes: manually detecting instances of collusion between a clerk and a service provider is a daunting task. The call-center processes

[1] Evidence collected from interviews with management. Interview transcripts are available upon request to the author.

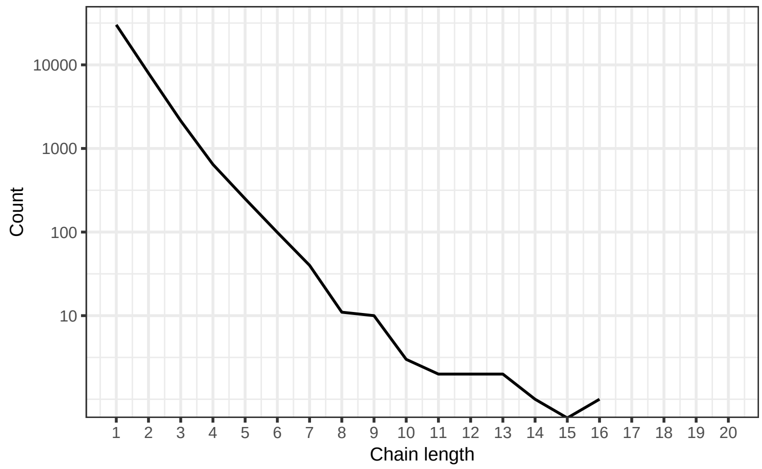

Figure 1: **Observed distribution of chain lengths under the rotation.** The distribution is skewed to the left: 92 percent of claims are allocated after at most two draws.

an average of 534 claims daily; monitoring their allocation would require listening to hours of recordings of phone conversations. Such lack of oversight leaves room for fraudulent behavior.

To address the problem, the company rolled out in November 2016 an update to its internal software, to limit clerks' discretion when selecting service providers. Under the new system, called the *rotation*, clerks request services in a given market–defined as a type of services in a city. The software randomly draws a provider within the market, using weights set by management. Clerks must call the selected provider and check whether it is available. If it is, then they should allocate the claim to that provider. If not, they must input a reason, and may *skip*–ask for another draw. The process ends when they reach an available provider. Reasons for unavailability include inability to reach the provider, the provider already operating at full capacity, or incapacity to service particular claims for technical reasons, such as lack of adequate equipment. Through this process, it takes an average of 1.4 draws to award a claim (Figure 1). Clerks may also *force* out of the rotation, i.e. terminate the random draw, and select a provider manually, which occurs in 23 percent claims. Clerks should force in three instances: when instructed to do so by management, for recurring claim where one customer always deals with the same provider, and when a customer jointly files two claims that could be serviced by the same provider. In this case, the provider should be selected through the rotation for the first claim, and is forced for the second claim.

Although the rotation does not completely eliminate discretion during provider selection, it makes collusion bear clear empirical implications. Because it is costly, management seldom verifies the reasons clerks invoke for skipping or forcing. Suppose that clerk i colludes with provider j. Then i may pretend that providers other than j were unavailable, and skip to j. Although it is more risky, i may also pretend that she was following an instruction from management, and force to j. So if i colludes with j, then i should skip non-j providers more than honest clerks, and/or i's rate of services forced to j should be higher than that of honest clerks.

The call-center operates 24-7 and employs about 177 clerks. As is typical of call-centers in developing countries, clerks are relatively young (28 years old on average), and well-educated,

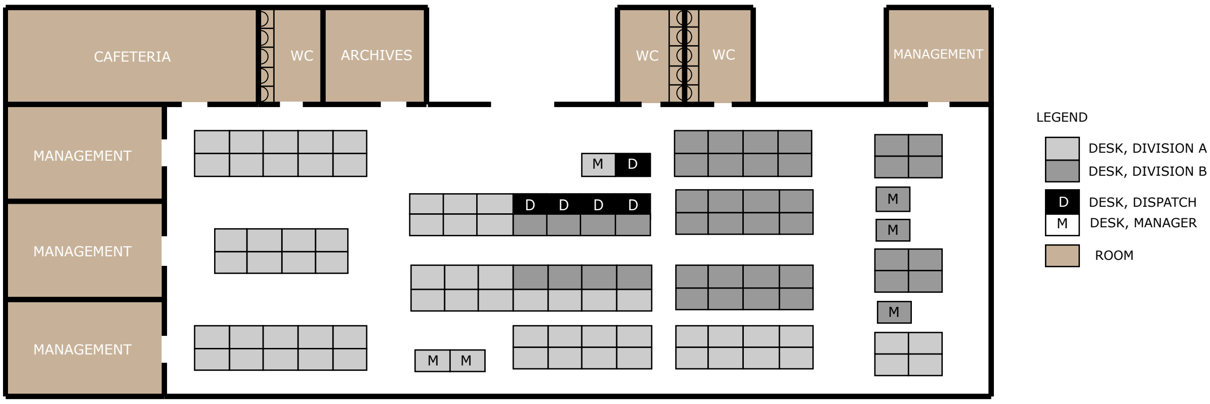

Figure 2: **Schematic floor plan of the call-center.** Desks sit one person. The space has an area of 4,725 squared feet. Clerks operate in a single room. They are separated in two functionally, and spatially distinct divisions.

with virtually all clerks having completed some tertiary education, and there is high turnover (70.3 percent in 2016). Clerks are paid an average of $424, which represents 1.7 times minimum wage.

Combined with features of the company, the internal software allows reconstructing interactions between clerks. Clerks all work in one room, on a single floor of 4,725 squared feet. (Figure 2), and can only access the company software from that room. Clerks log into the software as they start their shift, and close their session when they leave, for this information is used by management to monitor their presence, and pay them accordingly. As such, although the company software does not record the physical location of employees, interactions with the software allow deriving which clerks are in the room at any given time.

Since the goal is to assess how the presence of colleagues affects incidences of fraudulent behavior, I use interactions with the software to construct an indicator of *attendance*: the number of clerks that are present in the room at some point in time. Other features of the company allow making finer distinctions: clerks are separated in two divisions, that operate in two non-overlapping industries, and are also spatially distinct. While clerks rarely collaborate across divisions, they regularly do so within division, where clerks often grant claims within the same market. As such, I consider two finer measures of attendance, that capture closer professional relationships. While (total) attendance tallied the number of clerks on the floor, within-division attendance counts, for clerk i, the number of clerks currently operating in her division, and within-market attendance counts the number of clerks that are currently operating in her market.

Finally, I construct an indicator of tie strength (Granovetter, 1973). The software allows reconstructing the amount of time two clerks spent together in a given time-period. I use a cutoff to separate weak ties from stronger ones, and define as a weak tie two clerks that have spent less that some time operating on the same markets in the past 30 days. Because most clerks work series of 8-hour shifts, I use alternatively a cutoff of one or two shifts. While Granovetter's notion of tie strength referred to an emotional bond, separating "friends" from "acquaintances," I take a broader view of the concept, separating familiar colleagues from unfamiliar ones.

3 Descriptive evidence

This section devises a novel, simple test to show statistical evidence of fraud in the allocation of claims. The test considers only claims allocated through the rotation, and examines the distribution of chain lengths, ie the number of draws before awarding a claim. Assuming that all clerks are honest, I pin down the theoretical distribution of chain lengths, and compare it to the observed distribution. I show that overall, chains are longer than expected under the null distribution, which suggests that some dyads did behave dishonestly. I also identify the dyads that are most unlikely to conform to the null distribution, and hence more likely to be dishonest.

3.1 Theoretical distribution of chain lengths

A market is a pool of N companies. At each period, company i is drawn without replacement from the pool, with initial weight p_i, such that $\sum_{i\in N} p_i = 1$. Each company has a probability of being available $q_i \in (0,1)$. The process stops at the first occurrence of an available company. When the pool is empty, all companies are put back in the pool.

The distribution of L, the length of a chain, has discrete support from 1 to infinity, with parameters $p = (p_i)_{i=1}^N$, and $q = (q_i)_{i=1}^N$. I first pin down the distribution of L before the pool is empty; that is, for $L \leq |N|$.

This distribution is defined by recursion. Let $V_t \subseteq N$ be the (random) subset of companies that are in the pool at draw t. Suppose that the pool is full; that is, $V_t = N$. The probability that the process stops at the next draw is simply the probability of picking an available company: $\Pr(l = 1|V_t = N) = \sum_{i\in N} p_i q_i$. Deriving the probability that the process stops at the next draw given any pool amounts to rescaling the weights to the companies that are left in the pool: $\Pr(l = 1|V_t) = \frac{\sum_{i\in V_t} p_i q_i}{\sum_{i\in V_t} p_i}$. The probability that the process ends in more than one draw is defined recursively. Conditional on company i being drawn and unavailable at draw t, the pool contains $V_{t+1} = V_t \setminus \{i\}$ at draw $t+1$, and the probability that the process ends at draw $t+1$, $\Pr(l = 1|V_{t+1})$ is given by our previous result. Company i is unavailable given pool V_t with probability $\frac{p_i(1-q_i)}{\sum_{j\in V_t} p_j}$. As such, the probability of the process ending after two draws is $\Pr(l = 2|V_t) = \frac{\sum_{i\in V_t} p_i(1-q_i)\Pr(l=1|V_t\setminus\{i\})}{\sum_{i\in V} p_i}$. More generally, the probability of observing exactly $l \leq |V_t|$ additional draws is:

$$\Pr(l|V_t) = \begin{cases} \frac{\sum_{i\in V_t} p_i q_i}{\sum_{i\in V_t} p_i} & \text{if } l = 1 \\ \frac{\sum_{i\in V_t} p_i(1-q_i)\Pr(l-1|V_t\setminus\{i\})}{\sum_{i\in V_t} p_i} & \text{otherwise} \end{cases}$$

For $L > |N|$, it suffices to decompose the event in the number of times the pool has been refilled, and the number of additional draws after the last refill. The pool is refilled after $|N|$ successive failures. This event, R, occurs with probability $\Pr(R) = \prod_{i\in N}(1 - q_i)$. Let $L = m|N| + l$ be the Euclidean division of L by $|N|$; m is the number of times the pool needs be refilled, and l the number of additional draws to realize a chain of length L. Then, the distribution of L is:

$$\Pr(L) = \Pr(R)^m \Pr(l|N).$$

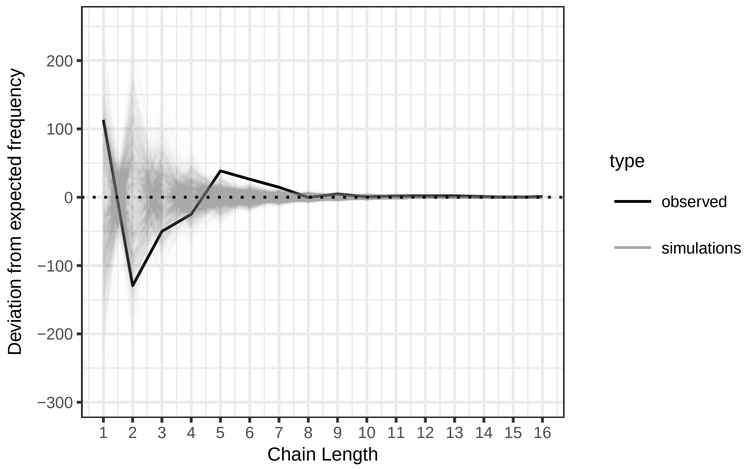

Figure 3: **Distribution of chain lengths as deviation from expected frequencies.** The thick black line represents observed frequencies of chain lengths as deviation from the expected frequencies. Each line of the grey shaded area represents one of 1000 simulation. Chains are longer than expected under honest behavior: there are significantly more chains of length 5-7 and significantly less chains of length 2-3.

While the weights p are directly available, the parameters q for a company being available need to be estimated. Assuming that all clerks are honest, we can easily estimate them from the data, since all skips then correspond to instances where the company was unavailable. I estimate q_i with $\hat{q}_i$, the mean skip rate on draws involving company i.

Estimating the unavailability rate q_i from the data makes the approach conservative. If some clerks are dishonest, then they have higher skip rates than their honest colleagues, which overestimates the true unavailability rate. As a result, the estimated distribution of chain lengths biased towards longer chains. Since only deviations towards *longer* chains are evidence of suspicious behavior, bias goes against the expected effect, making the approach more conservative.

3.2 Results

Using the procedure described above, I derive the null distribution of chain lengths for all markets. Suppose n claims were awarded in a market. In each market, I then derive the expected frequencies of chain lengths for n claims. To get a sense of uncertainty, I also simulate 1000 times n draws from the null distribution.

Aggregating all markets, Figure 3 shows that observed chains are longer than expected under the null, suggesting that some clerks did behave dishonestly by pretending that providers were unavailable. Chains of length 2 to 3 are less frequent than expected, while chains of length 5 to 7 are significantly more frequent. A Chi-squared test confirms the pattern, rejecting the that observed frequencies are drawn from the null distribution.

I then preform a comparable test within markets, to identify the clerks whose distribution of skips differs from the null. Suppose that clerk i awarded n claims in some market, with a mean chain length of $\bar{l}_i$. I examine the probability of obtaining a mean chain length $\bar{l} > \bar{l}_i$ after n draws from the null, to derive the following one-tailed p-value: $\Pr(\bar{l} > \bar{l}_i)$. I estimate this p-value by simulating 1000 times n draws from the null distribution. I obtain 582 such p-values, one

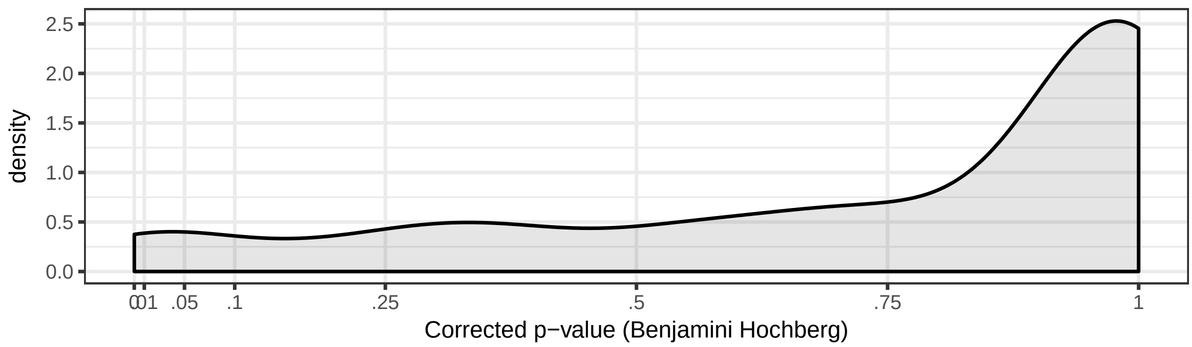

Figure 4: **Deviation from the null distribution at the individual level.** Distribution of Benjamini-Hocheberg corrected p-values for the probability of obtaining mean chain length larger than observed in the data under the null distribution for a given clerk-market. For a few clerk-markets, it is very unlikely that their mean chain length comes from the null distribution.

per clerk-market, that I correct for multiple testing using the Benjamini-Hocheberg procedure.

Figure 4 shows the distribution of such p-values for all clerk-markets. Most clerk-markets seem to conform to the null distribution: 92 percent of them have a p-value above 10 percent. However, 4 percent have a p-value below 1 percent. This result tells us that while the overwhelming majority of clerks seem to behave honestly, with chain lengths in line with what is to be expected if all clerks were honest, a small minority seem to behave dishonestly, with longer chains than expected.

Overall, the evidence in this section suggests that the null hypothesis of honest behavior fails to perfectly describe the observed distribution of chain lengths. Observed chains are significantly longer than expected under the null, suggesting that some clerks pretended that some providers were busy, in order to skip to their preferred provider. These deviations owe to a small number of clerk-markets, suggesting that some clerks may be regularly engaging in dishonest behavior in selected markets.

4 Causal evidence

Analyzing the distribution of the number of skips preceding the allocation of claims revealed significant deviations from would be expected if all clerks were honest, suggesting that some clerks may behave dishonestly in selected markets. In this section, I investigate how interpersonal relationships among clerks may cause such deviations.

4.1 Identification strategy and hypotheses

Identifying the causal effect of interpersonal relationships on dishonest behavior faces two main challenges. First, one need to credibly label behavior as honest or dishonest. Second, as with all studies of social influence, we face issues of endogeneity (Jackson, 2008): social relationships are not random. In this case, some markets feature more relationships than others, because the services awarded in those markets are more costly, and more complex. Such markets require more collaboration among clerks, but their increased complexity also opens up more opportunities for fraud. Furthermore, clerks might have a tendency to cluster by type, with dishonest types in one cluster, and honest types in another.

The identification strategy addresses the issue of measuring dishonest behavior by taking advantage of the highly structured nature of claim allocation. The claim allocation process makes dishonest behavior bear clear empirical implications that allow isolating suspicious behavior. Compared to their honest counterparts, dishonest clerks should have higher skipping and forcing rates, *ceteris paribus*. As such, I consider two outcomes for claim i: whether it was forced or not ($\text{force}_i = 1, 0$ respectively) and, if the claim was not forced, the number of draws leading to its allocation, $\text{skips}_i > 0$.

Ideally, to circumvent the endogeneity issue and identify the causal effect of social interaction on dishonest behavior, one would randomly assign social interactions during claim allocation, and observe the impact of such interactions on outcomes.

The identification strategy takes advantage of the size of the data to make comparisons within very small units, and consider only events where variation in social interactions is plausibly exogenous to dishonest behavior. I compare outcomes within dyad, within day of the week, and within hour, and control for the dollar value of the claim, as well as the level of business within market; that is, the number of claims awarded within that market within the past hour. Indeed, providers might be less available if there is more activity within the market, which would drive up skipping and forcing rates. I consider three events whose as-if-random nature is increasingly plausible, that I detail in the next subsection. The first is variation in attendance levels that owes to minor perturbations in how management schedules shifts and assigns clerks to markets. The second considers the same variation, but focuses on colleagues who do not know one another, which addresses concerns of selection by type. The third is recent hires, whose type is yet unknown to the rest of their colleagues.

Taking advantage of these events where variation in the number of interactions is plausibly exogenous, I first examine the average effect of the event. That is, for claim i in dyad j awarded at hour h of day d, I examine the following linear model:

$$Y_{ijhd} = \beta_{0j} + \beta_1 \text{ event}_i + \beta_2 \log(\text{amount}_i) + \beta_3 \text{ business}_i + \beta_{4h} + \beta_{5d}, \tag{1}$$

where $Y_{ijhd} \in \{\text{force}_i, \text{skips}_i\}$ is the outcome of interest.

Knowing that most clerks are honest, if they clerks deter dishonest behavior, then their presence should have no impact on other honest clerks, but decrease the incidence of fraudulent behavior by dishonest clerks, leading to lower skipping and forcing rates overall. That is:

Hypothesis 1. If honest clerks deter dishonest behavior, then $\beta_1 < 0$.

In our setting, however, the average effect may not be very indicative. If most clerks are honest, then the negative contribution of dishonest clerks to the average effect should be dwarfed by the null effect of honest clerks, leading us to observe $\beta_1 = 0$ instead.

As such, I look into heterogeneous effects, and examine a mixture model where dyad j can take one of two latent types $t_j \in \{H, L\}$, and let the intercept and coefficient on the event vary by type. Slightly adapting the previous approach, I now compare within provider instead of within dyad, and use some arbitrary provider as the reference category. That is, for claim i awarded by clerk j to company c on day d at hour h, with type $t_{jc} = t \in \{H, L\}$, I estimate the

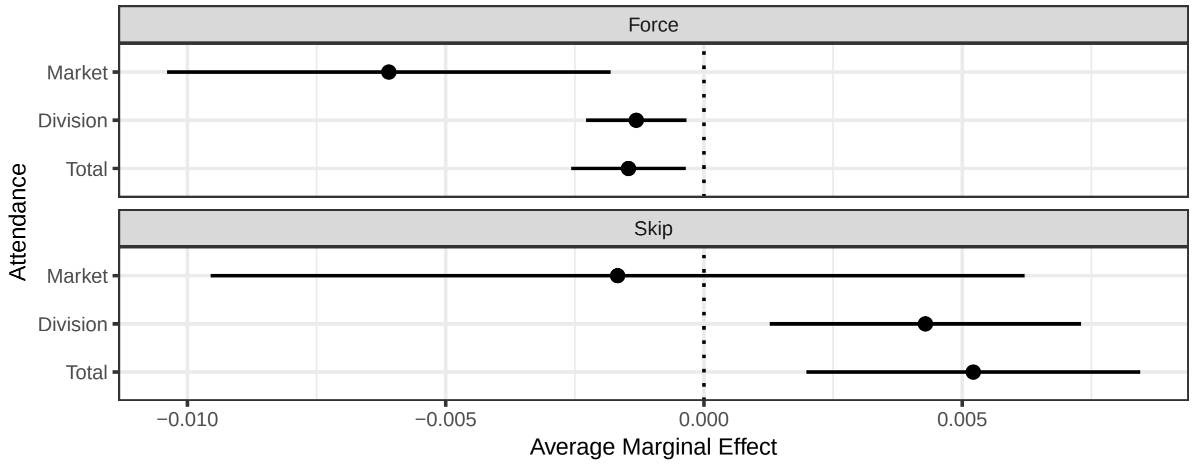

Figure 5: **Average marginal effect of attendance on force and skip rates.** Points are average marginal effects of one additional attendee, estimated using the specification in model 1. Bars represent 95 percent confidence intervals. Errors are clustered at the dyad level. On average, increased attendance decreases forcing behavior. It has no significant effect, or decreases skipping behavior.

following model:

$$Y_{ijchd} = \beta_{0t} + \beta_{1t}\ \text{event}_i + \beta_2 \log(\text{amount}_i) + \beta_3\ \text{business}_i + \beta_{4h} + \beta_{5d} + \beta_{6c}, \qquad (2)$$

where I set $\beta_{6c} = 0$ for some arbitrary company c that serves as a reference category. I estimate this model with the EM algorithm (Dempster, Laird and Rubin, 1977).

I identify the high type (H) as the one with the highest coefficient, and the low type (L) as the one with the lowest coefficient; that is, I define H and L such that $\beta_{1L} < \beta_{1H}$. Again, if honest clerks deter dishonest behavior, then their presence should have no impact on honest clerks, but decrease the incidence of fraudulent behavior by dishonest clerks, leading to lower skipping and forcing rates:

Hypothesis 2. If honest clerks deter dishonest behavior, then there are dyads of type L and $\beta_{1L} < \beta_{1H} = 0$.

4.2 Results

The first event I consider is variation in attendance levels: I use variation in the number of clerks present in the room when the claim was awarded. The strategy exploits small shocks in attendance levels that owe to minor variations in how management schedules shifts. Because the model includes hourly, and daily fixed effects, and the level of business, I control for the most obvious confounders; that is, temporality and business driving both attendance levels and our outcomes of interest. As detailed in section 2, I measure attendance at three levels, corresponding to increasingly close professional interactions. While total attendance is the number of clerks on the floor, I also consider attendance within division and and within market.

Figure 5 shows partial support for hypothesis 1: on average, increases in attendance decrease forcing behavior, suggesting that interpersonal interactions deter fraudulent behavior among

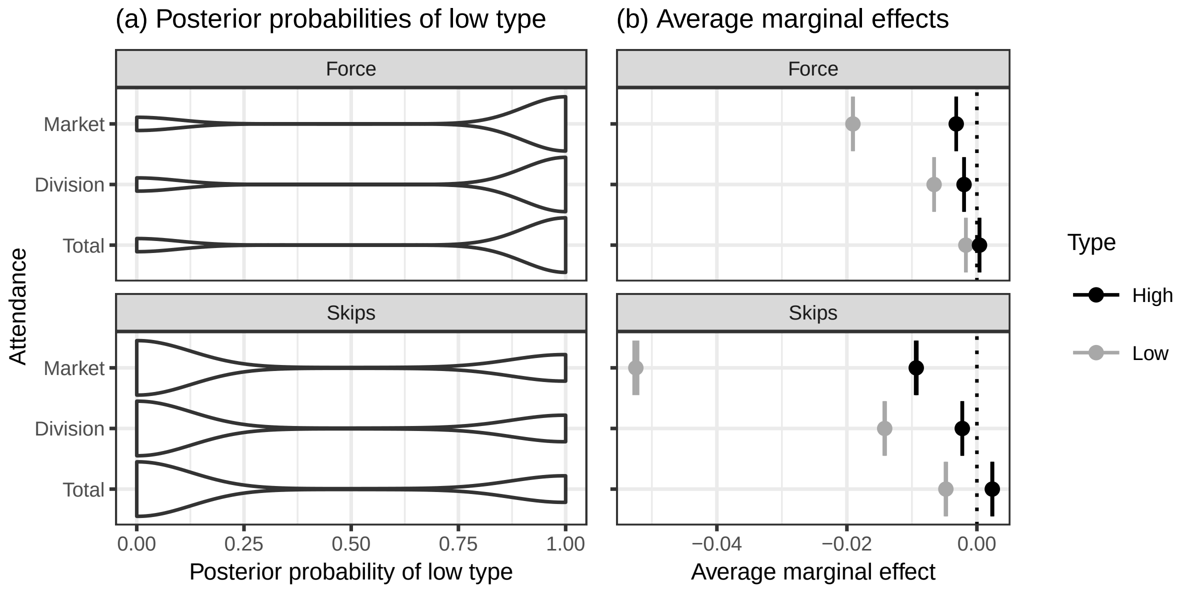

Figure 6: **Heterogeneous effects of attendance on force and skip rates.** Panel a represents the density of the posterior probability $\Pr(t_j = L|X, Y)$. The mixture models are separating. In panel b, points are average marginal effects of one additional attendee for each type, estimated using the specification in model 2. Bars are 95 percent confidence intervals. Increased attendance decreases forcing behavior for the low type, and has much smaller effects on the high type. The effect of increased attendance on skipping behavior is similar across types.

clerks. Furthermore, the effect is stronger for closer forms of professional interactions. In particular, market attendance has a markedly larger effect than division, or total attendance. However, such increases in attendance have either no effect, or increase skipping behavior.

Figure 6 considers heterogeneous treatment effects, and shows support for hypothesis 2. The mixture model shows that for the low type, increasing attendance decreases fraudulent behavior. In contrast, the high type has effect sizes that are statistically indistinguishable from zero, or have much smaller magnitude than the high type. As before, types become increasingly differentiated for closer forms of interactions, with the gap in effect size broadening as we move from total attendance to market attendance. Furthermore, a significant number of dyads compose the low type. The model classifies about a third of all dyads as low type for skipping behavior, and 80 percent as low type for forcing behavior. This suggests that undue forcing might be a milder, more widespread form of unethical behavior: instead of going through the cumbersome process of random drawing, clerks save time and force to some provider. When surrounded by colleagues, clerks quickly adjust their behavior and go back to normal. Finally, both models pick up the same dyads: for each kind of attendance, the correlation between the posterior probability of a dyad being of the low type for skipping and for forcing never goes below .23. This, in turn, suggests that the clerks that do engage in more serious forms of fraudulent behavior use both forcing and skipping.

Although this first result suggests that interpersonal relationships do reduce the suspicious behavior of some clerks, one challenge threatens identification: clerks may select into attendance. In particular, it might be that dishonest clerks cluster together on certain shifts, or on certain

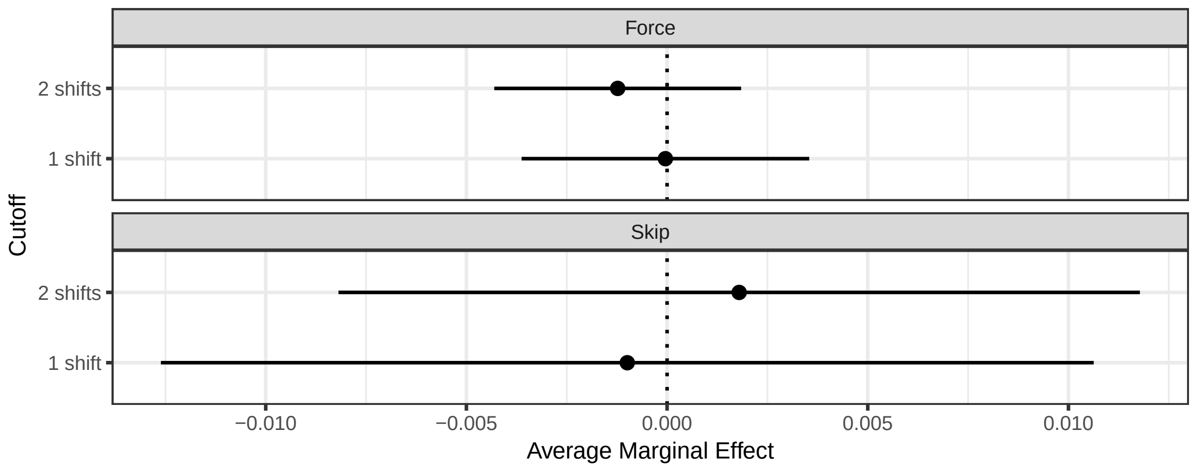

Figure 7: **Average marginal effect of unfamiliar colleagues on force and skip rates.** Points are average marginal effects of one additional attendee, estimated using the specification in model 1. Bars represent 95 percent confidence intervals. Errors are clustered at the dyad level. On average, unfamiliar colleagues have no statistically significant effect on forcing and skipping behavior.

markets.

To address the issue of selection, I consider another event: the within-market attendance of unfamiliar colleagues, defined as the colleagues with whom one has spent less than some threshold of time operating on the same markets in the past 30 days. Unfamiliar colleagues alleviate our concern of selection. Consider a dishonest clerk. Although unfamiliar colleagues may also be dishonest, their type is still unknown to the clerk under consideration, who will therefore behave as if those colleagues had been randomly selected. As detailed in section 2, I use two cutoffs, defining as unfamiliar the colleagues with whom one has spent less than 8, or 16 hours in the past 30 days.

Figure 7 considers average effects, and shows no support for hypothesis 1: the average marginal effect of unfamiliar colleagues is statistically indistinguishable from zero. As discussed in the preceding subsection, this likely owes to average effects pooling honest and dishonest types.

Conversely, Figure 8 considers heterogeneous effects, and shows support for hypothesis 2, with results that are very similar to considering both familiar, and unfamiliar colleagues. This confirms that there is a substantial fraction of clerks for which interpersonal relationships reduce the incidence of suspicious behavior, and that forcing behavior seems to be a milder, more widespread form of unethical behavior. Finally, results are stronger when using a cutoff of two shifts, and globally have a smaller magnitude than when considering both strong, and weak ties. This suggests that there is little concern for selection. Indeed, if selection was driving the results, then they should become weaker when considering stronger ties, who should be clustered together, hence exerting less monitoring on dishonest types.

Since there may still be concerns that our measure of familiarity does not measure perfectly relationships among colleagues–especially relationships outside the workplace–the third event considers the impact of new hires on the behavior of current employees. This event addresses

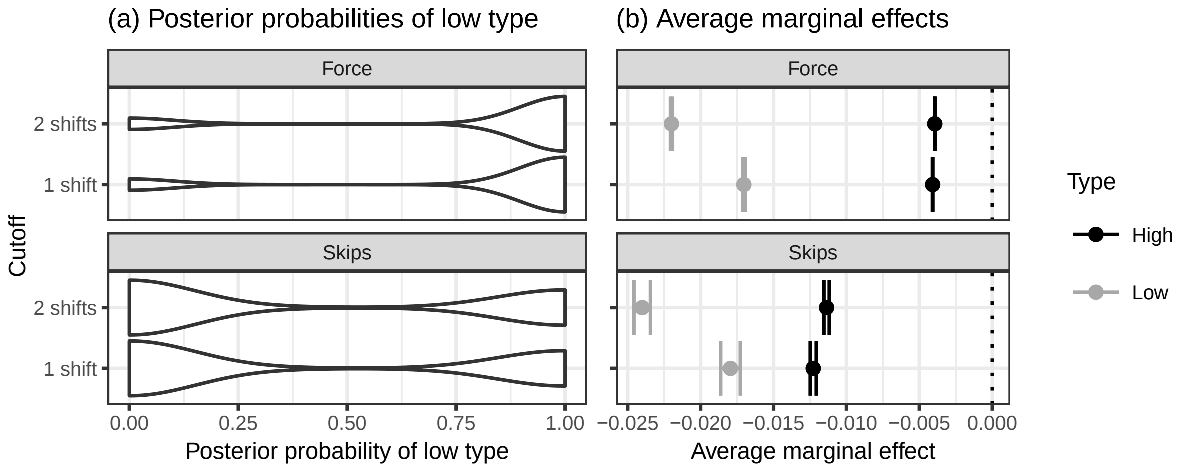

Figure 8: **Heterogeneous effects of unfamiliar colleagues on force and skip rates.** Panel a represents the density of the posterior probability $\Pr(t_j = L|X, Y)$. The mixture models are separating. In panel b, points are average marginal effects of one additional attendee for each type, estimated using the specification in model 2. Bars are 95 percent confidence intervals. Increased attendance decreases forcing and skipping behavior for the low type, and has smaller effects on the high type.

in a more convincing way the issue of selection: in this high-turnover environment, new hires are typically unknown of current employees. Because new hires may take some time to “settle in,” and become familiar to current employees, I use different windows to define the event of being a new hire, ranging from the first day of a clerk, to her first week. As for attendance, I define the event at three different scales: the entire call-center, the division, and the market, counting, respectively, the number of new hires currently operating in the room, the division, or the market of the claim under consideration.

However, this event comes with two major drawbacks. First, it is less frequent than the previous two, increasing our uncertainty about the estimates. Second, the behavioral implications are not as clear-cut as for the other events. By definition, new hires are inexperienced. As such, dishonest types may not alter their behavior in the presence of new hires, for the latter might be too inexperienced to notice suspicious behavior. When the new hire is experienced enough to notice this such behavior and have an impact on dishonest types, she has been working for long enough that her type becomes known to her colleagues.

Unsurprisingly, I find mixed results, that are reported in section A of the Appendix. On average, new hires have little effect on forcing and skipping behavior, and such effect is indistinguishable across types.

Overall, interpersonal relationships cause reductions in suspicious behavior for low-type dyads, increasingly so as we consider closer forms of monitoring: results are stronger for colleagues that interact on the same market, and that share a stronger tie. Furthermore, results suggest that a majority of clerks might engage in mildly unethical behavior: a majority of clerks reduce their forcing rates when other colleagues are present.

5 Structural evidence

The analysis in the previous section showed that interpersonal relationships cause reductions in suspicious behavior by dishonest clerks. It exploited on a reduced-form implication of the data generating process–dishonest clerks should have higher skipping and forcing rates than honest clerks, and reduce such dishonest behavior when surrounded with honest clerks.

This reduced-form approach, however, suffers from two major drawbacks. First, since this implication only holds *ceteris paribus*, it forced to make comparisons within very small units, potentially wasting a lot of power. Second, the analysis conflates identifying clerks' types with identifying properties about such types. In other words, the mixture models estimated in the preceding section define as dishonest the clerks that do reduce their forcing and skipping rates when surrounded with more clerks. A better analysis would first identify dishonest clerks, then examine how such clerks behave, and test whether their behavior is sensitive to monitoring from honest colleagues.

I estimate a structural model that takes advantage of the highly structured process defined by the rotation to recover this missing information–the probability that a provider is unavailable and the probability that a clerk is dishonest–and to describe the behavior of dishonest clerks. The model considers only claims allocated through the rotation, and ignores claims that were forced.

5.1 Model

Consider draw i from the rotation. The econometrician only observes the clerk $j[i]$ that requested the draw and the company $c[i]$ that was drawn, as well as the outcome $y_i = 0$ if the claim was awarded to $c[i]$, or $y_i = 1$ if the company was skipped. The clerk, on the other hand, knows her type, and observes whether the company was available ($b_i = 0$), or busy ($b_i = 0$).

To identify dishonest clerks, the model exploits the same implications of the data generating process as in the preceding section, but does so at the micro-level: a dishonest clerk may only unduly skip if the company was available. In other words, the model defines types at the dyad level. Consider the dyad between clerk j and company c. If that dyad is honest, it has type $t_{jc} = H$, and the clerk only skips when the company is busy. That is, $b_i = y_i$ for all draws from dyad jc. If the dyad is dishonest, it has type $t_{jc} = L$. In this case, the clerk may pretend that company $c[i]$ was unavailable and skips ($y_i = 1$), although it was not ($b_i = 0$). That is, $b_i \leq y_i$ for all draws from dyad jc.

Defining types at the dyad level has important implications for model estimation and interpretation. First, since the dependent variable y_i is binary, the model reduces to a mixture of binomials, which are generically not identified. Defining types at the dyad level pools claims within dyads, which identifies the model. Second, this approach impacts interpretation; in particular, it labels as honest the providers that dishonest clerks collude with, but labels as dishonest the providers that such clerks take revenue from. Suppose that clerk j is colluding with company c^*. Then j will only unduly skip when companies $c \neq c^*$ are drawn. That is, j will appear as type H for dyad jc^*, and as type L for dyads jc, $c \neq c^*$.

This suffices to define the model. Consider dyad jc and draw i from that dyad. Suppose that

jc is dishonest with probability π. Let p_i be the probability that the provider is available, and $1 - q_i$ the probability that an undue skip occurs, given that dyad jc is dishonest and provider c is available. I make p_i and q_i depend covariates using probit regressions. Let x_i and z_i be vectors of covariates that affect availability and dishonest skipping respectively, and β and γ associated vectors of parameters. With $\Phi(.)$ the cumulative density function of the standard normal distribution, we have the following model:

$$\begin{aligned}
\pi &= \Pr(t_{ij} = L) \\
p_i &= \Pr(b_i|x_i) = \Phi(x_i'\beta) \\
q_i &= \Pr(y_i|z_i, b_i, t_{j[i]c[i]}) = \begin{cases} \Phi(z_i'\gamma) \text{ if } t_{j[i]c[i]} = L \text{ and } b_i = 0 \\ b_i \text{ otherwise} \end{cases}
\end{aligned}$$

I estimate the model using a collapsed Gibbs sampler (Holmes and Held, 2006, see Appendix B for details) with the following conjugate priors:

$$\begin{aligned}
\pi &\sim \text{Beta}(\alpha_0, \alpha_1) \\
\beta &\sim N(0, \sigma_0^2) \\
\gamma &\sim N(0, \varsigma_0^2)
\end{aligned}$$

5.2 Specification and results

The model requires specifying covariates for the probability of being unavailable, $\Pr(b_i|x_i)$, and the probability of a dishonest type unduly skipping, $\Pr(y_i = 0|z_i, b_i = 0, t_{j[i]c[i]} = L)$. The chosen specification follows closely the specifications used for causal inference. I model the probability of being available using company, hour, and day of the week fixed effects, and control for claim value, as well as the level of business. Even without collusion, providers might be more willing to service high-value claims, and pretend they are unavailable for claims that are not profitable enough. Furthermore, at times of peak activity, all providers within a given market are more likely to be unavailable. As in the causal inference section, I also control for the level of business within the market. With $\Phi(.)$ the cdf of the standard normal distribution, the specification for being unavailable writes:

$$\Pr(b_i|x_i) = \Phi\left(\beta_{0c[i]} + \beta_{\text{hour}} + \beta_{\text{day}} + \beta_3 \log(\text{amount}_i) + \beta_4 \text{business}_i\right)$$

Dishonest clerks have a baseline probability of skipping unduly, but may also adjust their behavior according to circumstances. In particular, they should be more inclined to unduly skip high-value claims, since those claims bring more revenue to the provider they are colluding with. If interpersonal relationships affect dishonest behavior, they should also adjust their behavior when other clerks are present. As such, I control for the value of the claim, and the level of market attendance, which has been shown in the previous section to trigger the strongest

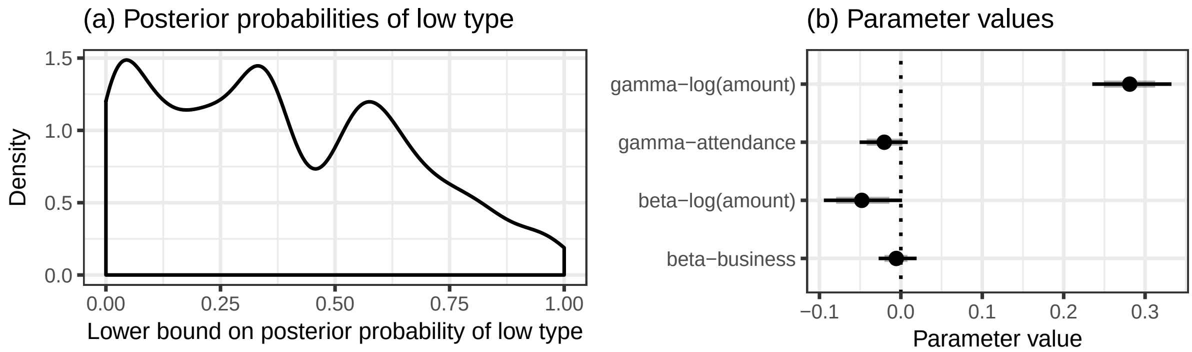

Figure 9: **Structural model, results.** Panel a represents the density of the lower bound of the 95 percent credible interval around the posterior probability of being a low type. The model isolates a few dyads that are very likely to be of the low type. In panel b, points are posterior mean parameter values; thick bars are 80 percent credible intervals, thin bars are 95 percent credible intervals. Low types skip more higher value claims, and skip less when there is high attendance. The model is estimated using 1000 iterations of the Gibbs sampler, with a burn-in of 100 iterations.

deterrence effect. I also include hour and day of the week fixed effects:

$$\Pr(y_i|z_i, b_i = 0, t_{j[i]c[i]} = L) = \Phi\left(\gamma_0 + \gamma_{\text{hour}} + \gamma_{\text{day}} + \gamma_1 \log(\text{amount}_i) + \gamma_2 \text{market attendance}_i\right)$$

The discussion yields two core hypotheses about the behavior of dishonest types. First, these types should be more inclined to behave dishonestly for high-value claims:

Hypothesis 3. There are dyads of type L and $\gamma_1 > 0$.

Second, if honest clerks deter dishonest behavior, then market attendance should reduce the extent of undue skipping.

Hypothesis 4. If honest clerks deter dishonest behavior, then there are dyads of type L and $\gamma_2 < 0$.

Empirical results confirm the hypotheses. Panel (a) of Figure 9 shows that while most dyads are likely honest, some have a very high probability of being dishonest. I use a conservative estimate for identifying dishonest types: instead of considering the mean posterior probability of a dyad being dishonest, I only consider the lower bound of the 95 percent credible interval. By this metric, .7 percent of dyads have a 99 percent chance of being dishonest, while 96 percent have less than 90 percent chance of being dishonest.

Panel (b) of Figure 9 confirms hypotheses relative to the behavior of low types. As expected, they are more likely to skip high-value claims, since those transfer more revenue to the provider they are colluding with. They are, however, less likely to behave dishonestly when more clerks are present, which implies that interpersonal relations within the organization deter dishonest behavior.

6 Conclusion

This paper answered one simple question: within an organization, do interactions between honest and dishonest colleagues spread, or deter corruption? I showed that these interactions deter corruption: dishonest colleagues reduce their suspicious behavior when other colleagues are present. I also showed that a milder form of unethical behavior may be widespread within the organization. Results show that most clerks reduce their forcing behavior when other colleagues are present, suggesting that they may unduly force when left to their own device, presumably to avoid the time-consuming process of random draws. I finally showed that dishonest types are sophisticated, in the sense that they are more likely to skip unduly for high value claims, since those transfer more revenue to the providers they are colluding with.

The findings have two important implications. First, they show that organizational structure affects the emergence and extent of corruption. By structuring the interpersonal relationships that take place within them, organizations give birth to patterns of informal monitoring that curtail corruption. The findings yield an important policy recommendation: in order to reduce corruption, it would be sensible to reduce the isolation of bureaucrats. They may explain why policies that combat corruption through major reforms of the organizational chart have proven so popular in recent years (eg. Bennet, 2012; Friedman, 2012; Hausman, 2011). Dishonest agents can circumvent norms: in our setting, dishonest agents exploit the smallest cracks in the claim allocation process to allocate claims to the providers they are colluding with. On the other hand, it is more difficult for such agents to circumvent monitoring from their peers. As a result, it should not be too surprising that one such popular policies is the creation of one-stop shops, a type of organizational unit that regroups previously disparate services into a single organization, therefore increasing the amount of peer monitoring between services.

Second, the findings force us to rethink existing macro-empirical work in light of these behavioral micro-foundations. Results show that peers deter dishonest behavior, including milder, more widespread forms of such behavior. As a result, instances of corruption spreading to an entire organization should be rare, and only occur in environments involving relatively fewer interpersonal relations. This suggests that arguments on corruption spreading to entire organizations in contexts of dense social ties may rely on some other behavioral foundation. Given the present set of results, it is unclear how overembedded organizations (Evans, 1995) may be more corrupt, and how dense social and cultural ties between the regulator and the industry may give birth to socio-cultural capture (Carpenter and Moss, 2013). As such, these results prompt for further research. Although the present study showed that bad apples can be deterred by good apples, lab evidence (Gino and Pierce, 2009) suggests that good apples can also be infected by bad ones. To understand those cases of organization-wide corruption, further research should examine the conditions under which the influence of honest agents trumps that of dishonest agents.

References

Aven, Brandy. 2015. "The Paradox of Corrupt Networks: An Analysis of Organizational Crime at Enron." *Organization Science* 26(4):980–996.

Bennet, Richard. 2012. "A Change Agent in the Tax Office: Nigeria's Federal Inland Revenue Service, 2004-2009." *Innovations for Successful Societies Case Study* .

Carpenter, Daniel and David A Moss. 2013. *Preventing Regulatory Capture.* New York, NY: Cambridge University Press.

Carpenter, Daniel P. 2001. *The forging of bureaucratic autonomy : reputations, networks, and policy innovation in executive agencies, 1862-1928.* Princeton, NJ: Princeton University Press.

Chandler, Alfred D. 1977. *The visible hand : the managerial revolution in American business.* Cambridge, MA: Harvard University Press.

Charron, Nicholas, Carl Dahlström, Mihaly Fazekas and Victor Lapuente. 2017. "Careers, Connections, and Corruption Risks: Investigating the Impact of Bureaucratic Meritocracy on Public Procurement Processes." *The Journal of Politics* 79(1):89–104.

Crozier, Michel. 2009. *The Bureaucratic Phenomenon (with a new introduction by Erhard Friedberg).* London: Transaction Publishers.

Dempster, A. P., N. M. Laird and D. B. Rubin. 1977. "Maximum Likelihood from Incomplete Data via the EM Algorithm." *Journal of the Royal Statistical Society. Series B (Methodological)* 39(1):1–38.

Evans, Peter B. 1995. *Embedded autonomy: states and industrial transformation.* Princeton, N.J.: Princeton University Press.

Friedman, Jonathan. 2012. "Saving a Sinking Agency: The National Port Authority of Liberia, 2006-2011." *Innovations for Successful Societies Case Study* .

Gambetta, Diego. 1996. *The Sicilian Mafia : the business of private protection.* Cambridge, Mass.: Harvard University Press.

Gino, Francesca and Adam D. Galinsky. 2012. "Vicarious dishonesty: When psychological closeness creates distance from one's moral compass." *Organizational Behavior and Human Decision Processes* 119(1):15–26.
URL: *http://dx.doi.org/10.1016/j.obhdp.2012.03.011*

Gino, Francesca and Lamar Pierce. 2009. "The abundance effect: Unethical behavior in the presence of wealth." *Organizational Behavior and Human Decision Processes* 109(2):142–155.

Gino, Francesca, Shahar Ayal and Dan Ariely. 2009. "Contagion and Differentiation in Unethical Behavior - The Effect of One Bad Apple on the Barrel." *Psychological Science* 20(3):393–398.

Gordon, Sanford C. 2011. "Politicizing Agency Spending Authority: Lessons from a Bush-era Scandal." *American Political Science Review* 105(04):717–734.

Granovetter, Mark. 1973. "The Strength of Weak Ties." *American Journal of Sociology* 78(6):1360–1380.

Granovetter, Mark. 1985. "Economic Action and Social Structure: The Problem of Embeddedness." *American Journal of Sociology* 91(3):481–510.

Haeder, Simon F and Susan W Yackee. 2015. "Influence and the Administrative Process: Lobbying the U.S. President's Office of Management and Budget." *American Political Science Review* 109(3):507–522.

Hausman, David. 2011. "Professionalization, Decentralization and a One-Stop Shop: Tax-Collection Reform in Ghana, 1986-2008." *Innovations for Successful Societies Case Study* .

Holmes, Chris C. and Leonhard Held. 2006. "Bayesian auxiliary variable models for binary and multinomial regression." *Bayesian Analysis* 1(1):145–168.

International Monetary Fund - Staff Team from the Fiscal Affairs Department and the Legal Department. 2016. Corruption: Costs and mitigating strategies - IMF Staff Discussion Note no. SDN/16/05. Technical report.

Jackson, Matthew O. 2008. *Social and economic networks.* Princeton, NJ: Princeton University Press.

Kaufman, Herbert. 1960. *The Forest Ranger: A Study in Administrative Behavior.* 2006 repri ed. New York, NY: Routledge.

Khanna, Vikramaditya, E Han Kim and Yao Lu. 2015. "CEO Connectedness and Corporate Fraud." *The Journal of Finance* 70(3):1203–1252.

Mauro, Paolo. 1995. "Corruption and Growth." *The Quarterly Journal of Economics* 110(3):681–712.

Pietryka, Matthew T and Donald A DeBats. 2017. "It's Not Just What You Have, but Who You Know: Networks, Social Proximity to Elites, and Voting in State and Local Elections." *American Political Science Review* 111(02):360–378.

Pitesa, Marko and Stefan Thau. 2013. "Compliant sinners, obstinate saints:How powerand self-focus determine the effectiveness of social influence in ethical decision making." *Academy of Management Journal* 56(3):635–658.

Rauch, James E and Peter B Evans. 2000. "Bureaucratic structure and bureaucratic performance in less developed countries." *Journal of Public Economics* 75(1):49–71.

Vannucci, Alberto and Donatella Della Porta. 2013. *The Hidden Order of Corruption: An Institutional Approach.* Farnham: Ashgate Publishing, Ltd.

Weber, Max. 1948. Bureaucracy. In *From Max Weber: Essays in Sociology*, ed. Hans Gerth and Charles W Mills. Oxford: Routledge Press pp. 196–244.

Zukin, Sharon and Paul DiMaggio. 1990. Introduction. In *Structures of capital : the social organization of the economy*, ed. Sharon Zukin and Paul DiMaggio. Cambridge England ; New York: Cambridge University Press pp. 1–36.

A Causal evidence

Figure 10 shows that on average, new hires have little effect on skipping and forcing behavior, and that this effect vanishes as the length of the window increases. The effect, however, seems more pronounced for forcing: new hires significantly decrease forcing behavior for windows of length 1, 5, 6, and 7. Yet, Figure 11 shows that such effect is largely indistinguishable across types, irrespective of the window length.

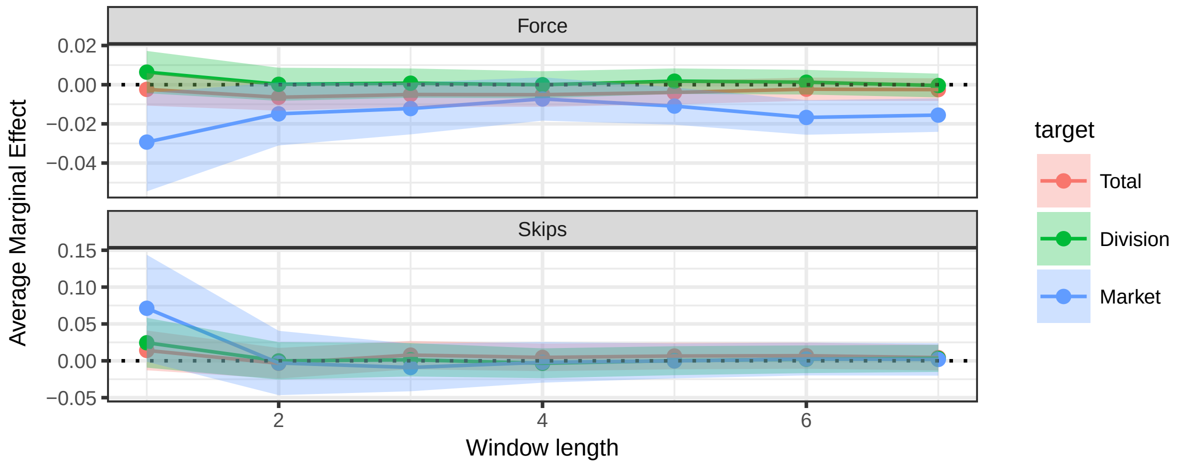

Figure 10: **Average marginal effect of attendance on force and skip rates.** Points are average marginal effects of one additional new hire, estimated using the specification in model 1. Shaded areas are 95 percent confidence intervals. Errors are clustered at the dyad level. On average, increased attendance decreases forcing behavior. It has no significant effect, or decreases skipping behavior.

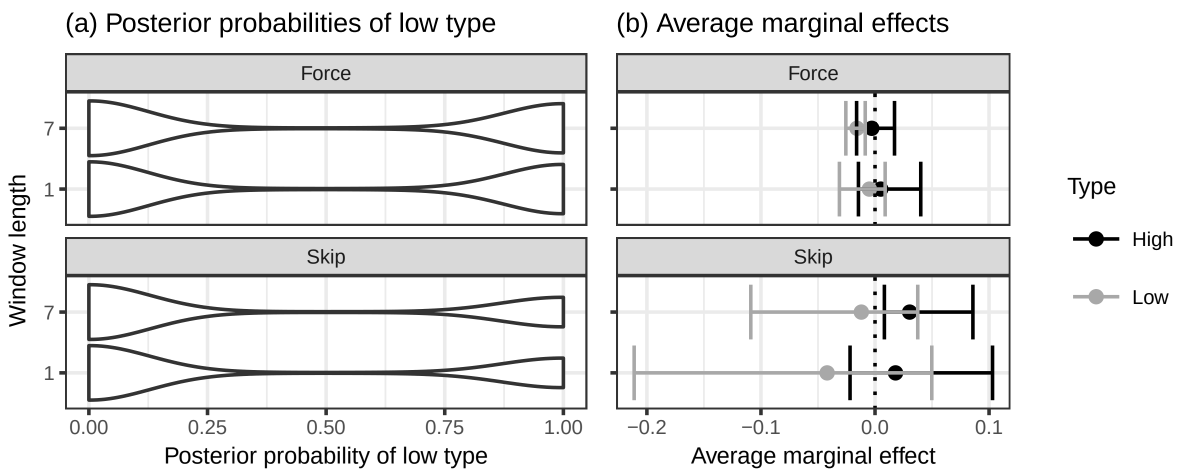

Figure 11: **Heterogeneous effects of new hires on force and skip rates.** Panel a represents the density of the posterior probability $\Pr(t_j = L|X, Y)$. The mixture models are separating. In panel b, points are average marginal effects of a new hire for each type, estimated using the specification in model 2. Bars are 95 percent confidence intervals. Irrespective of the window length, the effect of new hires is indistinguishable across types.

B Gibbs sampler

Note that p_i and q_i are probit regressions of b_i and y_i on x_i and z_i respectively. I use the standard latent utility representation, with

$$
\begin{aligned}
u_i &= x_i'\beta + \epsilon_i \\
v_i &= z_i'\gamma + \eta_i \\
\epsilon_i, \eta_i &\sim N(0,1) \\
u_i \geq 0 &\iff b_i = 1 \\
v_i \geq 0 &\iff y_i = 1 | t_{j[i]c[i]} = L, b_i = 0
\end{aligned}
$$

The Gibbs sampler proceeds as follows:

1. Update $p|\beta = \Phi(x\beta)$, and $q|\gamma = \Phi(z\gamma)$
2. Update $\tau_{ij} = \Pr(t_{ij}|\pi, p, q, y) = \frac{\pi \prod_{k \in C_{ij}} (p_k^L)^{y_k} (1-p_k^L)^{1-y_k}}{\pi \prod_{k \in C_{ij}} (p_k^L)^{y_k} (1-p_k^L)^{1-y_k} + (1-\pi) \prod_{k \in C_{ij}} p_k^{y_k} (1-p_k)^{1-y_k}}$
3. Update $\pi|p, q, \tau, y \sim \text{Beta}(\sum_{ij} \tau_{ij} + \alpha_0, D - \sum_{ij} \pi_{ij} + \alpha_1)$
4. Update $b_i|\pi, p_i, q_i, y_i \sim \text{Binom}\left(y_i \left[(1-\pi) + \pi \frac{p_i}{[p_i + (1-p_i)q_i]^{y_i} [(1-p_i)(1-q_i)]^{1-y_i}}\right]\right)$
5. Update $\beta, \sigma^2|b, x$ using a collapsed Gibbs sampler (Holmes and Held, 2006).
6. Update $\gamma, \varsigma^2|y, z, b, \tau$ using Holmes and Held (2006), with weights $w_i = \Pr(b_i = 0, t_{j[i]c[i]} = L|p_i, q_i, \tau_{j[i]c[i]}) = \tau_{j[i]c[i]} \frac{(1-p_i)q_i}{p_i + (1-p_i)q_i}$