The Unequal Diffusion of Honesty and Dishonesty in Workplace Networks

Romain Ferrali *

March 5, 2025

Abstract

Honest and dishonest behaviors may both diffuse among the members of an organization. Knowing which of the two spreads faster is important because it impacts the extent to which managers will need to resort to other, potentially more costly solutions to curb dishonest behavior. Assessing empirically which of honest or dishonest behavior spreads faster is challenging, because this requires field measurements of social relationships and dishonest behavior of individual members, which poses both measurement and inference problems. We examine an original, fine-grained dataset from a large company that allows for identifying agents likely to be dishonest and interactions among employees while offering a natural experiment that circumvents the inference problems associated with identifying peer-to-peer diffusion. We find (1) that dishonest behavior diffuses while honest behavior does not, (2) that diffusion likely operates through spreading information about opportunities for collusion, and (3) that policies that screen on dishonesty at hiring may be efficient to curb dishonest behavior in environments with high turnover.

1 Introduction

Unethical behavior is costly to individuals, organizations, and society (Mazar and Ariely, 2006; Rose-Ackerman and Palifka, 2016). Three categories of solutions have been identified for curbing such behavior in organizations (Treviño et al., 2014): (1) strengthening the organization's "ethical infrastructure" (through, e.g., ethical codes and training), (2) selecting ethical employees, and (3) fostering interpersonal relationships that promote ethical behavior at the leadership, employee-manager, and peer levels. Among these, peer relationships are particularly challenging. Indeed, they may facilitate the diffusion of both ethical (Abbink, 2004; Ariely et al., 2009) *and* unethical behavior (Gino et al., 2009; Weisel and Shalvi, 2015), and of other elements of interest to the organization, such as information and productive behavior (Mas and Moretti, 2009; Azoulay et al., 2010; Oettl, 2012; Song et al., 2018).

Knowing which of honest or dishonest behavior spreads faster has important implications for curbing unethical behavior. If honest behavior spreads faster, then peer interactions help. They reinforce the organization's ethical infrastructure, enable the diffusion of information and productive behavior, and reduce reliance on costly alternatives like screening and monitoring. If dishonest behavior spreads faster, then the other benefits of peer interactions come at the

*Aix-Marseille University, CNRS, AMSE, France, romain.ferrali@univ-amu.fr. I thank Olivier Bochet, Timothée Demont, Habiba Djebbari, Matias Iaryczower, Kosuke Imai, Jenn Larson, Horacio Larreguy, Adeline Lo, John Londregan, Kinga Makovi, Jean-Albert Margaine, Ameet Morjaria, Rebecca Morton, Malte Reichelt, Nikos Nikiforakis, Melina Platas, Talal Rahwan, Ernesto Reuben, Avner Seror, Smriti Sharma, Marie-Claire Villeval, Yves Zenou, as well as participants in the Princeton Political Economy seminar, the Imai Research Group, the IMEBESS, WESSI and WTFNS workshops, and the APSA, CSAE, and DIAL conferences for their helpful comments. I also thank Habibou Kassoum for his excellent research assistance. I gratefully acknowledge support from the Mamdouha S. Bobst Center for Peace and Justice, the French National Research Agency Grant ANR-17-EURE-0020, and the Excellence Initiative of Aix-Marseille University - A*MIDEX. All errors are mine.

price of spreading dishonest behavior which, in turn, increases the need to resort to other, costly dishonesty-reducing policies.

Experimental studies have shown that dishonest behavior diffuses more than honest behavior (Innes and Mitra, 2013; Dimant, 2019; Brunner and Ostermaier, 2019; Colzani et al., 2023), but these findings need confirmation in field settings for two reasons. First, lab results are sensitive to design features such as the size of the reward for cheating and the experimental task (Gerlach et al., 2019; Leib et al., 2021). While design dependence may not be problematic for testing how specific mechanisms impact the level of (dis)honest behavior, this becomes problematic when estimating the *relative* magnitudes of the diffusion of these behaviors. Field studies circumvent the issue by setting features of the environment such as the probability of detection, the norms surrounding dishonesty, and the incentives for cheating at naturally occurring levels.

Second, while the literature has identified a number of mechanisms supporting the diffusion of (dis)honest behavior, studies that examine the joint diffusion of honest and dishonest behavior consider designs that focus on one mechanism and control for the other. These mechanisms fall in three categories: (1) *information diffusion*, whereby social interactions enable learning about features of the environment such as the likelihood and severity of punishment (Palmer and Yenkey, 2015; Alm et al., 2017; Zhang et al., 2023) or the identity of potential partners in crime (Gross et al., 2018), (2) whether dishonest behavior is (i) *complementary*, e.g., when requiring collaboration (Weisel and Shalvi, 2015) or when collective dishonest behavior complicates punishment (Conrads et al., 2013) or (ii) *substitutable*, e.g., when too much dishonest behavior attracts attention and increases the likelihood of punishment disproportionately (Chan et al., 2021), and (3) *conformism*, when agents adopt the dominant norm of behavior, e.g., to uphold a good reputation among their peers and reap the benefits of long-term relationships (Banerjee and Duflo, 2000; Macchiavello and Morjaria, 2015), because of self-image concerns (Ariely et al., 2009; Dai et al., 2017), or perceptions of the acceptability of dishonest behavior (Chui et al., 2021). Evaluating which of these mechanisms most likely accounts for diffusion in field settings has important consequences for providing effective policy responses.

This paper examines how honest and dishonest behavior jointly diffuse in a real organization. It then identifies the set of mechanisms that account for the observed diffusion patterns. It finally conducts counterfactual exercises to identify policies that may curb the diffusion of dishonest behavior.

We consider the call center of a large roadside and health assistance company in the Middle East and North Africa region. Call-center clerks service customers' claims (e.g., a flat tire) by dispatching service providers registered in the firm's network (e.g., a tow truck). This unique setting provides a series of advantages in terms of measurement and inference. The environment leaves room for only one important form of dishonest behavior; namely, collusion between a clerk and a service provider (i.e., a clerk disproportionately dispatches the same tow-truck company in return for a kickback). Using a backup of the company's internal software, we derive a statistical measure of collusion and reconstruct interactions among clerks. Similar to other studies of peer effects inside the firm (Chan et al., 2014b,a; Hasan and Koning, 2019; Chan et al., 2021; Lindquist et al., 2022), as-if-random co-staffing offers a natural experiment that alleviates concerns of bias regarding the identification of peer effects.

We find that dishonest behavior diffuses across social ties, while honest behavior does not. When a clerk that behaves honestly interacts with a clerk that behaves dishonestly, her chances of behaving dishonestly the next month increase by 5.1 percentage points. Within the same interaction, the likelihood that the dishonest clerk behaves honestly the next month does not change meaningfully.

We then provide suggestive evidence of the mechanism that most likely supports observed peer effects in dishonest behavior, showing that such peer effects are more consistent with information transmission than with conformism or complementarities. More junior clerks learn about opportunities for collusion from the more senior clerks who behave dishonestly. A fraction of those junior clerks then leverage this information to engage in dishonest behavior.

We finally simulate counterfactual policies for curbing the diffusion of dishonest behavior. We consider a screening policy, where new hires are screened on honesty, and a network reassignment policy that isolates clerks behaving dishonestly in order to limit their diffusion potential. We find that limited screening is the most effective policy when networks are sparse and dishonest behavior is very prevalent at the onset. Indeed, in this high-turnover environment, limited investments in screening gradually increase the stock of honest behavior, while sparse networks prevent exposing employees who behave honestly to dishonest behavior. This suggests that, under specific conditions, limited investments in screening may strongly foster honest behavior.

This paper makes three main contributions to a growing literature on the diffusion of behavioral misconduct and unethical behavior.

Our first contribution is to provide what are, to the best of our knowledge, the first *field* estimates of the *relative* magnitude of peer effects in honest and dishonest behavior. Our results confirm extant experimental evidence (Innes and Mitra, 2013; Dimant, 2019; Brunner and Ostermaier, 2019; Colzani et al., 2023): dishonest behavior spreads, while honest behavior does not.

Our second contribution is to provide suggestive evidence of the mechanisms underlying the diffusion of dishonest behavior. We find that the observed patterns of influence are more consistent with information transmission than with conformism or complementarities. Experimental work studying the relative diffusion of honest and dishonest behavior has considered either conformism (Innes and Mitra, 2013; Dimant, 2019) or information (Brunner and Ostermaier, 2019; Colzani et al., 2023). Our field evidence lends external validity to the focus on informational mechanisms.

Our third contribution is to conduct a counterfactual analysis of policies to curb dishonest behavior, identifying a relatively feasible policy solution. Unlike lab studies, our field setting allows us to evaluate the effectiveness of interventions in real organizational contexts and to provide actionable insights. While prior research has explored how network structure[1] influences various organizational behaviors (e.g., Lindquist et al., 2022), its application to dishonesty remains underexamined. Unlike productivity, dishonest behavior is costly to measure, posing significant challenges for policy design. We address this gap by assessing the impact of a series of network and non-network interventions. A policy that screens new hires for honesty is relatively feasible and, as we demonstrate, effectively curbs dishonesty in high-turnover environments such as ours.

The remainder of the paper proceeds as follows. Section 2 presents a theoretical model that will guide our empirical exploration. Section 3 describes the context and data. Section 4 outlines the identification strategy. Section 5 presents our findings. Section 6 concludes.

2 Theory

This section introduces a simple formal model that will guide the empirical analysis. The model considers a population of n agents connected by the undirected graph $g = (G, N)$ where N is a set of nodes and G a set of links. Let $y_{it} = 1$ if agent i behaves dishonestly at time t, and $y_{it} = 0$ if otherwise. Let θ_{it} be the share of i's neighbors that behave dishonestly at time t and d_i the number of i's neighbors. Behaviors are randomly drawn with transition probabilities $p \equiv \Pr(y_{i,t+1} = 0 | y_{it} = 1, \theta_{it})$ and $q \equiv \Pr(y_{i,t+1} = 1 | y_{it} = 0, \theta_{it})$ which capture, respectively, the probability that an agent behaving dishonestly agent starts behaving honestly, and that an

[1] This paper does not distinguish between the formal and informal relationships occurring within the organization. We refer to both as a network structure.

agent behaving honestly starts behaving dishonestly. They read as follows:

$$p(\theta_{it}, d_i) = \alpha_p + \beta_p \overbrace{(1 - \theta_{it})d_i}^{\text{number of neighbors behaving honestly}} \quad (1)$$

$$q(\theta_{it}, d_i) = \alpha_q + \beta_q \underbrace{\theta_{it} d_i}_{\text{number of neighbors behaving dishonestly}} \quad (2)$$

Transition probabilities depend on an intrinsic propensity to switch behaviors $\alpha \geq 0$ and on the number of neighbors whose behavior is opposite to the agent's behavior, θd or $(1 - \theta)d$. The parameters β capture the rate of diffusion of honest and dishonest behavior; that is, how much the switching of behaviors depends upon one's peers. If $\beta = 0$, then behavior does not diffuse; that is, i's behavior does not depend on her neighbors'. Cases where $\beta > 0$ capture a notion of *diffusion*, wherein having more neighbors whose behavior is opposite fosters switching behavior. Conversely, cases where $\beta < 0$ capture a notion of *differentiation*, wherein having more neighbors whose behavior is opposite reinforces adherence to one's behavior. This simple setting expands upon Susceptible-Infected-Susceptible (SIS) models of the diffusion of diseases on networks (e.g., Jackson and Rogers, 2007) by introducing two innovations: first, diffusion may not only flow from dishonest/infected to honest/healthy agents but also from honest to dishonest agents. Second, the model not only allows for diffusion ($\beta > 0$) but also for differentiation ($\beta < 0$). Like these models, it allows considering arbitrarily complex networks and may be solved analytically under a set of standard assumptions.[2]

In the model, parameters α, which control agents' intrinsic propensity to switch behavior, may capture a wide array of mechanisms that are not related to peer-to-peer diffusion. They could, for instance, denote the organization's ethical infrastructure. Successful ethical training programs may result in increasing the baseline propensity towards honest behavior α_q and decreasing the baseline propensity towards dishonest behavior α_p. They may also capture other forms of monitoring, such as IT monitoring, which would operate similarly.

Parameters β, which control the extent to which (dis)honest behavior diffuses, subsume a wide array of mechanisms that have been shown to support the diffusion of (dis)honest behavior. As discussed in the Introduction, these mechanisms may fall into three broad categories: (1) information diffusion, (2) complementarity or substitution effects, and (3) conformism. All of these mechanisms imply a form of diffusion ($\beta > 0$), except for substitution effects, which imply differentiation ($\beta < 0$).

The model yields a series of intuitive predictions that allow discussing policies for curbing dishonest behavior. First, holding everything else constant, increasing the diffusion rate of honest behavior (β_p) or the baseline propensity towards honest behavior (α_p) will increase the share of honest individuals in the organization, while increasing the diffusion rate of dishonest behavior (β_q) or the baseline propensity towards dishonest behavior (α_p) will decrease that share (Online Appendix (OA), Proposition 3). As such, the faster dishonest behavior spreads relative to honest behavior (i.e., $\beta_p \gg \beta_q$), the larger the cost to the firm and the higher managerial willingness to pay for such policies. Furthermore, policies that improve the organizational context (i.e., increase α_p or decrease α_q), such as ethical training programs or IT monitoring also curb dishonest behavior.

Additionally, in an organization in which dishonest behavior diffuses while honest behavior is repellent (i.e., $\beta_p \leq 0, \beta_q \geq 0$), increasing network density (i.e., making the network more connected by increasing the number of links) will increase the share of dishonest individuals in the organization, while making the organization less connected will decrease that share

[2] Following Jackson and Rogers (2007); López-Pintado (2008); Lamberson (2010); Tarbush and Teytelboym (2017), we assume independence of degree across nodes and use a mean-field approximation. That is, we assume that the share of infected agents among one's neighbors matches the population average and approximate the variation of the share of dishonest agents in the population using a deterministic continuous-time process.

(OA, Propositions 4). In an organization in which both honest and dishonest behavior diffuse (i.e., $\beta_p, \beta_q \geq 0$), increasing (decreasing) density will decrease (increase) the share of dishonest individuals if, initially, there are sufficiently many honest individuals (OA, Proposition 5).

The model makes a series of simplifying assumptions: it assumes that honest and dishonest behaviors are binary and considers a linear functional form for the diffusion of both. The first assumption follows from our main question – which of honest or dishonest behavior spreads faster – which presumes binary forms of behavior. We discuss in Section 4 how this compounds with advantages in terms of empirical analysis. It also corresponds to a scope condition: the model considers the diffusion of a single (dis)honest action, a condition that matches our context reasonably well (see next section). A context in which two (dis)honest actions coexist would correspond to two instances of the same model, each with different parameter values. Assuming linear functional forms simplifies exposition and interpretation, allowing to neatly separate diffusion-related parameters β from non-diffusion-related parameters α. The specifics of the functional form matter, however, as different functional forms may capture different mechanisms of diffusion (Boucher et al., 2024), a fact we will leverage to characterize the mechanism underlying diffusion (Section 5.2).

3 Context and data

The data describes the daily operations of a call-center company based in the MENA region, between November 1st, 2016, and August 31st, 2018. It is a backup of the company's internal software.

The data is a backup from the company's internal software. For every 704,800 claims awarded between November 1, 2016, and August 31, 2018, the data record the timestamp of allocation, the clerk that processed the claim, the service provider who was awarded the claim, the monetary value of the claim which will be paid to the provider assigned to service the claim, and the *market* to which the claim belongs, where a market is defined as a type of service (say, tow-truck or ambulance) and a city in which the service is to be delivered. Our analysis is restricted to a subset of markets that are sufficiently large to lends themsselves to the construction of a statistical measure of dishonest behavior (Section 3.3). Most markets are small (the top 20% markets in terms of revenue account for 82% of the company's total revenue). Out of this subset, we consider the markets in which an automated claim allocation system was rolled out (Section 3.1) and at least two service providers operated in any given month. Table 1 provides a series of descriptive statistics on the entire sample and describes the subset considered for analysis.

The remainder of this section provides details about the company, including a description of dishonest behavior under consideration, the incentives underlying such behavior, and the natural experiment we leverage for identification. We then describe the construction of our two main quantities of interest – social ties and dishonest behavior, – discuss the limitations of these measures, and describe our validation procedures.

3.1 Context

The company and its industry. The roadside and health assistance industry specializes in delivering prompt and reliable assistance to individuals who encounter vehicle breakdowns or accidents, as well as health-related emergencies. As is standard in the industry, the company we consider maintains a network of service providers, including tow trucks, mechanics, ambulance services, and medical professionals, to ensure swift response times and comprehensive assistance in all markets. It offers subscription-based models whereby customers pay a monthly fee to gain access to their services. It operates a call center in which clerks handle incoming calls and coordinate assistance. Customers can reach out to the call center on the phone to request

	Full sample	Subset
Claim characteristics		
Opt-out rate*	.34	.3
Mean number of draws†	1.32	1.44
Number of claims	704,800	367,128
Revenue (m$)	35.78	15.18
Daily number of claims	1,054	549
Daily revenue (k$)	53.49	22.69
Market characteristics		
Number of markets	805	75
Mean monthly market revenue (k$)	2.78	9.23
Firm characteristics		
Number of firms	971	344
Mean monthly firm revenue (k$)	1.49	2.14
Clerk characteristics		
Number of clerks	423	406
Mean monthly clerk revenue (k$)	.7	1.26
Mean daily number of claims	9.79	7.09
Mean daily number of markets	6.87	4.98
Monthly wage ($)‡§	427	-
Percent females‡	.57	-
Age‡	28.29	-
Turnover¶	.48	-
Period		
Beginning	2016-11-01	2016-11-01
End	2018-08-31	2018-08-31

* Percentage of claims for which clerks opted out of the random draw.
† Average number of draws required to award claims when clerks did not opt out of the random draw.
‡ Communicated by management
§ Amounts to 1.6 times private sector minimum wage
¶ Computed for year 2017

Table 1: **Descriptive statistics.** The subset column considers markets pertaining to the top 20% in terms of revenue, in which the firm rolled out a new system for claim allocation, and where at least 2 firms operate on any month. Markets from the subset account for 42 percent of the company's total revenue over the period.

assistance. Clerks dispatch the appropriate service providers to the location of the incident.

A key feature of business operations is that the company compensates service providers for service provision. This creates a potential for dishonest behavior, as clerks may collude with service providers to transfer additional revenue to them in return for a kickback.

Call-center clerks and market assignment. Call-center clerks face a fast-paced environment with few possibilities for career progression and a highly irregular work schedule. Their salary, at 1.6 times the minimum wage (Table 1), is relatively low given that the position requires foreign language proficiency. As a result, turnover is high, reaching 48% in 2017.

Clerks operate on a single floor (Figure 1). The floor is divided into two divisions that handle services pertaining to health and roadside assistance, respectively. Clerks are assigned to one division, within which they typically handle several markets (7 per day on average).

Market assignment is primarily managed by software, with mid-management making marginal adjustments. The software assigns clerks to shifts and markets. Using historical data, it estimates staffing needs and allocates clerks based on skills and fairness, ensuring (1) that clerks are trained for their assigned market, and (2) equitable distribution of day/night

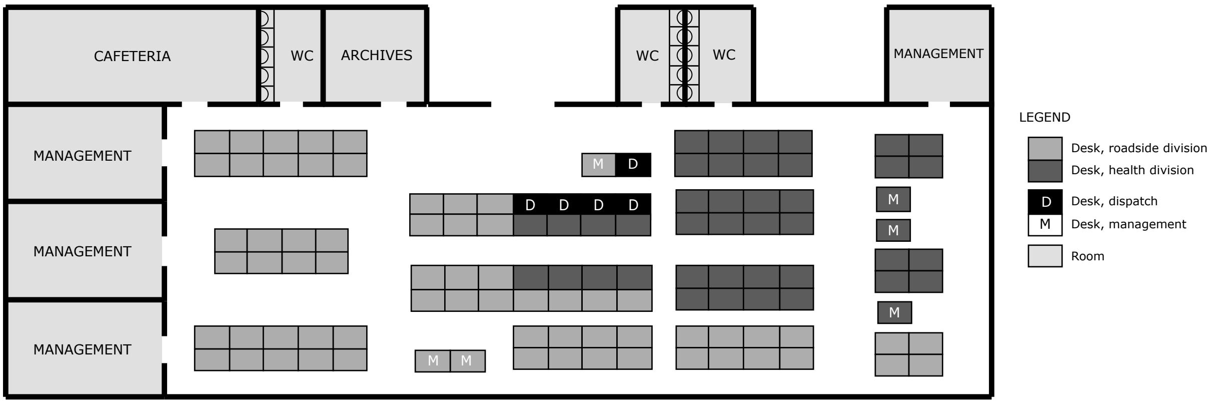

Figure 1: **Schematic floor plan of the call-center.** Desks sit one person. The space has an area of 4,725 squared feet. Clerks operate in a single room. They are separated in two functionally and spatially distinct divisions.

and week/weekend shifts while maintaining the required number of weekly hours. Management manually adjusts assignments for unforeseen events like severe weather or sick leaves. In Section 4, we discuss the implications of this policy for the identification of peer effects.

Assigning claims to providers. Maintaining a reliable network of service providers is crucial for the company, as its business model depends on swiftly servicing claims with high-quality service. This creates a tradeoff: prioritizing quality requires selecting top-performing providers, while ensuring availability demands redundancies, potentially lowering service quality. Management sets a monthly revenue distribution balancing quality and availability. More claims go to high-performing providers to maintain quality, while some revenue is allocated to lower-performing providers to prevent their defection and ensure their availability.

The revenue distribution is enforced by a software solution that governs claim allocation. The software guides clerks by randomly selecting a provider based on management-defined weights. Clerks must contact the provider; if unavailable, they log a reason and request another draw, averaging 1.4 draws per claim in the analyzed markets. In 30% of cases, clerks manually opt out of the random draw, e.g., due to specialized equipment needs that can only be fulfilled by specific providers. The analysis begins in November 2016, when this system was implemented.

Dishonest behavior. In 2018, we interviewed the call center's top manager to understand dishonest behavior within the call center. Having started as a clerk in the early 2000s, they[3] had witnessed only two detected cases of dishonesty, both predating the implementation of the random draw system.

One case, from the early 2010s, involved a clerk soliciting small loans from providers, later returning favors by allocating more claims. As the manager explained: "*[The culprit] needed money. So, they went on calling providers to ask 'Please, my mom is sick. I need money. Could you lend me $100?' They reached out to somewhere between 30 and 40 providers [and gathered about $4,000 in the process]. [Providers] wait for a few weeks, call them back, and ask 'Forget about the $100, you can keep them. That said, I have a problem. I noticed I'm getting very few claims. Can you help?' So, for $100, [the provider] gets more claims. [They] give [the clerk] another $100 and [collude]. And then, that clerk can help [the provider collude] with more clerks.*"

The anecdote illustrates that preferential claim allocation is the main form of dishonest

[3]We use the gender-neutral pronoun "they" to conceal the manager's identity.

behavior. Creating false claims is not a concern, as this behavior is heavily policed through technology, systematic auditing, and harsh sanctions. Management is, however, concerned about collusion through preferential claim allocation. The random draw system introduced in 2016 made this behavior more difficult, but not impossible. Indeed, clerks may still manipulate allocations by falsely marking providers as unavailable to favor colluding providers. Detecting such behavior is difficult, as verifying availability is costly, with over 1,000 daily claims and no automated solution.

The anecdote also underscores the role of social ties in both enabling and preventing collusion. Ties between management and providers facilitate enforcement. Indeed, providers often inform management about dishonest behavior they have observed in clerks or other providers. About 20 providers reported this incident, aiding in gathering evidence and sanctioning those involved. However, clerks rarely report misconduct by colleagues, presumably because clerks' dishonest behavior can be concealed from peers and supervisors.

Ties among clerks may also facilitate collusion through a variety of mechanisms. This incident suggests an informational mechanism: the manager feared that the dishonest clerk could have recruited others by sharing information about colluding providers. The second case, from the early 2000s, where three clerks fabricated claims for kickbacks,[4] suggests homophily, as these clerks had a romantic relationships. In Section 5.2, we show that the data align more with the informational mechanism.

3.2 Measuring social ties

We construct networks of social interactions among clerks using co-staffing patterns. A clerk is considered staffed on a market at a given hour if she awarded at least one claim there. Clerks i and j are co-staffed if both were staffed on the same market at the same hour (Figure 2).

Our baseline definition considers two clerks connected on a market-month if they co-staffed that market for at least three hours. Increasing this threshold captures stronger forms of interactions but reduces overall connectivity, limiting our ability to capture peer effects (Figure 3).

This measure understates interaction time, making networks appear sparse. Clerks handle 10 claims daily across seven markets in an 8-hour shift (Table 1). Since claims take time to process, two clerks may work on the same market for most of their shift yet rarely assign claims within the same hour, leading to lower measured interaction.

To ensure robustness, we explore alternative constructions of social ties. Our main result (Figure 5) is robust to varying the co-staffing threshold (Figure OA23). It is not robust to considering broader forms of interactions, such as being co-staffing on the same division or being co-staffed at all (Figure OA12). This suggests that markets are the relevant unit for social interactions.

We validate our measure against self-reported ties from an employee survey we conducted in January 2017. While co-staffing time weakly correlates with self-reported ties, the correlation is stronger for market- and division-level interactions (OA, Section 2). Our results are robust to replacing co-staffing time with self-reported ties (Figure OA5).

3.3 Measuring dishonest behavior

The random-draw system introduced in 2016 (see Section 3.1) provides empirical traces of collusion, which we leverage to derive a measure of dishonest behavior. This measure combines three necessary but individually insufficient conditions. Consider a scenario where clerk i colludes with provider k while most other clerks remain honest:

[4]False claims were not closely monitored at the time.

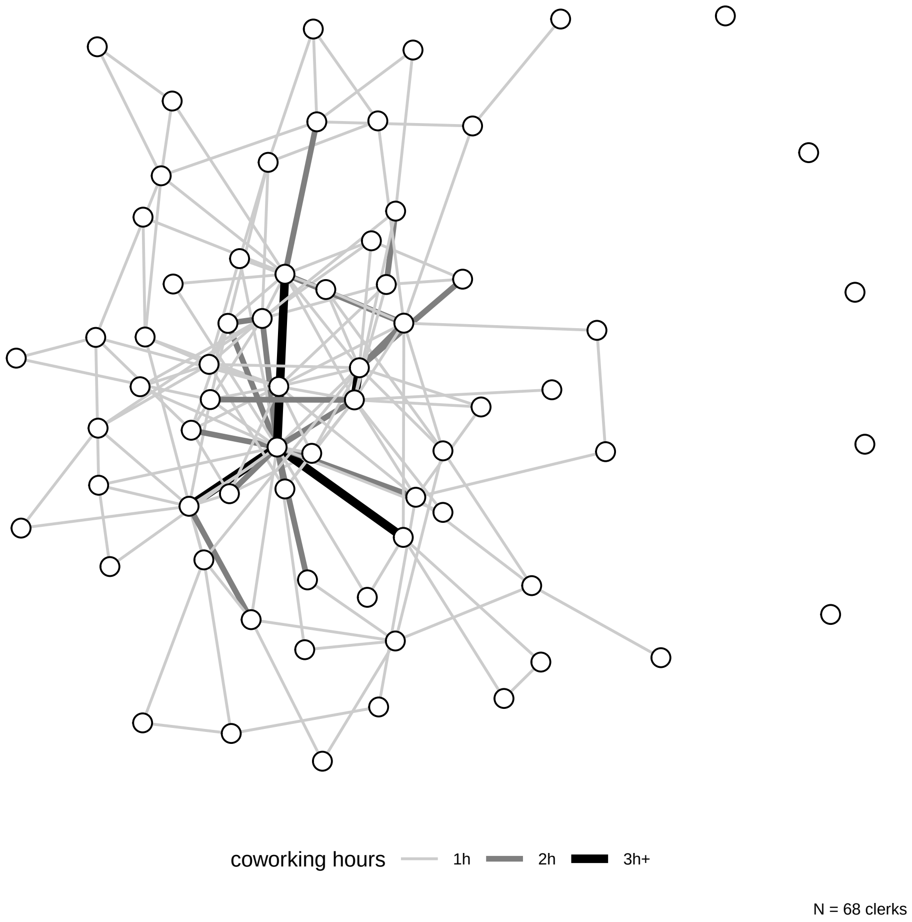

Figure 2: **Co-staffing network on a market-month.** Nodes represent clerks. Links indicate co-staffing relationships.

1. **Absolute revenue.** Clerk i must transfer a large amount of revenue to provider k to ensure sufficient kickbacks. However, high revenue alone does not indicate collusion, as i may simply be highly active or operating in a high-value market (e.g., replacement vehicles).

2. **Revenue distribution.** Clerk i must allocate disproportionately large shares of revenue to provider k compared to other clerks. This condition captures the core risk of collusion identified by management; i.e., preferential allocation of claims that deviates from company policy.[5] However, high variance in claim values (e.g., towage distance) can lead to excess revenue allocations for non-collusive reasons.

3. **Gaming of the random draw system.** Clerk i should be more likely than others to opt out of the random draw or show long sequences of such draws when allocating claims to provider k. This condition is necessary, as short sequences of random draws ensure compliance with company policy. However, opting out can occur for legitimate reasons, such as customer-specific needs (e.g., requiring specialized equipment).

We construct a measure of dishonest behavior $s_{ikmt} \in [0, 1]$ for the dyad between clerk i and provider k in market m during month t. This measure, the *s-score*, is defined as $s_{ikmt} =$

[5] A more direct approach would compare realized allocations to company policy, but deviations frequently occur for admissible reasons, such as reduced provider availability at night.

Figure 3: **Degree distribution under different tie thresholds.** Degree represents the number of links per individual. Each panel title indicates the co-staffing threshold used.

$s^1_{ikmt}s^2_{ikmt}s^3_{ikmt}$, where each component $s^j_{ikmt} \in [0, 1]$ operationalizes one of the three conditions. It equals 1 when all three conditions are fully met and 0 when at least one is fully not met.

For condition 1, we set $s^1_{ikmt} = 1$ if the revenue allocated by i to k exceeds \$427 (the average clerk wage) and 0 otherwise. Lower thresholds increase the rate of false positives, capturing revenues that are in fact too low to justify kickbacks, while higher thresholds increase the rate of false negatives and decrease sample size. Our threshold is conservative, as \$427 corresponds to the 96th percentile of the clerk-provider revenue distribution (Figure OA6). We vary this cutoff in robustness checks (see Section 4 for details).

We operationalize conditions 2 and 3 using null models that assume that all clerks are honest. We then look for upward deviations from the null using concepts related to the notion of p-values. Supplementary Materials Section 3.1 details model covariates and estimation procedures.

For condition 2, we estimate a null multinomial model for each market-month, predicting the probability of allocating a claim to provider k. This yields a null distribution for the revenue R_{ik} allocated by i to k. The probability $\Pr(R_{kimt} > r_{kimt})$ of observing a higher revenue than the actual revenue r_{ik} under the null is a one-tailed p-value. We define $s^2_{ikmt} = 1 - \Pr(R_{kimt} > r_{kimt})$. Higher values of s^2_{ikmt} indicate greater deviation from expected allocations, suggesting potential dishonest behavior.

For condition 3, we estimate a null logistic model for each provider-market-month predicting the probability of a suspicious allocation (i.e., after an opt-out or after three or more draws, knowing that over 90% of random-draw claims are allocated in two draws or fewer). With F_{ik} the fraction of suspicious allocations by i to k under the null and f_{ik} the observed fraction, we define $s^3_{ikmt} = 1 - \Pr(F_{ikmt} > f_{ikmt})$.

Figure 4 shows the s-score distribution for dyads meeting condition 1. The distribution is left-skewed, indicating that most dyads are honest (Panel a). Components s^2_{ikmt} and s^3_{ikmt} are weakly correlated ($\rho = .11$), confirming the necessity of both conditions (Panel b).

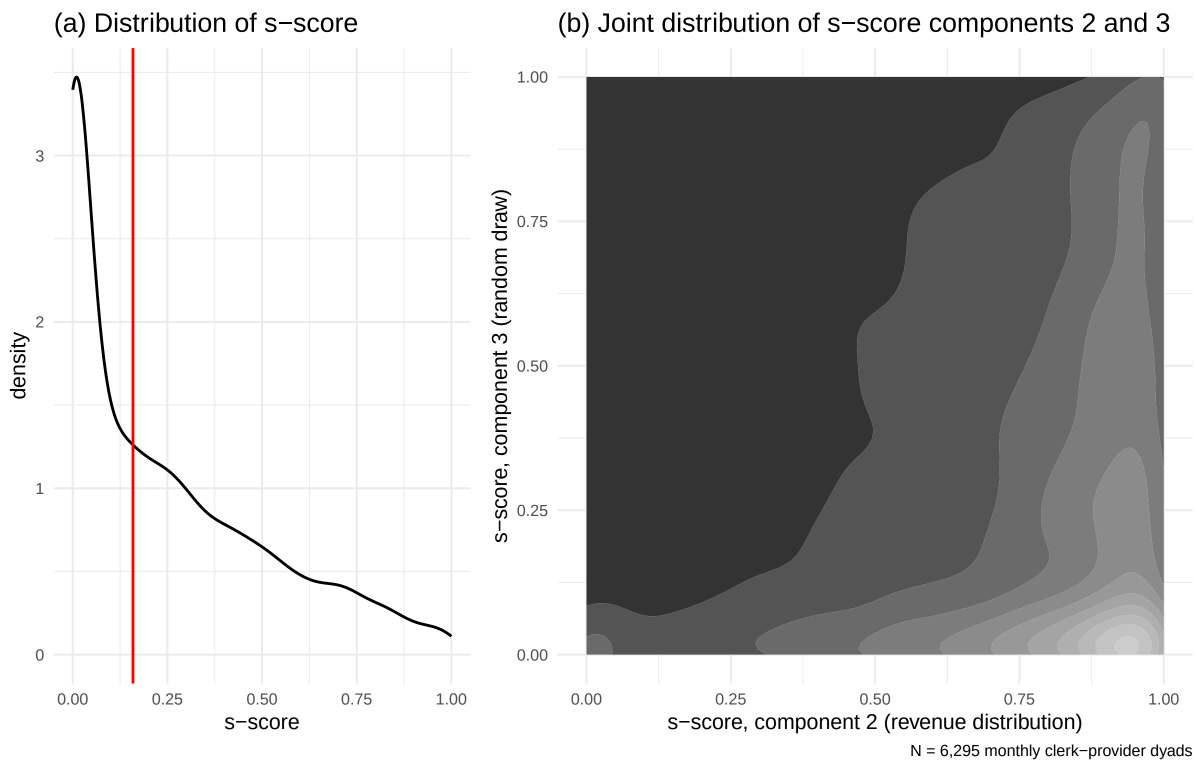

Figure 4: **Distribution of s-scores for dyads meeting condition 1.** Panel a shows the distribution of s_{ikmt}, with the solid line indicating the median. Panel b shows the joint distribution of s^2_{ikmt} (x-axis) and s^3_{ikmt} (y-axis), with lighter areas indicating higher mass.

The s-score has a series of advantages. The measure is likely comprehensive as it captures claim misallocation, which is the main risk of dishonest behavior identified by management (Section 3.1). Since claim allocation is clerks' only task, the measure picks up dishonest behavior consistently across clerks. The measure is non-invasive: as it was constructed ex-post, clerks did not know that their behavior was subjected to additional monitoring, which could have led them to alter their behavior. Finally, we can leverage the null models associated with the s-score to derive the monetary cost of misallocation (OA, Section 3.3).

The s-score's main limitation is that it understates dishonest behavior. In markets where most clerks behave dishonestly, the null behavior benchmark is itself dishonest, making it harder to detect deviations. This conservatism, however, increases confidence that high s-scores indicate genuine misconduct.

Validity checks confirm that the s-score captures dishonest behavior (OA, Section 3.2). Direct validation–i.e., linking s-scores to observed kickbacks–was not feasible. Instead, we first show that the metric captures behavior that deviates from the organizational norm. We then demonstrate that the rule-breaking associated with high s-scores is more consistent with private gains than with alternative explanations such as inexperience, "well-intentioned" rule-breaking, or clerk-provider complementarities that would facilitate cooperation.

4 Identification strategy

Slightly adapting notation from our theoretical model (Section 2), let $y_{imt} = 1$ if clerk i behaves dishonestly in market m during month t, and $y_{imt} = 0$ otherwise. We estimate the following auto-regressive linear probability model using ordinary least squares:

$$z_{imt+1} = \alpha_m + \alpha_t + \beta y_{imt} + \gamma_0 n_{imt} + \gamma_1 n_{imt} y_{imt} + \delta x'_{imt} + \epsilon_{imt}, \tag{3}$$

where $z_{imt+1} \equiv \mathbb{1}\{y_{imt+1} \neq y_{imt}\}$ is an indicator for whether i's behavior changes. We convert continuous s-scores s_{ikmt} into the binary measure y_{imt} by considering all providers i interacted with in m during t and setting $y_{imt} = 1$ if at least one provider's s-score exceeds 0.5. This threshold is chosen because it yields an estimated share of misallocated revenue close to 5%, the median global cost of fraud (Association of Certified Fraud Examiners, 2018, see OA, Section 3.3 for details).[6] The variable n_{imt} captures the number of i's neighbors in m whose behavior differs from i's (i.e., the number of honest neighbors if i is dishonest, and vice versa). The model includes market and month fixed effects α_m and α_t, a vector of control variables x_{imt},[7], and an error term ϵ_{imt} clustered at the market and month levels.

The key parameters are γ_0 and $\gamma_0 + \gamma_1$, capturing, respectively, the effect of a dishonest neighbor on an honest agent turning dishonest, and the effect of an honest neighbor on a dishonest agent turning honest. This model aligns closely with Equations (1) and (2) of the theoretical framework (Section 2): by definition, $n_{imt} = \theta_{imt} d_{imt}$ if $y_{imt} = 1$ and $n_{imt} = (1 - \theta_{imt}) d_{imt}$ otherwise, and we obtain $\gamma_0 = \beta_q$ and $\gamma_0 + \gamma_1 = \beta_p$. When control variables x_{imt} are omitted, the empirical model matches the theoretical model exactly, as we recover $\alpha_q = \alpha_m + \alpha_t$ and $\alpha_p = \alpha_m + \alpha_t + \beta$. Consistent with the theoretical framework, Equation (3) assumes that (dis)honest behavior is binary and spreads according to a linear function. Besides the advantages discussed in Section 2, considering binary behavior reduces noise in our measure of dishonest behavior.

We argue that, in our context, Equation (3) addresses the sources of bias that may prevent the identification of peer influence (VanderWeele and An, 2013). Market- and month-level fixed effects absorb contextual shocks that may affect the outcomes of i and her peers (Fowler et al., 2011). The use of lagged peer behavior addresses the reflection problem (Manski, 1993). Management's market-assignment policy (Section 3.1) should address homophily. Since staffing decisions are made algorithmically, without considering dishonest behavior, market-level social ties should be exogenous to individual behavior (Assumption 1). Thus, if peer influence flows through market-level social ties rather than other social ties (Assumption 2), then homophily should not bias estimates of parameters γ_0 and γ_1. We provide direct evidence supporting Assumptions 1 and 2 (OA, Section 4.2), showing that (i) the employment spells of dishonest clerks are not significantly shorter thant those of honest clerks, suggesting that management is unaware of dishonest behavior, (ii) that clerks enter new markets independently of their behavior, and (iii) that social influence flows through market-level social ties rather than broader forms of interaction. We also conduct statistical tests that show no support for homophily or reflection biases (OA, Sections, 4.3 and 4.4, respectively).

Another threat to identification is selection into dishonest providers. Equation (3) is evaluated at the clerk rather than the clerk-provider level. Indeed, Assumptions 1 and 2 guarantee that clerk-to-clerk, rather than clerk-to-provider interactions are independent of clerks' behavior. However, it could be that dishonest providers approach a set of clerks and collude with them separately, with between-clerks interactions playing no role. Thus, observed peer effects could reflect clerks selecting into the same set of providers rather than direct clerk-to-clerk influence. To rule this out, we reestimate equation (3) at the provider level and hold the provider fixed using provider-level fixed effects. We show that peer effects persist, suggesting

[6] By construction, if clerk i has no provider with revenue above \$427, then $y_{imt} = 0$. Thus, we estimate the model only for clerks with at least one such provider at $t + 1$.

[7] Controls include measures of i's behavior in market m (hours worked, log-revenue processed, number of peers), overall behavior (log-revenue across markets, total hours worked, maximum s-score), and a market-level control; namely, total log-revenue. See OA, Section 4.1 for further details about the constructions of these controls and a correlation matrix.

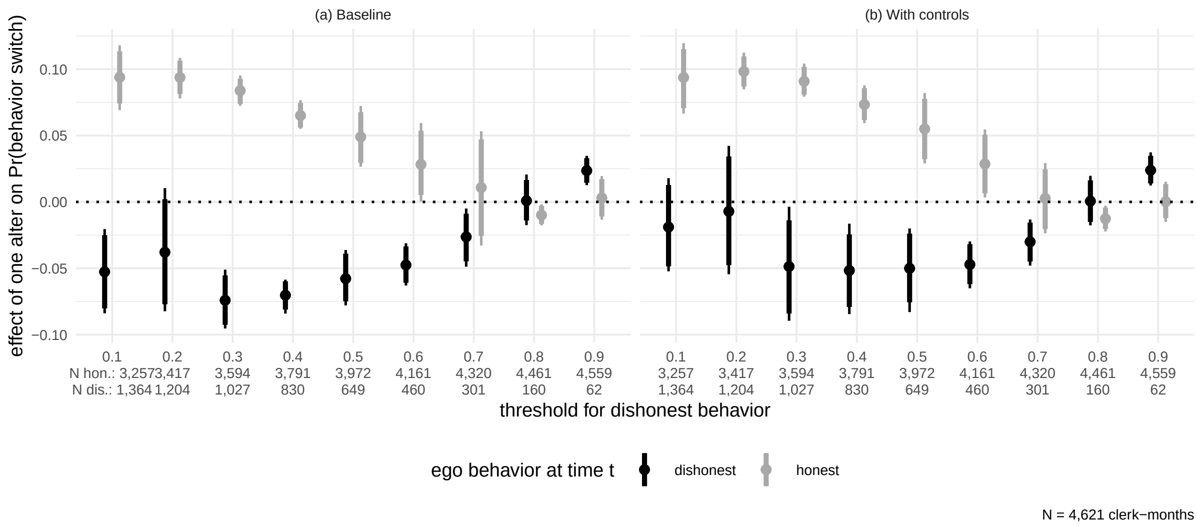

Figure 5: **Average marginal effect of a neighbor of opposite behavior on switching behaviors.** The x-axis represents the threshold in s-score used for behavior definition to estimate the model in equation (3), as well as the number of agents behaving honestly and dishonestly implied by such threshold. Points consider, for each model, agents that behave honestly (gray) or dishonestly (black) in period t and represent the average marginal effect of an additional neighbor with opposite behavior on the probability of the focal agent switching behavior; bars are 90 and 95% confidence intervals clustered at the month and market levels. All models include month and market fixed effects. Panel (a) reports models without controls, and panel (b) reports models with the controls discussed in footnote 7. Our preferred specification uses a cutoff of 0.5. Agents behaving honestly are influenced by agents behaving dishonestly, while agents behaving dishonestly weakly differentiate from agents behaving honestly. Results are qualitatively similar for a wide range of cutoffs.

that selection into provider cannot fully account for the result (OA, Section 4.5).

Finally, we test robustness to alternative data constructions. Our results remain consistent across different measures of dishonest behavior and of social ties, including considering each component of the s-score separately, varying the time threshold used to construct social ties, and considering strong and weak ties separately (OA, Section 5).

5 Results

This section presents estimates of peer effects in honest and dishonest behavior. We then provide evidence on the mechanism underlying peer effects. We finally conduct counterfactual exercises.

5.1 Peer effects in honest and dishonest behavior

We estimate peer effects in honest and dishonest behavior using Equation (3). Figure 5 reports parameters γ_0 and $\gamma_0 + \gamma_1$, corresponding to β_q and β_p in our theoretical model (Section 2). Parameter $\beta_q = \gamma_0$ (gray) measures the effect of peers behaving dishonestly on the probability that an agent who behaves honestly switches to dishonest behavior the next month. Parameter $\beta_p = \gamma_0 + \gamma_1$ (black) measures the effect of peers behaving honestly on the probability that an agent who behaves dishonestly switches to honest behavior the next month. Estimates are shown for varying s-score thresholds (0.1 to 0.9).

Using a 0.5 cutoff, we find that dishonest behavior spreads, while honest behavior is weakly repellent (i.e., $\beta_p \leq 0, \beta_q \geq 0$). Interacting with a peer that behaves dishonestly increases by 5.1 percentage points the likelihood that an agent that behaves honestly will behave dishonestly the following month. The same interaction reduces by 5.8 percentage points the likelihood that the agent behaving dishonestly will behave honestly the following month likelihood of behaving dishonestly (p-values < 0.01). Effect sizes remain stable across specifications (panels a and b) and s-score cutoffs, though results for thresholds above 0.7 suffer from power issues, as dishonest behavior is too rare using these cutoffs. While the diffusion of dishonest behavior is consistently observed in our robustness checks (Section 4), the repelling effect of honest behavior is weaker across tests.

Confronting these findings with our theoretical model implies that reducing network density (i.e., decreasing connectivity) should hinder the diffusion of dishonest behavior, thus increasing the prevalence of honest behavior (OA, Proposition 4). Note that this implication only holds if the linear count model used in the theoretical model is borne out empirically. Figures OA27 and OA28 show that the linearity assumption is reasonable. We show in the next section that the count model is reasonable.

5.2 Mechanism

We investigate the mechanism underlying the diffusion of dishonest behavior. Three potential explanations emerge from the literature (see Introduction for a brief review): (1) information diffusion, (2) complementarity/substitution effects, and (3) conformism.

First, results cannot be explained by substitution effects, as these would imply that increased exposure to peers behaving dishonestly would decrease the likelihood of dishonest behavior.

Second, we establish that the data is more consistent with complementarity that with conformism by comparing count models (based on the number of opposite-behavior peers) and percentage models (based on the share of opposite-behavior peers). Percentage models can capture both conformism and complementarity effects, while count models capture only complementarity effects (in line with Boucher et al. (2024); see OA, Section 6.1 for a discussion). Figure 6 shows that count models better fit the data, suggesting that complementarity effects explain diffusion better than conformism. However, our key insight that dishonest behavior diffuses remains robust to both specifications (Figure OA25).

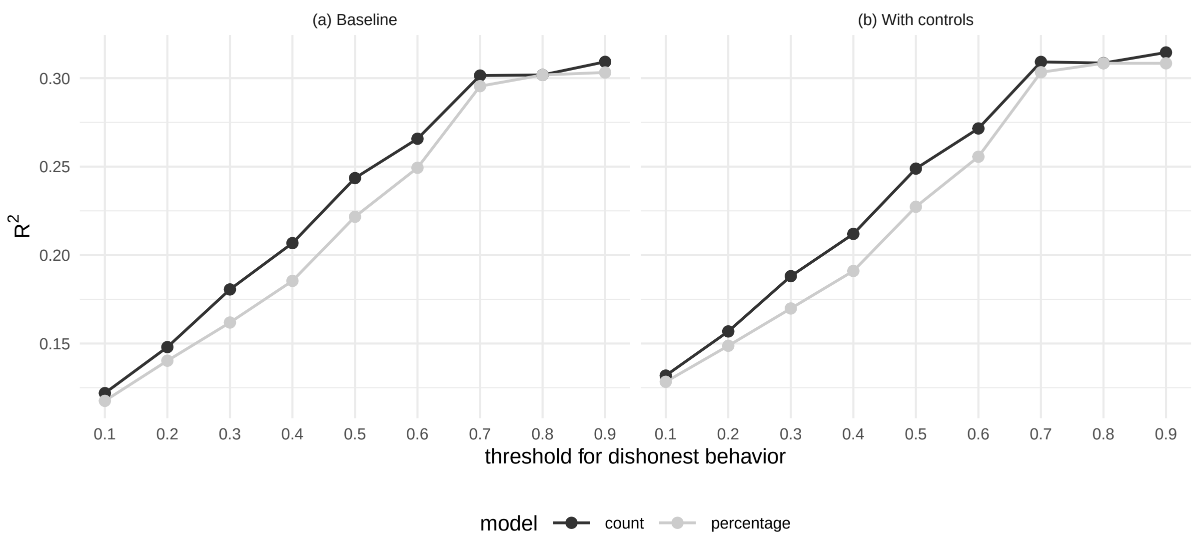

Figure 6: **Fit of count and percentage models.** This figure shows the R^2 of count models (Figure 5) and percentage models (Figure OA25). Count models have better fit, suggesting that complementarities better explain the observed patterns of diffusion than conformism.

We finally show support for information diffusion rather than complementarities. Figure 7 shows that clerks are most susceptible to peer influence in their second and third months in a new market. This suggests that clerks learn from their peers as they enter a market and may then decide to switch behavior. Their behavior then sticks with them over time. This is more consistent with information diffusion than with complementarities. Indeed, if peer effects were driven by complementarities, then they should remain stable over time.

Finally, we show evidence suggesting that social ties transmit information about providers' willingness to collude, echoing qualitative evidence from management interviews (Section 3.1). We estimate our main model at the provider level, distinguishing between providers that newly entered a market and established providers. We find evidence of peer effects in dishonest behavior among clerks only for providers that are already established (Figure OA26). This suggests that social ties transmit information about providers' willingness to collude: as new providers lack reputations, no information may transit.

5.3 Counterfactuals

We explore a series of interventions for curbing dishonest behavior. Given that our linear count model aligns well with the data (Sections 5.1 and 5.2), theoretical predictions (Section 2) apply directly, reducing the need for computational exploration.

Managers can intervene on the network (Valente, 2012) by altering (1) network nodes (e.g., screening clerks for honesty), (2) network links (e.g., modifying market-assignment policies to induce link formation), or (3) patterns of diffusion within links (i.e., influencing peer effects β_p and β_q). Managers can also intervene on the organizational environment (Palmer and Moore, 2016, 220-1, e.g., ethical training, IT monitoring), which would amount to adjusting intrinsic behavioral propensities (α_p and α_q). Simulations do not maniuplate parameters α and β from simulations, as finding sensible values for parameter calibration is not straightforward and the effects are intuitive: increasing the baseline propensity for (dis)honest behavior or its rate of diffusion increases its prevalence (OA, Proposition 3).

We consider two interventions that alter network nodes:

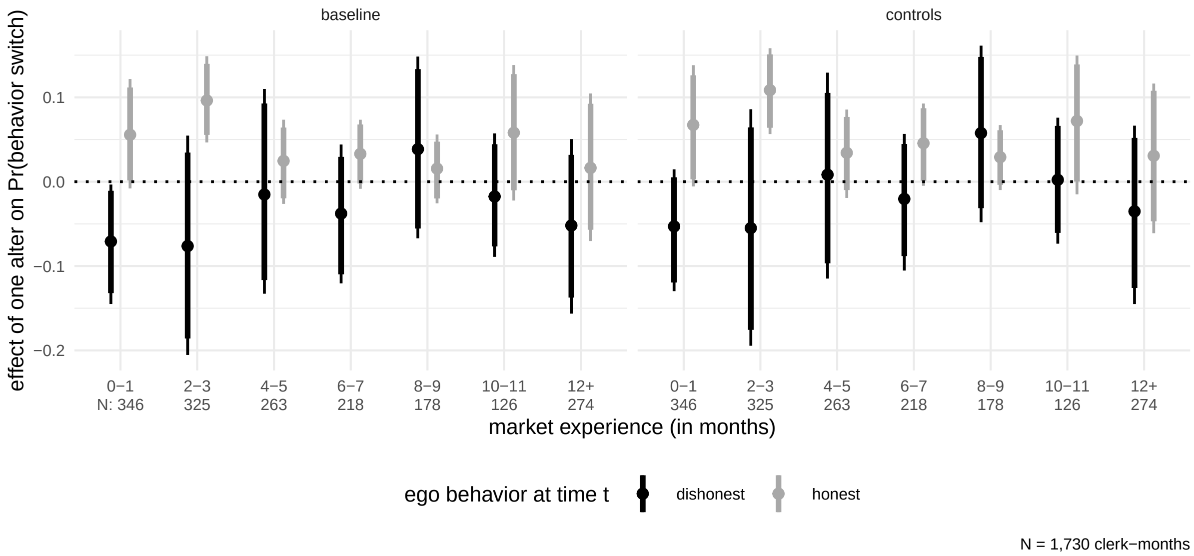

Figure 7: **Peer effects as a function of experience.** This figure amends our main specification (Equation (3)) by estimating peer effects separately as a function of market experience using interaction terms. We remove from the sample those clerks who entered markets prior to the beginning of our data, as their experience cannot be measured. The figure uses a threshold of 0.5 in s-scores to define dishonest behavior. Clerks are most sensitive to peer effects in their second and third months of experience.

1. **Reassignment policy.** At each time-period, dishonest clerks are shifted to the positions with the least number of links, similar to Lindquist et al. (2022). Note that this this policy requires complete knowledge of clerks' past (dis)honest behavior, which may prove costly and is not feasible in our setting.
2. **Limited screening policy.** New hires are screened for honesty before entering a market, implying that only honest clerks enter markets. This policy requires fewer information, as clerks' honesty is only measured once.

We compare these interventions against three benchmarks that manipulate network links:

3. **No interactions.** Remove all netowrk links. Since dishonest behavior diffuses while honest behavior is weakly repellent, this minimizes the prevalence of dishonest behavior (OA, Proposition 4).
4. **Full interactions.** Include all network links. This maximizes the prevalence of dishonest behavior (OA, Proposition 4).
5. **Business-as-usual.** Use the observed network structure.

We simulate these policies using OLS estimates derived from our specification without controls (Figure 5, left panel), ensuring alignment with theory. This has the additional advantage of avoiding having to set the endogenous control variables used in our specification with controls. We estimate the following transition probabilities:

$$\Pr(y_{imt+1} = 0 | y_{imt} = 1) = \underbrace{\hat{\alpha}_m + \hat{\alpha}_t + \hat{\beta}}_{\alpha_p} + \underbrace{(\hat{\gamma}_0 + \hat{\gamma}_1)}_{\beta_p} \underbrace{n_{imt}}_{(1-\theta_{it})d_i}$$

$$\Pr(y_{imt+1} = 1 | y_{imt} = 0) = \underbrace{\hat{\alpha}_m + \hat{\alpha}_t}_{\alpha_q} + \underbrace{\hat{\gamma}_0}_{\beta_q} \underbrace{n_{imt}}_{\theta_{it}d_i}$$

Similar to our reduced-form results (Section 5.1), we constrain clerks that handle too little revenue to be honest. All simulations are initialized using clerks' observed behavior at $t = 1$. Under limited screening, clerks entering markets at $t > 1$ are initialized with honest behavior, while other policies use their observed behavior. Our outcome of interest is the percentage of dishonest clerk-market-months. We estimate counterfactuals 100 times per policy and parameter configuration and report the median outcome as well as the 95% confidence interval. Our simulations vary network density by altering the threshold used to construct ties. We also vary the threshold for dishonest behavior using different cutoffs in s-scores. Using higher s-scores reduces the number of dishonest agents at initialization. Since our linear probability model may not imply well-defined predicted probabilities, we bound the predictions between 0 and 1 and show robustness to using logistic regression instead of OLS estimates (Figures OA29 and OA30).

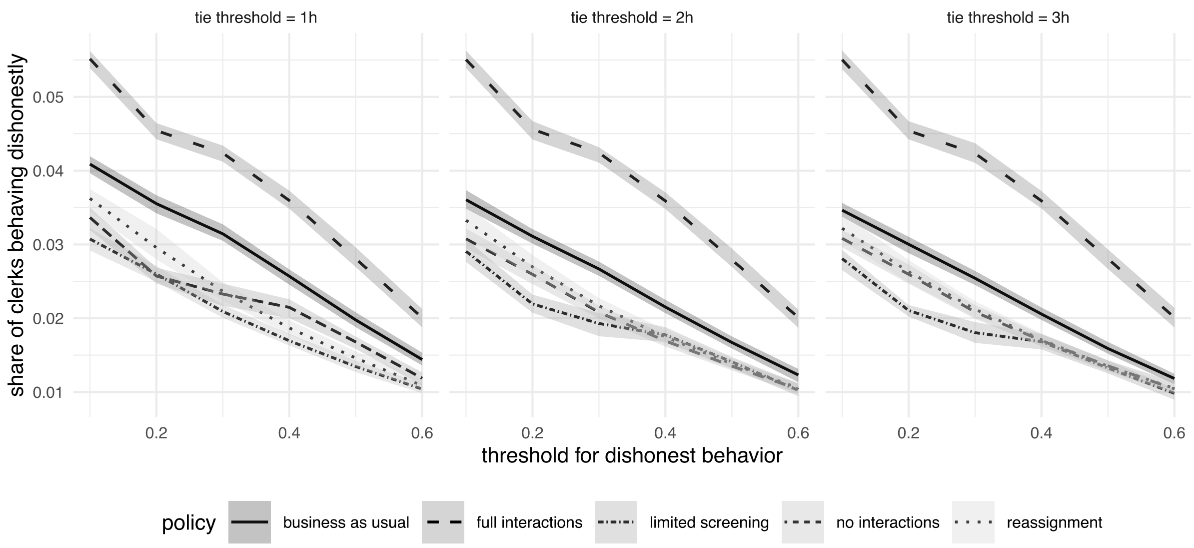

Figure 8: **Counterfactuals.** This figure shows the share of dishonest clerk-market-months in our counterfactual policies. Solid lines represent the median of 100 simulations. Ribbons represent the 95% confidence interval. The x-axis varies the threshold in s-scores used to define dishonest behavior at initialization. Panels vary the threshold in co-staffing time used to construct network links. The limited screening policy is most effective at curbing dishonest behavior when networks are sparse (high tie threshold) and dishonest behavior is more prevalent (low threshold in s-score). The reassignment policy is as effective as the no interactions policy when networks are sparse and less effective when networks are more dense.

Figure 8 shows the results. They confirm the ordering of our benchmarks: no interactions outperforms business-as-usual, which in turn outperforms full interactions. Using higher thresholds in s-scores reduces the prevalence of dishonest behavior, as simulations are initialized with fewer dishonest agents. The reassignment policy is nearly as effective as no interactions, except when networks are dense (tie threshold = 1h). In this case, isolating agents that behave dishonestly is not effective, as residual connectivity still enables diffusion. The limited screening policy performs best when networks are sparse (tie threshold > 1h) and dishonest behavior is initially high (low s-score threshold). Indeed, because of high turnover, dishonest agents are replaced by screened honest agents, while sparse networks limit contagion.

6 Discussion

We found that dishonest behavior diffuses while honest behavior does not (Section 5.1), echoing experimental evidence (Innes and Mitra, 2013; Dimant, 2019; Brunner and Ostermaier, 2019; Colzani et al., 2023) in field settings. The diffusion of dishonest behavior is best explained by information transmission (Section 5.2), aligning with prior field (Palmer and Yenkey, 2015; Alm et al., 2017) and experimental studies (Brunner and Ostermaier, 2019; Colzani et al., 2023).

We also highlight that screening for honesty can be effective in high-turnover environments with sparse networks. While high turnover has been shown to negatively affect firms (e.g., Li et al., 2022), this finding contributes to a growing body of evidence suggesting that it may also have positive effects, preventing the moral decay that may occur with longer tenure (Aven et al., 2021). Doing so, we complement recent work estimating the effect of counterfactual policies that alter network structure on other behaviors within the firm (Lindquist et al., 2022).

Managerial implications. The fact that dishonest behavior diffuses while honest behavior does not creates a tradeoff for organizations: enabling the transmission of productive behaviors and valuable information through peer relationships (Mas and Moretti, 2009; Azoulay et al., 2010; Oettl, 2012; Song et al., 2018) comes at the price of spreading dishonest behavior.

Our counterfactual analysis suggests that restructuring peer interactions to curb dishonest behavior is challenging. Since dishonest behavior is harder to track than productivity, policies requiring full behavioral monitoring may be impractical. Moreover, these interventions often require severing social ties, which may restrict the diffusion of beneficial behaviors.

This highlights the need for alternative strategies to mitigate dishonest behavior, such as strengthening ethical infrastructure, carefully screening new hires for honesty, leveraging supervisory and leadership relationships, and enhancing other forms of monitoring, such as IT-based solutions. These approaches can be further refined by acknowledging that dishonest behavior spreads faster than honest behavior. For example, ethical training programs could emphasize identifying collusion signals, reinforcing reporting mechanisms, and detecting early signs of dishonest behavior diffusion.

However, these insights may not generalize across settings. Managers should assess the extent and mechanisms underlying the diffusion of dishonest behavior within their organization to design tailored, effective policy responses.

Limitations and future directions. Our study faces four main limitations, which suggest avenues for future research:

1. **Measurement and inference constraints.** Dishonest behavior and social interactions are measured indirectly. While our natural experiment alleviates concerns regarding the identification of peer effects, bias cannot be fully ruled out.

2. **Generalizability concerns.** Most field evidence on the diffusion of dishonesty comes from Western contexts (e.g., Ichino and Maggi, 2000; Pierce and Snyder, 2008; Yenkey, 2018; Chan et al., 2021; Mohliver, 2019; Quispe-Torreblanca and Stewart, 2019; Frake and Harmon, 2024), raising questions about cross-cultural generalizability (Rahwan et al., 2019). Likewise, the mechanism underlying diffusion conditions the effectiveness of policy interventions. In our setting, decreasing network density should reduce dishonest behavior. Yet, when dishonest behavior poses a tradeoff between efficiency and secrecy, the same policy may instead increase dishonest behavior (Ferrali, 2020). Future studies should test whether our findings that dishonest behavior diffuses faster than honest behavior and that diffusion is supported by informational mechanisms generalize to other settings. We expect informational mechanisms to operate in settings where dishonest behavior requires

knowledge, secrecy imperatives are not too stringent, social interactions are dense enough, and the moral cost associated with the behavior is small enough.

3. **No network formation nor behavioral interactions.** Our study sidesteps the question of network formation and focuses on one form of dishonest behavior. Yet, evidence suggests link formation is endogenous to (dis)honest behavior, with different forms of dishonest behavior implying different patterns of link formation (Zhang and King, 2021; Zhang et al., 2023). Future research should explore how network self-selection influences the diffusion of dishonest behavior under multiple forms of dishonest behavior.

4. **Limited countrefactual policies.** Our exploration of counterfactual policies does not estimate the cost of alternative policies or their impact on other outcomes, such as the diffusion of productive behaviors, preventing from conducting accurate cost-benefit analyses. Likewise, we ignore individual-level heterogeneity in patters of diffusion. Recent research suggests that some agents act as "moral beacons" (Helzer et al., 2024) which, if strategically placed within the network, could act as firewalls. Future research could tackle these questions by considering settings that provide estimates of the cost of alternative policies, allow for measuring not only the rate of diffusion of (dis)honest behavior, but also of those other desirable outcomes, and for measuring individual-level heterogeneity.

References

Abbink, Klaus, "Staff rotation as an anti-corruption policy: an experimental study," *European Journal of Political Economy*, nov 2004, *20* (4), 887–906.

Alm, James, Kim M. Bloomquist, and Michael McKee, "When You Know Your Neighbour Pays Taxes: Information, Peer Effects and Tax Compliance," *Fiscal Studies*, dec 2017, *38* (4), 587–613.

Ariely, Dan, Anat Bracha, and Stephan Meier, "Doing Good or Doing Well? Image Motivation and Monetary Incentives in Behaving Prosocially," *American Economic Review*, feb 2009, *99* (1), 544–555.

Association of Certified Fraud Examiners, "Report to the Nations. 2018 global study on occupational fraud and abuse," Technical Report 2018.

Aven, Brandy, Lily Morse, and Alessandro Iorio, "The valley of trust: The effect of relational strength on monitoring quality," *Organizational Behavior and Human Decision Processes*, 2021, *166*, 179–193. Behavioral Field Evidence on Ethics and Misconduct.

Azoulay, Pierre, Joshua S. Graff Zivin, and Jialan Wang, "Superstar Extinction," *Quarterly Journal of Economics*, may 2010, *125* (2), 549–589.

Banerjee, Abhijit V. and Esther Duflo, "Reputation Effects and the Limits of Contracting: A Study of the Indian Software Industry," *The Quarterly Journal of Economics*, aug 2000, *115* (3), 989–1017.

Boucher, Vincent, Michelle Rendall, Philip Ushchev, and Yves Zenou, "Toward a general theory of peer effects," *Econometrica*, 2024, *92* (2), 543–565.

Brunner, Markus and Andreas Ostermaier, "Peer Influence on Managerial Honesty: The Role of Transparency and Expectations," *Journal of Business Ethics*, 2019, *154* (1), 127–145.

Chan, Tat Y., Jia Li, and Lamar Pierce, "Compensation and peer effects in competing sales teams," *Management Science*, 2014, *60* (8), 1965–1984.

_ , _ , and _ , "Learning from peers: Knowledge transfer and sales force productivity growth," *Marketing Science*, 2014, *33* (4), 463–484.

Chan, Tat Y, Yijun Chen, Lamar Pierce, and Daniel Snow, "The Influence of Peers in Worker Misconduct: Evidence From Restaurant Theft," *Manufacturing And Service Operations Management*, 2021, *23* (4), 952–973.

Chui, Celia, Maryam Kouchaki, and Francesca Gino, ""Many others are doing it, so why shouldn't I?": How being in larger competitions leads to more cheating," *Organizational Behavior and Human Decision Processes*, 2021, *164*, 102–115.

Colzani, Paola, Georgia Michailidou, and Luis Santos-Pinto, "Experimental evidence on the transmission of honesty and dishonesty: A stairway to heaven and a highway to hell," *Economics Letters*, 2023, *231*, 111257.

Conrads, Julian, Bernd Irlenbusch, Rainer Michael Rilke, and Gari Walkowitz, "Lying and team incentives," *Journal of Economic Psychology*, feb 2013, *34*, 1–7.

Dai, Zhixin, Fabio Galeotti, and Marie Claire Villeval, "Cheating in the Lab Predicts Fraud in the Field: An Experiment in Public Transportation," *Management Science*, jan 2017, *64* (3), 1081–1100.

Dimant, Eugen, "Contagion of pro- and anti-social behavior among peers and the role of social proximity," *Journal of Economic Psychology*, aug 2019, *73*, 66–88.

Ferrali, Romain, "Partners in crime? A theory of corruption as a criminal network," *Games and Economic Behavior*, 2020, *124*.

Fowler, James H., Michael T. Heaney, David W. Nickerson, John F. Padgett, and Betsy Sinclair, "Causality in Political Networks," *American Politics Research*, 2011, *39* (2), 437–480.

Frake, Justin and Derek Harmon, "Intergenerational transmission of organizational misconduct: Evidence from the Chicago Police Department," *Management Science*, 2024, *70* (6), 3856–3878.

Gerlach, Philipp, Kinneret Teodorescu, and Ralph Hertwig, "The truth about lies: A meta-analysis on dishonest behavior," *Psychological Bulletin*, jan 2019, *145* (1), 1–44.

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

Gross, Jörg, Margarita Leib, Theo Offerman, and Shaul Shalvi, "Ethical Free Riding: When Honest People Find Dishonest Partners," *Psychological Science*, dec 2018, *29* (12), 1956–1968.

Hasan, Sharique and Rembrand Koning, "Prior ties and the limits of peer effects on startup team performance," *Strategic Management Journal*, 2019, *40* (9), 1394–1416.

Helzer, Erik G, Taya R Cohen, Yeonjeong Kim, Alessandro Iorio, and Brandy Aven, "Moral beacons: Understanding moral character and moral influence," *Journal of Personality*, 2024, *92* (3), 735–752.

Ichino, Andrea and Giovanni Maggi, "Work Environment and Individual Background: Explaining Regional Shirking Differentials in a Large Italian Firm," *Quarterly Journal of Economics*, aug 2000, *115* (3), 1057–1090.

Innes, Robert and Arnab Mitra, "Is Dishonesty Contagious?," *Economic Inquiry*, 2013, *51* (1), 722–734.

Jackson, Matthew O and Brian W Rogers, "Relating Network Structure to Diffusion Properties through Stochastic Dominance," *The B.E. Journal of Theoretical Economics*, 2007, *7* (1), 1–13.

Lamberson, PJ, "Social Learning in Social Networks," *The B.E. Journal of Theoretical Economics*, 2010, *10* (1), 1935–1704.

Leib, Margarita, Nils Köbis, Ivan Soraperra, Ori Weisel, and Shaul Shalvi, "Collaborative dishonesty: A meta-analytic review.," *Psychological Bulletin*, 2021, *147* (12), 1241.

Li, Qin, Ben Lourie, Alexander Nekrasov, and Terry Shevlin, "Employee Turnover and Firm Performance: Large-Sample Archival Evidence," *Management Science*, 2022, *68* (8), 5667–5683.

Lindquist, Matthew J., Jan Sauermann, and Yves Zenou, "Peer Effects in the Workplace: A Network Approach," *SSRN working paper*, 2022.

López-Pintado, Dunia, "Diffusion in complex social networks," *Games and Economic Behavior*, 2008, *62* (2), 573–590.

Macchiavello, Rocco and Ameet Morjaria, "The Value of Relationships: Evidence from a Supply Shock to Kenyan Rose Exports," *American Economic Review*, sep 2015, *105* (9), 2911–45.

Manski, Charles F, "Identification of Endogenous Social Effects: The Reflection Problem," *The Review of Economic Studies*, jul 1993, *60* (3), 531.

Mas, Alexandre and Enrico Moretti, "Peers at work," *American Economic Review*, mar 2009, *99* (1), 112–145.

Mazar, Nina and Dan Ariely, "Dishonesty in Everyday Life and Its Policy Implications," *Journal of Public Policy & Marketing*, 2006, *25* (1), 117–126.

Mohliver, Aharon, "How Misconduct Spreads: Auditors' Role in the Diffusion of Stock-option Backdating," *Administrative Science Quarterly*, jun 2019, *64* (2), 310–336.

Oettl, Alexander, "Reconceptualizing stars: Scientist helpfulness and peer performance," *Management Science*, jun 2012, *58* (6), 1122–1140.

Palmer, Donald and Celia Moore, "Social networks and organizational wrongdoing in context," *Organizational wrongdoing: Key perspectives and new directions*, 2016, pp. 203–234.

_ **and Christopher B. Yenkey**, "Drugs, sweat, and gears: An organizational analysis of performance-enhancing drug use in the 2010 tour de France," *Social Forces*, dec 2015, *94* (2), 891–922.

Pierce, Lamar and Jason Snyder, "Ethical Spillovers in Firms: Evidence from Vehicle Emissions Testing," *Management Science*, nov 2008, *54* (11), 1891–1903.

Quispe-Torreblanca, Edika G. and Neil Stewart, "Causal peer effects in police misconduct," aug 2019.

Rahwan, Zoe, Erez Yoeli, and Barbara Fasolo, "Heterogeneity in banker culture and its influence on dishonesty," *Nature*, 2019, *575* (7782), 345–349.

Rose-Ackerman, Susan and Bonnie J Palifka, *Corruption and Government: Causes, Consequences, and Reform*, New York, NY: Cambridge University Press, 2016.

Song, Hummy, Anita L. Tucker, Karen L. Murrell, and David R. Vinsonc, "Closing the productivity gap: Improving worker productivity through public relative performance feedback and validation of best practices," *Management Science*, jun 2018, *64* (6), 2628–2649.

Tarbush, Bassel and Alexander Teytelboym, "Social groups and social network formation," *Games and Economic Behavior*, 2017, *103*, 286–312.

Treviño, Linda Klebe, Niki A den Nieuwenboer, and Jennifer J Kish-Gephart, "(Un)Ethical Behavior in Organizations," *Annual Review of Psychology*, 2014, *65* (1), 635–660.

Valente, Thomas W, "Network interventions," *science*, 2012, *337* (6090), 49–53.

VanderWeele, Tyler J and Weihua An, "Social Networks and Causal Inference," in Stephen L Morgan, ed., *Handbook of Causal Analysis for Social Research*, Springer, 2013, pp. 353–374.

Weisel, Ori and Shaul Shalvi, "The collaborative roots of corruption," *Proceedings of the National Academy of Sciences*, aug 2015, *112* (34), 10651–10656.

Yenkey, Christopher B., "Fraud and Market Participation: Social Relations as a Moderator of Organizational Misconduct," *Administrative Science Quarterly*, mar 2018, *63* (1), 43–84.

Zhang, Victoria, Aharon Cohen Mohliver, and Marissa King, "Where is all the deviance? Liminal prescribing and the social networks underlying the prescription drug crisis," *Administrative Science Quarterly*, 2023, *68* (1), 228–269.

_ **and Marissa D King**, "Tie decay and dissolution: Contentious prescribing practices in the prescription drug epidemic," *Organization Science*, 2021, *32* (5), 1149–1173.