Online Appendix for: "The Unequal Diffusion of Honesty and Dishonesty in Workplace Networks"

February 26, 2025

Contents

List of Figures

List of Tables

1 Theory

This section details the theoretical model mentioned in the paper.

1.1 Model

1.1.1 Additional assumptions about the setting

Since the graph g represents an organization, we assume that it is connected. Following standard practice, we also assume independence of degree[1] across nodes (Jackson and Rogers, 2007; López-Pintado, 2008; Lamberson, 2010; Tarbush and Teytelboym, 2017). This independence assumption has an important implication: it makes it easy to derive the *excess degree distribution*; that is, the degree distribution of a randomly chosen neighbor of a randomly chosen node. With $P(d)$ the probability that a randomly chosen node has degree d, $\bar{d}$ the mean degree in the network, and $\tilde{P}(d)$ the probability that a randomly chosen neighbor of a randomly chosen node has degree d, we have

$$\tilde{P}(d) \equiv \frac{P(d)d}{\bar{d}}$$

1.1.2 Steady states under the mean field approximation

To study the properties of this model, we make a standard simplification (Jackson and Rogers, 2007; López-Pintado, 2008; Lamberson, 2010; Tarbush and Teytelboym, 2017) and consider a *degree-based mean field approximation*. In other words, we solve the model assuming that all agents have the same neighbor infection rate at time t equal to mean neighbor infection rate in society: $\theta_{it} = \theta \equiv \mathbb{E}[\theta_{it}]$ for any $i \in N$. Let $\rho(d) \equiv \Pr(y_{it} = 1 | d_i = d)$ be the mean infection rate among agents of degree d. Note that ρ, the mean infection rate in society is $\rho = \sum_d P(d)\rho(d)$ and that

$$\theta = \sum_d \tilde{P}(d)\rho(d) \tag{1}$$

Using the mean field approximation allows analyzing the steady states of this dynamic process; that is, points at which the infection rate θ remains constant. At each time period t, a fraction $q(\theta, d)$ of non-infected agents of degree d becomes infected, and a fraction $p(\theta, d)$ of infected agents of degree d recovers. This pins down the law of motion of $\rho(d)$:

$$\frac{\partial \rho(d)}{\partial t} = [1 - \rho(d)]q(\theta, d) - \rho(d)p(\theta, d)$$

At the steady state, $\rho(d)$ remains constant. In other words, its law of motion satisfies

$$\frac{\partial \rho(d)}{\partial t} = 0 \iff \rho(d) = \frac{q(\theta, d)}{p(\theta, d) + q(\theta, d)} \tag{2}$$

Substituting this expression for $\rho(d)$ at the steady state into equation (1) gives $\theta = \sum_d \tilde{P}(d)\frac{q(\theta,d)}{p(\theta,d)+q(\theta,d)}$. Define $H(\theta) = \sum_d \tilde{P}(d)\frac{q(\theta,d)}{p(\theta,d)+q(\theta,d)}$. At the steady state, it must be that $H(\theta) = \theta$. In other words, steady states correspond to fixed points of H. Let $\mathcal{F} \equiv \{\theta : H(\theta) = \theta\}$ be the set of such fixed points. Since $H : [0, 1] \to [0, 1]$ is continuous, $\mathcal{F}$ is non-empty. Define $\overline{\theta} \equiv \max \theta \in \mathcal{F}$ and $\underline{\theta} \equiv \min \theta \in \mathcal{F}$ be the steady states that generate most and least infection, respectively.

Generically, there may be many steady states. However, some specific parameter values give us more traction. For instance, if there is no natural propensity to behave dishonestly, then a situation where everyone behave honestly is stable. In other words, if $\alpha_q = 0$, then there is a steady state at 0. The following proposition generalizes the intuition:

[1]The *degree* of node i on graph g is the number of neighbors of node i on graph g

Proposition 1 (Corner solutions)**.** $\underline{\theta} = 0 \iff \alpha_q = 0$ *and* $\overline{\theta} = 1 \iff \alpha_p = 0$.

Because situations in which all members of society behave similarly seem unrealistic, we focus on cases where $\alpha_p, \alpha_q \neq 0$. Additionally, note that when β_p and β_q have the same sign, our linear setup implies that H is monotonic, and its derivative is either concave or convex. In this case, there is a unique steady state. Formally:

Proposition 2 (Uniqueness)**.** *If* $\alpha_p, \alpha_q \neq 0$ *and* β_p *and* β_q *have the same sign, then* $H(\theta)$ *has a unique fixed point* $\theta^* = \underline{\theta} = \overline{\theta}$.

1.1.3 Comparative statics

Having a better sense of the properties of steady states, we can now analyze how they change across organizations. Let $f \equiv (g, \mu)$ be an organization, with $g = (G, N)$ its associated network, and $\mu = (\alpha_p, \alpha_q, \beta_p, \beta_q)$ a vector containing its associated parameters. We compare organizations f and f' that differ either with respect to their graph g or their parameter values μ. In what follows, notations use the $'$ symbol to refer to f'.

In order to make such comparisons, we must be able to compare across organizations that may each have multiple steady states. We say that an organization is more honest than another if both its maximal and minimal steady states sustain less dishonest behavior than the other. Formally:

Definition 1 (Honest organizations)**.** We say that organization f' is more *honest* than organization f and write $f' \succeq f$ when $\underline{\theta}' \leq \underline{\theta}$ and $\overline{\theta}' \leq \overline{\theta}$.

A simple lemma gives a lot of traction. If the H function of organization f' is above that of organization f, then f' is less honest than f. In other words, its rightmost and leftmost fixed points shift to the right. We introduce the result in the following lemma and provide an illustration in Supplementary Figure OA1:

Lemma 1 (Comparing organizations)**.** *If* $H_{f'}(\theta) \geq H_f(\theta)$ *for any* $\theta \in [0, 1]$*, then* $f' \preceq f$.

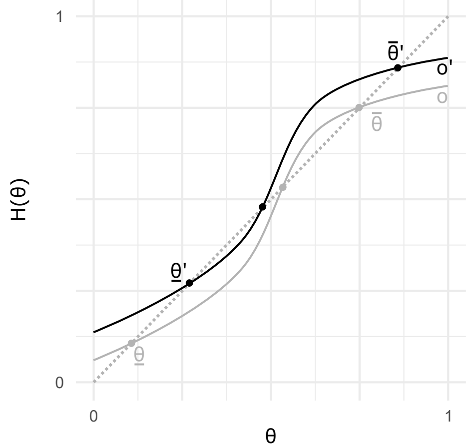

Figure OA1: **Illustration of lemma 1.** Because $H_{o'}$ is above H_o, its rightmost and leftmost fixed points are to the right of those of H_o.

A more natural definition would consider the percentage of agents behaving dishonestly at the steady state instead of the neighbor infection rate. Working with the more natural definition makes the problem less tractable. In the next section, we discuss the differences between the two definitions, explain how both definitions are similar for the kinds of networks considered in the empirical analysis, and show that the most important theoretical propositions are somewhat robust to the alternative, more natural definition.

Together, definition 1 and lemma 1 allow comparing organizations that differ only by one of their parameters $\phi \in \mu$. The following proposition gives the result:

Proposition 3. *Suppose organizations f and f' only differ by the value of one of their parameters $\phi \in \mu$. We have:*

$$\begin{aligned} \alpha_p' \geq \alpha_p &\iff f' \succeq f \\ \beta_p' \geq \beta_p &\iff f' \succeq f \\ \alpha_q' \geq \alpha_q &\iff f' \preceq f \\ \beta_q' \geq \beta_q &\iff f' \preceq f \end{aligned}$$

Finally, let's compare organizations that have the same parameter values μ, but differ in their network g. Two networks can differ in many ways. We consider one specific kind of variation, namely a shift in the degree distribution such that the degree distribution on f' first-order stochastically dominates (FOSD) the degree distribution on f,[2] implying that the mean degree on f' is higher than on f. The following two propositions give the results:

Proposition 4 (Increasing mean degree 1)**.** *Consider organizations o and f' with excess degree distributions $\tilde{P}$ and $\tilde{P}'$ respectively such that $\tilde{P}'$ FOSD $\tilde{P}$. If $\beta_p < 0$ and $\beta_q \geq 0$, then $f' \preceq f$. If $\beta_p \geq 0$ and $\beta_q < 0$, then $f' \succeq f$.*

Proposition 5 (Increasing mean degree 2)**.** *Consider organizations f and f' with excess degree distributions $\tilde{P}$ and $\tilde{P}'$ respectively such that $\tilde{P}'$ FOSD $\tilde{P}$. If $\beta_p \geq 0$ and $\beta_q \geq 0$, or $\beta_p < 0$ and $\beta_q < 0$, then there is a threshold $\hat{\theta} \in (0,1)$ such that*

$$\begin{cases} f' \succeq f, & \textit{if } \beta_p, \beta_q \geq 0 \textit{ and } \theta^* < \hat{\theta} \\ f' \succeq f, & \textit{if } \beta_p, \beta_q < 0 \textit{ and } \theta^* \geq \hat{\theta} \\ f' \preceq f, & \textit{if } \beta_p, \beta_q \geq 0 \textit{ and } \theta^* \geq \hat{\theta} \\ f' \preceq f, & \textit{if } \beta_p, \beta_q < 0 \textit{ and } \theta^* < \hat{\theta} \end{cases}$$

Together, propositions 4 and 5 show that the optimal organization differs widely depending on parameter values (see figure OA2 for a graphical summary). A maximally connected organization is optimal if (1) $\beta_p \geq 0$ and $\beta_q < 0$, or (2) $\beta_p \geq 0$, $\beta_q \geq 0$, and there is sufficiently few agents behaving dishonestly at the (unique) steady state, or (3) $\beta_p < 0$, $\beta_q \geq 0$, and there is many agents behaving dishonestly at the steady state. Supplementary Figure OA2 provides a graphical illustration.

1.2 Another definition of honesty

Results in the previous section define an honest organization using the steady state mean neighbor infection rate θ (Definition 1). A more natural definition would consider instead the population infection rate ρ associated with these steady states. Let $r : [0,1] \to [0,1]$ be the population infection rate associated with some steady state mean neighbor infection rate θ. Analogously to $\mathcal{F}$, the set of steady state mean infection rates in some organization, define $\mathcal{R}_f \equiv r_f(\mathcal{F})$, the set of population infection rates associated with those steady state mean neighbor infection rates in organization f. Finally, define $\overline{r} \equiv \max r \in \mathcal{R}_f$, and $\underline{r} = \min r \in \mathcal{R}_f$ to be, respectively, the maximum and minimum steady state population infection rates. The new definition says that an organization is more honest than another if both its maximal and minimal steady state population infection rates sustain less dishonest behavior than the other. This definition in this subsection and in the proofs of the results it introduces use the new definition, unless otherwise noted. We have:

[2] That is, f' has a degree distribution $P'(d)$ such that $\sum_d P'(d)u(d) \geq \sum_d P(d)u(d)$ for any non-decreasing function u.

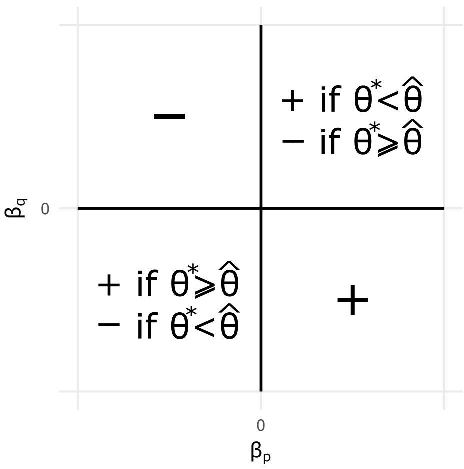

Figure OA2: **Illustration of propositions 4 and 5.** Consider o and o' such that $P'(d)$ FOSD $P(d)$. $+$ signs denote $o' \succeq o$, $-$ signs denote $o' \preceq o$. Increasing mean degree may help or hurt, depending on the values of β_p, β_q and the initial steady state θ^*.

Definition 2. We say that organization f' is more *honest* than organization f and write $f' \succeq f$ when $\underline{r}' \leq \underline{r}$ and $\overline{r}' \leq \overline{r}$.

The concepts of mean neighbor infection rate θ and population infection rate ρ can be at odds when the degree distribution has high variance. As an example, consider a star network of n nodes where the hub h is connected to the $n-1$ remaning spokes. Suppose that at time t, we have $y_{ht} = 1$ and $y_{it} = 0$ for all spokes $i \neq h$. The infection rate is $\rho = \frac{1}{n}$. The neighbor infection rate is $\theta_{ht} = 0$ for the hub, and $\theta_{it} = 1$ for all spokes $i \neq h$. As such, the mean infection rate is $\theta = \frac{n-1}{n}$. As n gets large, we have $\rho \to 0$ and $\theta \to 1$. Formally, recall that $\rho = \sum_d P(d)\rho(d)$, while equation (1) gives $\theta = \sum_d \tilde{P}(d)\rho(d) = \sum_d \frac{P(d)d}{\bar{d}}\rho(d)$. If $\mathbb{V}(d) = 0$, then $d = \bar{d}$, which implies $\theta = \rho$. As $\mathbb{V}(d)$ increases, the excess degree distribution $\tilde{P}(d)$ upweights the high-degree nodes relative to the degree distribution $P(d)$ since, by definition, many nodes are connected are connected to high-degree nodes. As such, θ upweights the infection rate of high-degree nodes relative to ρ.

The differences between the population infection rate and the mean neighbor infection rate imply that not all the regularities identified with Definition 1 are robust to using Definition 2. Indeed, the function $r_f(\theta)$ needs not be monotonic, which prevents extending claims on θ to claims on ρ. Specifically, the comparative statics on parameters (proposition 3) do not travel unless one imposes strong regularity conditions that guarantee that $r_f(\theta)$ is monotonic. Propositions 5 and 4, which hold parameter values constant but vary the degree distribution do travel, up to some additional restrictions that ensure that $r_f(\theta)$ is monotonic.

Proposition 6 (Increasing mean degree 1 - definition 2)**.** *Consider organizations f and f' with degree distributions P and P' respectively such that P' FOSD P and $\tilde{P}'$ FOSD $\tilde{P}$. If $\beta_p < 0$, $\beta_q \geq 0$, and $\left|\frac{\beta_q}{\beta_p}\right| \geq \frac{q(0,n-1)}{p(0,n-1)}$, then $f' \preceq f$. If $\beta_p \geq 0$, $\beta_q < 0$, and $\left|\frac{\beta_q}{\beta_p}\right| \leq \frac{q(0,n-1)}{p(0,n-1)}$, then $f' \succeq f$.*

Proposition 7 (Increasing mean degree 2 - definition 2)**.** *Consider organizations o and o' with degree distributions P and P' respectively such that P' FOSD P and $\tilde{P}'$ FOSD $\tilde{P}$. If $\alpha_p, \alpha_q \neq 0$ and β_p, β_q have the same sign, then there is a threshold $\hat{\theta} \in [0,1]$ such that*

$$\begin{cases} f' \succeq f, & \text{if } \beta_p, \beta_q \geq 0 \text{ and } \theta^* < \hat{\theta} \\ f' \preceq f, & \text{if } \beta_p, \beta_q \geq 0 \text{ and } \theta^* \geq \hat{\theta} \\ f' \preceq f, & \text{if } \beta_p, \beta_q < 0 \text{ and } \theta^* < \hat{\theta} \\ f' \succeq f, & \text{if } \beta_p, \beta_q < 0 \text{ and } \theta^* \geq \hat{\theta} \end{cases}$$

1.3 Proofs

Proof of proposition 1. Solving for $H(0) = 0$, we find $H(0) = 0 \iff \alpha_q = 0$. Similarly, solving for $H(1) = 1$, we find $H(1) = 1 \iff \alpha_p = 0$. □

Proof of proposition 2. Note that

$$\frac{\partial \rho(d)}{\partial \theta} = \frac{d[\beta_q p(0,d) + \beta_p q(0,d)]}{[p(\theta,d) + q(\theta,d)]^2}$$

Because $p(0,d), q(0,d) > 0$, if β_p and β_q have the same sign, then $\frac{\partial \rho(\theta,d)}{\partial \theta} \geq 0 \iff \frac{\partial \rho(\theta,d')}{\partial \theta} \geq 0$ for any $d, d' \in \{1, ..., n-1\}$. As such, H is monotonic. Furthermore,

$$\frac{\partial^2 \rho}{\partial \theta^2} = -2d \frac{[\beta_q p(0,d) + \beta_p q(0,d)](\beta_q - \beta_p)}{[p(\theta,d) + q(\theta,d)]^3}$$

A similar argument shows that H' is monotonic.

If H and H' are monotonic and $\alpha_p, \alpha_q \neq 0$, then H has a unique fixed point. □

Proof of lemma 1. It suffices to show that $\underline{\theta}' \geq \underline{\theta}$ and $\overline{\theta}' \geq \overline{\theta}$.

The cases where $\underline{\theta} = 0$ and $\overline{\theta} = 1$ are trivial. We have $\theta' \geq 0 = \underline{\theta}$, and since $\overline{\theta}' \leq 1$, it must be that $\overline{\theta}' = \overline{\theta} = 1$.

Consider $\underline{\theta} > 0$. We show that $\underline{\theta}' \geq \underline{\theta}$. Since $H_f(0) > 0$, it must be that $H_f(\theta) > \theta$ for any $x \in [0, \underline{\theta})$. So $H_{f'}(\theta) \geq H_f(\theta) > \theta$ for any $x \in [0, \underline{\theta})$. It must therefore be that $\underline{\theta}' \geq \underline{\theta}$.

Consider $\overline{\theta} < 1$. We show that $\overline{\theta}' \geq \overline{\theta}$. Suppose $\overline{\theta}' < \overline{\theta}$. Since $H_f(1) < 1$, it must be that $H_{f'}(\theta) < \theta$ for any $x \in (\overline{\theta}', 1]$. So $H_f(\overline{\theta}) \leq H_{f'}(\overline{\theta}) < \overline{\theta}$, a contradiction. □

Proof of proposition 3. Note that if $\frac{\partial \rho(d)}{\partial \phi} \geq 0$ for any (θ, d), then $\frac{\partial H_f}{\partial \phi} \geq 0$ for any θ. Using lemma 1, this implies $f' \preceq f$. Similarly, if $\frac{\partial \rho(d)}{\partial \phi} \leq 0$ for any (θ, d), then $\frac{\partial H_f}{\partial \phi} \leq 0$ for any θ, which implies $f' \succeq f$. We have:

$$\frac{\partial \rho(d)}{\partial \alpha_p} = -\frac{q(\theta,d)}{[p(\theta,d) + q(\theta,d)]^2} \leq 0$$
$$\frac{\partial \rho(d)}{\partial \beta_p} = -\frac{(1-\theta)dq(\theta,d)}{[p(\theta,d) + q(\theta,d)]^2} \leq 0$$
$$\frac{\partial \rho(d)}{\partial \alpha_q} = \frac{p(\theta,d)}{[p(\theta,d) + q(\theta,d)]^2} \geq 0$$
$$\frac{\partial \rho(d)}{\partial \beta_q} = \frac{\theta dp(\theta,d)}{[p(\theta,d) + q(\theta,d)]^2} \geq 0$$

□

Proof of proposition 5. Note that

$$\frac{\partial \rho(d)}{\partial \theta} = \frac{d[\beta_q p(0,d) + \beta_p q(0,d)]}{[p(\theta,d) + q(\theta,d)]^2}$$

Suppose that $\beta_p, \beta_q \geq 0$. Then, then we have $\frac{\partial \rho(d)}{\partial \theta} \geq 0$ for any d, which implies that $\frac{\partial H_f}{\partial \theta}, \frac{\partial H_{f'}}{\partial \theta} \geq 0$. Furthermore, note that

$$\frac{\partial \rho(d)}{\partial d} \geq 0 \iff \theta \geq \frac{\alpha_q \beta_p}{\alpha_q \beta_p + \alpha_p \beta_q} \equiv \hat{\theta}$$

This and $\tilde{P}'$ FOSD $\tilde{P}$ implies $H_{f'}(\theta) \geq H_f(\theta) \iff \theta \geq \hat{\theta}$.

Consider the case where $\theta^* \geq \hat{\theta}$. It suffices to show that $\theta^{*'} \geq \theta^*$. We prove this by contradiction. Suppose that $\theta^{*'} < \hat{\theta}$. Recall that $H_{f'}(0) > 0$. If $\theta^{*'} < \hat{\theta}$, then $H_{f'}(\hat{\theta}) = H_f(\hat{\theta}) \geq \hat{\theta}$ implies that $H_{f'}$ must have another fixed point in $(\theta^{*'}, \hat{\theta}]$, contradicting proposition 2.

Suppose that $\theta^{*'} \in [\hat{\theta}, \theta^*)$. Since H_f is increasing and has a unique fixed point at θ^*, it must be that $H_f(\theta) > \theta$ for any $\theta \in [\hat{\theta}, \theta^*)$. This implies $H_{f'}(\theta^{*'}) \geq H_f(\theta^{*'}) > \theta^{*'}$, a contradiction.

The three other cases prove similarly. □

Proof of proposition 4. Note that

$$\frac{\partial \rho(d)}{\partial d} = \frac{\theta \alpha_p \beta_q - (1-\theta)\alpha_q \beta_p}{[p(\theta, d) + q(\theta, d)]^2} \tag{3}$$

If $\beta_p < 0$ and $\beta_q \geq 0$, then $\frac{\partial \rho(d)}{\partial d} \geq 0$. Since ρ is non-decreasing, $\tilde{P}'$ FOSD $\tilde{P}$ implies

$$\sum_d \tilde{P}'(d)\rho(\theta, d) \geq \sum_d \tilde{P}(d)\rho(\theta, d),$$

that is $H_{f'}(\theta) \geq H_f(\theta)$ for any θ. By lemma 1, $f' \preceq f$.

If $\beta_p < 0$ and $\beta_q \geq 0$, then $\frac{\partial \rho(d)}{\partial d} \leq 0$. Since $-\rho$ is non-decreasing, FOSD implies

$$-\sum_d \tilde{P}'(d)\rho(\theta, d) \geq -\sum_d \tilde{P}(d)\rho(\theta, d),$$

that is $H_{f'}(\theta) \leq H_f(\theta)$ for any θ. By lemma 1, $f' \succeq f$. □

Proof of proposition 7. Suppose that $\beta_p, \beta_q \geq 0$. From the proof of proposition 5, define $\hat{\theta} \equiv \frac{\alpha_q \beta_p}{\alpha_q \beta_p + \alpha_p \beta_q}$, and recall that $\frac{\partial \rho(d)}{\partial d} \geq 0 \iff \theta \geq \hat{\theta}$. Consider the case where $\theta^* \geq \hat{\theta}$.

Proposition 5 implies that $\theta^{*'} \geq \theta^*$. We now show that $f' \preceq f$. Let $r^* \equiv r_f(\theta^*)$ and $r^{*'} \equiv r_{f'}(\theta^{*'})$ be the population infection rates at the steady state in f and f' respectively. It suffices to show that $r^{*'} \geq r^*$. Since $\frac{\partial \rho(d)}{\partial \theta} \geq 0$, we have $r^* = \sum_d P(d)\rho(d, \theta^*) \leq P(d)\sum_d \rho(d, \theta^{*'})$. Since $P'(d)$ FOSD $P(d)$ and $\rho(d, \theta)$ is non-decreasing in d, we have $\sum_d P(d)\rho(d, \theta^{*'}) \leq \sum_d P'(d)\rho(d, \theta^{*'}) = r^{*'}$.

The three other cases prove similarly. □

Proof of proposition 6. If $\beta_p < 0$ and $\beta_q \geq 0$, then proposition 4 implies that $\underline{\theta}' \geq \underline{\theta}$ and $\overline{\theta}' \geq \overline{\theta}$.

Let's show that $\underline{r}' \geq \underline{r}$. Note that $\left|\frac{\beta_q}{\beta_p}\right| \geq \frac{q(0,n-1)}{p(0,n-1)}$ implies that $\frac{\partial \rho(d)}{\partial \theta} \geq 0$ for any $d \in \{1, ..., n-1\}$. As such, we have $\frac{\partial \overline{r}(\theta)}{\partial \theta} \geq 0$. This implies that $\underline{r} = r_f(\underline{\theta})$ and $\underline{r}' = r_{f'}(\underline{\theta}')$. Furthermore, $\frac{\partial \rho(d)}{\partial \theta} \geq 0$ for any $d \in \{1, ..., n-1\}$ implies that $r_f(\underline{\theta}) = \sum_d P(d)\rho(d, \underline{\theta}) \leq \sum_d P(d)\rho(d, \underline{\theta}')$. Since P' FOSD P and $\rho(d)$ is non-decreasing in d, then $\sum_d P(d)\rho(d, \underline{\theta}') \leq \sum_d \tilde{P}(d)\rho(d, \underline{\theta}') = r_{f'}(\underline{\theta}')$. We show similarly that $\overline{r}' \geq \overline{r}$.

The case when $\beta_p \geq 0$ and $\beta_q < 0$ proves similarly. □

2 Measuring social ties

2.1 Descriptive statistics

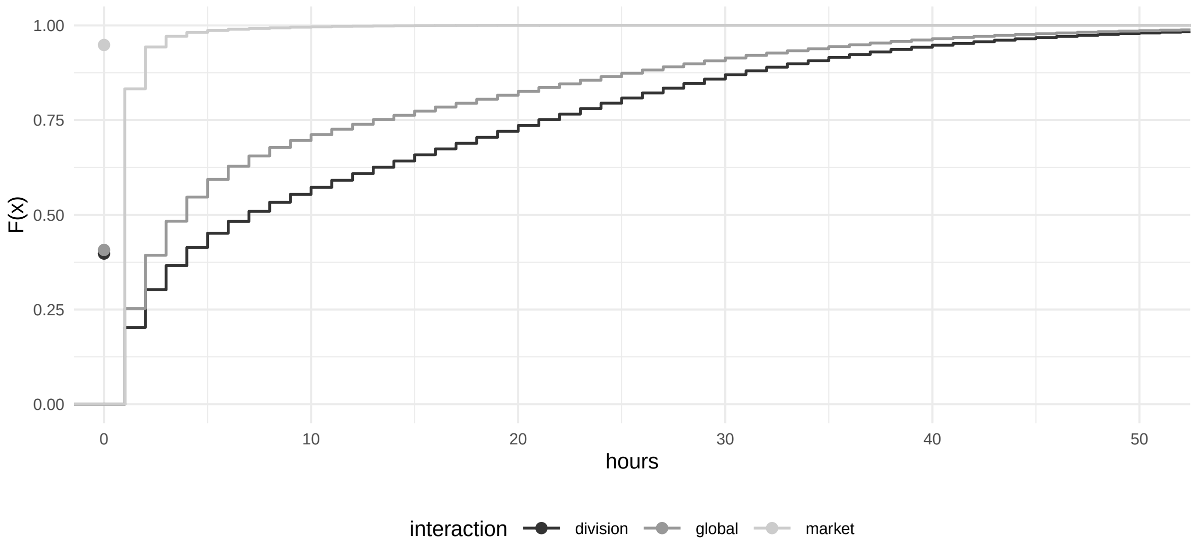

Figure OA3: **Cumulative distribution of coworking hours by interaction level.** The cumulative distribution conditions on dyads of 1+ coworking hours. Points at 0 indicate the percentage of 0 coworking hours dyads.

	Threshold				
Statistic	1h	2h	3h	4h	5h
betweenness	23.196	2.392	0.613	0.187	0.057
closeness	0.313	0.080	0.028	0.014	0.009
clustering	0.179	0.026	0.011	0.007	0.005
degree	3.038	0.509	0.173	0.087	0.057
eigenvector	0.184	0.434	0.758	0.885	0.934

Table OA1: **Network descriptive statistics.** This table reports the mean value of selected network statistics as a function of the threshold used to define ties. All statistics but eigenvector centrality decrease as the threshold increases, because networks become more sparse. Eigenvector centrality increases as the threshold increases, because the few nodes that have non-zero degree become more eigenvector central as networks become more sparse.

2.2 Validation with self-reported links

In January 2017, we surveyed all 174 clerks that had 3 months of tenure or more as of January 1st 2017. We administered two network surveys. In the first survey, we collected friendship and co-working networks using a standard name generator.

In the second survey, we randomly sampled three random alters about which we asked specific questions. Since most dyads interact infrequently, we divided the dyads in three categories of frequent, intermediate, and rare interactions using terciles of time spent together using punch-in data from the three months preceding the survey (October to December 2016), and sampled, for each respondent, one of each such dyads. The questions included self-reported measures of knowledge about and friendship with alter on a 1-5 scale, as well as responses to two quizzes designed to evaluate ego's knowledge about alter on the professional, and personal levels. Quizzes were incentivized. We report the questions below:

1. Name generator for the friendship network: "Who are the clerks with whom you discuss personal matters?"

2. Name generator for the coworker network: "On a typical work day, who are the clerks with whom you usually work?"

3. Self-reported measure of knowledge: "How much do you know X?" Answers: "Never heard that name before," "I know him/her by name," "A little," "Quite well," "Very well."

4. Self-reported measure of friendship: "How much do you discuss personal matters with X?" Answers: "Never," "Little," "Somewhat," "Much," "A great deal."

5. Personal knowledge quiz: "In which city was X born?", "At home, does X have a [CRT TV, flat-screen TV, desktop computer, laptop]?", "What is X's marital status?", "Does X have a sister?"

6. Professional knowledge quiz: "Is X a temporary or a permanent employee?", "What is X's division?", "What is X's salary?", "In what year did X join the company?"

We validate our measure of social interactions by investigating whether it correlates with survey measures. To do so, we construct dyad-level measures of interaction time that count the number of hours two clerks have spent over a period of time for a given level of interaction. We consider three periods that all end on January 31 2017, and have a length of 1, 2, and 3 months. We consider three increasingly weak level of interactions: "market," "division," and "global." "Market" counts the number of hours clerks i and j have spent operating in the same market. "Division" counts the number of hours clerks i and j have spent processing claims from markets that pertain to the same division, and "global" counts the number of hours clerks i and j have spent at the call-center. From these data, we derive, for each dyad, interaction level, and time period, the mean daily amount of time i and j have spent interacting.

Figure OA4 considers all our network data and estimates effect of interaction time on all the outcomes elicited. We find that although effect sizes are small, market- and division-level interactions significantly correlate with better knowledge on both the personal and professional level, while global-level interactions do not, which confirms that interactions occurring within market or within division capture stronger social relationships than interactions that occur more globally.

Figure OA5 reproduces Manuscript Figure 5 using self-reported friendship and coworking ties. Consistent with our main result, we find that dishonest behavior diffuses in all models. We also find that honest behavior does not significantly diffuse in the coworking model and, contrary to our main result, diffuses in the friendship model. This change when using self-reported friendship ties suggests that those ties are homophilous.

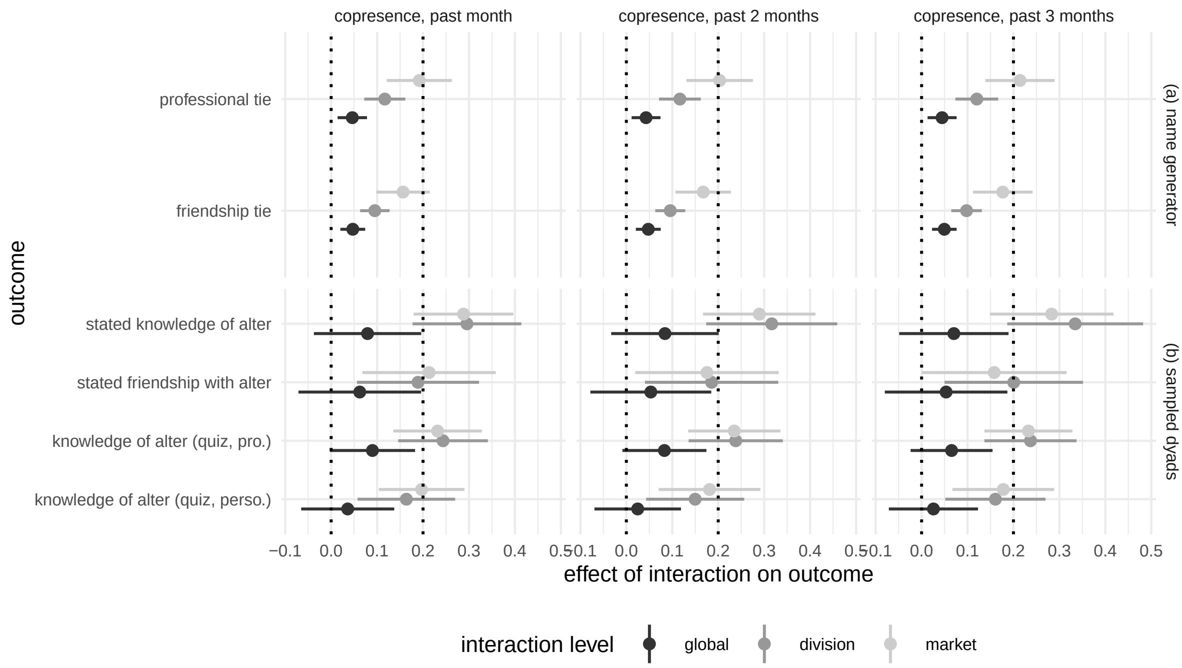

Figure OA4: **Correlation between interaction time and self-reported ties.** We report OLS estimates of models in which both the outcome and independent variable are standardized, with heteroskedastic-robust standard errors. Parameter values represent the effect of a one standard deviation increase of the independent variable on the outcome, in terms of standard deviations. As per Cohen's rule of thumb, a effect size below 0.2 corresponds to a small effect, while an effect size between 0.2 and 0.5 corresponds to a small to medium effect. Market- and division-level interactions correlate with better knowledge on both the personal and professional level, while global-level interactions do not.

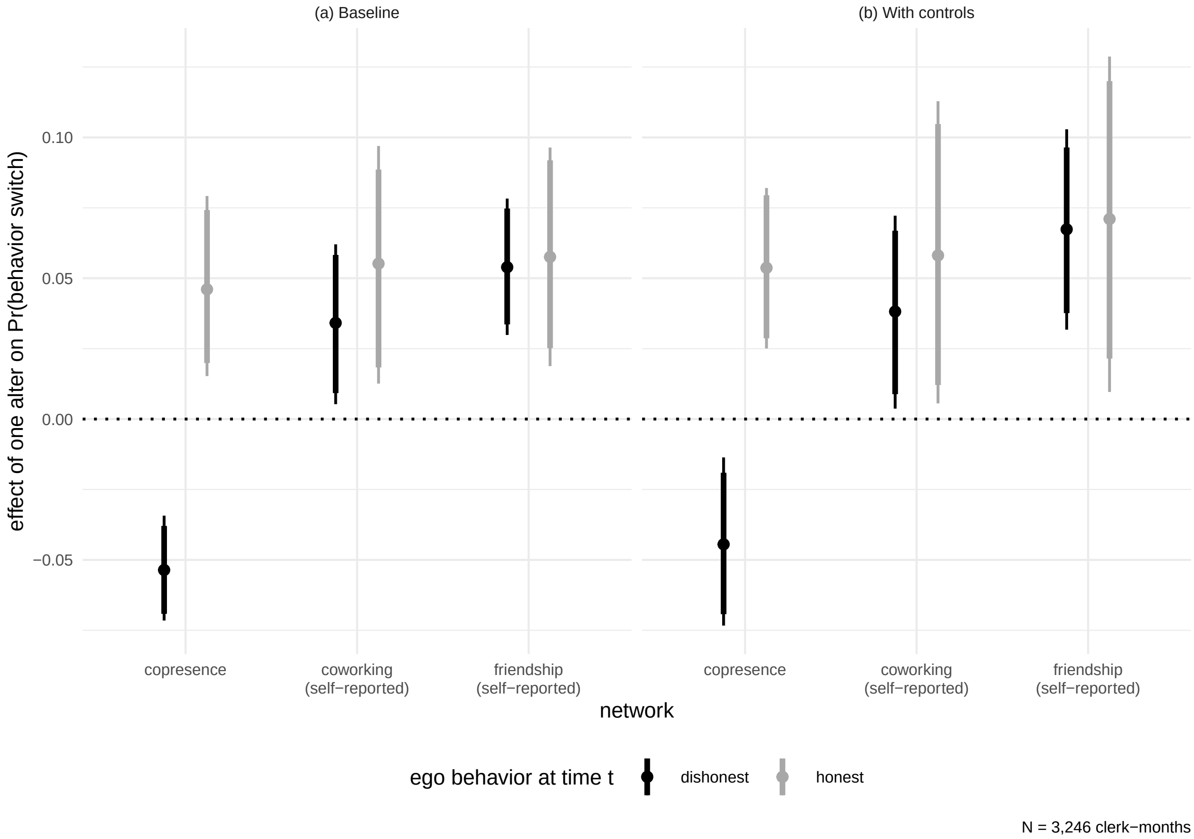

Figure OA5: This figure reproduces Manuscript Figure 5. All models subset to those clerks that took the survey. All models use a threshold of .5 in s-score to define dishonest behavior. The co-staffing model uses ties constructed through co-staffing for comparison. Other models use self-reported ties elicited through a name generator instead. Dishonest behavior diffuses in all models. Honest behavior is repellent in the co-staffing model, does not significantly diffuse in the coworking model, and diffuses in the friendship model. That results change when using self-reported friendship ties suggests that those ties are homophilous.

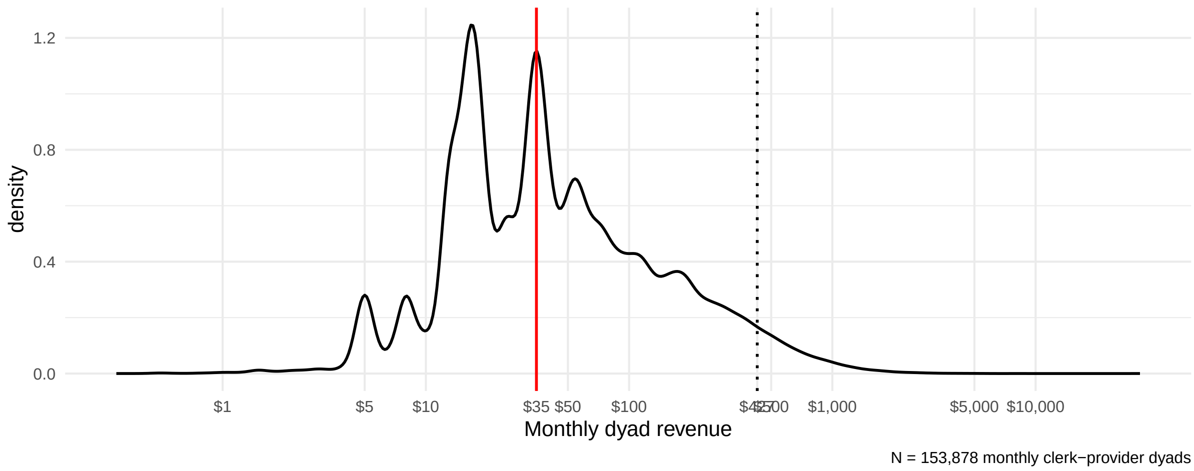

Figure OA6: **Distribution of dyad revenue.** The solid line represents the median, and the dotted line marks the \$427 threshold.

3 Measuring dishonest behavior

3.1 Constructing s-scores

Threshold for component 1. We set the threshold for component 1 at \$427, whcih matches clerks' average monthly wage (Table 1). Thus, with r_{kimt} the revenue transferred by clerk i to provider k during month t on market m, we define $s^1_{ikmt} \equiv \mathbb{1}\{r_{kimt} > 427\}$. Figure OA6 shows that the distribution of dyad revenue is largely skewed to the left, with \$427 corresponding to the 96th percentile.

Derivation of components 2 and 3. We now describe how we construct components 2 and 3 of the s-score. Consider first component 2. Let $Y_c = k \in \mathcal{K}$ be the provider to which claim c is allocated, with $\mathcal{K}$ the set of all providers operating on market m. Our null model M^2_{mt}, estimated for market m during month t estimates a multinomial logistic regression. With U_c be a vector of covariates of claim c (see description below), and β_{kmt} a vector of parameters for provider k, the model writes

$$M^2_{mt} : \Pr(Y_c = k|U_c) = \frac{\exp(U_c'\beta_{kmt})}{\sum_{k'\in\mathcal{K}} \exp(U_c'\beta_{k'mt})}$$

The model allows pinning down a null distribution of revenues across providers for any clerk. Let (R_{imt}) be a length-$\mathcal{K}$ random vector of revenues for clerk i, with R_{kimt} the revenue allocated to provider k. Under M^2_{mt} and with $\mathcal{C}_{imt}$ the set of claims handled by i on market m during month t and m_c the value of claim c, we have that $R_{imt} \sim \sum_{c\in\mathcal{C}_{imt}} m_c \text{ Multinom}(U_c, \beta_{mt})$. Recall that $s^2_{ikmt} \equiv 1 - \Pr(R_{kimt} > r_{kimt}|M^2_{mt})$.

Consider now component 3. Let $Z_c = 1$ if claim c was awarded suspiciously; that is, if it was allocated after three or more draws (more than 90% of claims allocated through random draws are allocated in two draws or fewer), or if it was allocated after opting out of the random draw. Let $Z_c = 0$ otherwise. Our null model M^3_{kmt}, estimated for provider k in market m during month t estimates a logistic regression. With V_c a vector of covariates of claim c (see description below), and γ_{mt} a vector of parameters, the model writes

$$M^3_{kmt} : \Pr(Z_c = 1|V_c) = \frac{\exp(V_c'\gamma_{kmt})}{1 + \exp(V_c'\gamma_{kmt})}$$

The model allows pinning down a null distribution of suspicious claim awards to provider k. Let F_{ikmt} be a random variable denoting the fraction of claims that were suspiciously awarded to k by clerk i, and f_{ikmt} be the observed value of such random variable. Under the null distribution and with $\mathcal{C}_{ikmt}$ the set of claims awarded by i to provider k on market m during month t, we have $F_{ikmt} \sim \frac{\sum_{c \in \mathcal{C}_{ikmt}} \text{Bernoulli}(V_c, \gamma_{kmt})}{|\mathcal{C}_{ikmt}|}$. Similar to component 2, we define $s^3_{ikmt} \equiv 1 - \Pr(F_{ikmt} > f_{ikmt} | M^3_{kmt})$.

The distribution of revenues R_{imt} and of suspicious claim awards F_{ikmt} are not tractable analytically. As such, we approximate them using Monte Carlo simulations taking 1,000 draws of R_{imt} and of F_{ikmt}.

Covariates. Estimating models M^2 and M^3 requires selecting covariates. Since months are a natural time unit in this setting, we estimate model parameters monthly. Selecting covariates poses a tradeoff. On the one hand, one would like to include many covariates, such that our null models account for as many sources of variation as possible across clerks. On the other hand, some of our markets have very few transactions, thereby preventing the estimation of overly complex models. Both models consider six time intervals that match known variations in activity and are used by management in staffing decisions (see Section 3.1 for details). We divide each day into three time intervals that correspond to the three shifts used by management (morning (7a-2p), afternoon (2p-9p), and evening (9p-7a)). We further divide the week into weekdays and weekends, for a total of 6 time intervals. Model M^2, which considers claim allocation, also controls for log-claim value, as some providers are often selected to handle more complex, more valuable claims.

3.2 Validation

We show that s-scores indeed pick up dishonest or deviant behavior, then that s-scores pick up dishonest, rather than deviant behavior.

We first show that s-scores indeed pick up dishonest or deviant behavior by showing that such behavior is rare, and therefore contrary to the norm. We have seen that most clerks have minuscule s-scores (Manuscript, Figure 4). We also show that clerks behaving dishonestly have a small number of dishonest partners (Figure OA7).

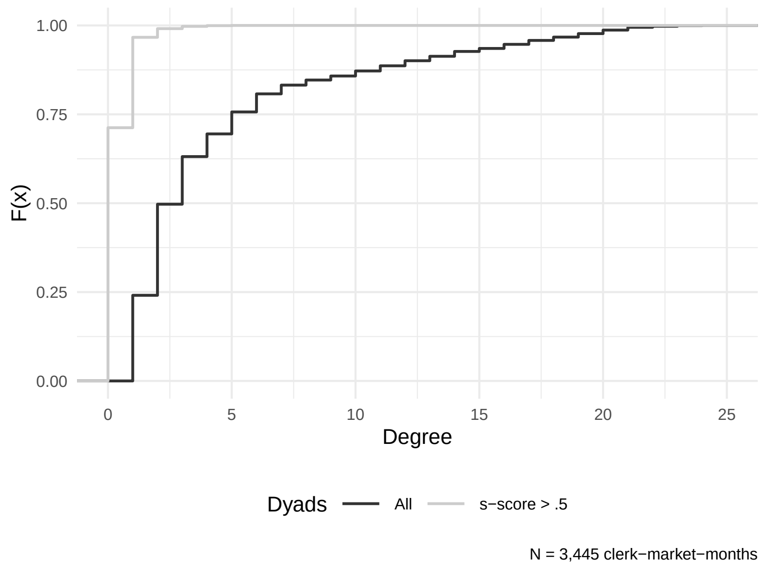

Figure OA7: **Cumulative degree distribution of dishonest clerk-provider ties.** This figure plots the degree distribution of clerk-provider ties for all those ties that satisfy condition 1 of the s-score (i.e., revenue greater than $427). The grey line plots the degree distributions of dishonest ties; i.e., ties with s-scores greater than .5. While clerks tend to have large neighborhoods (25% clerks interact with exactly 1 provider), they have few dishonest partners. 71% clerks have no dishonest partners and 3% clerks have more than 1 dishonest partner.

We then show that s-scores pick up dishonest, rather than deviant behavior. We provide evidence that the forms of rule-breaking captured by high s-scores underlie private gains instead of other motives such as inexperience, "well-intentioned" rule-breaking, or clerk-provider complementarities that facilitate cooperation.

We first show that engaging in dishonest behavior requires a modicum of experience, i.e. no clerks with less than a year of tenure have high s-scores, while some of the clerks with more than a year of tenure have high s-scores (Supplementary Figure OA8).

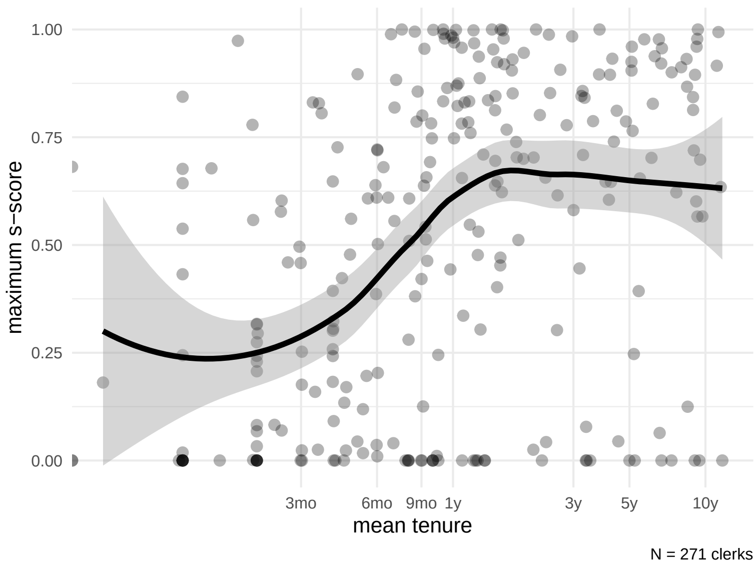

Figure OA8: **Correlation between tenure and dishonest behavior.** The x-axis represents the average tenure of a clerk over the period she has been observed, while the y-axis plots the maximum s-score of any clerk-provider dyad she has been involved in. The black line is a loess fit. It takes a modicum of experience to engage in dishonest behavior: clerks that have less than one year of tenure all have low s-scores, while some of the clerks that have more than one year of tenure have high s-scores.

We then rule out forms of "well-intentioned" rule-breaking by considering provider quality. By "well-intentioned" rule breaking, we mean breaking the rules for the company's gain; e.g., if the random draw system prevents fast allocation to better-performing providers. If s-scores mainly picked up well-intentioned rule-breaking, then we should observe that clerk-provider links with high s-scores should indicate high-quality providers. We show that such correlation exists but has an unclear direction and a minuscule size (Supplementary Table OA2). We then focus on those clerks that have high s-scores overall, as they should be more prone to selection based on quality criteria, and find similar inconclusive results (Supplementary Figure OA9).

Table OA2: Correlation between s-scores and provider quality

	yearly ratings			quarterly ratings		
	(1)	(2)	(3)	(4)	(5)	(6)
s-score	−0.002	−0.005**	−0.006**	0.005	0.009**	0.008*
	(0.004)	(0.002)	(0.003)	(0.006)	(0.004)	(0.004)
Num.Obs.	4805	4805	4805	5620	5620	5620
R2	0.000	0.204	0.256	0.000	0.090	0.106
Mean DV	0.756	0.756	0.756	0.714	0.714	0.714
SD DV	0.037	0.037	0.037	0.09	0.09	0.09
Market FE	-	✓	✓	-	✓	✓
Month FE	-	✓	✓	-	✓	✓
Clerk FE	-	-	✓	-	-	✓

Note: These models regress the quality rating of a provider on the s-score of each dyad with revenue above $427 she is involved in. Models 1 to 3 use objective, yearly quality ratings, models 4 to 6 use subjective, quarterly quality ratings that are both normalized to fall between 0 and 1. Objective quality ratings come from audits conducted yearly by the firm to verify whether service providers conform to a series of pre-defined quality standards. Subjective quality ratings use data from customer satisfaction surveys in which poor quality ratings are verified by the firm. We exclude data from dissatisfied customers whose complaints proved untrue. Standard errors are clustered at the month and market levels. There is little correlation between provider quality and s-scores, with effect size of about one tenth of a standard deviation of the dependent variable.

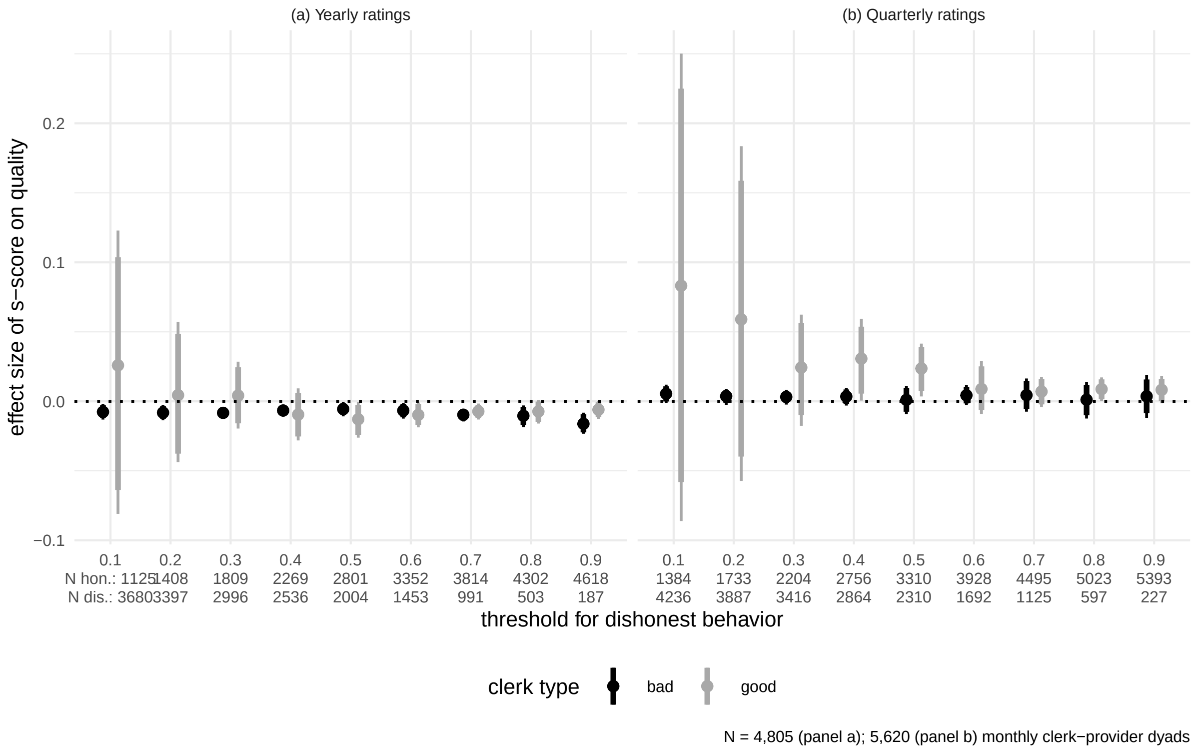

Figure OA9: **Correlation between s-scores and provider quality by clerk behavior.** We regress the quality rating of a provider on the s-score of each dyad with revenue above $427 she is involved in, adding an interaction term for clerk behavior (honest/dishonest). That is, we estimate the following linear model $r_{imt} = \alpha_j + \alpha_m + \alpha_t + \beta_0 y_{jmt} + \beta_1 s_{ijmt} + \beta_2 s_{ijmt} y_{jmt} + \epsilon_{ijmt}$. The dependent variable r_{imt} corresponds to the quality of provider i in market m during month t. It is measured through objective, yearly ratings (left panel) and subjective, quarterly ratings (right panel). Both are normalized to fall between 0 and 1. Objective quality ratings come from audits conducted yearly by the firm to verify whether service providers conform to a series of pre-defined quality standards. Subjective quality ratings use data from customer satisfaction surveys in which poor quality ratings are verified by the firm. We exclude data from dissatisfied customers whose complaints proved untrue. The independent variable s_{ijmt} corresponds to the s-score of the dyad between provider i and clerk j on market m during month t, and the variable $y_{jmt} \in \{0, 1\}$ corresponds to the behavior (honest, $y_{jmt} = 0$ vs. dishonest, $y_{jmt} = 1$) of clerk j on that same market-month. A clerk is defined to behave dishonestly if her maximum s-score on a given month is above some threshold, reported on the top row of the x-axis. Models include clerk, market, and month fixed effects ($\alpha_j, \alpha_m, \alpha_t$ respectively) and cluster standard errors at the month and market levels. Points report estimates of the effect of higher s-scores on provider quality for clerks behaving honestly and dishonestly; i.e. β_1 and $\beta_1 + \beta_2$ respectively. Bars represent 90 and 95 percent confidence intervals clustered at the month and market levels. The x-axis also reports the number of clerks behaving honestly and dishonestly for each model (second and third rows). There is little correlation between provider quality and s-scores. That correlation is minuscule for clerks behaving dishonestly (provider quality has a mean of 0.76 and 0.72 with a standard deviation of 0.04 and 0.09 for yearly and quarterly ratings, respectively), and comparable between clerks behaving honestly and dishonestly for thresholds above 0.5.

Our data do not allow ruling out that s-scores capture clerk-provider complementarities that facilitate cooperation, although limited contextual evidence casts doubt on this hypothesis. For instance, a clerk may find it easier to work with a provider because they speak the same language.

We do not have sufficient data on clerks' and providers' characteristics to test this hypothesis quantitatively. However, all clerks have similar backgrounds (college-educated individuals in their late twenties from the city in which the call center is located) and communicate with providers in the country's dialect of Arabic. Furthermore, any such communication advantage with a specific provider (e.g., clerk i and provider k speak the same regional dialect) should be shared with other providers on the market, since markets are defined geographically. Together, this suggests that such complementarities should be rare.

3.3 The cost of dishonest behavior

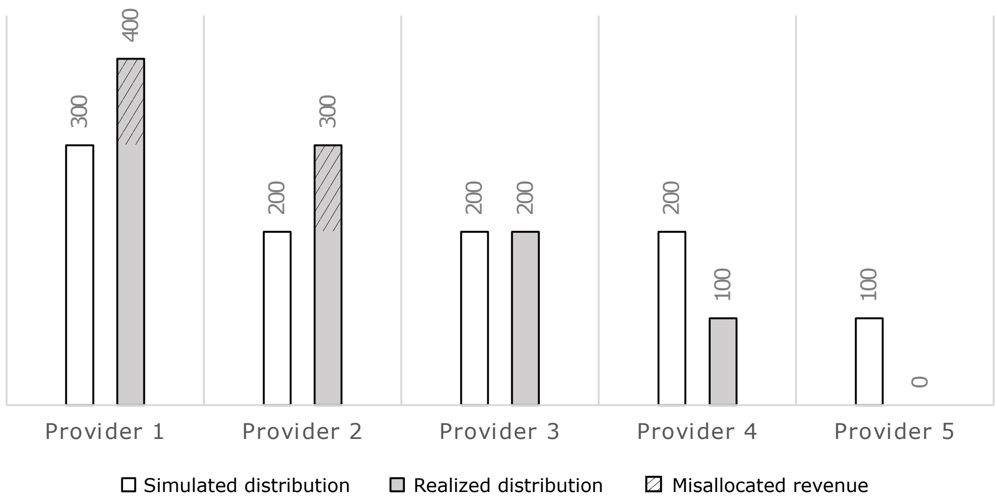

Figure OA10: **Deriving misallocated revenue.** In this example, white bars represent how a clerk behaving honestly would have allocated a total revenue of 1,000 among providers 1 to 5. Grey bars represent the observed allocation from a clerk behaving dishonestly. Misallocated revenue corresponds to the hatched portions, and equals 200; that is, 20% of total revenue.

Since dishonest behavior is the misallocation of claims following clerk-provider collusion (Manuscript, Section 3.1), we leverage the s-score to estimate the proportion of revenue that was misallocated over the period. Specifically, we single out dishonest behavior using a threshold in s-scores. Clerks that behave dishonestly on a market significantly deviate from the average allocation of revenue in that market (Condition 2, Manuscript Section 3.3). We use simulations to recover the claim allocation of a clerk that behaved dishonestly, had she behaved honestly. To do so, we leverage the null model M^2_{imt} used to estimate component 2 of the s-score, which provides the distribution of claims under honesty to obtain the distribution of revenues R_{imt} had they behaved honestly. We finally derive the amount of misallocated revenue by comparing, for each clerk-market, their simulated distribution to their observed distribution. The total amount misallocated corresponds to the sum of upwards deviations of the actual distribution of revenues as compared to the actual simulated distribution of revenue (Figure OA10 provides a graphical representation).

Formally, re-using notation from section 3.1, let K be the set of providers operating in market m, and r_{ik} be the revenue transfered by clerk i to provider k. Suppose dishonest behavior is defined by a threshold $\overline{s}$ in s-score. Suppose that as per threshold $\overline{s}$, clerk i behaved dishonestly. Suppose furthermore that she allocated the vector of revenues $r_i = (r_{ik})_{k\in K}$ on market m, but that draw n from M_2 yields the vector of revenues $\tilde{r}_{in} = (\tilde{r}_{ikn})_{k\in K}$. The total revenue misallocated by i from simulation n is therefore $\delta_{in} \equiv \sum_{k\in K}(r_{ik} - \tilde{r}_{ikn})1\{r_{ik} - \tilde{r}_{ikn} \geq 0\}$. We then average δ_{in} over 1,000 simulations to obtain the average revenue misallocated by i, $\bar{\delta}_i \equiv \frac{1}{1000}\sum_{n=1}^{1000}\delta_{in}$.

Table OA3: Misallocated revenue as a function of threshold in s-score

threshold in s-score	misallocated revenue (K$)	misallocated revenue (% revenue)	% dishonest clerks
0.1	1,414	9.68%	2.83%
0.2	1,304	8.93%	2.48%
0.3	1,161	7.94%	2.08%
0.4	990	6.78%	1.68%
0.5	803	5.50%	1.32%
0.6	609	4.17%	0.97%
0.7	423	2.89%	0.66%

Note: The last column reports the percentage of clerk-market-months with dishonest behavior.

Table OA3 shows that misallocation is substantial, ranging between 2.9 and 9.7% of the revenue generated by the markets analyzed in this study. For comparison, the third and second largest markets in this set of markets generated, respectively, K$514 and K$1,100 over the study period.

Results also show that fraud is committed by a minority, with even the lowest threshold value putting the percentage of clerk-market-months with dishonest behavior around 3%. An important reason for these low numbers is that most clerks manipulate little revenue in most of the markets they operate in, meaning that they do not have enough to gain from behaving dishonestly. As a result, their s-score is 0 (see Section 3.3).

Our preferred threshold value for the s-score is 0.5, as it is the estimate that matches most closely 5%, the estimated cost of fraud to the median organization worldwide (Association of Certified Fraud Examiners, 2018). We use this threshold as a baseline for estimation.

Finally, results remain fairly stable for a wide range of s-scores: using cutoffs between 0.3 and 0.7 makes the percentage of misallocated revenue vary by 1.5 percentage points relative to our baseline of 0.5.

4 Estimation strategy

4.1 Additional descriptive statistics

Figure OA11 below shows the correlation between all variables used in our main model (Equation (3)) for a threshold of 0.5 in s-scores. The variable `switch`, `dishonest behavior`, and `N alters with opposite behavior` correspond, respectively, to variables z_{imt+1}, y_{imt}, and n_{imt}. All other variables are controls included in the vector x_{imt}. We include:

- Indicators of activity on market m during month t:
  - `log revenue (market)`: revenue generated by market m during month t. This should correlate with higher revenue allocated by i to providers and thus, potentially, to higher s-scores through component s^1.
- Indicators of i's behavior on market m during month t. These indicators should correlate with higher revenue allocated by i to providers and thus, potentially, to higher s-scores through component s^1. These indicators are:
  - `log hours (clerk-market)`: the hours worked by clerk i on market m during month t. This should correlate with higher revenue allocated by i to providers and thus, potentially, to higher s-scores through component s^1.
  - `log revenue (clerk-market)`: the revenue generated by clerk i on market m during month t. This should correlate with higher revenue allocated by i to providers and thus, potentially, to higher s-scores through component s^1.
  - `log market experience`: the experience of clerk i on market m at month t, measured in months, to control for potential learning effects.
  - `degree`: clerk i's degree on market m at month t, measured using a threshold of 3 hours. We include this variable to compare clerks that share the same number of links within a market, hence the same potential for diffusion.
- Indicators of i's overall activity behavior during month t. These indicators should capture higher overall activity and potentially lower s-scores, as These should indicate higher overall activity, and potentially lower s-scores, as i's effort is spread across markets. These indicators are:
  - `log hours (clerk)`: hours worked by clerk i during month t on all markets. This should correlate with higher revenue allocated by i to providers and thus, potentially, to higher s-scores through component s^1.
  - `log revenue (clerk)`: revenue generated by clerk i during month t on all markets. This should correlate with higher revenue allocated by i to providers and thus, potentially, to higher s-scores through component s^1.
  - `clerk max. s-score`: maximum s-score of clerk i during month t on all markets. This capture i's general tendency towards dishonest behavior.

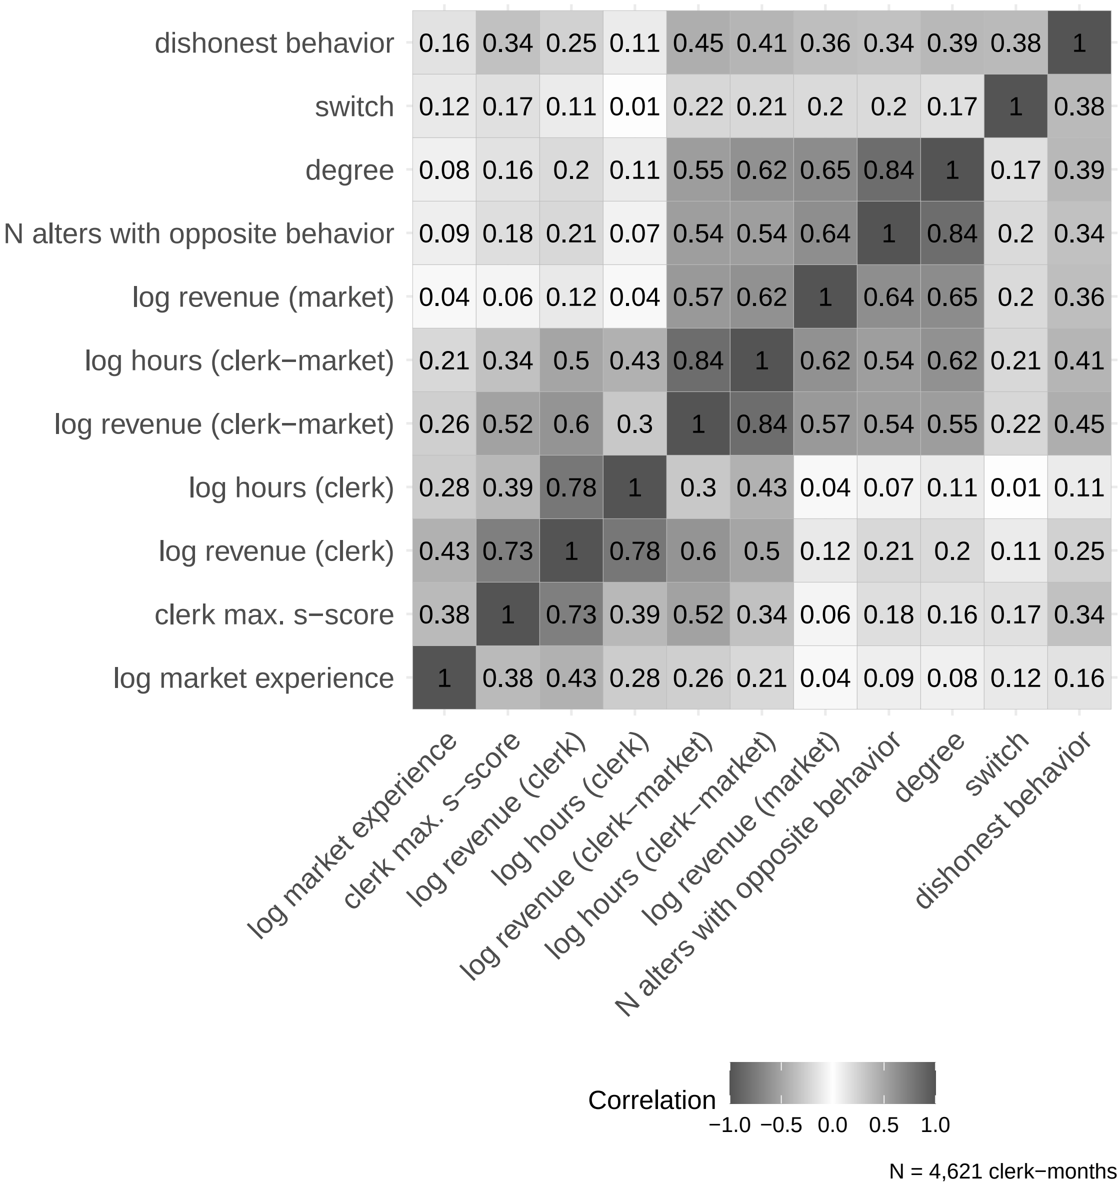

Figure OA11: Correlogram

4.2 Direct tests of the identifying assumptions

We provide direct evidence supporting identifying assumptions 1 and 2. We show that clerks are assigned to markets independently of their behavior (Assumption 1) through two tests. First, since management punishes dishonest behavior by dismissal (Section 3.1), then all else equal, a clerk behaving dishonestly should see a shorter employement spell, if management observes such behavior. Yet, we use survival models to show that dishonest behavior does not correlate with shorter employment spells end (Supplementary Table OA4). Second, we test directly that assignment to new markets is independent of one's behavior. Clerks who behave dishonestly are just as likely as clerks who behave honestly to be assigned to new markets and, conditional on being assigned to a new market, there is no correlation between one's own behavior and the behavior of clerks working in the destination market (Supplementary Table OA5).

We then show that peer influence operates indeed at the market level (Assumption 2): we find no meaningful evidence of influence when considering broader forms of social interactions

(Supplementary Figure OA12). To do so, we consider alternative measures of social ties that capture broader forms of social interactions. We consider division-level networks, in which two clerks are connected if they were co-staffed on the same division, and the global network, in which two clerks are connected if they were co-staffed at all during the same hour.

	(1)	(2)
s-score	−1.601***	−0.320
	(0.464)	(0.457)
log(revenue)		−0.346***
		(0.042)
experience		−0.012***
		(0.003)
Num.Obs.	3854	3854
AIC	1769.0	1695.3

Table OA4: **Identifying assumptions: Cox proportional hazard survival models for the duration of clerks' employment spells.** Time-periods are months, and all models use time-varying covariates, with s-score defined as a clerk's maximum s-score on a given month, experience defined in months from hiring date, and log revenue defined as the log of total revenue handled by a clerk on a given month. Standard errors are clustered at the clerk level. Without controls, clerks behaving dishonestly are less likely to see their employment spell end (model 1). Controlling for revenue and experience, the length of the employment spells of clerks behaving dishonestly is not significantly different from that of clerks behaving honestly (model 2). That clerks behaving dishonestly are no more likely to be dismissed suggests that management is unaware of their (dis)honest behavior.

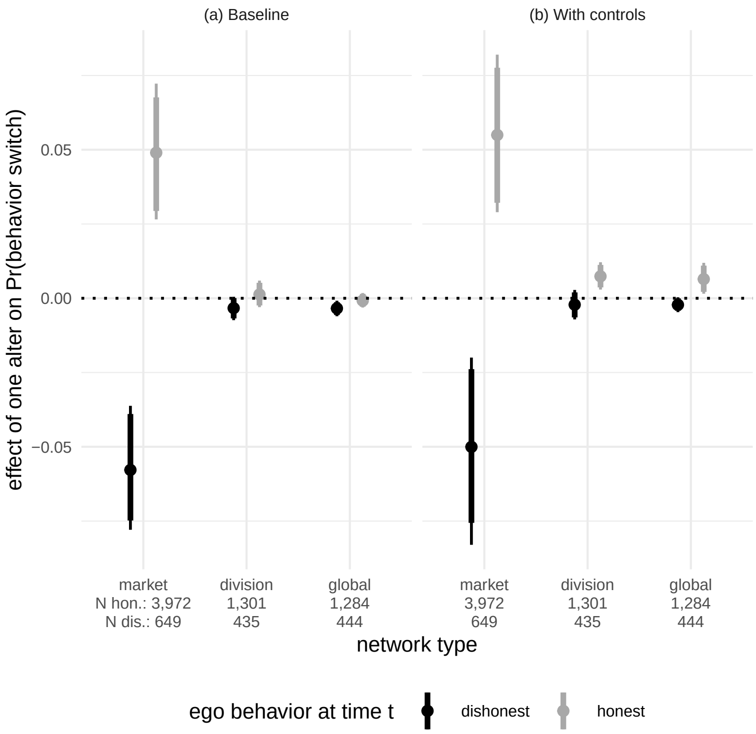

Figure OA12: **Identifying assumptions: varying the source of influence.** The x-axis represents the relevant network used to define ties; "market" reproduces Manuscript Figure 5 in the paper with a threshold of 0.5 in s-scores; "global" define a tie between clerks i and j during month t if they have spent at least one hour attending the workplace together; "division" defines that same tie if they have spent one hour operating in markets pertaining to the same same division. A clerk is defined as behaving dishonestly if her maximum s-score in all dyads pertaining to the relevant set of interactions (i.e., market, division, or all) is greater than 0.5. The x-axis reports the number clerks behaving (dis)honestly used in estimating the models corresponding to Manuscript equation (3). Models are estimated using only those clerks who have at least one dyad with revenue greater than \$427 in the relevant subset at time $t+1$. Points represent, for each model, the average marginal effect of an additional neighbor of opposite behavior on 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. In panel (b), controls are as follows. They are those discussed in footnote 7 of the Manuscript for "market." Controls for "division" are similar but consider i's degree within her division instead of within market, and the log revenue of i's division instead of the log revenue of market m. Controls for "global" are similar to thos discussed in Materials and Methods but consider i's degree within the workplace instead of within market, and discard the other controls defined for market m. Effect sizes are about 10 times smaller when considering interactions that occur within the division or globally than when considering interactions than occur within markets, suggesting that the relevant set of interactions do indeed occur within markets.

	Move$_{t+1}$		Market max s-score$_t$		
	(1)	(2)	(3)	(4)	(5)
Clerk max s-score$_t$	−0.124***	−0.019	0.031	0.021	0.004
	(0.037)	(0.028)	(0.022)	(0.018)	(0.012)
R2	0.284	0.335	0.031	0.056	0.325
Num.Obs.	1847	1847	7004	7004	7004
Month FE	✓	✓	✓	✓	✓
Clerk controls$_t$	–	✓	–	✓	✓
Market controls$_t$	–	–	–	–	✓

Table OA5: **Identifying assumptions: clerks entering new markets.** Clerk controls are measured at month t and include clerk experience in months, log revenue, and number of hours worked that month. Market controls are also measured at month t and include log market revenue, and the number of clerks and of firms operating in that market. All models cluster standard errors at the month and clerk level. Conditional on other determinants of moving, there is no significant correlation between clerk (dis)honest behavior and entering new markets (model 2). Conditional on moving, there is no significant correlation between clerk (dis)honest behavior and the amount of dishonest behavior in the destination market (models 3 to 5). That there is no correlation between clerk behavior and entering new markets suggests that clerks are assigned to markets independently of their (dis)honest behavior.

4.3 Statistical tests for homophily

We first implement a permutation test (LaFond and Neville, 2010) that leverages the dynamic nature of our network data. Under homophily, two clerks behaving similarly should be more (less) likely to form (sever) links than two clerks with opposite behaviors. The test shows no evidence of this pattern at conventional significance levels (see paragraph below for a description of the procedure and Figure OA13 for results). We also implement a procedure outlined by Shalizi and Thomas (2011). The procedure explicitly does not condition on existing, potentially homophilous links by randomly assigning clerks within a market-month to one of two equal-sized groups, and estimates the influence exerted by clerks of group i on clerks of group $-i$. If any influence flows through the links connecting both groups, we should observe a non-zero correlation between the outcomes of group i and those of group $-i$, which we verify empirically (Figure OA14). Finally, exposure to peer influence is endogenous to network position (Aronow and Samii, 2017). Individuals that are more central in the network are more likely to be exposed to peer influence since they have more neighbors, or neighbors who are themselves more central. Our main specification addresses this by controlling for degree. We show that our results are robust to more closely-defined comparisons, by controlling for degree more flexibly and for a variety of other centrality scores (Figure OA15).

A permutation test for homophily. We use a permutation test for homophily adapted from LaFond and Neville (2010). Their intuition is simple: if there is homophily, then over time, ties should form between individuals sharing similar values for some attribute – in our case, whether clerk i behaves (dis)honestly. Their permutation procedure decorrelates one's adoption pattern from that of her neighbors. Considering periods t and $t+1$, they hold constant the number of individuals that keep and drop the attribute between t and $t+1$, but permute their identity. In other words, they randomize, among the agents that had the attribute at t, the ones who drop it and randomize, among the agents that did not have the attribute at t, the ones who adopt it. The authors consider only one graph over multiple periods. As such, their test statistic is the variation between t and $t+1$ of the χ^2 statistic of the association table between whether the dyad i and j has a tie, and whether that dyad has the same attribute value. In order to use that test jointly for several networks, we use a parametric specification:

$$g_{ijmt} = \alpha_\tau + \alpha_m + \beta_0 t + \beta_1 y_{ijt} + \gamma t y_{ijt} + \epsilon_{ijmt},$$

with $g_{ijmt} = 1$ if there is a tie between clerks i and j in market m during period t, y_{ijt} a binary variable that equals 1 if i and j have the same (dis)honest behavior at period t, and $t \in \{0, 1\}$ a binary variable that equals 0 in the first period, and 1 the next period. The model compares within months and within markets. As such, we include month fixed-effects α_τ, and market fixed-effects α_m. The statistic of interest in this test is the parameter γ that captures whether the effect of having the same attribute value increases between the first and the second period. Note that there needs to be enough individuals capable of switching behaviors for there to be enough permutations. The main analysis only considered those clerks that had at least one dyad with revenue greater than \$427 at $t+1$. We now consider only those market-months that had 5 such clerks or more at $t+1$, and consider only those clerks within the market. The expectation is that the observed test statistic is not significantly different from the statistics generated by the permutation procedure. Supplementary Figure OA13 below shows the results.

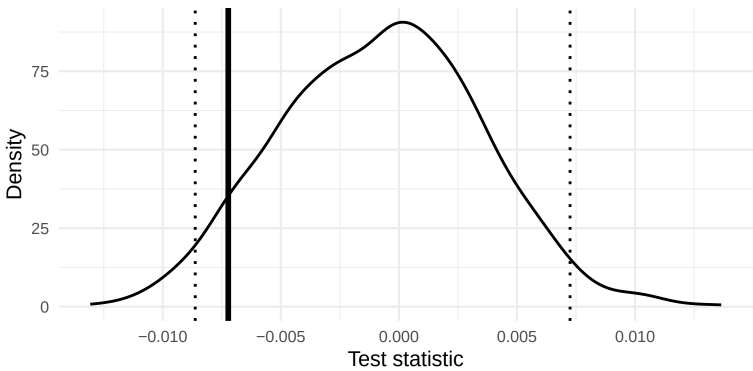

Figure OA13: **Identification - homophily: permutation test for homophily.** The figure reports the density of test statistics for a permutation test for homophily adapted from LaFond and Neville (2010), using 1,000 permutations. The black bar represents the observed test statistic, the dotted bars represent the 2.5 and 97.5 percentiles. Having similar attribute values has about no effect on tie creation, which is not significantly different from what is to be expected from a process where adoption decision is uncorrelated among neighbors.

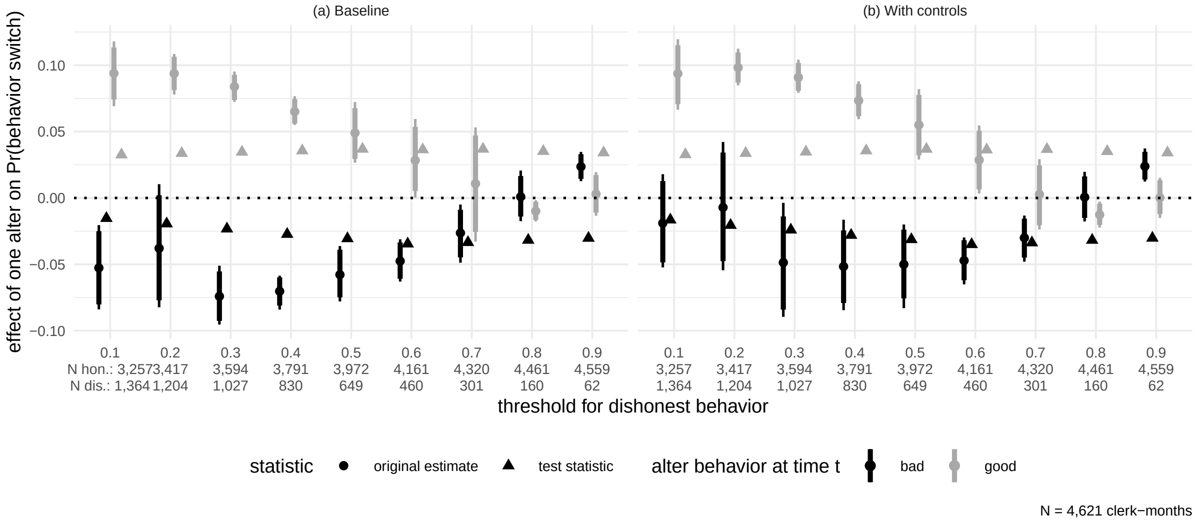

Figure OA14: **Identification - homophily: identifying diffusion from non-neighbors.** This figure reports the results of the random split procedure described in Shalizi and Thomas (2011). As per the procedure, we amend the model in Manuscript equation (3) as follows. Within each market-month, we randomly assign each clerk to one of two equal-sized groups. We then estimate the model in equation (3) considering as "peers" of clerk i those clerks that belong to the group that i does not belong to. In other words, with g_i the group of clerk i, we define n_i as the number of clerks in g_{-i} whose behavior at t is opposite to i's. We repeat the procedure 1,000 times, collecting parameters γ_0, γ_1 at each iteration. The figure reports the median γ_0, γ_1 (triangles), as well as the 95 and 90% confidence intervals (too small to show in this figure). The procedure explicitly does not condition on existing, potentially homophilous ties. That g_{-i} correlates with i indicates that some influence is at play in the network. For comparison, the Figure also reports estimates and 90 and 95% confidence intervals from our main models (circles, reproduced from Manuscript Figure 5). As expected by Shalizi and Thomas, the procedure understates the effects of peers. Influence as measured by the procedure is significantly different from zero, indicating that some influence is at play, and goes in the expected direction.

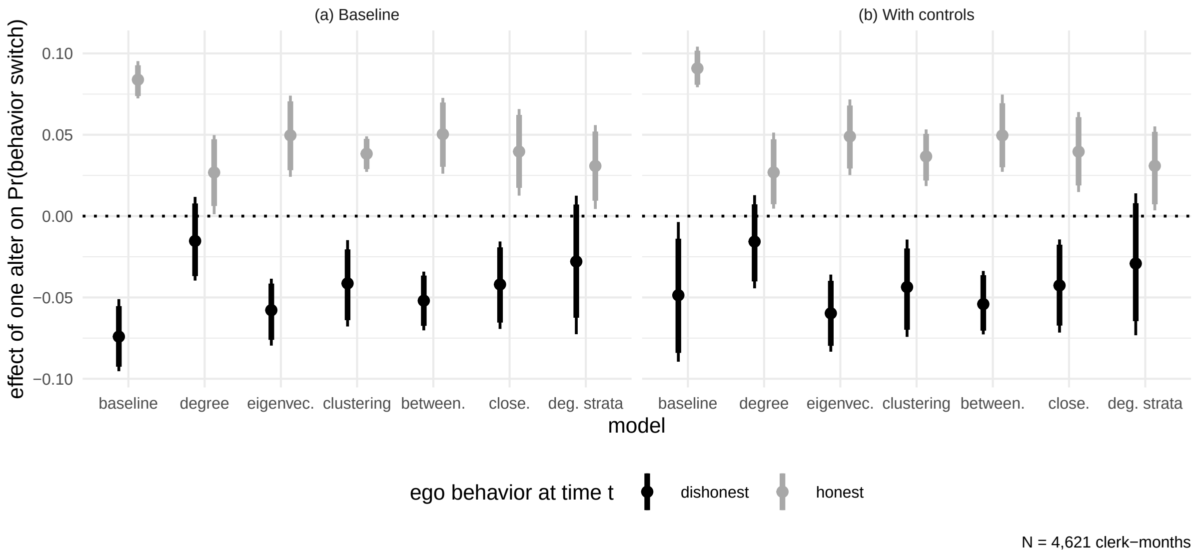

Figure OA15: **Identification: extended controls for network centrality.** This Figure reproduces Manuscript Figure 5 but introduces extended controls for network centrality. All models use a threshold of .5 in s-score to define dishonest behavior. The baseline model does not control for centrality scores. Other models control for the centrality score mentioned in the x-axis evaluated at t (degree, eigenvector centrality, local clustering, betweenness, closeness), as well as their interaction with (dis)honest behavior at t. The "degree strata" models control for degree using degree strata for degree = 0, degree = 1, and quintiles when degree > 1, as well as their interaction with (dis)honest behavior at t. Models in panel (b) further control for the set of controls mentioned in Manuscript footnote 7.

4.4 Statistical tests for the reflection problem

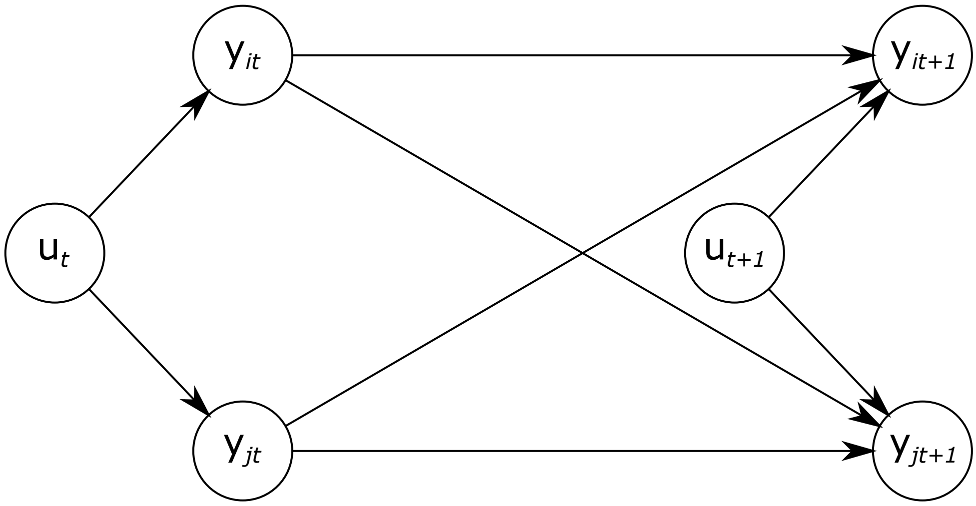

Figure OA16: **Causal directed acyclic graph of the influence process estimated by the model in equation** (3)**.** The figure considers two connected agents i, j and omits controls. Following the notation introduced in the theory, nodes $y \in \{0, 1\}$ represent the behaviors of clerks i, j at times $t, t+1$. Nodes u represent unobservables. Arrows represent causal relationships.

We discuss the rationale underlying our tests of the reflection problem. Figure OA16 represents the causal directed acyclic graph (Pearl, 2009) underlying the process of dynamic influence estimated by the model in equation (3) in the paper, omitting controls and only for two agents i, j that share a tie. As discussed in Section 4, identifying the causal effect of y_{jt} on y_{it+1} requires meeting three assumptions, which we describe formally:

1. No correlated shocks: nodes u do not contain (unobserved) shocks that may affect both i and j.

2. No homophily: that i and j share a link at t is independent of y_{it}, y_{jt}.

3. No reflection problem: the backdoor path $y_{jt} \rightarrow y_{jt+1} \leftarrow u_{t+1} \rightarrow y_{it+1}$ needs to be blocked.

Our first test changes time windows. Indeed, the backdoor path $y_{jt} \rightarrow y_{jt+1} \leftarrow u_{t+1} \rightarrow y_{it+1}$ would disappear if one removed $y_{jt+1} \leftarrow u_{t+1} \rightarrow y_{it+1}$. This last part captures the contemporaneous influence that i and j exert on one another. The magnitude of the bias should increase the more time i and j have to influence one another. Supplementary Figure OA17 below shows the results.

Our second test controls for peers' current outcomes. Doing so blocks the backdoor path from y_{jt} to y_{it+1} by controlling for y_{jt+1}. In other words, we augment the model in equation (3) with the number of peers of opposite behavior at $t+1$. Supplementary Figure OA18 below shows the results.

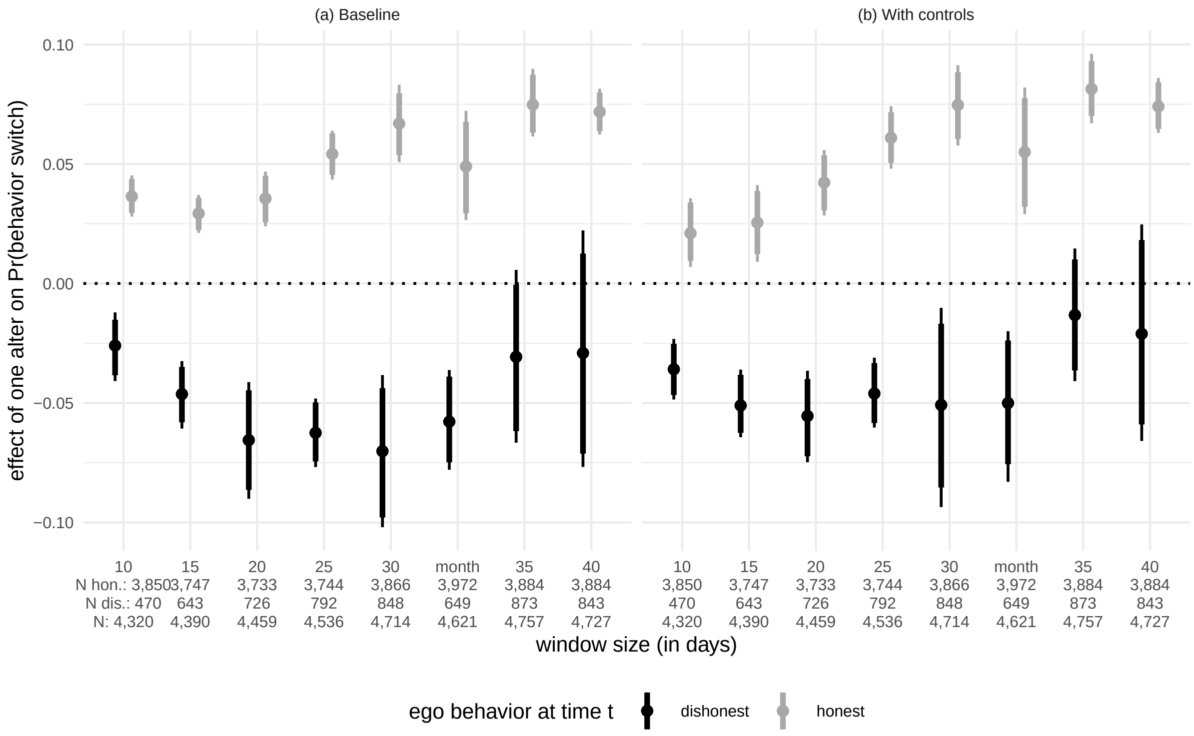

Figure OA17: **Identification - reflection problem: varying window size.** The x-axis represents the length in days used to define a time period to estimate the model in Manuscript equation (3), as well as the number of clerks behaving honestly and dishonestly implied by such threshold. Points represent, for each model, the average marginal effect of an additional neighbor of the opposite behavior on switching behavior; bars are 90 and 95% confidence intervals clustered at the period and market levels. All models include period and market fixed effects and use a threshold of .5 in s-score to define dishonest behavior. Panel (a) reports models without controls, panel (b) reports models with the controls discussed in Manuscript, footnote 7. Although magnitudes decrease as window sizes get smaller, results are qualitatively similar, suggesting that the reflection problem has little impact on results.

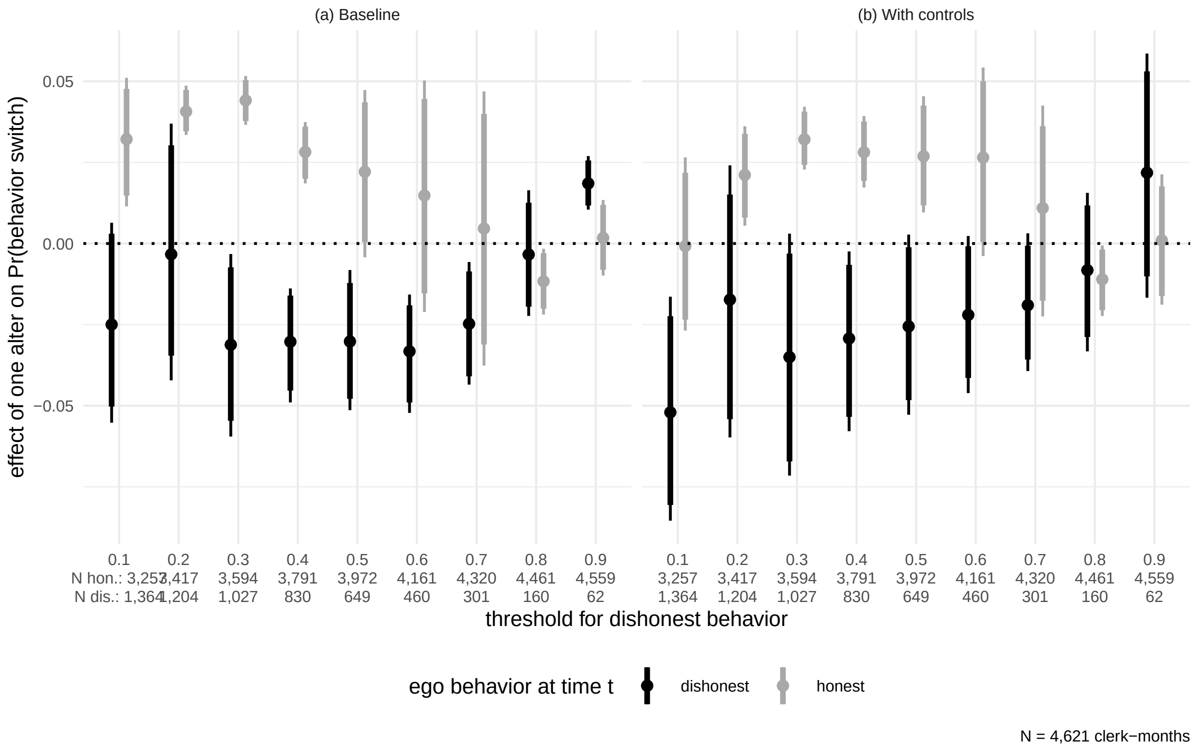

Figure OA18: **Identification - reflection problem: controlling for peers' behavior at $t+1$.** The figure reproduces Manuscript Figure 5 but controls for n_{it+1}, the number of peers from the opposite behavior at $t+1$ that were also operating in market m at time t. Additionally, models with controls (panel b) also control for x_{t+1}, the set of controls discussed in Materials and Methods in the paper, evaluated at $t+1$. Results are largely unchanged, suggesting that the reflection problem has little impact on results.

4.5 Test for selection into providers

We our concerned that our main specification, which aggregates clerk-provider interactions into a clerk-level indicator of dishonest behavior, may confound peer effects with selection into providers. It may be that clerks behaving dishonestly all collude with the same set of dishonest providers. In this environment, dishonest providers would, e.g., approach clerks and individually entice them into collusion, with between-clerks interaction playing no role. Thus, what looks like between-clerks influence would merely reflect selection into the same set of clerk-provider relationships. We rule out this possibility by testing whether between-clerk influence survives when fixing the provider of interest. To do so, we reestimate the specification in equation (3) in the Manuscript at the provider level, and show that between-clerk influence survives.

Considering clerk i operating with provider k in market m during month t, equation (3) becomes:

$$z_{ikmt+1} = \alpha_{km} + \alpha_t + \beta y_{ikmt} + \gamma_0 n_{ikmt} + \gamma_1 n_{ikmt} y_{ikmt} + \delta x'_{ikmt} + \epsilon_{ikmt} \tag{4}$$

Our dependent variable $z_{ikmt+1} \equiv 1\{y_{ikmt+1} \neq y_{ikmt}\}$ equals 1 if i's behavior *when interacting with provider k* changes between months t and $t+1$ and 0 otherwise. Similar to our main specification, we turn continuous s-scores into binary (dis)honest behaviors using a cutoff, considering this time only interactions with provider k; that is, we define $y_{ikmt} \equiv 1\{s_{ikmt} > \bar{s}\}$. We also amend our measure of social interactions to account for within-provider interactions; instead of considering co-staffing interactions, we consider co-working interactions. Consistently with our main specification, we consider that clerks i and j share a link through provider k if they have spent at least 3 hours assigning claims to the same provider within market m during month t. The variable n_{ikmt} captures the number of clerks that i shares a link with through co-working interactions with provider k and whose behavior at t is opposite to that of i. We finally perform within-provider comparisons by introducing a market-provider fixed effect α_{km}.[3]

Should results be driven by selection into providers, then we should observe no influence when considering within-provider, co-working interactions. In other words, it should be that parameters γ_0 and $\gamma_0+\gamma_1$ are not statistically different from zero. We show in Figure OA19 that this is not the case: we find that dishonest behavior diffuses while honest behavior is weakly repellent, even when considering within-provider interactions.

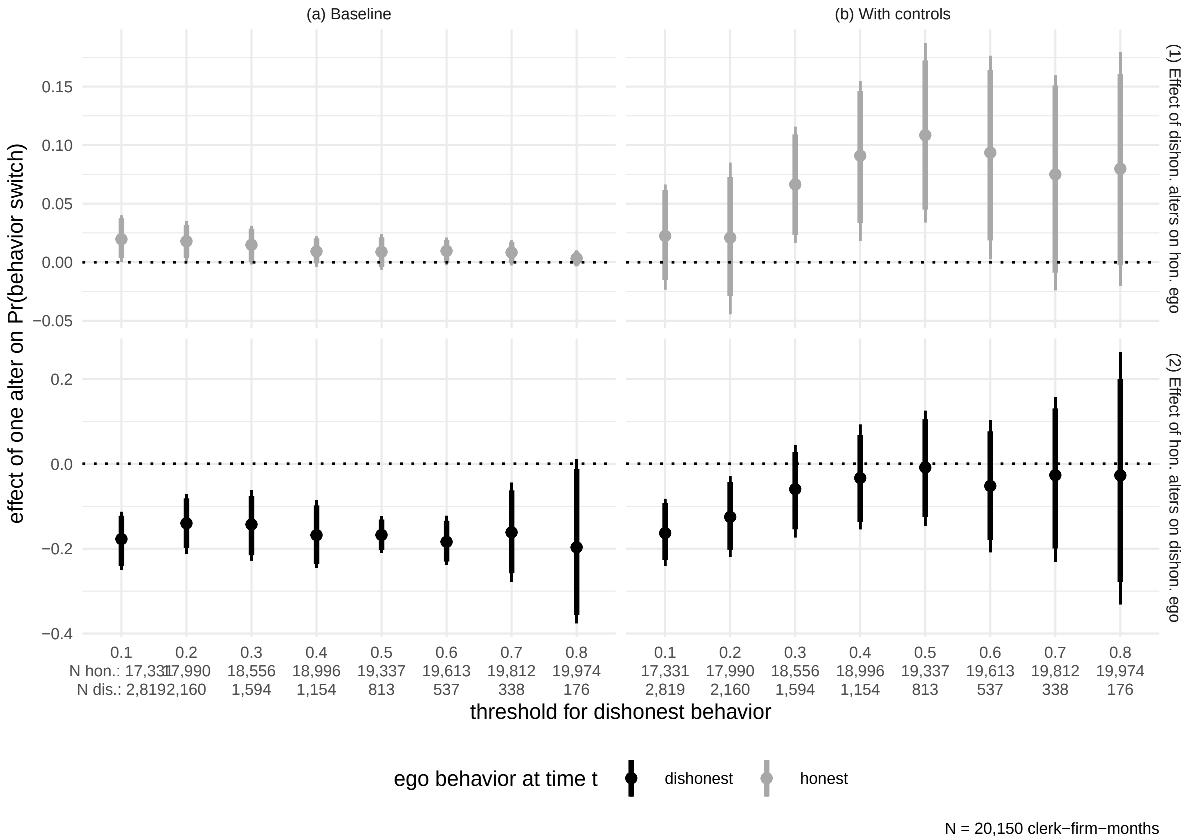

Figure OA19: **Identification: selection into providers.** This figure reports estimates of parameters γ_0 (gray) and $\gamma_0+\gamma_1$ (black) for the model reported in equation (4). It follows the same conventions as Manuscript Figure 5. Results are consistent with those reported in Manuscript Figure 5: dishonest behavior diffuses, while honest behavior is weakly repellent, even when considering within-provider interactions.

5 Robustness to alternative constructions of the data

Our main result reports estimates using different cutoffs for dishonest behavior (Manuscript, Figure 5). Here, we show robustness to considering only components 2 and 3 of the s-score (i.e., revenue distribution and gaming of the random draw system, respectively; Figures OA20 and OA21). We also show robustness to varying the threshold used to define component 1

[3]The vector of controls x_{ikmt} includes the controls discussed in Manuscript footnote 7 and further controls for provider-level degree (i.e., the number of co-working interactions) and clerk-provider revenue.

of the s-score (i.e., absolute revenue condition; Figure OA22). We also consider alternative constructions of links. While our baseline definition posits a link between i and j if they have spent at least three hours together, we show that results are robust to using cutoffs ranging from 1 to 5 hours (Figure OA23). We further leverage these cutoffs to separate weaker and stronger forms of interaction, showing that stronger forms of interaction (i.e., longer co-staffing times) are more influential (Figure OA24). As discussed above, we also examine whether social influence flows through more loosely defined forms of social interactions (Figure OA12) and consider alternative time units (Figure OA17).

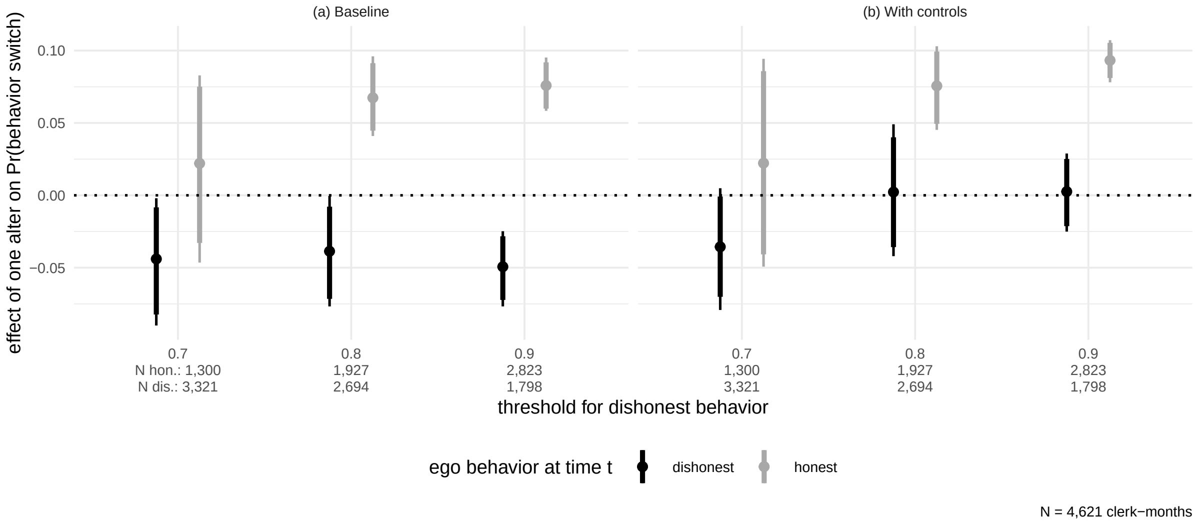

Figure OA20: **Robustness - data: using component 2 of the s-score.** This figure reproduces Manuscript Figure 5 but uses component 2 of the s-score (i.e., revenue distribution) as the dependent variable instead of the full s-score. This component is highly skewed to the right (Figure 4). For this reason, models using cutoffs smaller than .7 cannot be estimated, as they show too little variance. Results are qualitatively similar to Manuscript Figure 5.

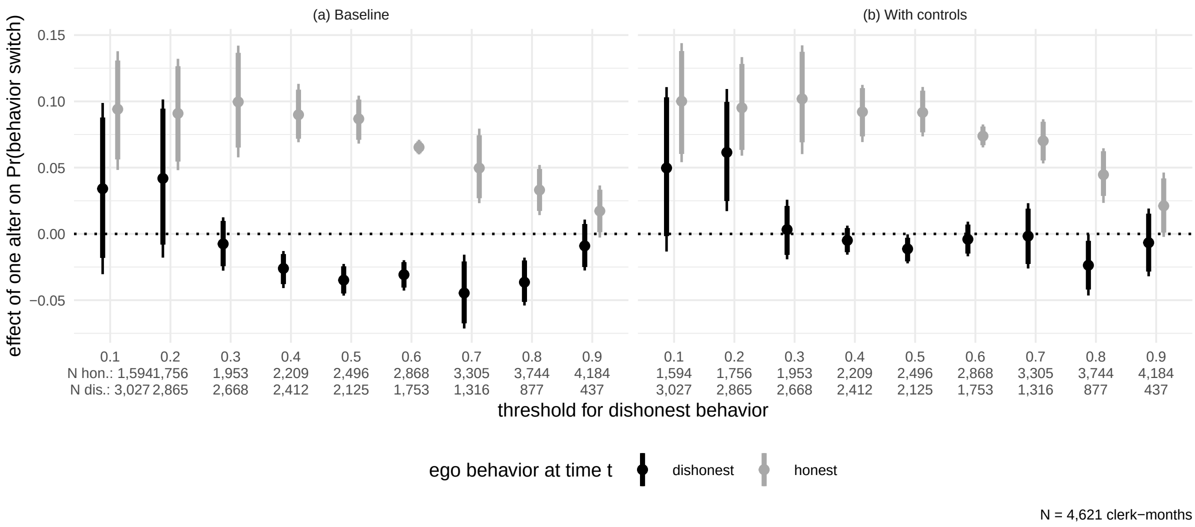

Figure OA21: **Robustness - data: using component 3 of the s-score.** This figure reproduces Manuscript Figure 5 but uses component 3 of the s-score (i.e., gaming the random draw system) as the dependent variable instead of the full s-score. Results are qualitatively similar to Manuscript Figure 5.

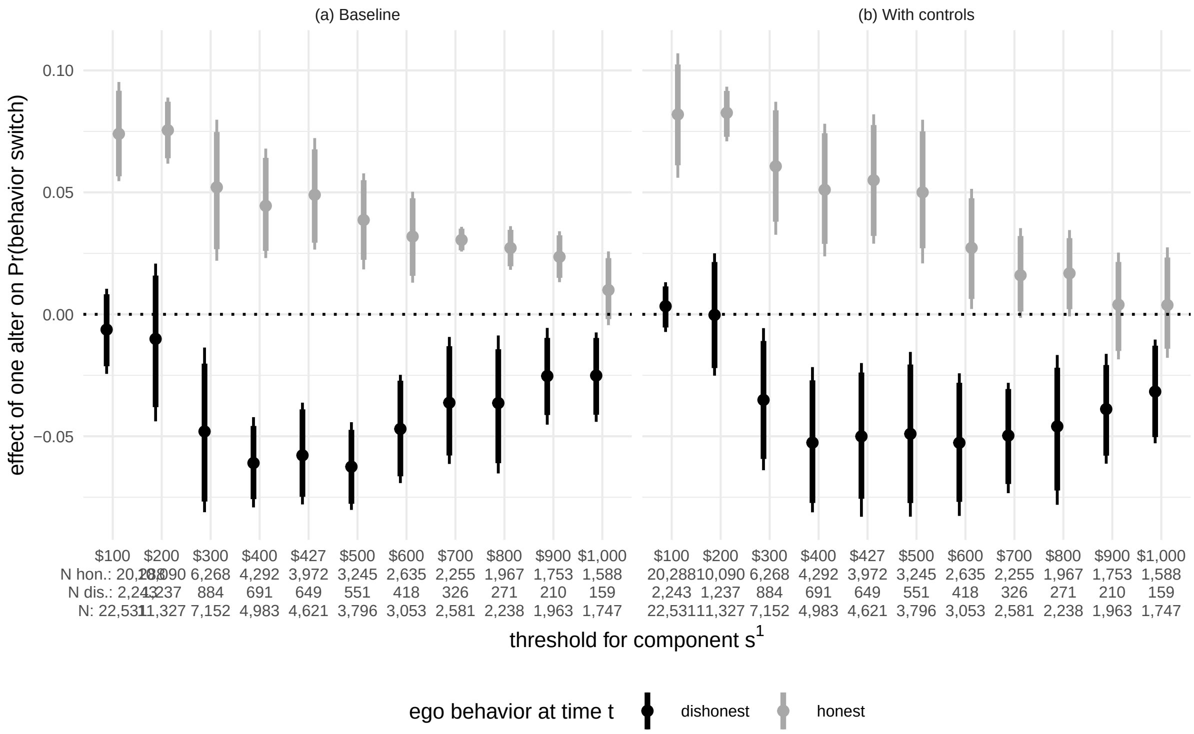

Figure OA22: **Robustness - data: varying component 1 of the s-score**. This figure reproduces Manuscript Figure 5 but varies the threshold used to define component 1 of the s-score (i.e., threshold in absolute revenue transferred to a provider). Note that using higher thresholds decreases sample size, as our estimates focus on those clerk-market-months that have at least one provider with revenue higher than the threshold. Results are qualitatively similar to Manuscript Figure 5.

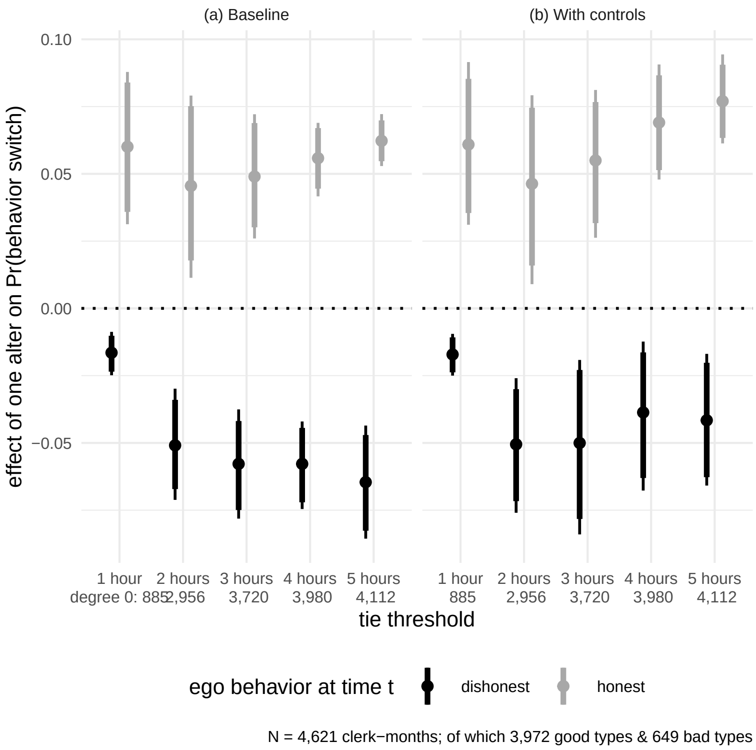

Figure OA23: **Robustness - data construction: tie threshold.** The x-axis represents the threshold in time spent operating jointly in a market used for tie definition to estimate the model in Manuscript equation (3), as well as the number of observations with no social ties implied by the definition. Points represent, for each model, the average marginal effect of an additional neighbor of the opposite behavior on switching behavior; bars are 90 and 95% confidence intervals clustered at the month and market levels. All models include month and market fixed effects, and use a threshold of .5 in s-score to define dishonest behavior. Panel (a) reports models without controls, panel (b) reports models with the controls discussed in Materials and Methods in the paper. With and without controls, effects have a similar magnitude irrespective of the threshold for dishonest behavior. For clerks behaving dishonestly, effects get stronger without controls, but are of comparable magnitude after introducing controls.

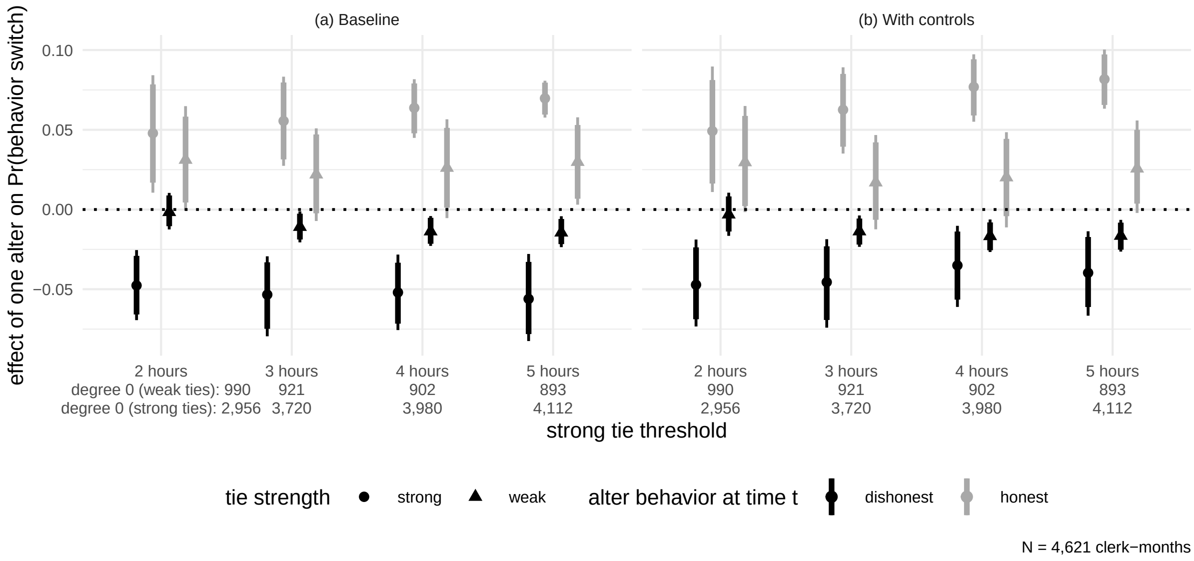

Figure OA24: **Robustness - data: strong and weak ties.** This figure augments our main model by separating interactions through weak ties and interactions through strong ties, where stronger ties capture longer monthly co-staffing time. The x-axis reports the threshold used to define strong ties (e.g., 3 hours means that co-staffing times ranging from 1 to 2 hours represent weak ties, while co-staffing times of 3 hours or more represent strong ties). Strong ties are more influential than weak ties. All models use a threshold of 0.5 in s-scores to define dishonest behavior.

6 Mechanism

6.1 Separating conformism and complementarities

In line with Boucher et al. (2024), we establish that percentage models can capture both conformism and complementarities, whereas count models do not account for complementarities. Suppose agent i undertakes a costly, continuous action y_i, with a baseline payoff function given by $v(y_i) = \alpha y_i - \frac{1}{2}y_i^2$, where the quadratic term represents convex costs. To incorporate peer influence, we define $\bar{y}_i$ as the mean behavior of i's peers and y_i^+ as the sum of their behaviors.

First, we show that percentage models capture both conformism and complementarities. Conformism can be modeled as agent i seeking to minimize the difference between her behavior and that of their peers, represented as $u(y_i, \bar{y}_i) = v(y_i) - \beta(y_i - \bar{y}_i)^2$. The quadratic term penalizes deviations from peer behavior, capturing a preference for conformity. Alternatively, complementarities can be captured by $u(y_i, \bar{y}_i) = v(y_i) + \beta y_i \bar{y}_i$, where an agent's payoff increases with the mean action of their peers. In both cases, solving the first-order conditions yields the best response function $y_i^* = \gamma + \delta \bar{y}_i$. Agents' best respones – which are what the econometrician observes empirically – are linear in mean peer behavior. Extending the reasoning to binary actions $y \in \{0, 1\}$, $\bar{y}$ is the percentage of i's neighbors taking action 1. Thus, $y_i^* = \gamma + \delta \bar{y}_i$ is a percentage model.

We then show how count models capture complementarities but not conformism. The payoff function $u(y_i, y_i^+) = v(y_i) + \beta y_i y_i^+$ introduces a term $\beta y_i y_i^+$, capturing complementarities. Solving for y_i yields a best response function of the form $y_i^* = \alpha + \beta y_i^+$, aligning with a count-based specification. However, payoffs dependent on counts $u(y_i, y_i^+)$ cannot easily incorporate distance-based utilities like $(y_i - \bar{y}_i)^2$, which are fundamental for modeling conformism. Thus, while count models accommodate complementarities, they fail to capture conformist behavior.

6.2 Additional evidence

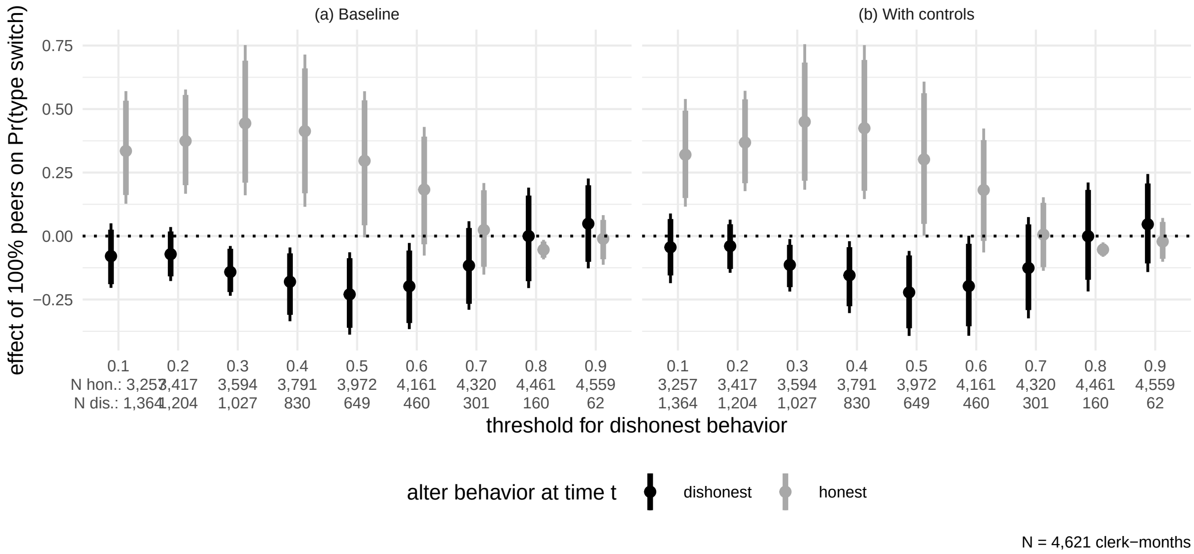

Figure OA25: **Mechanism: percentage models.** This figure reproduces Manuscript Figure 5 but uses the *percentage* of peers of opposite behavior instead of the *count* of peers of opposite behavior as the dependent variable. Results are robust to this alternative modeling approach.

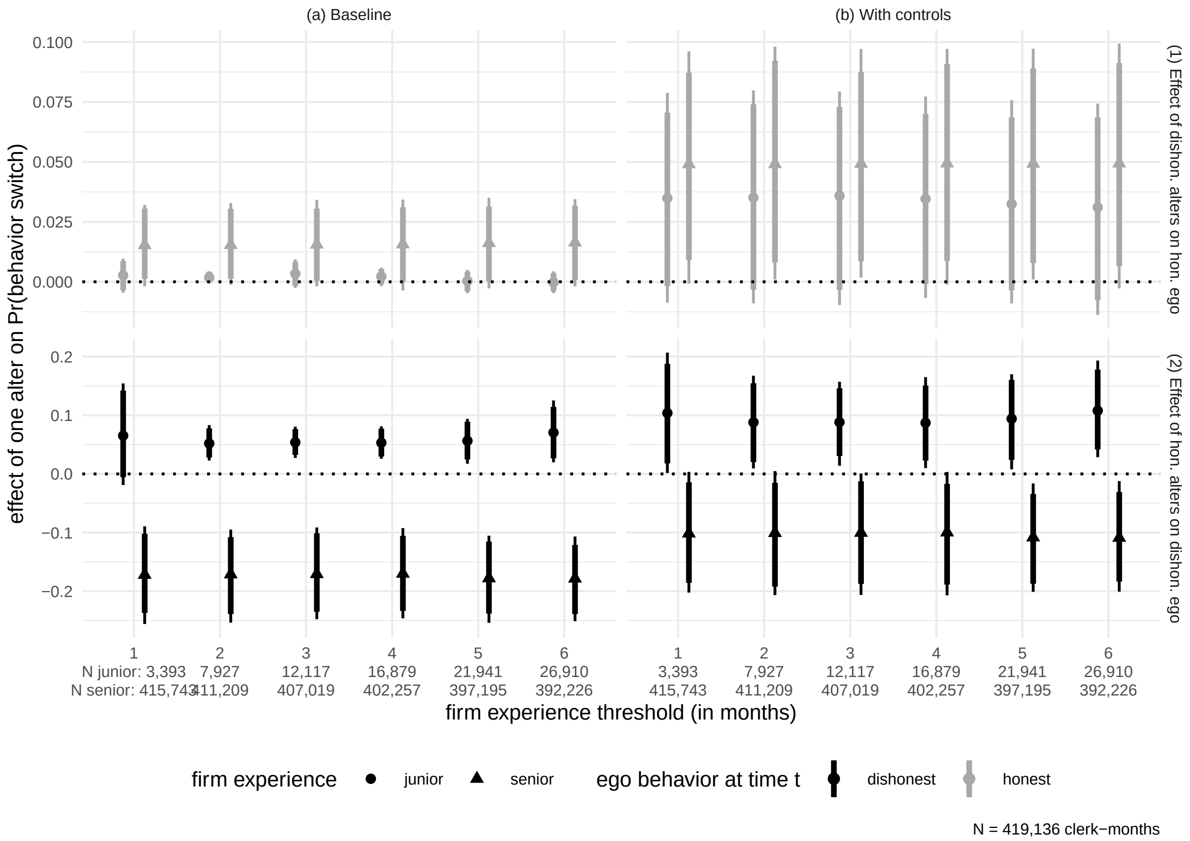

Figure OA26: **Mechanism: information - junior firms.** This figure reproduces Figure OA19 but estimates effects separately for junior and senior firms. The x-axis varies the threshold used to define junior and senior firms. All models use a threshold of .5 in s-score to define dishonest behavior. Influence is weaker for junior firms than for senior firms.

.

7 Counterfactuals

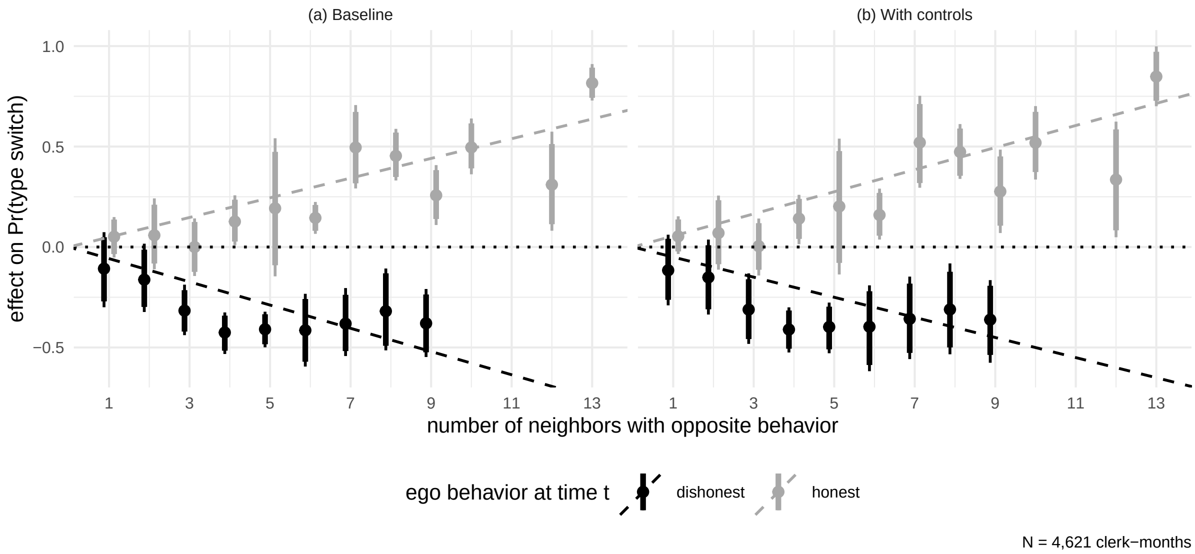

Figure OA27: **Robustness - modeling: flexible non-linear effects.** This figure estimates the model in Manuscript equation (3) but allows for flexible non-linear effects of the number of peers of opposite behavior. The x-axis represents the number of peers of opposite behavior, and the y-axis represents the effect of that number of peers of opposite behavior on the probability of switching behavior. We use a threshold of .5 in s-scores to define dishonest behavior. The dashed line represents the linear trend estimated in our main specification (Manuscript Figure 5). The fit is largely linear, although there is some evidence of concavities for the effect of alters behaving dishonestly on egoes behaving honestly at time t, and of convexities for the effect of alters behaving honestly on egoes behaving dishonestly at time t.

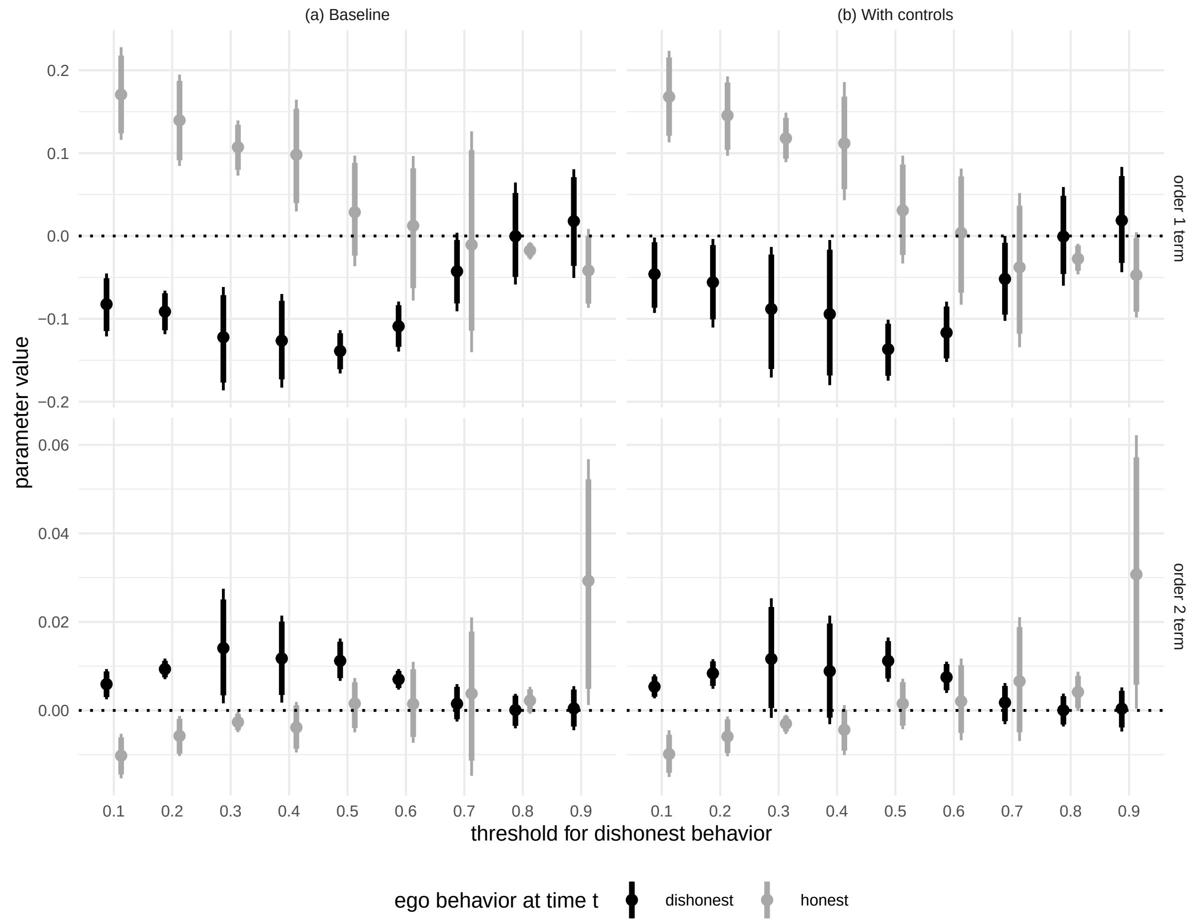

Figure OA28: **Robustness - modeling: polynomial non-linear effects.** This figure reproduces Manuscript Figure 5 but accounts for non-linear effects by estimating quadratic specifications on the number of peers of opposite behavior. Top panels report the parameter value of the first order term, bottom panels report the parameter value of the second order term. First-order effects are similar to those reported in Manuscript Figure 5. Second-order effects are not consistently different from zero. They suggest convexities for the effect of alters behaving dishonestly on egoes behaving honestly at time t, and concavities for the effect of alters behaving dishonestly on egoes behaving honestly at time t.

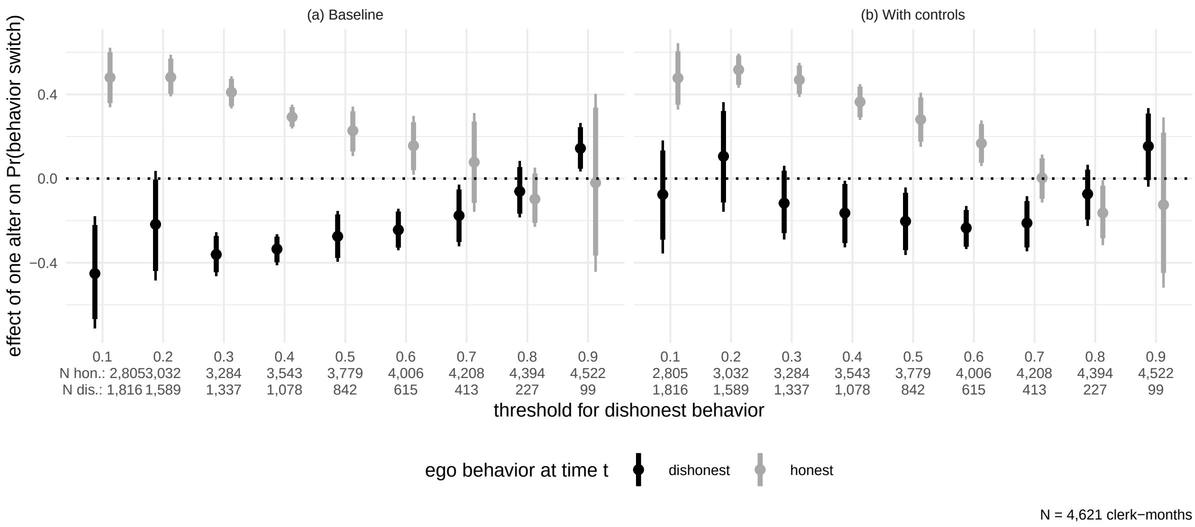

Figure OA29: **Robustness - modeling: logit models.** This figure reproduces Manuscript Figure 5 but reports logistic regression estimates of the model in equation (3) instead of ordinary least squares. Points report log-odds ratios and bars their associated 90 and 95% confidence intervals. Results are robust to this alternative modeling approach.

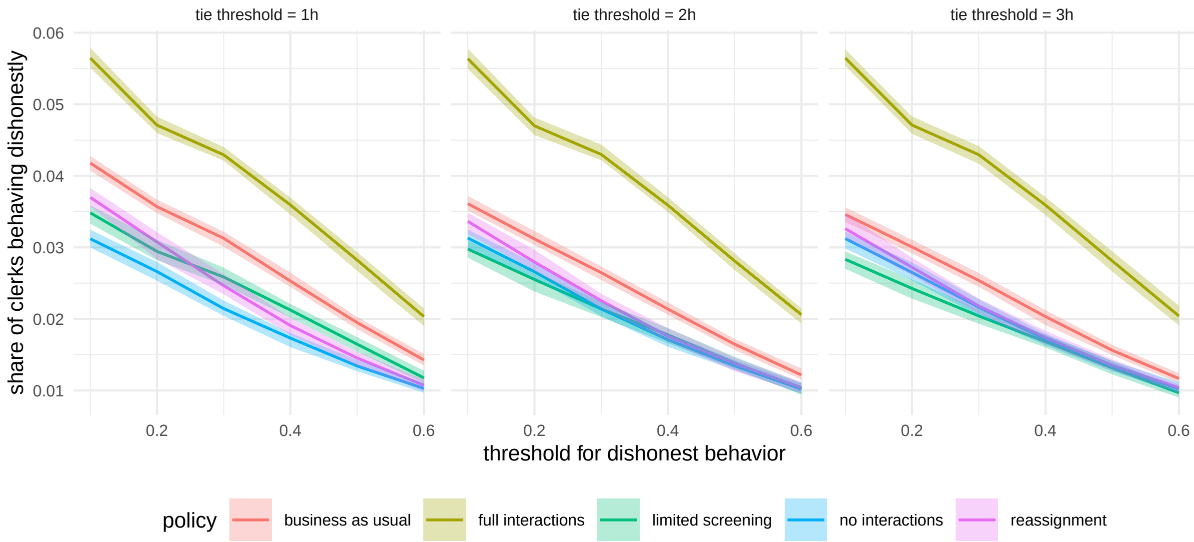

Figure OA30: **Counterfactuals: logit estimates.** This figure reproduces Figure 8 but uses estimates derived from logistic regression estimates (Supplementary Figure OA29) instead of OLS estimates. Results are robust to this change in specification.

References

Aronow, Peter M. and Cyrus Samii, "Estimating average causal effects under general interference, with application to a social network experiment," *The Annals of Applied Statistics*, 2017, *11* (4), 1912 – 1947.

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

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

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.

LaFond, Timothy and Jennifer Neville, "Randomization tests for distinguishing social influence and homophily effects," in "Proceedings of the International World Wide Web Conference (WWW)" 2010.

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

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

Pearl, Judea., *Causality : models, reasoning, and inference*, Cambridge University Press, 2009.

Shalizi, Cosma Rohilla and Andrew C. Thomas, "Homophily and Contagion Are Generically Confounded in Observational Social Network Studies," *Sociological Methods & Research*, 2011, *40* (2), 211–239. PMID: 22523436.

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