Corruption or incompetence? Disentangle through learning[1]

Evidence from a field experiment in the Mexico City Labor Court

Esteban Degetau[2]

This version: 2023-02-02

[1]This BA thesis was made possible by Joyce Sadka and Enrique Seira, who ran the Notifiers experiment, among other experimental interventions in the Mexico City Labor Court. I also want to thank Marco Medina, Sebastián García, Frederico Finan, and Ernesto dal Bó for helpful comments. All mistakes are my own.

[2]Instituto Tecnológico Autónomo de México, E-mail: esteban.degetau@itam.mx

Abstract

Labor justice in Mexico City is notoriously slow. *De jure*, labor disputes should take no longer than 3 months to be resolved. However, in practice, the average dispute will take a year before it can move to hearings stage. Since the hearings stage cannot begin until defendants have been personally notified by a court official (a notifier), hearings lag so much as notifiers are either corrupt or incompetent. This work shows preliminary evidence of a randomized control trial in the Mexico City Labor Court that seeks to evaluate whether a rotation scheme of notifiers across geographical regions can reduce corruption and increase success rates, by introducing competition among bureaucrats. Despite the increased difficulty that rotating notifiers face, they proved to be no less successful than fixed notifiers. Evidence for corruption reduction is suggested by the fact that rotating notifiers are 12.6 percentage points more successful on repeated interactions with defendants, relative to fixed notifiers on the same repeated interactions, from a baseline of 41.4%. Under the assumption that notifiers cannot become (differentially) more/less competent on posterior interactions, differences in learning rates (i.e. success on repeated interactions) may be interpreted as evidence for corruption reduction. Altogether, these results suggest the statistical zero treatment effect on success is driven by gains of corruption reduction and losses of increased difficulty.

Table of contents

Chapter 1

Introduction

It is often said that corruption is one of the most pressing issues that developing countries face. Indeed, corruption is a first order problem with political economy and development implications in Mexico. In the context of labor justice, corruption determines the speed at which suits move through procedural stages, and it also determines both *de jure* and *de facto* outcomes[1]. Hence, corruption in justice procurement has certainly redistributive implications, but it may also have efficiency implications in the real economy (Ponticelli and Alencar 2016; Boehm and Oberfield 2020; Chemin 2020).

A suitable bureaucracy for rent seeking activities is that of personal notifications in the Mexico City Labor Court. Labor law requires defendants to be personally notified by a bureaucrat (a notifier) before any hearing can be held. Often, notifiers will fail to notify because defendants pay them not to, or because they can sit on case files with notification orders to extract payments from plaintiffs, who may be willing to pay to get their cases notified and moving (Sadka et al. 2019).

[1] Kaplan and Sadka (2011) found that only 40% of convictions were materially collected by plaintiffs of unfair dismissal labor disputes in a sample of 332 finalized suits in the Mexico City Labor Court.

Notifiers get away with such malfeasance because they face no *competition*. Under the status quo, a notifier is a monopolist over a defined region of Mexico City. Since she can commit to never notify a case file, she can extract rents from either party effectively. Because case files cannot move to the next procedural stage until all defendants are persionally notified, swift labor justice is directly correlated with notification effectiveness, and the state of affairs is one of slow justice and ineffective notifiers.

In an effort to increase notification effectiveness, Sadka et al. (2019) have designed and tested a rotation scheme that seeks to reduce corruption and increase success rates by introducing competition among notifiers. The Notifiers experiment tests whether the rotation of individual notifiers across geographical regions can increase success rates and reduce corruption.

The Notifiers experiment hypothesizes that rotating notifiers will no longer be able to commit to never notify, because over the next rotation, some other notifier will have control over the case file. Hence, one could expect the rotation scheme to increase success rates and reduce corruption.

Some countries have legally required the use of rotation schemes since at least the 1990s in an effort to subvert corruption in especially susceptible public offices, such as procurement. However, no rigorous field experiment that I have knowledge of has been able to show whether rotating schemes effectively reduce corruption. Such field evidence is yet to be produced because of the difficulty that measuring corruption entails.[2]

[2] Evidence of corruption reduction under rotation schemes is only classroom experimental, but has yielded some interesting insights. Bühren (2020) found that students portraying procurement bureaucrats under rotating schemes do collect bribes, but do not reciprocate to the bribes collected and allocate contracts efficiently. Whether would-be contractors foresee such lack of reciprocity depends culturally, as noted by the cross-culture study by Bühren (2020) noted. Consistently, Barr and Serra (2010) find that corruption reduction under rotation schemes is prone to underlying societal corruption perceptions, but that these may be altered effectively by changes in the environment. Moreover, Abbink (2004) shows that rotating schemes in fact reduce corruption by half, but finds that corruption may still happen in one-shot scenarios, where bureaucrats do not benefit from reciprocating.

Often, what seems like corrupt behavior cannot be reasonably distinguished from sheer incompetence. In scenarios where a powerful party may benefit from a bureaucrat not doing her job, it is especially true that a researcher cannot call on her not doing her job as corrupt behavior, because she may well be just incompetent. In societies where corruption is rampant and may also take the form of nepotism like Mexico, this is very likely to happen as well, since there is no shortage of incompetent bureaucrats (Casar, n.d.; Olken and Pande 2012).

Measurement of corruption is such a challenge not only as parties engaging in corruption actively try to hide their behavior, but also because they might change their involvement altogether if they think they may be a subject of corruption measurement. Hence, Banerjee, Mullainathan, and Hanna (n.d.) have proposed a tasks based approach on corruption measurement. This tasks based approaches seeks to answer questions in the form of —Does policy A increase or reduce success on task X?—. This is the approach that the Notifiers experiment has carried through its evaluation and that I will try to reproduce, with a slight twist.

The rotation treatment displays two contrary effects on repeated interactions with defendants, relative to the fixed control group; (i) rotating notifiers face greater difficulty, as they must relearn how to approach defendants after every rotation, thus decreasing their effectiveness, and (ii) rotating notifiers cannot commit to stop all notification attempts of a case file after collecting a bribe from a defendant to not notify, increasing effectiveness and decreasing the effects of corruption on success. The empirical question of which of these two effects dominates success rates in the rotating setting is the main motivation for the Notifiers experiment.

Results from 67,739 notification orders in 22,990 case files processed during Phase 4 of the Notifiers experiment show no significant effect of the rotation scheme on success. Comparing within geographical regions, fixed notifiers succeeded on 42.5% of their attempts, while rotating notifiers on 41% with a non-significant difference. Does the null treatment effect on success imply that it had no effect on corruption? I argue it need not.

Indeed, it may be the case that corruption reduction in the rotation scheme is happening, but we cannot see it through success rates as the negative effect from the increased difficulty is equating the positive effect of corruption reduction. However, any measure of corruption I claim must be reasonably distinguishable from sheer incompetence for it to be valid.

A way to test whether the rotating scheme caused corruption reduction distinctly from sheer incompetence is through notifier learning. Under the assumption that notifiers cannot become differentially more/less successful upon repeated interactions with defendants as a result of differences in incompetence, differences in learning rates may be interpreted as differences in corruption.

Because corruption requires trust among participants, as its agreements are non-enforceable (Bühren 2020; LaPorta et al. 1997), defendants who have been acquainted to a notifier may be more eager to trust them to follow through their commitment not to notify in exchange for a payment. Rotating notifiers should find it harder to establish personal relations with defendants that allow them to collect bribes, as they do not visit the same region every day. Since fixed notifiers can build corruption networks more easily in their defined region, they may stop notifying defendants they have been acquainted to, while rotating notifiers keep attempting to notify.

The validity of this comparison is argued as follows. Some defendants are firms large enough to have several suits filed against them during the course of Phase 4, and indeed large enough to be visited by fixed and rotating notifiers more than once. Moreover, because the beginning of Phase 4 saw all notifiers assigned to new regions, fixed notifiers had to build trust with their defendants as much as rotating notifiers had to upon rotation. Because time will make fixed notifiers more familiarized with the large firms in their region, they will stop notifying these firms while rotating notifiers will not, as a result of differences in trust building with defendants.

Results show notifiers in both treatment groups were more successful on later interactions with defendants. However, rotating notifiers become more successful more rapidly, i.e. they learnt more rapidly than fixed notifiers. The differential learning rates across treatment arms reflect a *wedge of missed opportunities* that may be interpreted as corruption reduction in the rotating scheme of 12.6 percentage points, from a baseline of 41.4%. This final statement rests under the assumption that notifiers cannot learn differentialy as a result of differences in incompetence.

Although the wedge of missed opportunities did not hold to some robustness checks, notifiers in the rotating condition learnt at no slower pace relative to notifiers in the fixed condition. Since rotating notifiers also showed behaviors consistent with corruption reduction, as increased traceability and in-proximity rates, as well as reduction of illegal success on repeated interactions, the losses of the wedge under robustness checks arguably may be driven by lack of power, as rotating notifiers indeed had very few chances to visit the same defendant more than once.

Altogether, results suggest that any corruption reduction in the rotating arm was completely offset by the increased difficulty that rotating notifiers face, resulting in a null net effect on success. Analysis of repeated interactions under rotation schemes proved to be a limited tool for corruption measurement, as indeed rotating bureaucrats simply do not have the chance to engage in trust-related corruption.

This work is organized as follows: Section 2 describes the data used and the setting of the Notifiers experiment; Section 3 shows main results as well as several robustness checks; finally, Section 4 concludes.

Chapter 2

Data

The main source of data used in this work comes from Phase 4 of the Notifiers experiment conducted by Sadka et al. (2019). Phase 4 ran from November 2021 to July 2022 and consisted of 50 notifier assigned to 25 geographical regions, half in the rotating scheme and half in the fixed comparison scheme. Data from the Notifies experiment is best described as the administrative data of the Central Notifications Office of Mexico City Labor Court. Data were recorded automatically through the Notifiers App developed for the Notifiers experiment.

A secondary source of data is the National Firm Directory (INEGI 2021) and the National Firm Census (INEGI 2019). The secondary data source was necessary to gather defendant level (firm level) features to include as covariates in some specifications and was matched to administrative data by defendant name, or address. Appendix 1: DENUE Match expands on matching procedures and shows summary statistics about firm features.

This Chapter is organized as follows: Section 2.1 describes the Notifiers experiment design and implementation; Section 2.2 briefly describes a theoretical model of corruption in the notifiers setting; Section 2.3 shows preliminary results of the Notifiers experiment to motivate the analysis in this work; finally, Section 2.4 shows summary statistics of the variables that I leverage to show corruption in the Notifiers experiment.

2.1 Setting

The Notifiers experiment was conducted in the Mexico City Local Labor Courts to test whether the rotation of individual notifiers across geographical regions and the compulsory use of a notifications app would increase success rates and decrease corruption.[1] To test this hypothesis, Sadka et al. (2019) ran a randomized control trial based at the Central de Diligencias Actuariales (roughly, Central Notifications Office).

Several phases of the experiment have taken place thus far. A phase in the context of the Notifiers experiment is defined as a period of time through which no notifiers changed conditions and where geographical regions suffered no important changes. Phase 4 of the Notifier experiment transpired from November 2021 to July 2022 and it consisted of 50 notifiers assigned to one of two treatment arms and to one of 25 geographical regions.

Experimental design

Case file population. Every day some 120 lawsuits are filled in the Court, out of which 95% claim unfair dismissal. The experiment focused on the subset of case files claiming unfair dismissal and that had at least one first notification order pending.[2]

[1] A notification order is considered successful if the defendant it is addressed to is properly notified of the trial. Notification is legally required to be delivered physically at the defendant's address by a notifier who has to make sure the defendant is there at the time of notification.

[2] First notifications are required by law before a suit can proceed to hearings stage. The average case file will take a full year before it has finalized all its first notifications and can move to the hearings stage.

Additionally, the Notifiers experiment included older case files with at least one first notification pending. By the end of Phase 4, the Notifications Office had processed 22,990 case files.

Notifier population. During Phase 4, the Notifications Offices operated with 50 notifiers, 25 were assigned to the treatment condition and 25 to the fixed condition. Assignment between rotation and fixed conditions was done as follows:

1. The Labor Court's jurisdiction (Mexico City) was divided into 25 geographical regions, considering their area, case file load and number of addresses, according to data on previous phases. The process through which the jurisdiction was split is beyond the scope of this work. The resulting region division can be seen in Figure 4.2.
2. Make pairs of two similar notifiers. Similarity within pairs was determined by a survey at phase baseline. Notifier pairs were then randomly assigned to each of the 25 regions.
3. Through a public lottery, randomly assign pair members to a rotation condition. So, within a pair of similar notifiers, one notifier would be assigned to the rotating condition and the other pair member to the fixed condition.
4. Verify notifiers features are balanced across rotation conditions. See Table 4.1.

Treatment arms. To test the hypothesis that the rotation of individual notifiers across geographical regions will increase success rates and decrease corruption, notifiers in the Notification Office were assigned to one of two conditions.

1. **Fixed condition:** Notifiers in the fixed condition are called "control notifiers" and they are fixed to their assigned geographical region. Within Phase 4, no fixed notifiers switched regions. The fixed condition resembles the status quo of notification operations in the Court, where a single notifier is the monopolist of a given region.

2. **Rotating condition:** Notifiers in the rotating condition are called "treatment notifiers" and they rotate across geographical regions. The particular rotation of treatment notifiers across regions is assigned randomly every other day.

Notice that under the Notifiers experiment, two notifiers are assigned to every region on any given work day; the region's fixed notifier and a randomly selected rotating notifier.

Additionally, although the fixed condition is set to simulate the status quo of the Court, some non-trivial differences may arise by the sheer implementation of the Notifiers experiment. Indeed, the fixed condition may not perfectly resemble the Court's status quo since the notifiers under the experiment face additional monitoring and administrative control by the research team and by the Notifiers app in the Notifications Office.

Implementation

Case file level treatment assignment. Case files arriving to the Notifications Office went through the following process.

1. *Eligibility screening.* Eligible case files had at least one first notification pending within the Court's jurisdiction.
2. *Automatic random assignment to either treatment arm.* A case file assigned to, say the treatment arm, remained in the treatment arm for its entire procedural life. Notice in Table 4.2 that casefiles are balanced across treatment arms.
3. *Automatic assignment to geographical regions.* For each notification order in a case file[3] a region was assigned automatically to match addresses by zip code.

[3] A case file may have several first notification orders if more than defendant is sued or if a defendant is sued in more than one address.

4. *Storage by treatment, region and hearing date.* Eligible case files get stored by their treatment arm, region and hearing date, so that the most urgent case files would be sent to route first.

Routes. Every business day, the research team assembled 25 sets of case files to be executed by notifiers in either treatment arm over the following days. The routes team included the most urgent notification orders from storage, added some case files in the vicinity of urgent orders, and delivered the route to its corresponding notifier. One day, routes are assembled for rotating notifiers and the next day for fixed notifiers. Routes are defined within regions and constitute the main mechanism through which rotating notifiers are rotated across regions while fixed notifiers remain fixed.

Figure 4.1 shows that rotating notifiers indeed were randomly rotated across geographical regions. Fixed notifiers visited their regions up to 60 times throughout the Phase 4, while rotating notifiers visited at most 7 times the same region. Notice, however, that fixed notifiers occasionally visited different regions. This is the case because Covid related personnel shortages were adverted by adding case files to routes from adjacent regions.[4]

Table 4.3 shows that routes assigned to notifiers are balanced across treatment arms. It also shows that the average route had 26 notification orders in 10 case files to 24 defendants in 10 addresses.

2.2 Corruption in the Labor Notifications setting

Having established what the Notifiers experiment is and how it was designed and implemented, this section turns to describe the behaviors that the Notifiers exper-

[4] To guarantee anonymity of notifiers, Region IDs showed in Figure 4.1 were randomized and need not match the division numbers in Figure 4.2.

iment seeks to change. Namely, corruption. Corruption in the labor notifications setting can generally have two forms.

1. **Avoidance bribes** are paid by defendants to notifiers to avoid a notification. Because case files cannot begin the hearings process until all defendants have been properly notified, defendants will try to avoid notifications to delay trails through bribes to notifiers.
2. **Encouragement bribes** are paid by plaintiffs to notifiers to encourage notification efforts. It is common for them to be referred to as “apoyo” or “ayuda” (roughly, assistance) because they represent a payment for the notifier’s troubles in physically attempting to notify[5].

It is useful to think about notifiers as auctioning whether they attempt a notification. If a notifier can collect a larger payment on an avoidance bribe, she will not notify. If she can collect a larger payment on an encouragement bribe, she will attempt to notify.

Transmission mechanisms

How does the rotation scheme cause success rates to increase while decreasing corruption? Two main transmission mechanisms are considered.

1. **Notification avoidance costs:** Under the rotation scheme, notification avoidance becomes costlier to defendants. Relative to notifiers in the fixed condition who can commit to never notify after collecting avoidance bribes, rotating notifiers may not stop trying to extract avoidance bribes after some other notifier previously collected avoidance bribes. Because defendants need to keep bribing rotating notifiers every so often to ensure avoidance, this increased cost may deter them from paying further avoidance bribes.

[5]Notifiers have told me how costly it is to be out and about notifying. In many cases, they pay for transport costs out of their own pockets.

2. **Control over case files**: Under rotation, notifiers have less control of case files to extract bribes. Rotating notifiers receive routes to a random region every other day. As such, they cannot expect to return to the same region to execute orders that they had not executed from earlier routes. Relative to notifiers in the fixed condition, they have fewer time to conduct the auction to extract bribes from either side.

It is important to note that success will be the main observable variable through which we will determine the results of the Notifiers experiment. Because corruption will not be observed, the next best impact evaluation measure will be success rates. However, the description of corruption is helpful to think clearly about the link between corruption and success rates. If notifiers under the rotation scheme show higher success rates relative to notifiers under the status quo, we will be able to say that that increase was driven by corruption reduction in the terms outlined above.

2.3 Preliminary results from the Notifiers experiment

Now that the reader has the background of the Notifiers experiment and the theoretical models of bribe collection in the notifiers context, I will show some preliminary results of the notifiers experiment and discuss the difficulty of reaching a fulfilling answer on whether the rotation scheme caused corruption reduction. Note, however, that this work is about corruption measurement and as such, I will only show preliminary results of the Notifiers experiment as part of the motivation of this work.

Results are shown for the average treatment effect on the main outcome variables defined as:

- **Success:** a dummy that takes a value of 1 if the notifier was successful in notifying a defendant in a particular address and 0 otherwise.

- **Traceable:** a dummy that takes a value of 1 if the notifier had her phone's location services turned on during the time she attempted to notify, and 0 if she was not traceable through her phone's location services.
- **In proximity:** a dummy that takes a value of 1 if the notifier was traced to be in a vicinity during the attempt and takes avalue of 0 otherwise.
- **Has photo:** a dummy that takes a value of 1 if the notifier attached a photo to the notification order and 0 otherwise.
- **Illegal success:** a dummy that takes a value of 1 if the attempt was successful while the notifier was traced far away of the address vicinity during the time she notified. It is indeed illegal to notify a defendant without personal attendance to the address.

Table 4.4 shows the main results of Phase 4 of the Notifiers experiment. Most shockingly, the rotation scheme did not increase success rates. Fixed notifiers succeeded on 42.5% of attempts while rotating notifiers only on 41%. The difference in success rates is not statistically significant.

Despite the null treatment effect on success, rotating notifiers indeed behaved differently than notifiers in the fixed condition. They were traceable more often, reached the addresses more often, and attached photos more often, while having no greater illegal success rate. Overall, these results are consistent with more honest behavior. What meaning can be attributed to the null treatment effect on success as a corruption outcome?

The zero treatment effect on success may be explained by one of two stories, either (i) the rotation scheme had no effect on corruption and the effect from increased difficulty is negligible, or (ii) there is a positive effect on success from fewer notification avoidance bribes but it is offset completely by the increased difficulty treatment notifiers face of attempting notifications on different geographical regions every other day.

Suppose the success of a notification order i is determined heterogeneously depending on the assignment of notification order i to treatment arm $T_i \in \{\text{Treatment}, \text{Control}\}$ as in Equation 2.1.

$$
\begin{aligned}
Y_i = \; & \alpha + \beta \text{Difficulty}_i + \gamma \text{Avoidance}_i + \phi \text{Encouragement} + \eta T_i \\
& + \delta \text{Difficulty}_i \times T_i + \lambda \text{Avoidance}_i \times T_i + \psi \text{Encouragement}_i \times T_i + \epsilon_i
\end{aligned}
\tag{2.1}
$$

A hypothesis that may explain the null treatment effect on success is that neither difficulty, avoidance or encouragement bribes are heterogeneously determining success by treatment ($\delta = \lambda = \psi = 0$). An alternative hypothesis consistent with the observed null treatment effect is that the rotation scheme indeed decreases the negative impact of avoidance bribes on success ($\lambda < \gamma;\ \lambda,\ \gamma > 0$) but this effect is offset by the increased difficulty of notification under the rutation scheme ($|\delta| > |\beta|;\ \delta,\ \beta < 0$). How can we distinguish between these two alternative hypotheses without direct observation of avoidance bribes?

2.4 Repeated interactions with defendants

Let learning rates be defined as the change of success rates on repeated interactions to address-defendants.[6] Suppose success Y_i of notification order i is determined by the number of visits V_i a notifier has paid to a defendant before, as in Equation 2.2.

$$
Y_i = \alpha + \beta V_i + \epsilon_i \tag{2.2}
$$

[6] The notation "address-defendant" symbolizes the destination of a notification order, which is composed by a defendant (a firm or an individual) to be notified at a particular address. Filed suits may sue several defendants at several addresses each.

Then α shows average success rates of notifiers on their first visit to defendants and β represents the average learning rate of notifiers in the experiment.

A clear piece of evidence for widespread avoidance bribes collection and enforcement would be for $\beta < 0$. This would mean that notifiers succeed at a lower rate on later visits, implying that after getting acquainted with defendants, notifiers stop notifying defendants that they previously notified. Because notifiers previously notified successfully, we may say that the reduction of success rates is attributed to avoidance bribes rather than incompetence.

A way to assess whether the rotation scheme indeed caused corruption reduction while maintaining success rates is to examine heterogeneous learning rates. Because defendants may be more eager to pay avoidance bribes to notifiers they know well,[7] and because fixed notifiers are more likely to be known in their region, differences on learning rates may be driven by differences in avoidance bribe collection across treatment arms.

Summary statistics

Despite the fact that rotation schemes are designed to discourage trust building between bureaucrats and users, Phase 4 of the notifiers experiment lasted long enough for rotating notifiers to visit defendants up to three times, as can be seen in Table 4.5. Although most notification orders are located in the first visit, 4,114 notification orders happened between the second and third visits. These are the observations that will identify repeated interactions with defendants and are the primary source of variation that I leverage to show corruption reduction in the rotating scheme.

An important fact to note about of the repeated interaction approach to corruption measurement is that posterior visits happened on average later for rotating notifiers.

[7] Bühren (2020) and LaPorta et al. (1997) have stressed the importance of trust in corruption incidence. Because corrupt contracts are non-enforceable, parties can only expect their contracts to be fulfilled if they trust their partner. Trust among parties is the key mechanism that rotation schemes exploit, since users cannot punish opportunistic bureaucrats in future interactions under rotation.

Figure 4.3 shows the modes for interactions with defendants are farther apart from each other in the rotating scheme than in the fixed. Figure 4.4 shows that indeed control notifiers are more likely to revisit a defendant after fewer days had passed.

Additionally, Appendix 2: Address-defendant identification elaborates on the identification of address-defendant pairs and includes descriptive statistics on the distribution of notification orders by address-defendant pairs and visits.

Chapter 3

Results

To show whether the rotating treatment had an effect on corruption that can be reasonably differentiated from sheer incompetence, I first estimate the probability of success conditional on the number of visits a notifier has paid to specific address-defendants. This exercise is included to show the reader differences in learning rates visually. I estimate the conditional probability with a parsimonious linear Probit model as follows:

$$\mathbb{P}(\mathrm{Y}_i = 1 | \mathrm{V}_i = v_i) = \Phi(\alpha_{\mathrm{T}} + \beta_{\mathrm{T}} \mathrm{V}_i + u_i) \tag{3.1}$$

Where Y_i is the main outcome variable, V_i is a variable that takes the value of 0 the first time a notifier attempts to execute an order on an address-defendant, and a subsequent integer on subsequent attempts. Hence, α_{T} is the probability of success on the first visit for notifiers in the treatment arm $\mathrm{T} = \{\text{Treatment, Control}\}$. β_{T} is the average change in success rates when executing posterior orders to address-defendants. Thus, β_{T} is the average learning rate of notifiers in the treatment arm T when visiting address-defendant pairs.

Figure 4.5 shows results of specification in Equation 3.1. Rotating notifiers are as likely to succeed an attempt the first time they visit an address-defendant as fixed notifiers. However, on posterior visits, rotating notifiers are more likely to succeed relative to control notifiers, i.e. they learn at a higher rate ($\hat{\beta}_{\mathrm{T}} > \hat{\beta}_{\mathrm{C}}$) on repeated interactions with address-defendant pairs.

The differential learning rates across treatment arms shown in Figure 4.5 shows a *wedge of missed opportunities* that portrays notification attempts where control notifiers could have succeeded at, but failed. Because notifiers are balanced across treatment groups (see Table 4.1), it should be reasonable to suggest that the wedge of missed opportunities is not driven by differential incompetence. Rather, the wedge of missed opportunities must be driven by differential behaviors unrelated to incompetence. The next section will show evidence consistent with avoidance bribe reduction in the rotating scheme.

3.1 Heterogeneous learning rates

As illustrative as Figure 4.5 can be, specification in Equation 3.1 has room for refinement. I estimate the heterogeneous treatment effect of success on repeated visits (heterogeneous learning rates) with the following specification:

$$\mathrm{Y}_i = \alpha_R + \beta \mathrm{T}_i + \gamma \mathrm{V}_i + \delta \mathrm{T}_i \times \mathrm{V}_i + \epsilon_i \tag{3.2}$$

Where Y_i is an outcome variable at notification order level. T_i is a dummy that takes a value of 1 if observation i belongs to a rotating case file, and 0 if it belongs to a fixed case file. α_R is a region level fixed effect. V_i is defined as above. $\hat{\alpha}_R$ represents the success rate that the control notifier assigned to region R has during her first visit to an address-defendant.

Because control notifiers are fixed to one of 25 geographical regions, including region level fixed effects means comparing the fixed notifier in each region to all the rotating notifiers who visited their region. Estimating the slope of the interaction term, δ shows whether treatment notifiers learn at a different rate vis-a-vis control notifiers, thus accounting for region level, time-invariant unobservable variables.

Nonetheless, rotating notifiers only visit a particular region a few times over the duration of the experiment. Hence, the range of repeated visits to address-defendants they show is much shorter than what control notifiers show (see Table 4.5). Because notifiers cannot learn linearly infinitely, diminishing learning rates may subdue learning coefficients for the longer ranges that fixed notifiers show. This concern is addressed by including only observations with repeated visits where both treatment arms have registered attempts, i.e. comparable visits.

Table 4.6 shows consistent results with those shown in Figure 4.5. Success rate during the first attempt to address-defendants was 41.3% regardless of the treatment arm. Although notifiers in both treatment arms are more successful on later visits to the same defendant in the same address, rotating notifiers learn 10.1 percentage points faster than their fixed peers (see column 1).

Table 4.6 also shows learning rates on other outcome variables. Note in column 3 that rotating notifiers not only reach proximity more often during their first attempt on any address-defendant, but also increase the rate at which they reach proximity on later attempts. This is a stark comparison to fixed notifiers who indeed reach proximity less often on later attempts.

Under the assumptions of linear learning rates and no differential incompetence, the results shown in this section present strong evidence for corruption reduction caused by the rotating scheme. Relative learning slowness in the control group can be interpreted as failures driven by greater collection of avoidance bribes, as fixed notifiers can be seen as more trustworthy to fulfill their commitment to avoid notification on later notification orders. Supporting evidence for this thesis is that rotating notifiers are able to show up more often on later attempts, while their fixed peers indeed stop showing up to notify, as they have reached agreements with the defendants in their regions. In the next subsection I relax the assumption of linear learning rates, and address some other concerns with specification 3.2, namely firm-level heterogeneity, actual visits and days between visits.

3.2 Robustness checks

Although the methodological strategy leveraged above is powerful enough to show heterogeneous learning rates across treatment arms on repeated interactions with defendants, its corruption reduction interpretation rests on some assumptions. Thus far, I have assumed (i) that notifiers learn linearly on repeated interactions with defendants; (ii) that heterogeneous learning rates show differences in notifier-defendant behavior, rather than heterogeneous firm composition; (iii) that notifiers learn on repeated interactions regardless of whether they actually visited defendants; and (iv) that notifiers do not learn differentially as a sheer result of time elapsed during the experiment.

In this section, I will test whether the heterogeneous learning rates hold while relaxing each of these assumptions. Note, however, an assumption I will not be able to relax is that of *non-differential incompetence*. I cannot show that there does not exist another mechanism that explains the heterogeneous learning rates other than corruption reduction in the treatment arm. Thus, the corruption reduction interpretation of the differential learning rates rests on the relatively strong assumption that notifiers cannot become *differentially* more/less competent on posterior interactions with defendants.

Nonlinear learning rates

In this section, I relax the assumption of linear learning rates. One may worry about the consistency of the heterogeneous learning rates on each visit, and that notifiers cannot learn linearly indefinitely upon repeated interactions with defendants. Consider the following specification:

$$\mathrm{Y}_i = \alpha_R + \beta \mathrm{T}_i + \gamma_v \mathbb{1}(\mathrm{V}_i = v_i) + \delta_v \mathrm{T}_i \times \mathbb{1}(\mathrm{V}_i = v_i) + \epsilon_i \tag{3.3}$$

Where the visits variable V_i is taken as a factor under the indicator $\mathbb{1}(\mathrm{V}_i = v_i)$. Then, γ_v and δ_v show the learning rates during visit v for fixed and rotating notifiers, respectively.

Table 4.7 shows estimations for nonlinear learning rates on our outcome variables, as specified in 3.3. Overall, results show notifier behavior under rotation is consistent with corruption reduction and linear learning rates was not such a strong assumption.

Some interesting details are worth noting from results in Table 4.7, though. Rotating notifiers learnt at a higher rate on every posterior interaction with defendants, relative to fixed notifiers, while succeeding at no lower rate during the first interaction. However, during the third interaction, learning estimates for rotating notifiers are positive and higher than for fixed notifiers, but not statistically significant. Such is the case as not enough power is left from the 14 observations left in this interaction.

Learning estimates in other outcome variables are consistent with corruption reduction caused by the rotation scheme. Rotating notifiers were traceable more often, and thus they were seen in proximity from defendants more often on repeated interaction, vis-a-vis fixed notifiers. They also attached photos to their notification orders more often upon later attempts.

An interesting story can be told about the effects of the rotation scheme on *illegal success*. Note that fixed notifiers consistently increase their rates of illegal success over time, while rotating notifiers do not. This finding is consistent with corruption reduction under rotation, since fixed notifiers may be increasingly notifying illegally from afar, as they have become acquainted with defendants in their region, while rotating notifiers find it more difficult to build and maintain relations with defendants.

Firm level heterogeneity

A source of concern from specification 3.2 has to do with firm level heterogeneity. Consider the firms that faced several unfair dismissal claims during the course of the experiment. They ought to be different in size and knowledge about the labor process, relative to firms that have only seen one suit filed against them. Hence, results in Table 4.6 may only reflect firm level heterogeneity, rather than showing heterogeneous learning rates associated with corruption.

A way to address this concern is to include firm features as treatment interacted covariates in the regressions and see if notifiers still exhibit heterogeneous learning rates. Firms' features included in the regressions are employment stratum (0 to 5 employees, and so on) matched from the National Firm Directory (INEGI 2021) at the firm level, and the average revenue at the sector level matched from the National Firm Census (INEGI 2019).[1]

Additionally, I include plaintiff features and notification order distance measures interacted with the treatment dummy as in Equation 3.4 to test the robustness of the results shown above.[2]

[1] Appendix 1: DENUE Match expands on matching procedures and shows firms' features are balanced across treatment arms.

[2] Appendix 3: Difficulty measurement shows summary statistics of (linear) distance related difficulty measures.

$$Y_i = \alpha_R + \beta T_i + \gamma V_i + \lambda X_d + \eta X_p + \tau D_i + T_i \times (\delta V_i + \phi X_d + \xi X_p + \rho D_i) + \epsilon_i \quad (3.4)$$

Where X_d are defendant (firm) level features, X_p are case file level plaintiff features and D_i distance related difficulty measures of notification order i. Then, δ should show an estimation for the heterogeneous learning rates that is not driven by firm level heterogeneity.

Table 4.8 shows the wedge of missed opportunities holds when including treatment interactions with case file features and notification distance measures, but loses statistical significane when firm features are included as interacted covariates. Although not significant, the magnitude of the point estimate is comparable to the model without interactions, and may point towards some corruption reduction, despite that not enough observations are left in later interactions to hold power.

Note that including all the levels of interactions together, as in column 5, about 20,000 observations are lost as they have missing values on either itneracted covariate. Column 6 drops the same observations dropped in Column 5 but does not include treatment interacted covariates, and it shows that the lost wedge effect on Column 5 is due to the loss of observations and not necessarily due to the addition of interacted covariates. Also note that treatment notifiers learnt at no lesser rate relative to control notifiers on every specification in Table 4.8.

Another way to rule out the firm heterogeneity hypothesis is to include address/defendant fixed effects to account for their time-invariant unobservable features (namely, size and knowledge of the labor process). Table 4.9 uses specification 3.2 but includes several levels of fixed effects to test whether the wedge of missed opportunities holds when comparing within each level.

Table 4.9 shows the wedge of missed opportunities holds when including region and notifier fixed effects but does not hold when including address, defendant or address-defendant level fixed effects.

These results may have two interpretations. On one hand, the wedge of missed opportunities is not capturing corruption reduction in the rotation scheme because it is capturing firm level heterogeneity. Or, on the other hand, specifications with thousand of clusters lack the statistical power to reach a meaningful result. The reader may interpret results at face value.

Note, however, columns 3 through 5 compare within address, defendants, and address-defendant pairs, and show treatment notifiers learn at no higher rate vis-a-vis control notifiers. Although with these specifications the wedge of missed opportunities is no larger, the rotation scheme proves no less successful relative to the fixed status quo, despite the increased difficulty that the rotation scheme poses to notifiers. Appendix 3: Difficulty measurement shows that rotating notifiers find it harder to notify addresses far away from the Labor Court, but also far away from the clusters of common addresses.

Actual visits

Another source of concern from 3.2 is that the *Visits* variables used thus far do not take into account whether notifiers actually went to the notification order address. Therefore, the wedge of missed opportunities shown in Table 4.6 may only reflect the fact that on a posterior date, rotating notifiers are (differentially) more eager to do their job.

A way to address the actual visits problem is through the location services that the Notifications app leverages. Based on notifiers geotrackings, an "In proximity" variable was constructed and it assesses whether a notifier was in proximity of a notification order address during a time window relative to the time they claim they went. The In proximity variable takes values as follows:

$$\text{In proximity}_i = \begin{cases} 1 & if \quad \text{Notifier seen in proximity} \\ 0 & if \quad \text{Notifier seen far away or cannot be seen} \end{cases}$$

A way to address the actual visits concern is to split the notification orders by their In proximity status and regress specification 3.2 conditional on each status. If we see the wedge holds for each value of In proximity, we will know learning is adequately approximated by the *Visits* variables.

This approach is not *kosher* since traceability and in proximity status are both outcomes that have significant and non-trivial treatment effects (see Table 4.4). Thus, any attempt to condition by their status is going to lead to biased estimators. The best way I can think of addressing this issue is to show robustness on either state of In proximity.

Table 4.10 shows the wedge of missed opportunities holds regardless of in proximity status, just not significantly. Because the In proximity variable splits the sample, the fact that the wedge does not hold on either In proximity status but it does in the whole sample suggests that the problem with this specifications is lack of power, rather than a problem with the corruption interpretation.

Another way to address the actual visits concern is to construct an *Actual Visits* variable that only counts a visit if the notifier was In proximity of the notification address. This approach is going to add precision relative to the conditional sampling in Table 4.11, but is going to remain biased since traceability and in proximity are both outcomes with non-trivial treatment effects.

Table 4.11 indeed shows a very similar result to Table 4.10. Although the wedge is not statistically distinct from zero, the addition of precision from counting only visits in proximity in the Actual Visits variable increased the magnitude of such estimate, relative to the wedge shown in Table 4.10.

Altogether, despite the that wedge of opportunities was not robust to In proximity status and despite the difficulty of assessing actual visits without biasing results, rotating notifiers proved to be no less successful than their fixed peers.

Days between visits

Because rotating notifiers revisit address-defendants on later dates on average, relative to fixed notifiers, the wedge of missed opportunities may not reflect corruption reduction in the rotation scheme, but may be driven entirely by time correlated variables (see Figure 4.3 and Figure 4.4). If success is positively correlated to time through a variable that is not learning, then that variable could explain the wedge of missed opportunities and its corruption interpretation would not be valid.

However, success is not positively correlated with time. Figure 4.6 shows weekly outcomes by treatment arm. There seems to be no positive linear correlation of time with success and no heterogeneous treatment time effect on success. Thus, the wedge of missed opportunities could not be explained by time correlated unobservable variables simply because later attempts are no more likely to be successful on either treatment arm, while attempts on repeated visits are differentially more likely to be successful.

Figure 4.6 shows success declines over time on both treatment arms. This may be explained by internal Court management changes that saw new leaderships beginning in May 2022. New leadership appointed was notoriously incompetent and did not manage notifiers as thoroughly as past leadership did.

Note that, although all notifiers became more traceable over time[3] they did not achieve higher rates of in proximity attempts.

[3] Over winter holidays, notifiers were required to return their phones to the research team. On January, their phones' location services were manually turned on by the research team. That exogenous variation in traceablity explains the jump seen on Figure 4.6 starting January 2022.

3.3 Discussion

Overall, the heterogeneous learning rates portrayed in the wedge of missed opportunities proved to be quite sensitive to the addition of robustness checks. Additionally, its corruption reduction interpretation rests on the strong assumption that notifiers cannot become differenctially more/less competent on repeated interactions with defendants.

However, there are not many stories as compelling as the corruption reduction interpretation that fit the heterogeneous learning rates. A likely second best candidate may be that, as fixed notifiers acquire notoriety in their region, defendants may avoid notification through other means that are not corruption, while rotating notifiers leverage anonymity to notify.

Chapter 4

Conclusions

Despite that the Mexican Labor Law requires all disputes to be resolved in under three months, the average unfair dismissal claim will take a year before its first hearing can be held; after which a long and tedious process still awaits. The basic human right to swift justice is far from being fulfilled in practice, as rent seeking bureaucrats mismanage its provision. And indeed, swift justice is just one of many aspects of justice provision that courts ought to manage to ensure market efficiency.

If firms know that labor justice is not only paralyzed, but that it can serve to their needs through bribery, they have no incentive to comply with labor regulations. And, although labor regulation compliance is of particular interest to many parties, including Mexico's trade partners, the most vulnerable party to firms' illegal and opportunistic behavior are workers, especially unskilled and marginalized workers.

In the specific context of labor law, Notifiers are the bureaucrats who hold the first door to swift justice, as complete and personal notifications are required before any hearing can be held. Although the rotation scheme did not increase notification success, some evidence for corruption reduction was found.

Under the relatively strong assumption that notifiers cannot become differentially more/less competent on repeated interactions with defendants, the rotation scheme caused corruption reduction equivalent to a 12.6 percentage points increase in successful notification rates on second visits to defendants in the same address, from a baseline of 41.4%.

Overall, results suggest that any gains in success driven by corruption reduction were completely offset by an opposite force, namely the increased difficulty of re-learning geographical and defendants' characteristics after every rotation.

Although the results shown are consistent with corruption reduction under rotation, its causal identification relies on the strong and untestable assumption of homogeneous competence. Furthermore, the analysis of repeated interactions with users proved to be a rather limited tool for corruption measurement, as indeed rotation schemes prevent such repeated interactions from happening.

References

Abbink, Klaus. 2004. "Staff Rotation as an Anti-Corruption Policy: An Experimental Study." *European Journal of Political Economy* 20 (4): 887–906. https://doi.org/10.1016/j.ejpoleco.2003.10.008.

Banerjee, Abhijit, Sendhil Mullainathan, and Rema Hanna. n.d. "Corruption." https://doi.org/10.3386/w17968.

Barr, Abigail, and Danila Serra. 2010. "Corruption and Culture: An Experimental Analysis." *Journal of Public Economics* 94 (11): 862–69. https://doi.org/10.1016/j.jpubeco.2010.07.006.

Becker, Gary S., and George J. Stigler. 1974. "Law Enforcement, Malfeasance, and Compensation of Enforcers." *The Journal of Legal Studies* 3 (1): 1–18. https://www.jstor.org/stable/724119.

Bertrand, Marianne, Simeon Djankov, Rema Hanna, and Sendhil Mullainathan. 2007. "Obtaining a Driver's License in India: An Experimental Approach to Studying Corruption*." *The Quarterly Journal of Economics* 122 (4): 1639–76. https://doi.org/10.1162/qjec.2007.122.4.1639.

Boehm, Johannes, and Ezra Oberfield. 2020. "Misallocation in the Market for Inputs: Enforcement and the Organization of Production*." *The Quarterly Journal of Economics* 135 (4): 2007–58. https://doi.org/10.1093/qje/qjaa020.

Bühren, Christoph. 2020. "Staff Rotation as an Anti-Corruption Policy in China and in Germany: An Experimental Comparison." *Jahrbücher Für Nationalökonomie*

Und Statistik 240 (1): 1–18. https://doi.org/10.1515/jbnst-2018-0036.

Casar, María Amparo. n.d. "Anatomía de la Corrupción." 115.

Chemin, Matthieu. 2020. "Judicial Efficiency and Firm Productivity: Evidence from a World Database of Judicial Reforms." *The Review of Economics and Statistics* 102 (1): 49–64. https://doi.org/10.1162/rest_a_00799.

Cunningham, Scott. 2021. *Causal Inference: The Mixtape.* Yale University Press. https://doi.org/10.2307/j.ctv1c29t27.

Djankov, Simeon, Rafael La Porta, Florencio Lopez-de-Silanes, and Andrei Shleifer. 2003. "Courts*." *The Quarterly Journal of Economics* 118 (2): 453–517. https://doi.org/10.1162/003355303321675437.

INEGI. 2021. "Directorio Estadístico Nacional de Unidades Económicas." https://www.inegi.org.mx/app/mapa/denue/?ag=17.

INEGI, Instituto Nacional de Estadística y Geografía. 2019. "Censos Económicos 2019. CE." https://www.inegi.org.mx/programas/ce/2019/.

Kaplan, David S., and Joyce Sadka. 2011. "The Plaintiff's role in enforcing a court ruling: Evidence from a labor court in Mexico." https://www.econstor.eu/handle/10419/115395.

Kaplan, David S., Joyce Sadka, and Jorge Luis Silva-Mendez. 2008. "Litigation and Settlement: New Evidence from Labor Courts in Mexico." *Journal of Empirical Legal Studies* 5 (2): 309–50. https://doi.org/10.1111/j.1740-1461.2008.00126.x.

Khan, Adnan Q., Asim I. Khwaja, and Benjamin A. Olken. 2016. "Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors *." *The Quarterly Journal of Economics* 131 (1): 219–71. https://doi.org/10.1093/qje/qjv042.

LaPorta, Rafael, Florencio Lopez-de-Silanes, Andrei Shleifer, and Robert Vishny. 1997. "Trust in Large Organizations." *American Economic Review Papers and Proceedings* 87 (2): 333–38.

Olken, Benjamin A., and Rohini Pande. 2012. "Corruption in Developing Countries." *Annual Review of Economics* 4 (1): 479–509. https://doi.org/10.1146/annurev-economics-080511-110917.

Ponticelli, Jacopo, and Leonardo S. Alencar. 2016. "Court Enforcement, Bank Loans, and Firm Investment: Evidence from a Bankruptcy Reform in Brazil *.” *The Quarterly Journal of Economics* 131 (3): 1365–1413. https://doi.org/10.1093/qje/qjw015.

Robinson, Amanda Lea, and Brigitte Seim. 2018. “Who Is Targeted in Corruption? Disentangling the Effects of Wealth and Power on Exposure to Bribery.” *Quarterly Journal of Political Science* 13 (3): 313–31. https://doi.org/10.1561/100.00017067.

Sadka, Joyce, Enrique Seira, Frederico Finan, and Ernesto Dal Bó. 2019. “Improving Labor Courts: Fighting Corruption with Competition.” https://www.povertyactionlab.org/initiative-project/improving-labor-courts-fighting-corruption-competition.

Sadka, Joyce, Enrique Seira, and Christopher Woodruff. n.d. “Information and Bargaining Through Agents: Experimental Evidence from Mexico’s Labor Courts.” https://doi.org/10.3386/w25137.

Schmelzer, Christoph Hanck, Martin Arnold, Alexander Gerber, and Martin. n.d. *Introduction to Econometrics with r*. https://www.econometrics-with-r.org/.

Shleifer, Andrei, and Robert W. Vishny. 1993. “Corruption*.” *The Quarterly Journal of Economics* 108 (3): 599–617. https://doi.org/10.2307/2118402.

Tables

Table 4.1: Balance Table - Notifiers

Characteristic	**Control**, N = 25	**Treatment**, N = 25	**p-value**
Female	0.5 (0.5)	0.5 (0.5)	>0.9
Year of birth	1,978.0 (6.3)	1,979.6 (9.9)	0.2
Lawyer since	2,010.8 (5.4)	2,012.2 (5.9)	0.2
Notifier since	2,011.7 (5.0)	2,011.8 (6.7)	0.6
Weekly contact with Plaintiffs	3.7 (5.0)	3.5 (4.7)	>0.9
Weekly contact with Defendants	1.8 (2.5)	2.2 (3.2)	0.7
Weekly encouragement bribes offered	0.6 (1.3)	0.9 (1.8)	0.4
Weekly avoidance bribes offered	0.3 (0.7)	0.5 (1.0)	0.5
Weekly successful attepmts	26.8 (12.3)	21.9 (8.1)	0.2
City knowledge (0-10)	6.5 (2.1)	5.2 (2.2)	0.024
Would like a notifications app	0.4 (0.5)	0.2 (0.4)	0.14
Monthly family expenditure (MXN)	13,509.3 (6,917.8)	15,370.0 (5,516.4)	0.091
Share in family expenditure (%)	73.0 (26.5)	68.8 (21.5)	0.4

[1] Mean (SD)

[2] Wilcoxon rank sum test

Table 4.2: Balance Table - Case files

Characteristic	**Control**, N = 11,668	**Treatment**, N = 11,322	**p-value**
Addresses	1.16 (0.45)	1.15 (0.44)	0.4
Distance from Court (km)	7.06 (4.36)	6.99 (4.30)	0.2
Defendants sued	2.66 (2.39)	2.56 (2.16)	<0.001
Plaintiffs	1.09 (0.68)	1.09 (0.63)	0.7
Filed via SIREDE	0.87 (0.33)	0.88 (0.32)	0.028
Age	39.88 (12.28)	40.03 (12.30)	0.4
Weekly worked hours	57.68 (87.22)	57.12 (74.57)	0.6
Daily wage (MXN)	324.24 (406.11)	332.67 (406.72)	0.10
Vulnerable plaintiffs	0.08 (0.27)	0.08 (0.27)	0.8
Day shift	0.76 (0.43)	0.75 (0.43)	0.8
Paid biweekly	0.45 (0.50)	0.43 (0.50)	0.13
Power of attorney	0.44 (0.50)	0.43 (0.49)	0.022
Social Security	0.18 (0.38)	0.18 (0.38)	0.3
Services industry	0.59 (0.49)	0.59 (0.49)	>0.9
Claim reinstallation	0.44 (0.50)	0.46 (0.50)	0.13

[1] Mean (SD)

[2] Wilcoxon rank sum test

Table 4.3: Balance Table - Routes

Characteristic	**Control**, N = 1,067	**Treatment**, N = 1,229	**p-value**
Case files	10.66 (4.44)	10.38 (4.54)	0.4
Notification orders	27.31 (14.90)	25.44 (14.93)	0.002
Addresses	10.24 (4.17)	9.90 (4.24)	0.2
Defendants	25.62 (13.49)	23.85 (13.28)	0.003
Distance from Court (km)	8.11 (4.61)	7.67 (4.55)	0.024
Route dispersion (km)	1.51 (0.91)	1.49 (0.82)	0.4

[1] Mean (SD)

[2] Wilcoxon rank sum test

Table 4.4: Treatment effect on main outcomes

Dependent Variables:	Success	Traceable	In proximity	Has photo	Illegal success
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Treatment	-0.015*	0.087***	0.035***	0.045***	0.002
	(0.008)	(0.005)	(0.008)	(0.008)	(0.004)
Intercept	0.425	0.850	0.532	0.477	0.078
Fit statistics					
Dependent variable mean	0.414	0.881	0.545	0.498	0.079
Observations	67,739	67,739	67,739	67,739	67,739

Clustered (Case file) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*
OLS estimates with region fixed effects. Observations at notification order level. Intercept shows fixed effects simple mean.

Table 4.5: Summary statistics of visits to address-defendants

Characteristic	**Control**, N = 34,932	**Treatment**, N = 32,807
Visits		
1	30,339 (86.85%)	32,452 (98.92%)
2	3,024 (8.66%)	338 (1.03%)
3	735 (2.10%)	17 (0.05%)
4	308 (0.88%)	0 (0.00%)
5	173 (0.50%)	0 (0.00%)
6	114 (0.33%)	0 (0.00%)
7	75 (0.21%)	0 (0.00%)
8	42 (0.12%)	0 (0.00%)
9	32 (0.09%)	0 (0.00%)
10	28 (0.08%)	0 (0.00%)
11	18 (0.05%)	0 (0.00%)
12	11 (0.03%)	0 (0.00%)
13	15 (0.04%)	0 (0.00%)
14	7 (0.02%)	0 (0.00%)
15	6 (0.02%)	0 (0.00%)
16	3 (0.01%)	0 (0.00%)
17	2 (0.01%)	0 (0.00%)
Days between visits		
Median (IQR)	51 (16, 105)	68 (33, 114)
Range	2, 237	2, 235

[1] n (%)

Table 4.6: Heterogeneous learning rates on main outcomes

Dependent Variables:	Success	Traceable	In proximity	Has photo	Illegal success
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Treatment	-0.005	0.089^{***}	0.028^{***}	0.039^{***}	0.010^{**}
	(0.008)	(0.005)	(0.008)	(0.008)	(0.004)
Visits	0.050^{***}	0.015^{***}	-0.039^{***}	-0.022^{*}	0.043^{***}
	(0.011)	(0.006)	(0.011)	(0.012)	(0.007)
Treatment × Visits	0.101^{***}	0.019	0.083^{**}	-0.004	0.017
	(0.035)	(0.013)	(0.037)	(0.038)	(0.027)
Intercept	0.413	0.847	0.538	0.483	0.069
Fit statistics					
Observations	66,905	66,905	66,905	66,905	66,905
Dependent variable mean	0.411	0.882	0.546	0.500	0.078

Clustered (Case file) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Notes: OLS estimates with region fixed effects. Sample of comparable visits. Observations at notification order level. Intercept shows outcome rate during the first visit of a particular notifier to a particular address-defendant pair.

Table 4.7: Visits as dummies

Dependent Variables:	Success	Traceable	In proximity	Has photo	Illegal success
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Intercept	0.414	0.847	0.539	0.484	0.069
Treatment	-0.006	0.089***	0.027***	0.038***	0.010**
	(0.008)	(0.005)	(0.008)	(0.008)	(0.004)
2nd visit	0.023	0.016*	-0.056***	-0.033*	0.042***
	(0.016)	(0.009)	(0.016)	(0.017)	(0.010)
Treatment × 2nd visit	0.126***	0.017	0.080*	-0.024	0.038
	(0.040)	(0.016)	(0.043)	(0.044)	(0.032)
3rd visit	0.151***	0.028**	-0.046*	-0.024	0.088***
	(0.029)	(0.013)	(0.027)	(0.033)	(0.020)
Treatment × 3rd visit	0.166	0.045*	0.328***	0.275***	-0.164***
	(0.123)	(0.027)	(0.098)	(0.106)	(0.024)
Fit statistics					
Observations	66,905	66,905	66,905	66,905	66,905
Dependent variable mean	0.411	0.882	0.546	0.500	0.078

Clustered (Case file) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: OLS estimates with region fixed effects. Samples of comparable number of visits. Observations at notification order level. Intercept shows outcome mean during first visit.

Table 4.8: Interacted covariates

Dependent Variable:	Success					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Variables						
Intercept	0.413	0.393	0.473	0.402	0.418	0.429
Treatment	-0.005	0.013	0.048	1.06***	1.18***	0.010
	(0.008)	(0.023)	(0.044)	(0.128)	(0.164)	(0.010)
Visits	0.050***	0.054***	0.046***	0.050***	0.050***	0.050***
	(0.011)	(0.012)	(0.012)	(0.012)	(0.013)	(0.013)
Treatment × Visits	0.101***	0.071*	0.082**	0.094***	0.046	0.058
	(0.035)	(0.042)	(0.036)	(0.035)	(0.044)	(0.045)
Interactions: firm features		Yes			Yes	
Interactions: case file features			Yes		Yes	
Interactions: notification difficulty				Yes	Yes	
Fit statistics						
Observations	66,905	53,643	59,117	63,373	44,972	44,972
R^2	0.021	0.028	0.024	0.029	0.039	0.027
Dependent variable mean	0.411	0.415	0.425	0.415	0.432	0.432

Clustered (Case file) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: OLS estimates with region fixed effects. Observations at notification order level. Sample of comparable number of visits. Firm features are categorical variables for employed staff and firm's economic sector mean revenue. Case file features are number of plaintiffs, their mean wage, mean age, % of women, and % of vulnerable plaintiffs. Notification difficulty variables are log distance from labor court, log distance from region centroid, log distance from route centroid and log route dispersion.

Table 4.9: Several levels of fixed effects

Dependent Variable:			Success		
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Intercept	0.413	0.412	0.410	0.356	0.376
Treatment	-0.005		0.001	-0.009	-0.015
	(0.008)		(0.010)	(0.011)	(0.014)
Visits	0.050***	0.055***	0.052***	0.061***	0.044***
	(0.011)	(0.011)	(0.009)	(0.011)	(0.011)
Treatment × Visits	0.101***	0.108***	0.031	0.006	-0.001
	(0.035)	(0.034)	(0.031)	(0.037)	(0.037)
Fixed-effects					
Region (25)	Yes				
Notifier (50)		Yes			
Address (10,179)			Yes		
Defendant (10,494)				Yes	
Address-defendant (7,625)					Yes
Fit statistics					
Observations	66,905	66,905	62,617	34,062	21,002
R^2	0.021	0.053	0.624	0.524	0.637
Within R^2	0.002	0.002	0.003	0.005	0.005
Dependent variable mean	0.411	0.411	0.412	0.385	0.415

Clustered (Case file) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: OLS estimates. Observations at notification order level. Sample of comparable number of visits. Singletons removed.

Table 4.10: In proximity status

Dependent Variable:	Success		
Sample:	Full sample	Far away	In proximity
Model:	(1)	(2)	(3)
Variables			
Intercept	0.413	0.257	0.550
Treatment	-0.005	-0.025**	0.003
	(0.008)	(0.011)	(0.012)
Visits	0.050***	0.062***	0.062***
	(0.011)	(0.015)	(0.016)
Treatment × Visits	0.101***	0.092	0.066
	(0.035)	(0.061)	(0.041)
Fit statistics			
Observations	66,905	30,379	36,526
R^2	0.021	0.020	0.030
Dependent variable mean	0.411	0.250	0.546

Clustered (Case file) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: OLS estimates with region fixed effects. Observations at notification order level. Sample of comparable number of visits.

Table 4.11: Actual visits

Dependent Variable:	Success	
Model:	(1)	(2)
Variables		
Intercept	0.413	0.550
Treatment	-0.005	0.003
	(0.008)	(0.012)
Visits	0.050***	
	(0.011)	
Treatment × Visits	0.101***	
	(0.035)	
Actual visits		0.094***
		(0.016)
Treatment × Actual visits		0.071
		(0.049)
Fit statistics		
Observations	66,905	36,526
R^2	0.021	0.030
Dependent variable mean	0.411	0.546

Clustered (Case file) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: OLS estimates with region fixed effects. Observations at notification order level. Sample of comparable number of visits. Actual visits variable only counts visits if the notifier was traced in proximity of the order address, while Visits variable counts visits regardless of the in proximity status.

Figures

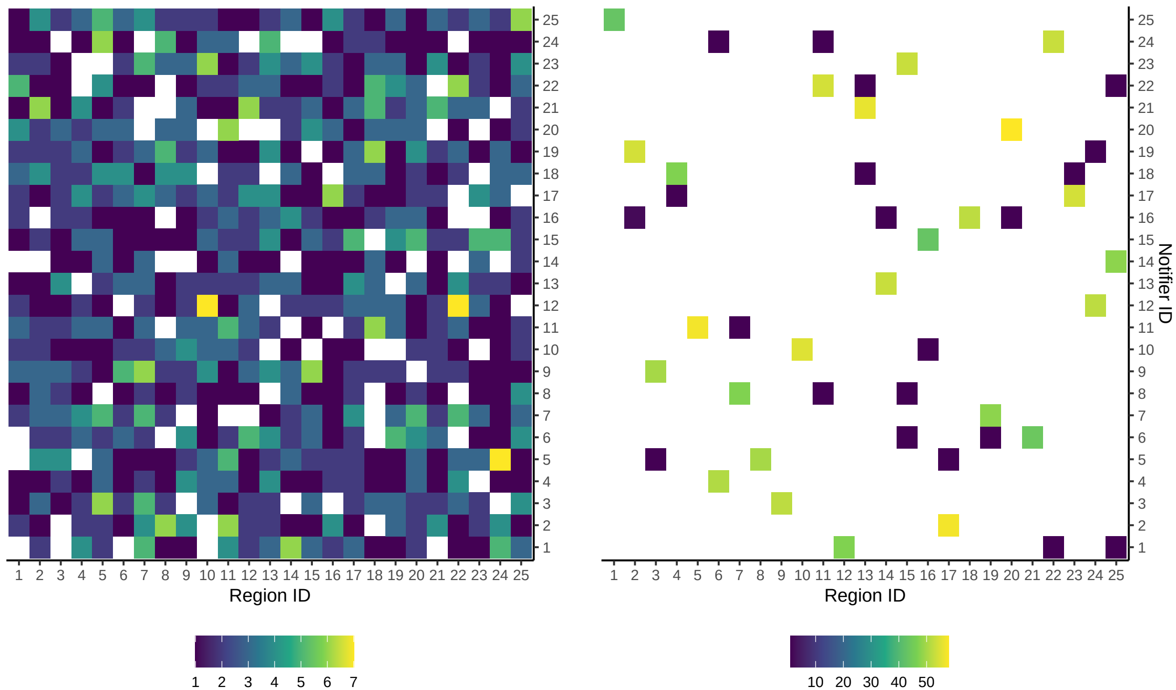

(a) Rotating notifiers

(b) Fixed notifiers

Figure 4.1: Visits to regions by treatment arm

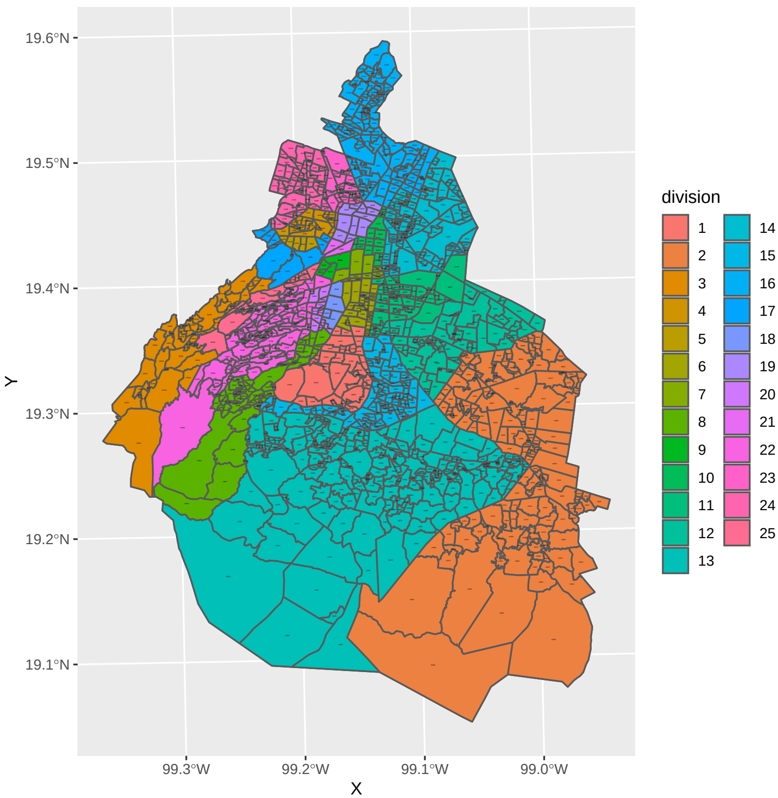

Figure 4.2: Mexico City split into 25 regions

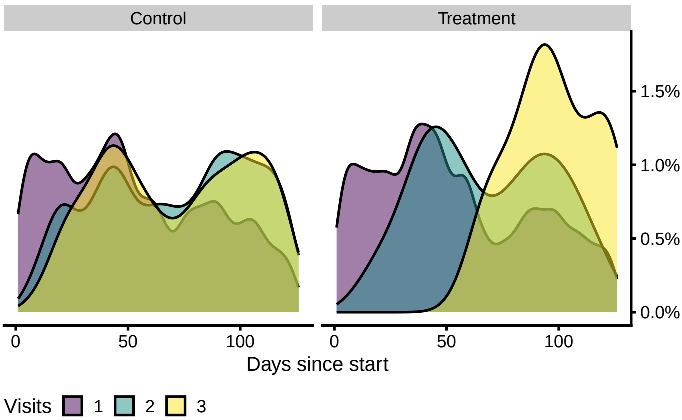

Figure 4.3: Density of comparable visits to address-defendants across days since start, by treatment arm.

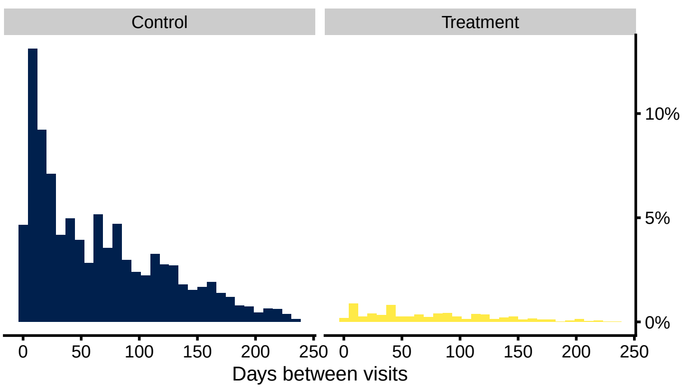

Figure 4.4: Histogram of days between visits to address-defendants by treatment arms

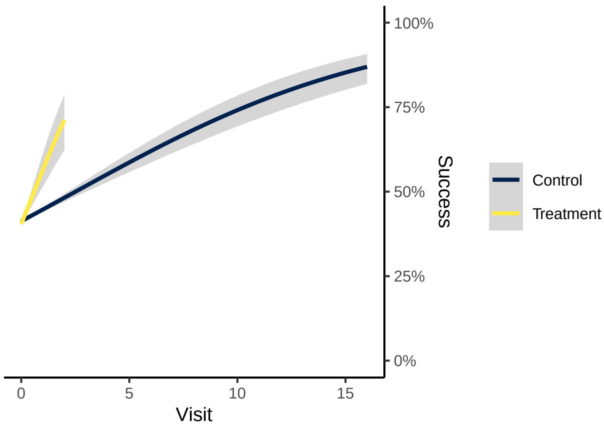

Figure 4.5: Probability of success conditional on number of visits to address-defendants

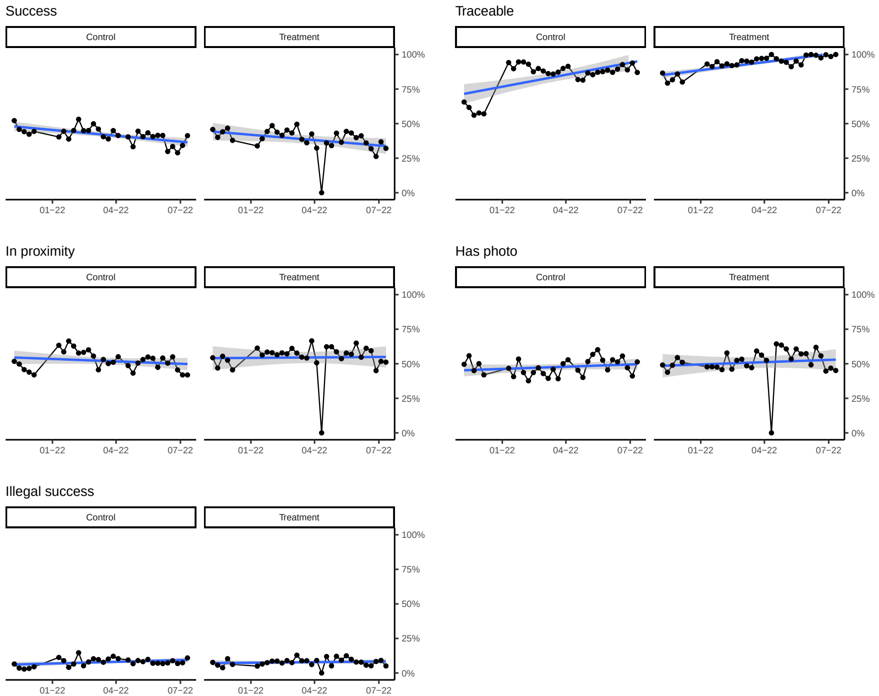

Figure 4.6: Weekly outcomes by treatment arm

Appendix

DENUE Match

Firms' features were obtained through the National Firm Directory (DENUE). Address/defendant ids were computed to match those computed in the Notifiers data set. Because DENUE is public information, it does not include firms' revenues or exact number of employees. It does include, however, firms' employment category (see Table 8) and it includes firms' economic sector (via SCIAN nomenclature at 6 digits, the most granular level). After obtaining firms' SCIAN code, I matched through it to the Economic Census to obtain sector mean revenue. The match was done through several steps.

1. Find direct matches through address-defendant id, then through defendant id, and finally through address id.
2. Fill repeatedly by casefile docket, address-defendant id, defendant id and finally by address id.
3. Fuzzy match through address-defendant id and defendant id.
4. Fill repeatedly by casefile docket, address-defendant id, defendant id and finally by address id.

Distribution of matches across type and link looks as follows:

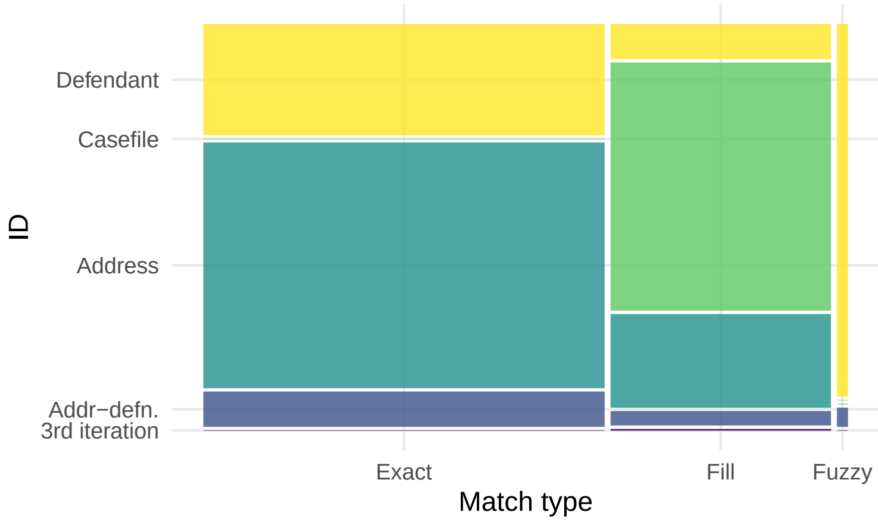

Figure 4.7: Match with DENUE

Balance across treatment arms looks as follows:

Table 4.12: Balance table - Defendants features

Characteristic	**Control**, N = 34,932	**Treatment**, N = 32,807
DENUE labor category		
0 a 5 personas	3,794 (13%)	3,664 (13%)
101 a 250 personas	3,209 (11%)	2,890 (10%)
11 a 30 personas	5,421 (18%)	4,705 (17%)
251 y más personas	7,659 (25%)	6,955 (25%)
31 a 50 personas	2,918 (9.7%)	2,697 (9.6%)
51 a 100 personas	3,974 (13%)	3,928 (14%)

(continued)

Characteristic	**Control**, N = 34,932	**Treatment**, N = 32,807
6 a 10 personas	3,208 (11%)	3,316 (12%)
Unknown	4,749	4,652
SCIAN code		
11	13 (<0.1%)	4 (<0.1%)
22	16 (<0.1%)	30 (0.1%)
23	1,557 (5.2%)	1,148 (4.1%)
31	973 (3.2%)	887 (3.2%)
32	750 (2.5%)	771 (2.7%)
33	868 (2.9%)	819 (2.9%)
43	3,129 (10%)	2,862 (10%)
46	4,231 (14%)	3,706 (13%)
48	653 (2.2%)	645 (2.3%)
49	270 (0.9%)	211 (0.7%)
51	780 (2.6%)	944 (3.4%)
52	882 (2.9%)	871 (3.1%)
53	634 (2.1%)	779 (2.8%)
54	3,159 (10%)	2,916 (10%)
55	154 (0.5%)	141 (0.5%)
56	5,881 (19%)	5,340 (19%)
61	925 (3.1%)	726 (2.6%)
62	693 (2.3%)	584 (2.1%)
71	404 (1.3%)	474 (1.7%)
72	2,598 (8.6%)	2,693 (9.6%)
81	1,146 (3.8%)	1,117 (4.0%)
93	467 (1.5%)	487 (1.7%)
Unknown	4,749	4,652
Revenue (Millions MXN)	24 (5, 67)	23 (5, 62)

(continued)

Characteristic	**Control**, N = 34,932	**Treatment**, N = 32,807
Unknown	6,852	6,557

[1] n (%); Median (IQR)

Address-defendant identification

Address-defendant pairs are characterized by an address-defendant id constructed as follows:

1. Keep only street name and street number without white spaces or non-alphanumerical characters as an address id.
2. Strip defendants names from white spaces, or non-alphanumerical characters as a defendant id.
3. Concatenate address id and defendant id, remove accents and capitalize characters.
4. Assign randomly a unique numerical id to preserve anonymity of defendants in their address-defendant ids.

The following tables show the number of unique address-defendants in the Notifiers data to reach identification of heterogeneous learning rates.

Table 4.13: Address-defendants with more than 1 casefile

	Control	Treatment
1 casefile	25,097	23,249
More than 1 casefile	4,212	4,073

Table 4.14: Address/defendants visited by both arms

	Addresses	Defendants	Address-defendants
Both arms	2,394	4,146	3,103
Only 1 arm	12,073	39,191	50,425

Table 4.15: Address-defendants by total number of visits.

Visits	Control	Treatment	Both arms
1	26,509	27,082	2,976
2	2,150	230	118
3	387	10	9
4	117		
5	55		
6	35		
7	24		
8	7		
9	9		
10	3		
11	3		
13	3		
14	2		
15	2		
16	1		
17	2		

Table 4.16: Address-defendants by total number of visits by treated (columns) and control (rows) notifiers.

	0	1	2	3
0		24,137	82	
1	24,413	2,065	31	
2	1,611	507	31	1
3	155	210	20	2
4	23	78	16	
5	3	37	15	
6	1	23	10	1
7		15	8	1
8		2	3	2
9		5	4	
10		1	2	
11		2	1	
13			2	1
14			2	
15			1	1
16			1	
17			1	1

Table 4.17: Address-defendants by total number of visits and legal entity

Visits	Individual	Firm
1	26,915	23,974
2	925	1,350
3	100	296
4	16	102
5	2	55
6	9	27
7		24
8		7
9		9
10		3
11		3
13		3
14		2
15		2
16		1
17		2

Table 4.18: Balance Table - Difficulty as distance measures

Characteristic	**Control**, N = 34,932	**Treatment**, N = 32,807	**p-value**
Distance from Court	7,152 (4,447)	7,134 (4,403)	0.6
Distance from region centroid	2,013 (1,608)	2,007 (1,615)	0.7
Distance from route centroid	1,259 (1,046)	1,261 (1,020)	0.005
Route dispersion	1,399 (785)	1,402 (741)	<0.001

[1] Mean (SD)

[2] Wilcoxon rank sum test

Difficulty measurement

Difficulty measures of notification orders was computed as several distance measures from notification orders. Although some distance measures are statistically significantly different across treatment arms, the magnitude of the differences are not worrisome.

Table 4.19: Heterogeneous difficulty effect on success

Dependent Variable:	Success			
Model:	(1)	(2)	(3)	(4)
Variables				
Treatment	0.622***	0.393***	0.419***	0.869***
	(0.098)	(0.078)	(0.077)	(0.120)
Log dist. from Court	-0.018			
	(0.014)			
Treatment × Log dist. from Court	-0.073***			
	(0.011)			
Log dist. region centroid		0.021***		
		(0.008)		
Treatment × Log dist. region centroid		-0.056***		
		(0.011)		
Log dist. route centroid			0.031***	
			(0.007)	
Treatment × Log dist. route centroid			-0.063***	
			(0.011)	
Log route dispersion				0.054***
				(0.013)
Treatment × Log route dispersion				-0.124***
				(0.017)
Intercept	0.587	0.276	0.213	0.042
Fit statistics				
Observations	65,363	65,363	64,169	64,169
R^2	0.025	0.023	0.024	0.026
Dependent variable mean	0.417	0.417	0.417	0.417

Clustered (Case file) standard-errors in parentheses
*Signif. Codes: ***: 0.01, **: 0.05, *: 0.1*

Note: OLS estimation with region fixed effects. Observations at notification order level.