

Compliance Behavior in Networks: Evidence from a Field Experiment[†]

By FRANCESCO DRAGO, FRIEDERIKE MENGEL, AND CHRISTIAN TRAXLER*

This paper studies the spread of compliance behavior in neighborhood networks in Austria. We exploit a field experiment that varied the content of mailings sent to potential evaders of TV license fees. The data reveal a strong treatment spillover: untreated households are more likely to switch from evasion to compliance in response to mailings received by their network neighbors. Digging deeper into the properties of the spillover, we find that it is concentrated among close neighbors of the targets and increases with the treated households' diffusion centrality. Local concentration of equally treated households implies a lower spillover. (JEL C93, D12, D85, L82, L88, Z13)

Interacting with individuals and firms that are suspected of violating laws is an important challenge for many governmental agencies and regulators. Over the past decade, a significant body of research has experimentally tested numerous strategies to improve efficacy in such interactions (see, e.g., Kleven et al. 2011; Dwenger et al. 2016; Shimeles, Gurara, and Woldeyes 2017 on mailing campaigns to enforce taxes and other payments). While progress has been made in understanding what type of enforcement strategies work, it is less understood if and how such interventions generate spillovers on untreated actors within a given social or economic network (Rincke and Traxler 2011; Pomeranz 2015; Brollo, Kaufmann, and La Ferrara 2017). Gaining more knowledge of how treatment-induced information diffuses among friends, neighbors, or coworkers—thus influencing the behavior among a broader population—is important from a policy perspective; it enables authorities to target enforcement interventions to maximize their overall impact (i.e., direct

*Drago: Department of Economics, University of Catania, Corso Italia 55, 95131 Catania, Italy, CSEF and CEPR (email: francesco.drago@unict.it); Mengel: Department of Economics, University of Essex, Wivenhoe Park, Colchester CO4 3SQ, United Kingdom, Department of Economics, Lund University (email: fmengel@essex.ac.uk); Traxler: Hertie School, Friedrichstraße 180, 10117 Berlin, Germany (email: traxler@hertie-school.org). Benjamin Olken was coeditor for this article. We thank two anonymous referees, Ciro Avitabile, Alexandra Avdeenko, Abhijit Banerjee, Will Boning, Esther Duflo, Florian Ederer, Ronny Freier, Andrea Galeotti, Ben Golub, Matt Jackson, Adam Szeidl, and Noam Yuchtman for many helpful comments and suggestions. The paper also benefited from discussions at conferences/seminars/workshops in Barcelona, Berlin, Bonn, Bristol, Capri, Copenhagen, Essex, Frankfurt, Gothenburg, Harvard, Kent, Linz, London, Lugano, Maastricht, Mannheim, Munich, Naples, Oxford, Paris, St. Louis, and Stanford. Kyle Ott, Nicolai Vitt, and Andreas Winkler provided outstanding research assistance. Francesco Drago thanks the Compagnia SanPaolo Foundation and Friederike Mengel the Dutch Science Foundation (NWO) for financial support (Veni grant 451-11-020).

[†]Go to <https://doi.org/10.1257/app.20170690> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

effects plus spillovers on the untreated population) and offers, as we are going to show, a new perspective on the debate about “concentrated” versus “widely spread” enforcement actions.

This paper exploits a large administrative dataset and a large-scale field experiment to study spillovers from enforcement mailings in neighborhood networks in Austria. We explore how communication among neighbors affects the compliance decision of untreated households. Our analysis addresses the policy questions from above and contributes to the growing literature on information transmission in networks (e.g., Banerjee et al. 2013, 2019; Alatas et al. 2016; BenYishay and Mobarak 2019).

We build on an experiment that tested different strategies to enforce compliance with TV license fees (Fellner, Sausgruber, and Traxler 2013). The experiment introduced exogenous variation in the treatment of 50,000 potential license fee evaders. In a *baseline* treatment, households received a mailing that asked them why they were not paying fees. In a *threat* treatment, the mailing communicated an imminent inspection and emphasized possible financial and legal consequences from noncompliance. Relative to a *control* group that did not receive any mailing, the two mailing treatments significantly increased compliance. Mediated by a higher perceived detection risk, the threat triggered the largest effect (Fellner, Sausgruber, and Traxler 2013).

This paper now studies the treatments’ impact on the compliance behavior of the *untreated* population. Since neither receiving a mailing nor compliance is observable to the other neighbors, behavior can only spread via treatment-induced communication. We first explore communication patterns in a large online survey. The survey documents high communication frequencies among neighbors, especially in rural areas. The intensity of communication declines with the distance to the next neighbor—a pattern that matches with other studies documenting the role of geographic proximity for social interaction (e.g., Marmaros and Sacerdote 2006). Our survey further highlights people’s willingness to share information on TV license fee enforcement with their neighbors.

To test whether communication generates a spillover, we use precise microdata and geocoded information on the full population of small Austrian municipalities to compute neighborhood networks based on geographic distance. Motivated by the results from our survey, we assume two households are linked if they live within a distance of 50 meters.¹ A network is composed of all households that are directly or indirectly linked. Identification of the treatment effects on untreated neighbors is achieved by the fact that conditional on the number of households covered by the experiment, the treatment of these “experimental households” varies exogenously. Several pieces of evidence offer ample support for this conditional independence assumption.

We find a pronounced spillover effect: untreated households, which were not part of the experimental sample, are more likely to switch from evasion to compliance in response to mailings received by neighbors in the same network. Our estimates

¹ We document that our main results are qualitatively robust to distance cutoffs up to 500 meters.

suggest that sending one additional threat (baseline) mailing into a network increases each untreated evader's propensity to comply by 7 (5) percentage points. A back-of-the-envelope calculation implies that 1,000 additional threat (baseline) mailings spread over 3,764 neighborhood networks would induce 68 (48) untreated households to start complying. While the comparison between direct and indirect treatment effects is complicated by different sample compositions, it is worth stressing that the overall spillover is similar in magnitude to the direct treatment effect. Hence, the implied *social multiplier* is roughly two (and thus comparable to findings on social interaction in other domains; Glaeser, Sacerdote, and Scheinkman 2003; Bruhin et al. 2014), meaning that the diffusion process doubles the impact of the enforcement intervention.

The evidence on treatment spillovers turns out to be very robust. Only if we increase the distance threshold that defines a network link to above 500 meters or if we analyze treatment spillovers at the level of (fairly small) municipalities, do the effects vanish. The same holds for placebo tests that ignore the geographic structure within municipalities: allocating households from the *same* municipality to randomly generated networks, we again obtain a null result. Our basic findings therefore highlight the crucial role neighborhood networks play in shaping the diffusion that underlies the spread of compliance. Moreover, we show that behavioral changes among targeted households are not necessary to induce compliance spillovers. While behavioral interdependencies (related, e.g., to the strength of social norms) do play some role, informational channels (e.g., the updating of risk perceptions) are crucial in shaping the indirect treatment effects.

To shed more light on the properties of the communication-induced spillover and the role of the network in the behavioral adaptation process, we approach three questions of significant policy relevance: (i) Which households are reached by the communicated information? (ii) Which households in the networks are "best targeted" to maximize the spillovers? and (iii) Are spillovers higher when treatments are locally concentrated rather than broadly spread?

Regarding the first question, we find that the spillovers from the baseline treatment are mainly limited to first-order (i.e., directly linked) neighbors of treated households. For the threat treatment, the spillovers reach further into the networks. To address the second question, we study how network centrality of targeted households (the "injection points") amplifies the spillovers. Consistent with our first finding and with theoretical models of diffusion (Banerjee et al. 2019), the data reveal strong, positive interactions with the treated households' diffusion centrality. Hence, the network structure matters for the spillover beyond mere geographic proximity. From a policy perspective, this result means that targeting a network's most "diffusion central" households will, *ceteris paribus*, maximize the intervention's indirect effects.

As for the third question, treatment concentration allows fewer households to learn about a mailing, but those that do are more likely to hear about it repeatedly. Which of these two effects dominates is ultimately an empirical question. The data show that local treatment concentration is associated with substantially *smaller* treatment spillovers. The potential gains from repeated exposure to treatment-induced information are thus dominated by the "loss" from reaching fewer households.

Again, this finding has a straightforward policy implication: to maximize the spillovers, mailing campaigns should avoid local concentration.

Our study contributes to several important strands of literature. First, we advance the research on enforcement and compliance by linking it to the networks literature (Jackson 2008). Our results document that communication can mediate evasion and avoidance decisions not only within firm (Pomeranz 2015) or family networks (Alstadsæter, Kopczuk, and Telle 2018) but also within neighborhood networks. The key innovation compared to earlier studies (e.g., Rincke and Traxler 2011; Brollo, Kaufmann, and La Ferrara 2017)—and research on spillovers on geographically proximate agents more generally (e.g., Bayer, Ross, and Topa 2008; Kuhn et al. 2011)—is that we analyze different network structures to assess how the diffusion process induces the spillovers in an experimental setting. The network analysis, as well as the identification strategy, distinguishes the present paper from Rincke and Traxler (2011), who rely on an instrumental variable approach to identify enforcement spillovers from door-to-door audits. Moreover, the focus on indirect treatment effects differs fundamentally from Fellner, Sausgruber, and Traxler (2013), who solely explore the direct effects of the mailing intervention.

Our findings point out how geographic information, which is readily available in many applications, could be incorporated in algorithms that are used to target audits or inspections. Rather than relying solely on a node's isolated characteristics, one could account for the nodes' positions within a network (e.g., to put more weight on highly diffusion-central actors while avoiding local concentration of targets). Geographic information should thereby be particularly relevant when enforcement activities are local and geographically correlated (as is the case with many door-to-door inspections at households or firms; e.g., Olken 2007; Khan, Khwaja, and Olken 2016). The present paper has already spurred tax authorities and public finance researchers to start examining how nongeographic networks contribute to enforcement spillovers (see, e.g., Boning et al. 2018).

The results from this paper also speak to a much broader set of applications that might exploit neighborhood communication, e.g., to effectively seed fundraising (Landry et al. 2006), technologies (Bollinger and Gillingham 2012), and health programs (Miguel and Kremer 2004) or to improve the effectiveness of marketing campaigns (Aral and Walker 2011). The relevance of geographic proximity will certainly depend on the type of communities (proximity tends to be more important in smaller municipalities) and on the types of issues considered (whether the issue is a relevant topic of conversation among geographic neighbors). Geographic networks have been shown to matter in such diverse domains as households' energy consumption (Allcott 2011), blood donations (Bruhin et al. 2014), or the diffusion of knowledge of the tax code (Chetty, Friedman, and Saez 2013). Beaman et al. (2015), who study technology adoption, show that seeding based on geographic networks works fairly well. While seeding based on a complex model of elicited social networks increases spillovers, the geographic network approach is cheaper and easier to implement.

Finally, our results on local treatment concentration, which builds on an inbreeding homophily index (Coleman 1958; Currarini, Jackson, and Pin 2009)

to capture whether neighbors are “similar” in terms of receiving the same treatment, relate to the literature on homophily (e.g., Jackson, Rodriguez-Barraquer, and Tan 2012; Jackson 2016). We discuss that a local concentration of a given treatment translates into a higher level of congruence of, e.g., neighbors’ posttreatment beliefs. Random treatment assignment thereby ensures that, conditional on having neighbors in the experiment sample, their treatments (and thus local treatment concentration) vary exogenously. Unlike other dimensions of similarity, this measure is exogenous to unobserved household characteristics. We can therefore isolate the effect from similarity in treatment on the size of the spillovers. Our estimates show that local treatment concentration reduces the spillovers. The finding is in line with the idea that similarity of neighbors’ beliefs hampers social learning (Golub and Jackson 2012).

The remainder of the paper is organized as follows. Section I provides further information on the institutional background and the field experiment. Section II reports survey results on communication patterns among neighbors. Our main data are described in Section III. In Section IV, we present our basic results. Section V discusses additional results on the nature of diffusion in our setting. Section VI concludes.

I. Background of the Field Experiment

A. License Fees

Obligatory radio and television license fees are a common tool to fund public service broadcasters. A typical license fee system is operated by Fee Info Service (henceforth FIS), a subsidiary of the Austrian public broadcasting company. In Austria, the Broadcasting License Fee Act prescribes that all “households” (including apartment-sharing communities, etc.) owning a TV or a radio must register their broadcasting equipment with FIS. The authority then collects an annual license fee of roughly €230 per household.² Households face an incentive to evade the fee because public broadcasting programs can be received without paying.

FIS takes several actions to enforce compliance. Using official data from residents’ registration offices, they match the universe of residents with data on those paying license fees. Taking into account that 99 percent of all Austrian households are equipped with a radio or a TV, each resident not paying fees is flagged as a potential evader (unless another household member has been identified as paying). Potential evaders are then contacted by mail and asked to clarify why they have not registered any broadcasting equipment. Data on those who do not respond are handed over to FIS’s enforcement division. Members from this division personally approach households and make door-to-door inspections (see Rincke and Traxler 2011). A detected evader is registered and typically has to pay the evaded fees for up to several past months. In addition, FIS can impose a fine of

²The fee is independent of the number of household members and varies between states. In 2005, the year covered by our data, the fee ranged between €206 and €263.

up to €2,180. If someone does not comply with the payment duty, legal proceedings will be initiated.

The enforcement efforts are reflected in the compliance rate: in 2005, around 90 percent of all Austrian households had registered broadcasting equipment and paid a total of €650 million (0.3 percent of GDP; see Berger et al. 2016). The number of registered households is in constant flux. New registrations emerge from mailing campaigns, door-to-door inspections, as well as from unsolicited registrations. The latter originate from households that register, for instance, using a web form or by calling a hotline.

B. Field Experiment

Fellner, Sausgruber, and Traxler (2013) tested different enforcement strategies in a field experiment. In cooperation with FIS, they randomly assigned more than 50,000 potential evaders, who were selected following the procedures described above, to an untreated *control* group or to different mailing treatments. All mailings, which were sent out during September and October 2005, included a cover letter and a response form with a prepaid envelope. The experiment varied whether or not the cover letter included a threat. The cover letter in the *baseline* mailing treatment simply clarified the legal nature of the interaction and asked why there was no registered broadcasting equipment at this household. In the *threat* treatment, the letter included an additional paragraph that communicated a significant risk of detection and emphasized possible financial and legal consequences from noncompliance (see the online Appendix D).

Fellner, Sausgruber, and Traxler (2013) found that the mailings had a significant impact on compliance. Most of the treatment responses occurred during the first weeks: within the first 50 days of the experiment, only 0.8 percent registered their broadcasting equipment in the control group. In the baseline mailing treatment, the fraction was 6.5pp higher. The threat treatment raised the registration rate by one additional percentage point. Beyond 50 days, there were no observable differences in registration rates. Complementary survey evidence suggested that in comparison to the control group, all mailings had a strong positive impact on the expected detection risk. Relative to the baseline, the threat mailing further increased the expected sanction risk. This pattern is consistent with the larger effect of the threat treatment.

The present paper studies whether the treatments triggered any spillover effects on the untreated population that were not covered by the experiment. More specifically, we exploit the experimental variation to analyze whether the mailing interventions affected untreated neighbors of those who were targeted. Given that neither the intervention itself (receiving a mailing³) nor the behavioral response (registering with FIS and starting to pay license fees) is observable to neighbors, communication is necessary for any spillover from treated households to untreated

³ Similar as in other countries, the privacy of correspondence is a constitutional right in Austria. Violations are punished according to the penal code (§118).

neighbors. In a first step, we will therefore discuss survey evidence on communication patterns.

II. Communication among Neighbors

To study communication among neighbors, we ran a survey with a professional online survey provider. The company maintains a sample that is representative of Austria's adult population. From this pool, we surveyed a subsample of almost 2,000 individuals. Participants were asked about the geographic distance to and the communication frequency with their first-, second-, and third-closest neighbors in terms of geographical (door-to-door) distance. We also elicited the relevance of TV license fees in the communication among neighbors. Details of the survey are relegated to online Appendix A.

The main findings from the survey are the following. First, the average intensity of communication among neighbors is fairly high, averaging about 60 percent of the respondents' communication intensities with their best friends from work/school. This finding is in line with other evidence suggesting that neighbors form an important part of people's social capital. The International Social Survey Programme's 2001 survey, for instance, shows that 11.2 percent of Austrians would turn to their neighbors as a first or second choice to ask for help in case they had the flu and had to stay in bed for a few days. Similar rates are observed for other central and northern European countries (e.g., Switzerland: 16.0 percent, Germany: 9.4 percent, Great Britain: 10.6 percent). For southern European countries (e.g., Italy: 4.7 percent) and the United States (6.3 percent), the data document lower rates.

Second, the intensity of communication declines with geographic distance. This result, which is again consistent with other research documenting the important role of geographic proximity for social interaction (e.g., Marmaros and Sacerdote 2006), is depicted in Figure 1. Communication frequencies monotonically drop from first-, to second-, and third-closest neighbors. A similar correlation is observed when we explore variation in the door-to-door distance to the closest neighbor; the farther away this neighbor lives, the lower is the reported communication frequency. Once the distance surpasses 200 meters, communication levels drop. We will return to this point below.

Third, the positive link between geographic proximity and communication intensity is systematically violated in larger, more urban municipalities. The survey evidence indicates that this is due to households living in apartment buildings. By definition, these households live very close to each other but, at the same time, communicate fairly infrequently with their neighbors.⁴ This problem does not seem to occur in more rural areas. The survey data show that in small municipalities, where apartment buildings tend to be smaller and less anonymous, the "closeness"

⁴For details, see the corresponding discussion of Figure A.3 in the online Appendix A. It is worth noting that the evidence supports arguments made by Jacobs (1961), who criticized the urban planning policy of the 1950s/60s with its emphasis on large apartment blocks—precisely because it prevents many types of social interaction common in smaller municipalities.

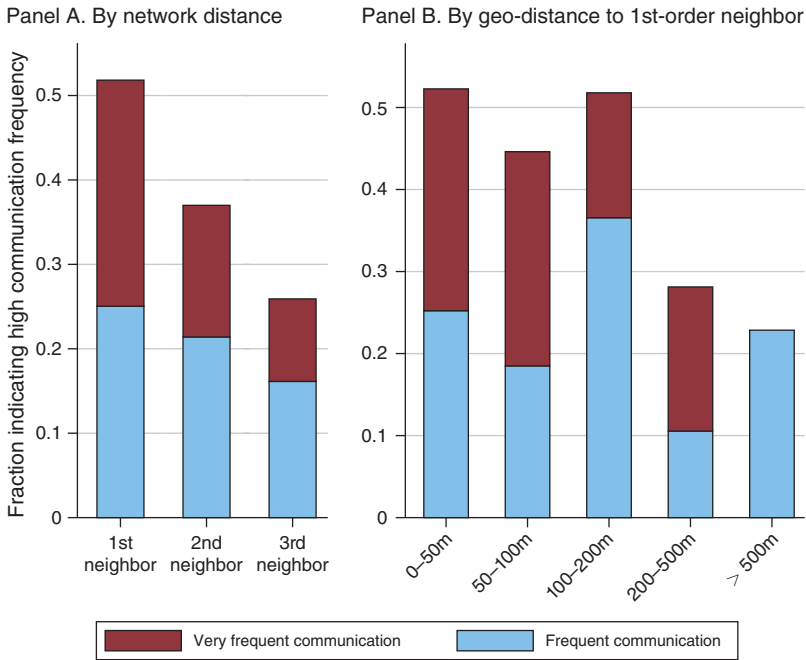


FIGURE 1. SURVEY RESULTS: COMMUNICATION FREQUENCY WITH NEIGHBORS

Notes: Panel A presents frequencies of communication with the first-, second-, and third-closest neighbor (in geographical terms), respectively. Panel B illustrates communication frequencies with the first-closest neighbor for different distance ranges to this neighbor.

of neighbors in apartment buildings is not aligned with lower communication frequencies (see online Appendix A).

Fourth, concerning the content of communication among neighbors, we observe that, in general, TV license fees are a relatively uncommon topic, similar to neighbors talking about job offers or financial opportunities (see Figure A.5 in online Appendix A). However, the survey reveals that people are willing to pass on license-fee-related information to their neighbors once some relevant news arrives. For a scenario where a household receives a FIS mailing, which indicates a possible inspection, almost two out of three respondents say that they would share this information with their neighbor and “warn” them (see Figure A.6 in online Appendix A). This seems reasonable, as field inspections are locally correlated. The evidence thus suggests that households are willing to initiate communication with their neighbors after receiving a mailing.

III. Data

To evaluate the impact of the experiment on the nonexperimental population, we build on several unique sets of data provided by FIS. The first data cover the universe of all Austrian households and include precise address information from official residency data together with FIS’s assessment of the households’ compliance before

the implementation of the field experiment. FIS derives this information—compliant or not (and thus potentially evading)—from data on all households paying license fees, data on past mailing campaigns and field inspections, as well as administrative data from the residents' registration office.

A second dataset covers the population from the field experiment (a subset of the first data) and indicates the treatment condition to which households were assigned. The third dataset contains information on all incoming registrations—unsolicited registrations, responses to mailings, and detections in door-to-door inspections—after the experiment. Using these data, we can observe behavioral changes in compliance, in particular registrations among the population from the field experiment and unsolicited registrations among those not covered by the field experiment. Our analysis will focus on the latter population.

A. Sample

The survey documents that geographic proximity is positively correlated with communication frequencies among neighbors in *small* but not necessarily in large municipalities (see Section II). In line with this finding, we focus on municipalities with less than 2,000 households (corresponding to a population size of approximately 5,000—the cutoff for small municipalities in the survey). The restriction is further motivated by the fact that these jurisdictions are predominantly characterized by detached, single-family houses. Less than 20 percent (5 percent) of households in these municipalities live in buildings with 3 (10) or more apartment units.⁵ For the geographical network approach introduced below, this is an important attribute.

Full Sample.—The sample restriction leaves us with 2,112 municipalities (out of 2,380) with an average of 1,700 inhabitants, covering almost half of the Austrian population. We geocoded the location of each single household from these municipalities. In a few cases, we failed to assign sufficiently precise geographic coordinates; we then excluded the affected parish (“*Zählsprenkel*”). With this procedure, we arrive at a sample of 576,373 households. Among these, one can distinguish three types: (I) potential evaders from the experimental sample, (II) potential evaders who were not covered by the experiment, and (III) compliant households (not part of the experiment). Figure 3, which is further discussed below, gives an overview and illustrates the different types' role in our research design.

Type I: Experimental Participants.—Our sample includes 23,626 households that were part of the field experiment. Summary statistics for these type I households, which will serve as “injection points” in our analysis of indirect treatment effects, are provided in Table 1. The table splits the experimental sample according to the three treatment groups: 1,371 households were in the untreated control group, 11,117 in

⁵ Among municipalities with 2,000–3,000 households, the share jumps to 39 percent [15 percent].

TABLE 1—SUMMARY STATISTICS FOR TYPE I HOUSEHOLDS

	Experimental condition		
	Control (no mailing)	Baseline mailing	Threat mailing
Registration rate	0.0109 (0.1041)	0.0701 (0.2553)	0.0765 (0.2659)
Ex ante compliance	0.3610 (0.4805)	0.3628 (0.4808)	0.3752 (0.4842)
Male	0.7644 (0.4245)	0.7472 (0.4346)	0.7588 (0.4278)
Age ^a	40.04 (13.00)	39.67 (12.47)	39.28 (12.11)
Degree	19.88 (52.09)	21.67 (55.39)	21.42 (55.83)
Eigenvector centrality	0.3075 (0.2399)	0.3031 (0.2415)	0.3051 (0.2431)
Diffusion centrality ^{T=1}	0.6443 (0.3776)	0.6379 (0.3790)	0.6372 (0.3798)
Diffusion centrality ^{T=10}	0.4392 (0.4785)	0.4349 (0.4771)	0.4337 (0.4775)
Number of households	1,371	11,177	11,078

Note: The columns present sample means (and standard deviations in parentheses) for the experimental sample, i.e., all type I households from the three treatment arms of the experiment: Control, Baseline, and Threat Mailing.

^a Information on age is only available for a subset of 2,778 households.

the baseline mailing, and 11,078 in the threat-mailing treatment. Consistent with Fellner, Sausgruber, and Traxler (2013), we observe three patterns:

- The observables are balanced across the treatments; this holds for age, gender, as well as network characteristics (degree and Eigenvector and Diffusion Centrality; see below).⁶
- The registration rates for the mailing treatments are significantly higher than in the untreated group. After the first 50 days of the experiment, 1.09 percent of all households in the control group registered for license fees. For the baseline mailing treatment, it was 7.01 percent.
- The threat mailing has a stronger effect. Table 1 indicates a registration rate of 7.65 percent.

Type II: Potential Evaders Not Covered by the Experiment.—In addition to the experimental participants, the sample includes 128,059 type II households that were classified as potential evaders at the time of the experiment. There are at least three reasons why these households were not part of the experimental sample. First, FIS excludes those that were “unsuccessfully” contacted with mailings in the past from future mailing campaigns. Second, all households that first appeared in the official

⁶Table 1 does not include any point estimates for the between-treatment-group difference. However, as is clear from the summary statistics, no variable turns out to be statistically different across the three groups. Note further that the high share of males is due to FIS’s procedure treating male individuals as household heads.

residents' registration record during the experiment's setup time could not be included in the experiment (e.g., recently formed households). Hence, some type II households might be longtime evaders, others short-term evaders. Third, the classification of potential evaders is also based on information that was not available to FIS during the experiment's setup phase (see below).

It is worth noting that types I and II households together account for a fourth of the total sample. This high fraction, which is well above the overall rate of noncompliance, reflects the fact that FIS's method to identify potential evaders is imperfect and delivers many "false positives": compliant households that are wrongly flagged as evaders.⁷ This point is also reflected in Table 1, which shows that the ex ante compliance rate (*before* the experimental intervention) among type I households was roughly 36 percent. A nonnegligible fraction of the mailing targets could therefore not respond by switching from evasion to compliance—a fact that we will exploit below (see Section IVE). Finally, note that the classification of potential evaders in the nonexperimental sample should be more accurate. For the geocoding procedure, we thoroughly cleaned the address data, which allowed us to eliminate many false positives within FIS's database (see footnote 7). As a consequence, the ex ante compliance rate in the type II sample should be considerably lower than in the type I sample.⁸

B. Geographical Networks

Our analysis studies whether potential evaders who were *not* covered by the experiment (type II households) start to comply with license fees in response to experimental interventions (the treatment of type I households) in their geographical network of neighbors. We therefore focus on networks that cover at least one type I and at least one type II household. We call these the *relevant* networks. To derive geographical networks, we first compute Euclidean distances between all households in each municipality. Whenever the distance between two households i and j is below an exogenous threshold z , we say there is a link between i and j . A network then consists of all households that are either directly or indirectly linked. Households that are directly linked to i are referred to as i 's first-order neighbors (FONs). Households one link farther away are referred to as second-order neighbors (SONs). Figure 2 illustrates this approach and shows how it produces disjoint networks.

A reasonable choice for the threshold z can be motivated by the survey evidence, which suggests that communication frequencies with FONs decline sharply once the geographical distance exceeds 200 meters (see Figure 1). This suggests $z \leq 200$ meters. Note further that larger thresholds leave us with fewer but larger networks. This point is illustrated in Table B.1 in online Appendix B. The table displays the number of relevant networks as well as the number of different household types per

⁷For instance, FIS's database frequently fails to account that households are already complying as they omit compliant household members (i.e., an individual who pays and one who does not are treated as different households), mainly due to typos in names and addresses.

⁸Note further that cheaters, who were detected in field inspections *before* the experiment, are classified as type III households.



FIGURE 2. ILLUSTRATION OF GEOGRAPHICAL NETWORKS

Note: The figure presents an example of two disjoint networks for a distance $z = 50$ meters.

network for different thresholds z . For $z = 50$, we observe the largest number of relevant networks. Since this will facilitate any between-network analysis, we will use a threshold of 50 meters as a benchmark for our analysis. To assess the robustness of our findings with respect to z , we rerun all our main estimations for networks based on thresholds between 25 and 2,000 meters.

With a 50-meter threshold, we arrive at 3,764 relevant networks that were covered by the experiment. The networks come from 771 different municipalities and include about 68,000 households (of types I, II, and III; see Table B.1 in online Appendix B). Among these, there are 14,987 type II households. Summary statistics for this group are provided in panel A of Table 2. The variable degree shows that the median (mean) type II household is linked to 6 (11) FONs who live within 50 meters distance. Note, 33.04 percent (33.34 percent) of type II households have a FON treated with a baseline (threat) mailing in the experiment. Finally, the variable registration rate indicates that 8 percent of type II households unsolicitedly registered within 50 days after the experiment. This is more than twice the average registration rate among *all* nonexperimental potential evaders (i.e., type II households inside and outside of networks covered by the experiment; in this population, the registration rate is 3 percent).⁹ Below we will show that the higher registration rate can be explained by the presence of spillover effects from experimental to nonexperimental households in these networks.

Panel B reports descriptive statistics at the network level. The network size, in the following denoted by N_k , has a median (mean) of 6 (18) households. For each network k , we computed variables that measure the treatment coverage: $Total_k$ captures the rate of other households in the experiment sample divided by $N_k - 1$. Similarly, $Base_k$, $Threat_k$, and $Control_k$ indicate the ratios of other

⁹One cannot directly compare these registration rates to those observed among type I households. As pointed out in Section IIIA, the latter sample contains a high fraction of households that were already complying before the experiment.

TABLE 2—SUMMARY STATISTICS FOR RELEVANT NETWORKS, $z = 50$ METERS

	Mean	SD	Median	1st quarterile	3rd quarterile
<i>Panel A. Household level (14,987 type II households)</i>					
Degree	11.11	22.22	6	4	11
FON in experiment	0.54	0.50	1	0	1
FON w/base	0.33	0.47	0	0	1
FON w/threat	0.33	0.47	0	0	1
FON w/control	0.08	0.26	0	0	0
SON in experiment	0.12	0.33	0	0	0
SON w/base	0.06	0.25	0	0	0
SON w/threat	0.07	0.25	0	0	0
SON w/control	0.01	0.09	0	0	0
Spatial distance to closest neighbor...					
... in experiment	95.94	143.86	45.06	2.27	113.18
... w/base	291.01	734.07	84.54	32.27	240.34
... w/threat	315.89	739.32	90.67	34.15	270.87
... w/control	678.53	1,081.36	295.58	114.66	732.02
Registration rate	0.08	0.27	0	0	0
<i>Panel B. Network level (3,764 networks)</i>					
Network size (N_k)	17.98	45.48	6	3	14
Treatment rates:					
<i>Total_k</i>	0.45	0.33	0.40	0.39	0.67
<i>Base_k</i>	0.21	0.28	0.10	0.00	0.33
<i>Threat_k</i>	0.22	0.29	0.10	0.00	0.33
<i>Control_k</i>	0.03	0.11	0.00	0.00	0.00
Network w/ $Control_k > 0$	0.14	0.35	0.00	0.00	0.00
<i>Panel C. Municipality level (771 municipalities)</i>					
Population	1,790	1,006	1,570	1,060	2,365
Labor income	27,250	2,172	26,936	25,977	28,333
1- or 2-family dwellings	0.82	0.14	0.84	0.75	0.92
Average age	47.95	1.55	47.92	46.85	48.97
Non-Austrian citizens	0.05	0.04	0.04	0.02	0.07
Catholic	0.89	0.10	0.91	0.86	0.95
Voter turnout	0.76	0.07	0.77	0.72	0.81

Notes: Panel A presents summary statistics at the household level for all potential evaders in the non-experimental sample (type II households) that are located in networks with at least one experimental participant (type I household). Panel B reports (unweighted) network level statistics for these relevant networks. Panel C considers several municipality characteristics (unweighted, at municipality level). FON and SON abbreviate First- and Second-Order Neighbors, respectively. Spatial distance is measured in meters.

households targeted with a baseline, a threat mailing, and untreated experimental households, respectively. (For a numerical example, see Figure 3.) Using $(N_k - 1)$ as denominator assures that the treatment rates vary between zero and one.¹⁰ Table 2 shows that from the perspective of a type II household in an average network, 45 percent of the other households in a network were covered by the experiment; 21, 22, and 3 percent of the other network members were in the baseline, threat, or control treatment, respectively.¹¹ While these numbers also reflected the small control group from the original experiment, it is important to note that 528 networks (14 percent of all networks) contain at least one untreated experimental household

¹⁰ Our estimates focus on the responses of type II households. Computing treatment rates relative to N_k would impose an upper bound (at $(N_k - 1)/N_k$) which mechanically varies with the network size.

¹¹ The high treatment ratios reflect our focus on (relevant) networks with at least one experimental household. The small share of households in the control group is a result of FIS's pressure to keep the untreated group fairly small.

(i.e., $Control_k > 0$). As we further discuss below, these networks are important for the identification of our main model.

Panel C of Table 2 presents summary statistics for census data at the municipality level. An average municipality (with a relevant network) is populated by 1,790 inhabitants with a mean labor income (wages and salaries) of €27,250; 82 percent of households live in single- and two-family homes. Household heads are on average 48 years old. The fraction of non-Austrian citizens is low (5 percent), and a large majority of the population is Catholic (88 percent). We also observe a high voter turnout at the 2006 national elections (77 percent on average).

IV. Spillover Effects

This section studies the indirect effects from the experiment on the nonexperimental population. We want to identify whether and how a type II household's probability to register for license fees changes in response to the experimental interventions in their neighborhood networks. Recall that in our setting neither the treatment nor compliance behavior is publicly observable (see Section IB). Any spillovers must therefore stem from communication rather than from direct observability of others' behavior.

A. Conceptual Framework

Each household i has a latent propensity to comply p_i , which depends on numerous factors governing compliance (including, for instance, the perceived sanction risk or the strength of social norms). We will discuss these different factors in Section IVE. For the moment, we focus on how propensity p_i evolves over time as a function of information acquired via communication in networks.

Each household is located in one network defined as a collection of nodes $\mathcal{N} = \{1, \dots, N\}$ and a set of edges (links between the nodes) defined as $\Xi \subseteq \{(i, j) | i \neq j \in \mathcal{N}\}$, where an element (i, j) indicates that i and j are linked. The set of i 's first-order neighbors (FONs) is denoted by $\mathcal{N}_i = \{j \in \mathcal{N} | (i, j) \in \Xi\}$. In each round of communication t , households i transmit to their FONs some information I_i^t with probability q independently across neighbors. The first round of communication, $t = 1$, takes place *after* treatment has occurred. Receiving a mailing directly increases a treated household's p_i (i.e., there is a direct treatment effect). In addition, the mailings affect the information that is passed on to a treated household's neighbors.¹²

After receiving information, households update their propensities and start to comply as soon as p_i^t exceeds a given, household-specific threshold \hat{p}_i . In each further round of communication t , households communicate with probability q with *their* FONs, again pass on information I_i^t , and update p . Indirect treatment effects

¹²We do not specify what information is transmitted, as our institutional setting does not constrain I_i^t . For the moment, one may simply consider that I_i^t refers to statements such as "I received a mailing" or "household j received a mailing." However, I_i^t could contain further information, e.g., about i 's propensity to comply, her perceived sanction risks, or any other dimension affected by the mailings. Note further that all households can communicate, but one may find it natural to think of communication being initiated by treated households.

then arise as the different mailing treatments can (differentially) alter the communicated information, which in turn influences the compliance decision of untreated households within the network.

In the following, we will first establish the existence of indirect treatment effects, i.e., compliance spillovers within networks (Sections IVB, IVC, and IVD). We then discuss evidence on what type of information I_i^t is communicated (Section IVE). Section V investigates the role of the network (\mathcal{N}, Ξ) in structuring the communication process.

B. Identifying Indirect Effects from the Experiment

A graphical illustration of our research design and how it relates to the different types introduced in Section IIIA is provided in Figure 3. The figure depicts three networks of equal size and structure, which all include three experimental households (type I), two potential evaders (type II), and two compliant (type III) households. For any given rate of experimental households, our strategy exploits variation in the randomly assigned treatments among experimental households.

The spillover effects from the experiment are estimated with the following model:

$$(1) \quad y_{ik} = \delta^{Total_k} + \beta_1 Base_k + \beta_2 Threat_k + \epsilon_{ik},$$

where y_{ik} indicates if a type II household i from network k starts to comply with license fees within 50 days after the intervention. The specification non-parametrically controls for the total rate of experimental participants in the networks by adding δ^{Total_k} , a set of fixed effects for each level of $Total_k$ (with $Total_k = Control_k + Base_k + Threat_k$; see Section IIIB). The key regressors measure the treatment rates at the network level, i.e., the fraction of type I households in the network that were in the baseline ($Base_k$) or in the threat treatment ($Threat_k$).¹³ For the examples from the top, middle, and bottom of Figure 3, we would obtain $Base_k [Threat_k]$ equal to 1/6, 2/6, and 1/6 [1/6, 1/6, and 2/6], respectively. For all three cases, $Total_k$ equals 1/2. Clearly, equation (1) also exploits variation in rates related to variation in network size. This is not captured in Figure 3, which keeps the size of the networks and the number of experimental participants constant.

As long as communication takes place among all households, the treatment rates should be positively associated with the information that household i receives. If i 's propensity to comply is increasing in the obtained information, the spillover effects should then be captured by a positive β_1 and β_2 in model 1. The two coefficients of interest measure the effects on type II households' propensity to start paying fees, in response to an increase in the network's rate of baseline and threat treatments, while keeping constant $Total_k$. Put differently, these are the effects from moving experimental households from the control group to one of the mailing treatments.

¹³ Alternative specifications in levels yield almost identical results to those reported below. Estimations in levels, however, turn out to be more sensitive to outliers related to a few very large networks.

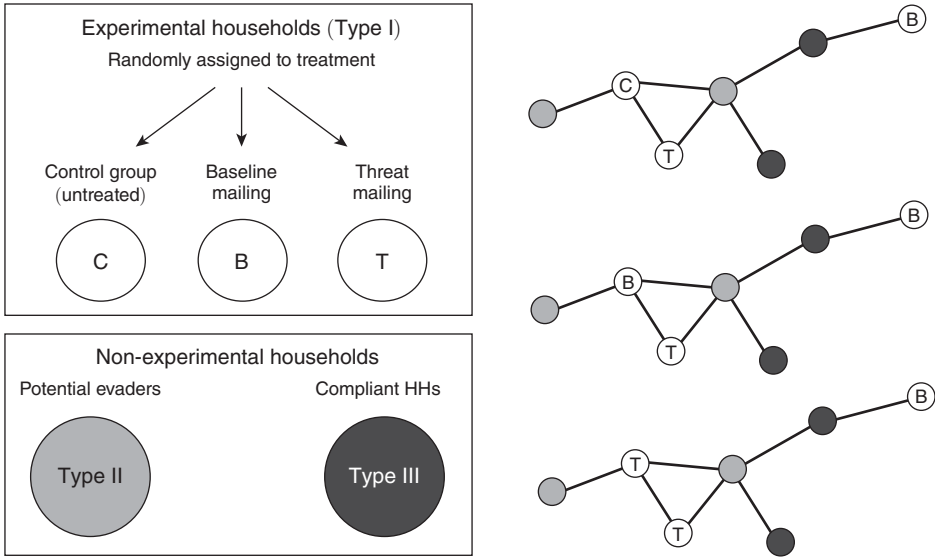


FIGURE 3. ILLUSTRATION OF RESEARCH DESIGN: HOUSEHOLD TYPES AND EXPERIMENTAL TREATMENTS

Notes: The figure provides an overview of the different household types (see Section IIIA) and illustrates the type of variation across networks that is explored in Section IV. The illustrated networks have the same size and structure. All networks include three type I households. For this given rate of experimental households (as captured by $Total_k$), our main specification exploits the variation in the randomly assigned treatments. The three examples cover the case of (i) one Control, one Baseline, and one Threat treatment; (ii) two B and one T; and finally (iii) one B and two T. Treatment rates are computed—from the perspective of type II households—relative to all other households, i.e., we divide the number of treated households by $N - 1 = 6$. For the examples, all networks have a $Total_k$ equal to $1/2$; $Control_k$ equal to $1/6$, 0 , and 0 ; $Base_k$ equal to $1/6$, $2/6$, and $1/6$; and $Threat_k$ equal to $1/6$, $1/6$, and $2/6$, respectively.

Identification.—As pointed out above, β_1 and β_2 are identified from network-level variation in the treatment rates for a given level of experimental coverage ($Total_k$). The identifying assumption is that conditional on $Total_k$, variation in $Base_k$ and $Threat_k$ is exogenous. Given our experimental setup, this assumption seems plausible: between networks with the same coverage, the assignment of the experimental households to the different treatments varies randomly and thus supports this conditional independence assumption (CIA).

Note that equation (1) accounts for selection into the experiment by including dummies δ^{Total_k} . Networks that differ in experimental coverage are therefore allowed to be different in unobservables that might affect the propensity to comply. One might expect, for instance, to see more experimental participants in networks where the average propensity to comply is lower. There might also be more “tougher” evaders among the untreated households in such networks. By (nonparametrically) controlling for $Total_k$, specification (1) accounts for such selection effects. Despite the small control group in the original experiment, model (1) is identified because we have enough networks with a positive share of control households within the experimental sample. (As reported in Table 2, 14 percent of all networks have $Control_k > 0$.) In fact, identification rests on the cross-sectional variation of all networks (even those with $Control_k = 0$) with the same rate of experimental

coverage but different baseline and threat rates (see Figure 3 for simple examples with a fixed network size).

To illustrate why the experiment supports our CIA, note first that randomization guarantees the balance of type I households' characteristics \mathbf{x} across treatments:¹⁴

$$(2) \quad E(\mathbf{x}|Base) = E(\mathbf{x}|Threat) = E(\mathbf{x}|Control).$$

Evidence in Fellner, Sausgruber, and Traxler (2013) and in our Table 1 is consistent with orthogonality, suggesting that randomization in the experiment was successful. For our CIA to be violated, it would take that, after partitioning the experimental households into networks, some mailings be (for a given level of $Total_k$) disproportionately sent to certain types of networks rather than others. This case can be excluded if we still have orthogonality after condition on $Total_k$:

$$(3) \quad E(\mathbf{x}|Base, Total_k) = E(\mathbf{x}|Threat, Total_k) = E(\mathbf{x}|Control, Total_k).$$

Note further that the following identities hold:

$$E(\mathbf{x}|Base) = E(\mathbf{x}|Base, Total_k) \cdot f(Total_k|Base),$$

$$E(\mathbf{x}|Threat) = E(\mathbf{x}|Threat, Total_k) \cdot f(Total_k|Threat),$$

$$E(\mathbf{x}|Control) = E(\mathbf{x}|Control, Total_k) \cdot f(Total_k|Control).$$

If the probability density functions $f(Total_k|\cdot)$ are equal across the three treatments and if (2) holds, it immediately follows that (3) must hold too. Support for this case is provided in Figure B.1 in online Appendix B, which indicates that the probability densities for the three groups are indeed almost identical.

To provide additional evidence in support of our identifying assumption, we estimate models of the following structure:

$$(4) \quad Base_k = \delta^{Total_k} + \mu^{Base} x_k + \epsilon_k^{Base} \quad \text{and} \quad Threat_k = \delta^{Total_k} + \mu^{Threat} x_k + \epsilon_k^{Threat},$$

where x_k is an observable characteristic that varies at the network level. Our conditional independence assumption implies that, controlling for $Total_k$ fixed effects, we should not find any correlation between observable network characteristics and our key regressors: neither μ^{Base} nor μ^{Threat} should be statistically different from zero. This is what we observe in Table B.2 in online Appendix B where we report the estimated μ^{Base} (columns 1 and 2) and μ^{Threat} (columns 3 and 4). We do not obtain any economically or statistically significant coefficients for network characteristics or characteristics of injection points (type I households, averaged at the network level). The results support the notion that

¹⁴In the following, the variables *Base*, *Threat*, and *Control* without subindex are dummies indicating which experimental group a type I household was assigned to.

TABLE 3—BASIC RESULTS: IMPACT OF MAILINGS ON COMPLIANCE OF NON-EXPERIMENTAL HOUSEHOLDS

	Networks with $z = 50m$		Municipality level	Probit	Municipality fixed effects	Field inspections	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Base_k$	0.2431 (0.0381)	0.2473 (0.0383)	-0.5672 (0.8102)	0.1972 (0.0554)	0.2219 (0.0451)	0.2440 (0.0381)	0.2486 (0.0382)
$Threat_k$	0.3496 (0.0390)	0.3532 (0.0391)	0.1157 (0.7995)	0.2423 (0.0554)	0.3376 (0.0464)	0.3499 (0.0390)	0.3566 (0.0392)
$Total_k$	—	-0.0594 (0.0332)	0.3825 (0.7395)	-0.0569 (0.0545)	—	—	—
$Total_k$ fixed effects	Yes	No	No	No	Yes	Yes	Yes
F -test: $Base_k = Threat_k$	13.44	14.01	5.16	13.32	14.34	13.31	13.58
Observations	14,987	14,987	14,987	14,987	14,987	14,987	13,395
Networks	3,764	3,764	771	3,764	3,764	3,764	3,588
R^2	0.0946	0.0553	0.0144	0.0770	0.1873	0.0966	0.0994

Notes: This table estimates the impact of the mailings on the compliance of households outside of the experimental sample (type II households). Column 1 presents the results from a linear probability (LPM) estimation of equation (1). Column 2 replaces δ^{Total_k} (i.e., the 451 dummies for each value of $Total_k$) with a linear control for $Total_k$. Column 3 uses the sample from the first two specifications but assumes that the network is defined by the municipality. Column 4 replicates the specification used in column 2 using a Probit model. It reports marginal effects evaluated at the mean of the independent variables. The remaining output comes again from LPM estimations including $Total_k$ fixed effects. Column 5 adds fixed effects at the municipality level. Column 6 controls for the enforcement rate (i.e., the rate of households detected in field inspections in the period before the start of experiment); the estimated coefficient on the enforcement rate is 0.387 (SE 0.120). Column 7 excludes all networks with an enforcement rate greater than zero. The table further reports F -statistics, testing $Base_k = Threat_k$. Standard errors, clustered at the network (all columns except 3) or municipality level (column 3), are reported in parentheses.

after controlling for experimental coverage, there is no selection on observables. Hence, both exercises lend support to the validity of our identifying assumption.

C. Basic Results

Using a linear probability model, we estimate model (1) for all potential evaders from the nonexperimental population (type II households) in relevant networks with $Total_k > 0$.¹⁵ The results, together with standard errors clustered at the network level, are reported in column 1 of Table 3. The coefficients of interest are both positive and precisely estimated. The point estimates imply that a 1 percentage point increase in the rate of the baseline (threat) treatment increases the likelihood that an untreated potential evader registers by 0.24 (0.35) percentage points, respectively. An F -test on the equality of the coefficients on the baseline and threat treatment rates rejects the null that the two effects are equal. The estimates hardly change when we switch from a nonparametric to a parametric control for the networks' experimental coverage. In fact, specification (2) shows that the coefficient on $Total_k$ is small and only weakly significant. The negative sign indicates that a higher experimental coverage in a network is correlated with a lower probability of unsolicited registration among type II households.

¹⁵Including networks with $Total_k = 0$ does not change the results reported below.

To illustrate the size of the spillover effect, consider the thought experiment in which we move one experimental household from the control to the threat treatment. For a median network with $N = 6$, the additional mailing implies a 20 percentage point increase in the threat treatment rate ($1/(6 - 1) = 0.2$). Our estimates imply that the additional threat mailing would increase the type II households' probability to register by 7 percentage points ($0.2 \times 0.35 = 0.07$). On average, type II households account for roughly 20 percent of the network population (see Table B.1 in online Appendix B). Hence, in a network with $N = 6$, only one will be of type II (20 percent of $N - 1 = 5$). We would therefore expect a total spillover of 0.07×1 unsolicited registrations for license fees.¹⁶ Keeping constant the experimental coverage, one additional threat (baseline) mailing thus increases the probability of observing one additional registration among the untreated evaders in the network by 7 (5) percentage points. Although the comparison of registration rates between the treated and the untreated sample is complicated (see footnote 9), it is worth noting that the total spillover effect seems to be of similar magnitude as the direct treatment effects on type I households (5.9 and 6.6 percentage points for the baseline and threat treatment, respectively; see Table 1). The implied *social multiplier* is thus around two, meaning that the overall effect of the intervention on compliance is twice the one observed by only looking at the targeted households.

Having detected a significant spillover from the experiment onto the nonexperimental population, let us point out that one would miss this indirect treatment effect if one estimates equation (1) at the municipality rather than the network level. This is documented in column 3 of Table 3. For this specification, we assigned the sample from columns 1 and 2 (i.e., all type II households from geographic networks with $Total_k > 0$) into one network per municipality and recomputed the treatment rates. Despite the fact that these municipalities are still fairly small observational units (with an average population of 1,790 individuals), we obtain estimated coefficients on the baseline and threat treatment rates that are both statistically insignificant. Column 4 reports the marginal effects from a probit estimation of the specification from column 2. The spillover effects are similar to those indicated by the linear probability model estimates, and the effect size differs again significantly between the two mailing treatments. In column 5, we augment the basic model (1) by adding fixed effects at the municipality level. As expected, this leaves our results unchanged.

To account for possible effects from local enforcement activities (see Rincke and Traxler 2011), column 6 controls for enforcement activities during the months before the experiment. More specifically, we compute an enforcement rate at the network level that measures the number of households detected in field inspections before the experiment (during the third quarter of 2005) relative to the network size. Consistent with Rincke and Traxler (2011), enforcement has a significantly positive relationship with the propensity to register (coefficient reported in the notes to Table 3). However, the point estimates for our coefficients of interest are essentially

¹⁶For larger networks, the effect of one additional mailing on the treatment rates would be smaller, but the spillovers would spread to a larger number of potential evaders. It is straightforward to show that equation (1) implies that the *total* spillover from sending one additional mailing into a network is independent of the network size.

identical to those from column 1, suggesting that pretrial enforcement was not correlated with the treatment variation. In an additional step, we run our basic model excluding networks in which at least one household was detected in field inspections before the trial.¹⁷ The results reported in column 7 show that the estimated coefficients remain again almost unchanged.

We also examined heterogeneous spillover effects according to different municipality characteristics. The analysis reveals significant interaction effects for only two variables: dwelling structure and voter turnout. The magnitude of the spillover increases with the fraction of people living in single- or two-family dwellings (as compared to multifamily homes) as well as with the turnout (see Table B.3 in online Appendix B). The latter interaction could reflect different levels of social capital and communication intensities or stronger norms supporting prosocial behavior.

D. Additional Results and Robustness Checks

Different Network Assumptions.—To understand the sensitivity of our findings with respect to z , we computed geographical networks based on distance thresholds that vary between 25 and 2,000 meters.¹⁸ We then replicate our estimates for the different samples of relevant networks. Panel A in Table 4 reports the results from this exercise. The estimated coefficients turn out to be fairly stable for $z < 500$ meters. For larger values of z , the coefficient on the baseline treatment starts to decline, whereas the one on the threat treatment remains substantial (but the standard error becomes large).

It is important to keep in mind that any change in z also varies the number of relevant networks, the average network size, as well as the number of type II households (see online Appendix Table B.1). This clearly complicates the comparison and interpretation of the different point estimates as well as the estimation of the model from (1). The last point is due to the fact that the number of networks approaches the number of $Total_k$ fixed effects as z becomes larger. For large networks (with $z > 500$), the δ^{Total_k} dummies then absorb almost all variation across networks. For this reason, Table 4 is based on the specification that linearly controls for experimental coverage (as in column 2 of Table 3). Estimations from our preferred model (with nonparametric controls for $Total_k$) yield almost identical results for networks with $z < 500$.

A first attempt to facilitate a meaningful comparison across samples is provided in panel B of Table 4. Here we normalize the estimated coefficients relative to the mean network size. With this normalization, we get the effect from sending *one* additional mailing to *each* relevant network on the register probability of an average type II household (in an average-sized network). Panel B shows that the effect from one mailing monotonically declines with z . Given the results from panel A, this pattern is due to the fact that the average network size increases monotonically

¹⁷This drops 176 networks from the sample. The number of nonexperimental households drops by a larger share because the excluded networks tend to be larger ones. This is due to the fact that, ceteris paribus, the probability to have at least one household detected by a field inspector is increasing with the network size.

¹⁸Employing within-municipality distance matrices, we exclude links between networks from different municipalities. This restriction becomes relevant for $z \geq 500$ but affects only a small part of the sample.

TABLE 4—SENSITIVITY ANALYSIS: DIFFERENT NETWORK (DISTANCE) ASSUMPTIONS

Threshold $z =$	25	50	75	100	250	500	1,000	1,500	2,000
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>Panel A. Estimation results</i>									
<i>Base_k</i>	0.2454 (0.0327)	0.2473 (0.0383)	0.2692 (0.0573)	0.2761 (0.0544)	0.2720 (0.0796)	0.2641 (0.1354)	0.2210 (0.2564)	0.0658 (0.3650)	-0.0731 (0.4771)
<i>Threat_k</i>	0.3144 (0.0333)	0.3532 (0.0391)	0.3954 (0.0590)	0.3904 (0.0560)	0.3564 (0.0823)	0.5038 (0.1360)	0.4816 (0.2572)	0.4196 (0.3555)	0.2713 (0.4507)
<i>Total_k</i>	0.0160 (0.0304)	-0.0594 (0.0332)	-0.1095 (0.0523)	-0.1231 (0.0475)	-0.1237 (0.0713)	-0.2050 (0.1226)	-0.1867 (0.2359)	-0.0863 (0.3315)	0.0511 (0.4295)
Constant	0.0195 (0.0078)	0.0306 (0.0027)	0.0345 (0.0021)	0.0384 (0.0020)	0.0420 (0.0018)	0.0430 (0.0016)	0.0446 (0.0015)	0.0446 (0.0015)	0.0444 (0.0015)
Observations	5,337	14,987	23,673	29,212	41,547	50,688	58,498	60,520	61,551
Networks	3,243	3,764	3,319	2,990	2,113	1,554	1,169	1,073	1,020
R ²	0.0775	0.0553	0.0363	0.0258	0.0146	0.0110	0.0078	0.0075	0.0075
<i>Panel B. Average individual effect from one additional mailing into each network</i>									
Base	0.0533	0.0111	0.0050	0.0035	0.0017	0.0004	0.0001	-0.0001	-0.0001
Threat	0.0673	0.0173	0.0090	0.0061	0.0026	0.0020	0.0013	0.0013	0.0012
<i>Panel C. Total spillover from 1,000 additional mailings, spread over all networks</i>									
Base	87.68	44.04	35.68	33.89	32.60	12.76	7.41	-4.45	-4.77
Threat	110.80	68.87	63.86	59.19	51.18	64.55	63.63	72.10	69.84

Notes: Panel A reports the results from LPM estimations of the equation $y_k = \beta_0 Total_k + \beta_1 Base_k + \beta_2 Threat_k + \epsilon_k$ (see specification 2 in Table 3) for different samples that emerge for different distance thresholds z defining the networks. Standard errors, clustered at the network level, are in parentheses. Panel B presents the effect of sending *one* additional baseline [threat] mailing into each network on a type II household's probability to register. The effect is derived from the point estimates from panel A and given by $(\hat{\beta}_0 + \hat{\beta}_1)/(N - 1) [(\hat{\beta}_0 + \hat{\beta}_2)/(N - 1)]$, where N indicates the average network size for a given threshold z . Panel C computes the expected spillover from sending a fixed number of 1,000 additional baseline [threat] mailings into the networks covered by the respective sample; i.e., for the baseline treatment, the total number of expected spillover registrations is computed as follows: Number of Observations \times (1,000/Number of Networks) \times $(\hat{\beta}_0 + \hat{\beta}_1)/(N - 1)$. The effect is weighted with the total number of observations to account for the fact that the spillover applies to all type II households in these networks.

with z . One additional mailing thus implies a smaller increase in treatment rates in the larger networks that are obtained from higher values of z . However, the metric from panel B neglects that the average spillover applies to a larger population.

The latter point is accounted for in panel C, which reports the results from a different thought experiment. It considers sending a fixed amount of 1,000 additional baseline or threat mailings to relevant networks. Based on our estimates and the network properties, we then compute the *total* number of additional registrations that we expect to be induced by the spillover effects from these mailings.¹⁹ For $z = 50$, for instance, the number of expected spillover registrations that *indirectly* emerges from 1,000 additional baseline (threat) mailings adds up to 44 (69), respectively. As pointed out above, this is very similar to the *direct* compliance effect of 1,000 additional mailings, implying a social multiplier of approximately two, which is not uncommon in the literature on social interactions

¹⁹For the baseline mailings, this number is computed as follows: Number of Observations \times (1,000/Number of Networks) \times $(\hat{\beta}_0 + \hat{\beta}_1)/(N - 1)$, where $\hat{\beta}_0$ is the coefficient on *Total_k* and N is the *mean* network size (for each different z). The effect is weighted with the total number of observations to account for the fact that the spillover applies to all type II households in the networks.

(Glaeser, Sacerdote, and Scheinkman 2003).²⁰ Finally, panel C indicates that the overall spillover from the baseline mailing shrinks for $z > 250$. This observation fits the survey evidence, which showed that communication frequencies among FONs sharply decline in this range.

Spillovers within the Experimental Sample.—Given our main results from above, it seems natural to ask whether there are also spillovers *within* the experimental sample. If a type I household's treatment response depended on the treatment of other households in the network, this would imply a violation of the stable unit treatment value assumption (SUTVA)—see Imbens and Wooldridge 2009—for evaluating the direct effect of the experiment. To explore this case, we focus on type I households and analyze whether the behavior of a treated household depends on the treatment rates in its network in addition to its own treatment. Our analysis does not yield evidence that treatment responses of type I households are influenced by the treatment of their neighbors: controlling for the baseline and threat mailing rates does not alter the estimates for the direct treatment effects of the mailings (see Table B.4 in online Appendix B). The results suggest that, for the experimental sample, the direct treatment effect dominates any indirect effects from the experiment.

Permutation Test on Networks.—In principle, columns 1–3 from Table 3 could be interpreted in support of the idea that geographical networks of neighbors are a key unit for the information transmission that shapes the spillover effects. A concern with this interpretation is that we might simply have too little variation to detect any spillover when we estimate the regressions at the municipality level (see column 3). To address this concern and provide further evidence that geographical networks are crucial in determining the spillovers, we perform the following permutation test.

Within each of the 771 municipalities covered by the sample from our main specification from Table 3, we randomly allocate all (types I, II, and III) households into networks of size $N = 10$.²¹ With this procedure, households remain in their “true” municipalities, but they are randomly grouped in different networks, irrespective of their geographic location within the municipality. For such randomly generated networks, we then compute our regressors and estimate equation (1). The results from 1,000 iterations of this exercise are illustrated in Figure 4.

The figure displays the cumulative distribution functions of the estimated coefficients for the baseline (β_1 , panel A) and the threat mailing (β_2 , panel B) and the point estimates from Table 3. We obtain one single case (out of 1,000 iterations) for which an estimate from the permutation tests is larger than the estimates from above. This suggests that the results from Table 3 are not simply driven by partitioning municipalities into smaller units. Instead, the networks based on geographical

²⁰ Based on Table 1, the average *direct* treatment effect would add up to 59 [66] registrations from 1,000 additional baseline [threat] mailings. Hence, relative to this benchmark, diffusion and social interaction double the impact of the mailings. Reassuringly, this ratio does not change that much if we consider lower ($z = 25$) or higher (e.g., $z = 250$) distance thresholds for computing the networks. For $z = 25$ [$z = 250$] the implied social multiplier would be approximately 2.5 [1.75].

²¹ For a municipality with, say, 1,017 households, we would randomly form 100 networks with 10, and 1 network with the remaining 17 households. Note further that our analysis yields very similar results when we use $N = 5$ (close to the median network size, see Table 2) or $N = 15$ (close to the third quartile).

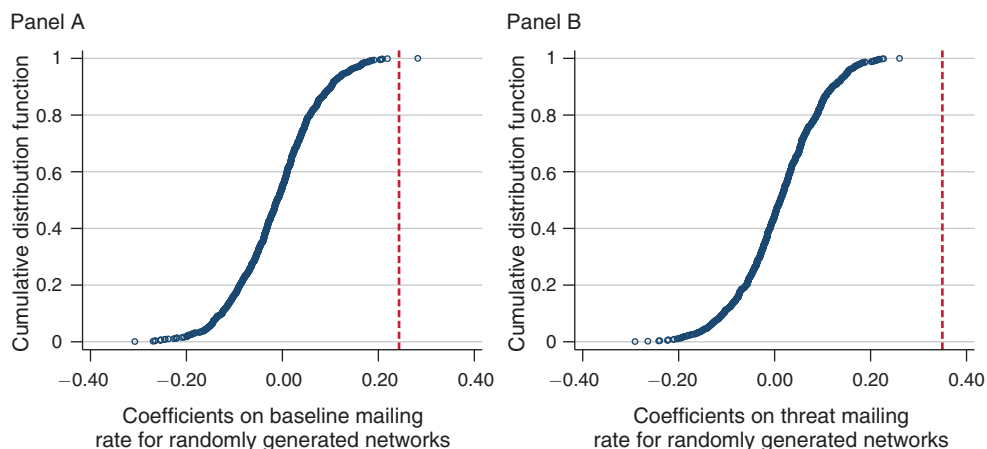


FIGURE 4. PERMUTATION TESTS: DISTRIBUTION OF ESTIMATED COEFFICIENTS

Notes: The figures present the cumulative distribution function of the coefficients on the baseline (panel A) and threat mailing rates (panel B), obtained from the permutation test described in Section IVD. The red vertical lines represent the coefficients on the baseline and the threat mailing rates obtained from estimating model (1).

distance seem to pick up a systematic spillover effect that is shaped by communication within these networks. Section V will study the microstructure of the spillovers in more detail.

E. Discussion of Channels

As pointed out in Section IB, neither receiving a mailing nor compliance is observable to one's neighbors. The compliance spillovers must therefore stem from the dispersion of information via communication (as, e.g., in the context of job referrals, see Beaman and Magruder 2012, Dustmann et al. 2016). This raises the question of which type of information makes households update their compliance propensity and ultimately comply. As in other studies on information diffusion, pinning down the precise piece of information that triggers behavioral responses is complex and not the focus of this paper. Nevertheless, our data offer several insights on the micromechanisms related to different dimensions of information.

In one broad class of mechanisms, the crucial information content is compliance behavior. After receiving a mailing, a type I household might communicate a switch to compliance. Communication could then reinforce, among other things, conformity pressure or social norms about compliance (Akerlof 1980, Bernheim 1994), implying a potential spillover on the compliance of type II households, which later communicate their compliance to their FONs, etc. Several pieces of evidence suggest that the spillovers do *not* (solely) rest on these mechanisms but rather on the diffusion of information that increases the perceived risk of detection.

If treatment spillovers were contingent on the compliance responses of treated households, the indirect effects should be closely aligned with the direct treatment effects. Our results reject this case. Table 1 shows that the threat produces roughly

12 percent more direct registrations than the baseline mailing. By contrast, the estimates from Table 3 indicate that the indirect effect from the threat is 45 percent larger than the one from the baseline mailings ($\beta_2/\beta_1 = 0.35/0.24$). The latter finding points to the role of more “traditional” enforcement channels: communication about the mailings may change the perceived sanction risk. After all, learning about an enforcement activity that was targeted at neighbor j might alter the subjective detection risk of an untreated cheater (see, e.g., Sah 1991). This risk channel appears particularly relevant for threat mailings, which announce possible door-to-door inspections. Given the high spatial correlation in the inspections of FIS’s enforcement division (Rincke and Traxler 2011), it seems plausible that type II households would update their risk perceptions. This case is further supported by the evidence from Fellner, Sausgruber, and Traxler (2013), who document that the larger direct treatment effect of the threat relative to a baseline is driven by a higher perceived sanction risk. Additional support comes from our online survey, which shows that people are willing “to warn” their neighbors (see Section II).

To further assess the role of direct treatment responses for the emergence of spillovers, we make use of the fact that many type I households were actually complying with license fees at the time of the experiment (see Section IIIA). Hence, these treated households could, by definition, not switch from evasion to compliance. If behavioral changes were necessary to induce an indirect effect, the mailings sent to compliant households should not produce any spillover. To test this hypothesis, we rerun our basic regression model on the sample of networks where *all* mailing targets were already complying with license fees before the treatment. The results from this exercise (which are reported in the first column of Table B.5 in online Appendix B) show that we do observe spillovers in networks where the compliance rate of type I households is 100 percent. We also find, though, that the spillovers, in particular those from the baseline mailing, decrease with the ex ante compliance rate among the injection points.²² Changes in compliance behavior thus seem to contribute to larger indirect treatment effects; however, they are not necessary to induce the spillovers. Obviously, this empirical design does not allow rejecting the case in which the shared information is not a *change* in compliance but rather the (previously unknown) compliance *status*. Even if all treated households were compliant, communicating their status might be sufficient to trigger spillovers among neighbors. Conceptually, however, this communication content would represent an informational rather than a pure behavioral imitation channel.

Summing up, our results indicate that while behavioral interdependencies do play some role, informational channels clearly contribute to the spillover. Moreover, the evidence is unambiguous in that any channel must operate via communication.

²²This mirrors a result from Banerjee et al. (2013), who find that adopters of new technologies are crucial in the diffusion of the technology. Note, however, that variation in the ex ante compliance rate might be correlated with relevant unobserved factors.

V. Communication and Diffusion in Networks

This section explores the role of the geographic network in diffusing information. We ask how far the spillover travels, i.e., whether the information of the mailing reaches only first-order neighbors or whether it diffuses farther (Section VA), which measures (if any) of injection points' centrality can be used for targeting (Section VB), and whether it is more effective to locally concentrate mailings or to spread treatments broadly within a network (Section VC). While all questions are informative about the role of the network in shaping the diffusion process, the last two questions are particularly important from a policy perspective.

A. How Local Are Spillovers?

Let us first study how far the spillover reaches into the network. If new information does not travel very far in the network, then it can only affect (postcommunication) propensities to comply of households located “near” the injection points in terms of network distance.²³ In this case, spillovers would be limited to households that are “close” to a treated household (compare the conceptual framework from Section IVA).

Figure 5, which presents a hypothetical network, illustrates a case of local spillovers. In this network, only household k receives a mailing. Since only k can start to spread the word, it will always first reach k 's FONs (panel B in Figure 5). If the FONs do *not* pass on the information, then any spillover will be limited to the treated household's FONs. If they do pass it on, then the news will reach k 's SONs (panel C in Figure 5). Roughly speaking, if a treatment induces more widespread communication, then more distant households will be reached via the spillover.

To analyze how far the spillovers reach, we compute treatment rates that distinguish between the treatments of a household's FONs, SONs, or higher-order neighbors (HONs). More specifically, we count the number of i 's FONs, SONs, and HONs who received a given treatment. Normalizing these numbers by $N_k - 1$, we obtain the treatment rates $Base_{ik}^h$ and $Threat_{ik}^h$ for $h \in \{FON, SON, HON\}$. This approach ensures the identities $Base_k = \sum_h Base_{ik}^h$ and $Threat_k = \sum_h Threat_{ik}^h$. We then estimate a refined version of equation (1), which now exploits variation between and within networks:

$$(5) \quad y_{ik} = \delta^{Total_k} + \sum_h \beta_1^h Base_{ik}^h + \sum_h \beta_2^h Threat_{ik}^h + \epsilon_{ik}.$$

The more locally confined diffusion is, the more sharply β_1^h and β_2^h should decline with the order of vicinity of the treated neighbors.

Note that the estimated coefficients β_1^h and β_2^h are identified from the experimental variation in treatment rates for a given level of $Total_k$. Conditional on $Total_k$, we argue that whether and how a type II's close (FON) or more distant neighbor (SON or HON) is treated varies exogenously. However, one might be concerned that the

²³ Network distance, often also referred to as geodesic distance, refers to the length of the shortest path between two agents (nodes).




FIGURE 5. ILLUSTRATION: HOW LOCAL ARE SPILLOVERS?

Notes: The figure shows targeted household k with its first-, second-, third-, and fourth-order neighbors (indicated by numbers 1, 2, 3, and 4) in a hypothetical network. Panel A represents the time before communication starts where only k knows about its letter, in (panel B) the news has spread to k 's FONS, in (panel C) to her SONs, etc. Note that each time the message spreads its impact becomes smaller: in (panel C) for instance, the SON of k learns about the letters but also about the fact that k 's FONS—with whom she communicates—did not receive a letter.

location of experimental households (relative to type II households) is not random within networks. Accounting for this point, we also consider an augmented specification which estimates the effects conditional on $Total_k^{FON}$, $Total_k^{SON}$, and $Total_k^{HON}$:

$$(6) \quad y_{ik} = \sum_h \alpha^h Total_k^h + \sum_h \beta_1^h Base_{ik}^h + \sum_h \beta_2^h Threat_{ik}^h + \epsilon_{ik}.$$

The identification of the β coefficients in this model rests again on the random variation in treatment rates among FONS, SONs, and HONS—but now we condition on the rate of experimental households, $Total_k^h$ for each $h \in \{FON, SON, HON\}$.²⁴ Note that this augmented specification will absorb much variation in the data.

Linear probability model estimates of equations (5) and (6) are presented in columns 1 and 3 of Table 5, respectively. The results show consistently that both mailing treatments trigger sizable spillovers on the treated households' FONS. For

²⁴We linearly control for $Total_k^h$ as there would be a too-large number of fixed effects for each rate $Total_k^h$ in $h \in \{FON, SON, HON\}$.

TABLE 5—IMPACT BY NETWORK DISTANCE: HOW FAR-REACHING ARE THE SPILLOVERS?

	(1)	(2)	(3)	(4)
$Base^{FON}$	0.2568 (0.0387)	0.2540 (0.0390)	0.2541 (0.0399)	0.2505 (0.0400)
$Base^{SON}$	0.0922 (0.0686)	0.0908 (0.0690)	0.1426 (0.1302)	0.1393 (0.1303)
$Base^{HON}$	0.0450 (0.0878)	0.0378 (0.0886)	0.2995 (0.1922)	0.2946 (0.1927)
$Threat^{FON}$	0.3615 (0.0397)	0.3648 (0.0400)	0.3580 (0.0407)	0.3609 (0.0409)
$Threat^{SON}$	0.1396 (0.0681)	0.1473 (0.0686)	0.1550 (0.1247)	0.1606 (0.1307)
$Threat^{HON}$	0.3884 (0.0948)	0.3998 (0.0955)	0.5889 (0.2044)	0.6073 (0.2107)
$Total^{FON}$			-0.0391 (0.0341)	-0.0397 (0.0341)
$Total^{SON}$			-0.0568 (0.1162)	-0.0590 (0.1162)
$Total^{HON}$			-0.3222 (0.1792)	-0.3304 (0.1796)
Distance to nearest base	-	-0.0000 (0.0000)	-	-0.0000 (0.0000)
Distance to nearest threat	-	0.0000 (0.0000)	-	0.0000 (0.0000)
<i>F</i> -tests: (see table notes)				
Base:				
FON versus SON	6.65	6.51	4.65	4.73
FON versus HON	6.60	6.70	14.57	15.33
SON versus HON	0.26	0.32	1.82	2.08
Threat:				
FON versus SON	12.16	11.67	12.48	12.34
FON versus HON	0.09	0.15	0.54	0.38
SON versus HON	5.76	5.92	3.42	3.69
R^2	0.0967	0.0974	0.0605	0.0610

Notes: FON, SON, and HON abbreviates First-, Second-, and Higher-Order Neighbors, respectively. Column 1 and 3 present results from LPM estimations of equations (5) and (6), respectively. Columns 2 and 4 add controls for the Euclidean distance to the nearest household treated with a baseline or a threat mailing, respectively. The *F*-tests from columns 1 and 2 are based on equation (5) and report *F*-statistics for the $H_0: \beta_j^h = \beta_j^\ell$ for $h \neq \ell \in \{FON, SON, HON\}$ and $j = 1$ (Base) and $j = 2$ (Threat), respectively. Hence, we test $Base^{FON} = Base^{SON}$, etc. Columns 3 and 4 are based on equation (6) and test the $H_0: \alpha^h + \beta_j^h = \alpha^\ell + \beta_j^\ell$ for $h \neq \ell$ and $j = 1$ (Base) and $j = 2$ (Threat), respectively. Put differently, we are testing $Total^{FON} + Base^{FON} = Total^{SON} + Base^{SON}$, etc. Adding up coefficients in columns 3 and 4 accounts for the fact that $Total^h$ differs across h . Number of observations: 14,987; number of networks: 3,764. Standard errors, clustered at the network level, are in parentheses.

the baseline treatment, the spillover seems mainly limited to FONs. To compare the size of coefficients in column 3 across FON, SON, and HON neighborhoods, we need to sum the coefficients for $Total_k^h$ and $Base_k^h$ for each $h \in \{FON, SON, HON\}$ as now $Total_k^h$ also differs across h . Doing so, we see that the estimates indicate spillovers from sending an additional mailing into each of these neighborhoods of 0.215 for FON ($= Total_k^{FON} + Base_k^{FON}$) as opposed to 0.086 for SON and -0.023

for HON, respectively. (Except for the comparison of SON with HON, all spillovers are significantly different from each other; see the F -statistics reported in Table 5.)

For the threat treatment, the spillovers are not limited to FONs or SONs but also extend to HONs. The estimated β_2^h coefficients from column 2 are statistically significant at the 1 percent—or 5 percent—level for all types of neighborhoods (FON, SON, HON), and the F -tests indicate that we cannot reject the hypothesis that spillovers are equally large for FONs as they are for HONs. All these findings are robust when we control for a household's Euclidean distances to the nearest household in a given treatment (columns 2 and 4).

The evidence from Table 5 indicates that the geographic network does shape diffusion. The spillover is more concentrated among households with a smaller network distance to the targeted nodes. At the same time, the spillover is not only confined to FONs: for the threat treatment, the indirect effect clearly has a larger scope.

The result also offers insights about two basic reasons why the threat treatment triggers a larger spillover. On the one hand, the threat might induce more communication. On the other hand, the treatment could—conditional on the level of communication—simply increase the propensities to comply by a larger amount. (Note that our basic estimates did not allow us to distinguish between these two interpretations; see Section IVE.) If the difference in spillovers from the base and the threat mailing was entirely due to the threat's stronger impact on compliance propensities, then we should see smaller spillovers for the baseline treatment—but those smaller spillovers should be equally distributed among FONs, SONs, and HONs, the groups reached by the threat. The fact that the spillover generated from the base mailing seems to be largely concentrated among FONs suggests, though, that the threat induces indeed more communication.

This finding also has implications for which households to optimally target. If the spillovers were limited to FONs, the indirect effect of the intervention would be maximized by simply targeting those households with the highest number of FONs (households with the highest degree). If spillovers travel beyond FONs, as is the case for the threat mailing, then other centrality measures may be more relevant for “optimal” targeting. The next subsection explores this point in more detail.

B. Which Households to Target?

To study optimal targeting, we focus on centrality measures that have been shown to play an important role for diffusion. The first such measure is diffusion centrality, which captures a household's ability to spread information in a finite amount of time (Banerjee et al. 2019). A special case of diffusion centrality is degree, which simply counts how many network neighbors (FONs) a household has. The second measure is eigenvector centrality, which is focused on infinite (or very long lasting) communication (DeGroot 1974; DeMarzo, Vayanos, and Zwiebel 2003; Golub and Jackson 2010).²⁵

²⁵These measures are more closely discussed in online Appendix C.




FIGURE 6. ILLUSTRATION: DEGREE, DIFFUSION CENTRALITY, AND EIGENVECTOR CENTRALITY

Notes: The figure shows a stylized representation of a typical network in our sample. Four nodes (indicated with a star) have the highest degree: they all have three network neighbors. The black node uniquely has the highest diffusion centrality after two rounds of communication ($DC^{T=2}$). By targeting the black node the entire network can be “reached” after two rounds of communication (assuming that the probability to passing on information is equal to 1). This is not the case if the patterned nodes are targeted. The latter nodes have the highest eigenvector centrality (EC), i.e., those are linked to nodes who are linked to “more important” nodes than the black node, thus contributing to their higher eigenvector centrality.

Figure 6, which shows a typical network from our sample, illustrates how different households would serve as optimal injection points, depending on which centrality measure is key (see figure notes). While these centrality measures tend to be strongly correlated, the variation among the large number of networks in our sample allows us to differentially assess their role for mediating the spillovers. We do so by exploring variation in the injection points’ (type I households’) centrality.

First, we compute the median centrality of all type I households in our relevant networks for each centrality measure c introduced below. Based on this median, we then calculate new treatment rates, $Base_k^{c-H}$ and $Base_k^{c-L}$, which measure a network k ’s fraction of type I households in the baseline mailing treatment with an above- or below-median level of centrality, respectively. The rates for the threat treatment, $Threat_k^{c-H}$ and $Threat_k^{c-L}$, are defined accordingly. We then estimate the following model:

$$(7) \quad y_{ik} = \delta^{Total_k} + \beta_1^H Base_k^{c-H} + \beta_1^L Base_k^{c-L} + \beta_2^H Threat_k^{c-H} + \beta_2^L Threat_k^{c-L} + \epsilon_{ik}.$$

Treatment rates are again defined relative to $N_k - 1$. This ensures the identity $Base_k^{c-H} + Base_k^{c-L} = Base_k$, in which $Base_k$ is the baseline rate from our basic model (1). The analogue holds for $Threat_k$. Equation (7) is thus nested in (1).

Intuitively, model (7) distinguishes the effect from treating nodes with relatively high or low centrality (as captured by β_1^H versus β_1^L , and β_2^H versus β_2^L for the baseline and threat treatment, respectively). For a given level of $Total_k$, the model exploits experimental variation in the treatment of more or less central nodes. Still, one might argue that the treated households’ centrality is correlated with unobserved network properties. In particular, having experimental households with a relatively high or low centrality might correlate with the nonexperimental households’ propensity to

comply. Hence, we also consider an augmented specification that responds to this concern.

Following the same strategy as in Section VA, we replace the $Total_k$ fixed effects from (7) and instead condition on $Total_k^{c-H}$ and $Total_k^{c-L}$, the rate of all experimental households in a given network k that have an above- or below-median level of centrality.²⁶ We then estimate

$$(8) \quad y_{ik} = \alpha^H Total_k^{c-H} + \alpha^L Total_k^{c-L} + \beta_1^H Base_k^{c-H} + \beta_1^L Base_k^{c-L} + \beta_2^H Threat_k^{c-H} + \beta_2^L Threat_k^{c-L} + \epsilon_{ik}.$$

Controlling for both $Total_k^{c-H}$ and $Total_k^{c-L}$ ensures that the β coefficients from this model are identified from the random treatment assignment, even if the potential injection points' centrality in network k was indeed correlated with unobservables.

We estimate (7) and (8) for three different measures of diffusion centrality (DC) as well as eigenvector centrality (EC). Note that DC is a function of T , the time horizon of communication (see equation (1) in the online Appendix C). Our main analysis uses DC for $T = 1$ and $T = 10$ rounds of communication, $DC^{T=1}$ and $DC^{T=10}$. For both measures, the second parameter of DC , the probability to pass on information (see Section IVA), is fixed at $q = 1$.²⁷ In addition, we also examine the DC benchmark proposed in Banerjee et al. (2019), with T equal to the network's diameter and the probability q equal to the inverse of the first eigenvalue of the network's adjacency matrix. This measure, DC_q^T , is thus based on network specific parameters q_k and T_k . Finally, note that all four centrality variables, $c = \{DC^{T=1}, DC^{T=10}, DC_q^T, EC\}$, are normalized.²⁸

Linear probability models estimates of equations (7) and (8) for each of the four centrality measures c are presented in Table 6. The results for the diffusion centrality measures are straightforward. For almost all combinations of treatments and DC measures, the point estimates for β^H are higher than those for β^L . The point estimates from column 1, for instance, suggest that the spillovers from the baseline mailing might double if one targets injection points with relatively high $DC^{T=1}$. At the same time, sending baseline mailings to nodes with high $DC^{T=10}$ does not necessarily increase spillovers (the difference between β_j^H and β_j^L is not statistically significant in column 3). This pattern also indicates that the diffusion process driving the spillover from the baseline mailing is limited—which is consistent with the results described in Section VA. (Recall that the indirect effects of the baseline mailing are mainly confined to first-order neighbors.)

²⁶Note that the two rates obviously add up to $Total_k$ from equation (1). Equation (8) is thus nested in specification 2 from Table 3. Due to a too-large number of fixed effects for each level of $Total_k^{c-H}$ and $Total_k^{c-L}$, we linearly control for the two rates (as in model (6) above).

²⁷In online Appendix B, we provide results for alternative values of $0 < q < 1$ (see online Appendix Table B.8). As mentioned above, the concept of degree centrality (the number of first-order neighbors) is nested in diffusion centrality. Specifically, $DC^{T=1}$ (with $q = 1$) corresponds to the degree (up to normalization).

²⁸For the diffusion centrality, we use the normalization suggested in Banerjee et al. (2019), i.e., we compare $DC^T(\ell)$ to the hypothetical case of complete networks (for which one obtains the maximum possible entry for $DC^T(\ell)$ for each node ℓ and any given T).

TABLE 6—IMPACT BY CENTRALITY OF INJECTION POINTS

Centrality measure c :	$DC^{T=1}$		$DC^{T=10}$		DC_q^T		EC	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Base_k^{c-L}$	0.1233 (0.0547)	0.1777 (0.1160)	0.2150 (0.0646)	0.1937 (0.1491)	0.1763 (0.0662)	0.2738 (0.1453)	0.1759 (0.1188)	0.4187 (0.2326)
$Base_k^{c-H}$	0.2571 (0.0390)	0.2601 (0.0401)	0.2459 (0.0386)	0.2552 (0.0394)	0.2474 (0.0385)	0.2486 (0.0392)	0.2441 (0.0382)	0.2444 (0.0382)
$Threat_k^{c-L}$	0.2620 (0.0607)	0.2813 (0.1177)	0.2782 (0.0691)	0.2261 (0.1482)	0.2985 (0.0729)	0.3421 (0.1460)	0.4415 (0.1183)	0.6225 (0.2403)
$Threat_k^{c-H}$	0.3598 (0.0397)	0.3632 (0.0408)	0.3549 (0.0394)	0.3651 (0.0402)	0.3533 (0.0394)	0.3554 (0.0400)	0.3479 (0.0391)	0.3468 (0.0390)
$Total_k^{c-L}$		-0.0970 (0.1055)		-0.0604 (0.1346)		-0.1552 (0.1313)		-0.3774 (0.2130)
$Total_k^{c-H}$		-0.0463 (0.0342)		-0.0520 (0.0337)		-0.0444 (0.0335)		-0.0378 (0.0327)
$Total_k$ fixed effects	Yes	No	Yes	No	Yes	No	Yes	No
F -tests: (see table notes)								
Base:	7.662	10.92	0.279	2.255	1.397	3.184	0.346	5.152
High versus low								
Threat:	3.188	9.051	1.461	8.871	0.662	5.747	0.668	0.596
High versus low								
R^2	0.096	0.059	0.095	0.058	0.095	0.058	0.095	0.058

Notes: The table presents results from LPM estimations of equations (7) and (8). The estimates explore how the impact of the mailings on the compliance of non-experimental households varies with the centrality of the injection points. We differentiate treatment rates according to the injection points' having above or below media measures of centrality c for: diffusion centrality for one ($DC^{T=1}$, columns 1 and 2) and 10 rounds of communication ($DC^{T=10}$, columns 3 and 4), diffusion centrality with T equal to the network's diameter and probability q equal to the inverse of the first eigenvalue of the network's adjacency matrix (DC_q^T , columns 5 and 6), and finally eigenvector centrality (EC , columns 7 and 8). The F -tests reported in columns 1, 3, 5, and 7 are based on equation (7) and test the $H_0: \beta_j^L = \beta_j^H$ for $j = 1$ (Base) and $j = 2$ (Threat), respectively. We thus test $Base_k^{c-H} = Base_k^{c-L}$ for the different centrality measures c (and analogously for the threat treatment rates). The tests from columns 2, 4, 6, and 8 are based on the augmented equation (8) and test $H_0: \alpha^L + \beta_j^L = \alpha^H + \beta_j^H$ for $j = 1$ (Base) and $j = 2$ (Threat), respectively. Hence, we test $Total_k^{c-L} + Base_k^{c-L} = Total_k^{c-H} + Base_k^{c-H}$ for the different centrality measures c (and analogously for the threat treatment rates). Adding up coefficients here accounts for the fact that $Total_k^{c-L}$ differs from $Total_k^{c-H}$ (compare Table 5). Number of observations: 14,987; number of networks: 3,764. Standard errors, clustered at the network level, are in parentheses.

For the threat mailing, both $DC^{T=1}$ and $DC^{T=10}$ seem to matter, suggesting that treatment-induced diffusion is less limited and happens over multiple rounds of communication. Again, this is consistent with the broader reach of the spillover documented in section VA. The results for the network-specific DC_q^T are in between those for $DC^{T=1}$ and $DC^{T=10}$.²⁹ Finally, for eigenvector centrality, there is no clear pattern indicating that a higher EC of the injection points would increase the spillover. Note that EC would be the decisive variable for targeting, if the diffusion process would be unlimited ($T \rightarrow \infty$). The results from columns 7 and 8 do not support this case.

The estimates from Table 6 are further corroborated by models in which the treatment rates linearly interact with the treated households' centrality c (see Table B.6 in online Appendix B). Consistent with the evidence from above, the estimates reveal

²⁹The measure DC_q^T can be considered as an intermediate case as the (network-specific) diameter, which determines T , is on average roughly 4.

significant interactions of the baseline treatment with $DC^{T=1}$; for the threat, instead, one obtains significant interactions with $DC^{T=10}$. In additional robustness exercises, we ran specifications that further augment equations (7) and (8) by including controls for the injection points' centrality. The estimates (which are reported in Table B.7) are almost indistinguishable from those reported in Table 6. Finally, we also replicated results for different DC measures with $T \in \{1, \dots, 10\}$ and $q < 1$. The results show again qualitatively similar patterns (see Table B.8).

To wrap up, the results from this and the previous section consistently indicate that the diffusion process driving the spillovers is (i) strongly limited for the baseline mailing but (ii) less limited for the threat mailing. For the baseline treatment, mostly close neighbors of targeted nodes respond to the intervention. The threat treatment, by contrast, induces compliance spillovers also on households with higher network distance. In terms of which households an authority should target to maximize spillovers, our results indicate that diffusion centrality serves as a good indicator. To operationalize such a strategy, however, an authority would have to determine a value of T . Which value to use will depend on the specific communication or enforcement strategy. For a rather innocent intervention that is unlikely to lead to much communication (as our baseline mailing), a low value of T might be best. If the intervention involves a message that is likely to induce much more communication (as the threat mailing), a higher value of T seems more reasonable.

C. Local Treatment Concentration

Another important policy question concerns the effect of treatment concentration: is it more effective to locally concentrate mailings or to spread treatments broadly within a network? The problem is illustrated in Figure 7, which displays a case with high local treatment concentration, i.e., where treated households are FONs (panel A), and another example where concentration is low (panel B). Where should we expect larger spillovers? In the context of limited diffusion, both targeting strategies could be reasonable. On the one hand, a low concentration *ceteris paribus* means that more households will hear about the treatment. Each of these households, however, is likely to hear about only *one* mailing and to talk to many other untreated households. Hence, while many households will hear about a mailing, their propensities to comply might only increase marginally. High concentration resolves the latter issue at the cost of reaching fewer households. High concentration should hence be particularly desirable if households need to hear sufficiently often about a mailing in order to change behavior, as is the case in, e.g., threshold models (Granovetter 1978, Centola and Macy 2007, Beaman et al. 2015).

To measure the extent to which mailings are locally concentrated, we define a measure called local treatment concentration, which measures the extent to which households in the same treatment are directly linked to each other in the network. This concept is reminiscent of the concept of homophily that has received a lot of attention in recent years across a variety of fields, including economics (Benhabib, Bisin, and Jackson 2010), sociology (McPherson, Smith-Lovin, and Cook 2001), and management (Borgatti and Foster 2003). There are two key differences, however, between our local concentration measure and homophily. First, the notion

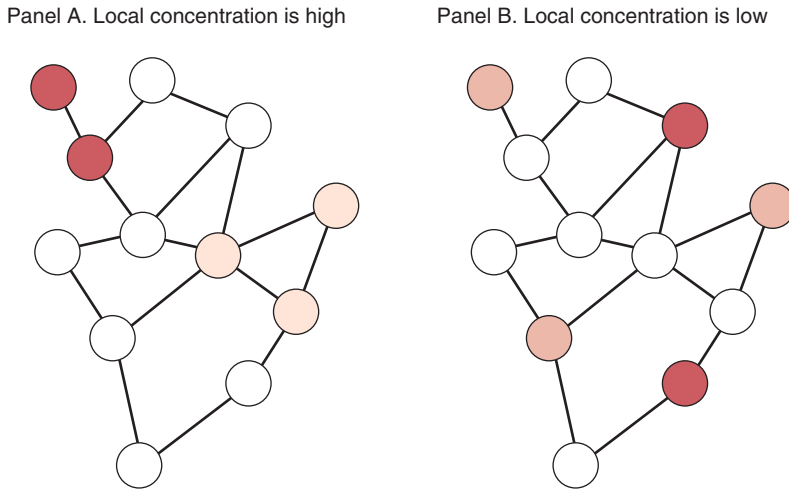


FIGURE 7. ILLUSTRATION: LOCAL TREATMENT CONCENTRATION

Notes: The figure shows two networks with the same structure but different degrees of local treatment concentration. There are 12 households of three types represented in this picture. Let the dark nodes represent threat-treated households, the lighter nodes base-treated and the empty nodes untreated households (type II or III). In panel A, households of the same type tend to be neighbors, while households in panel B tend to be neighbors with households of a different type. Based on the formula from the online Appendix C, it is straightforward to compute the inbreeding homophily (IH) index for the different cases. In panel A, the IH index for the threat treatment is given by $IH^{Threat} = \left(\frac{2}{4} - \frac{11}{12} \frac{1}{6}\right) / \left(1 - \frac{11}{12} \frac{1}{6}\right) \approx 0.41$, reflecting the fact that (i) there are two neighbors of a threat-treated node who are also threat-treated, (ii) threat-treated nodes have four neighbors overall, (iii) there is a total of twelve nodes in the network, and (iv) the threat-rate is $1/6$. For the base treatment, one analogously obtains $IH^{Base} = \left(\frac{6}{10} - \frac{11}{12} \frac{1}{4}\right) / \left(1 - \frac{11}{12} \frac{1}{4}\right) \approx 0.48$. In panel B, there are no neighbors of threat- (base-) treated households who are threat- (base-) treated. For this network, one gets $IH^{Threat} \approx -0.18$ and $IH^{Base} \approx -0.30$.

of local treatment concentration considers two neighbors as “similar” if they are in the same treatment, whereas homophily usually defines similarity in terms of race, age, or other characteristics. Second, the random assignment of treatments assures that, conditional on being experimental participants, local concentration will not reflect other household properties. In particular, whether or not experimental households are linked is exogenous to the assigned treatment condition. By contrast, with homophily whether two households are linked can be a *result* of the dimension of similarity considered.³⁰ This is why homophily can be thought of as a *preference* for being linked with similar others.³¹ This interpretation does not apply to our concept of local treatment concentration. However, both the endogenous concept of homophily and the exogenous concept of local treatment concentration measure to which extent neighbors are similar to each other. Exogeneity in this sense can be seen as a strength compared to previous empirical work, as it allows us to isolate the

³⁰ If, e.g., a single (white, old, etc.) person chooses to move next to other single (white, old, etc.) people.

³¹ Many empirical studies, though, consider homophily as the mere statistical fact that similar people are more likely to be linked to each other in the network (McPherson, Smith-Lovin, and Cook 2001; Currarini, Jackson, and Pin 2009).

effects of similarity (in treatment and posttreatment beliefs) on behavioral outcomes within networks.

We use the inbreeding homophily (IH) index to measure local concentration (Coleman 1958; Currarini, Jackson, and Pin 2009, 2010).³² The IH index is defined at the network level for each treatment group. It ranges in $[-1, 1]$. A value of zero indicates that on average within a given network the share of a household's neighbors in the same treatment group as the household itself corresponds to what we would expect from random assignment. A positive value indicates that there is homophily (high local concentration), i.e., households are linked overproportionately with households from the same treatment. A negative value indicates heterophily (low local concentration), i.e., households are linked overproportionately with households that are in a different treatment group.

Figure 7 shows a case of high (panel A) and low local treatment concentration (panel B). As detailed in the caption, the IH indices would be positive for the one (left) and negative in the other (right) case (see the caption to Figure 7 for computational details). To account for a potential downside of the IH index (it is slightly biased downward in small networks; see online Appendix C), we also consider a dummy indicating whether the IH index is positive. In 17 percent of networks, we observe at least 1 pair of FONs who receive the same mailing treatment. There is a lot of variation in the IH index, with a majority of networks having a negative inbreeding homophily for the two mailing treatments but also a substantial share of networks with a positive IH index.

We estimate the model

$$(9) \quad y_{ik} = \delta^{Total_k} + \beta_1 Base_k + \beta_2 Threat_k + \delta_1 IH_k^{Base} + \delta_2 IH_k^{Threat} \\ + \delta_3 (IH_k^{Base} \times Base_k) + \delta_4 (IH_k^{Threat} \times Threat_k) + \epsilon_{ik},$$

where IH_k^{Base} and IH_k^{Threat} capture one of the two local concentration measures (the IH index or the IH dummy). Conditional on $Total_k$ fixed effects, random treatment assignment implies that the variation in the treatment-specific concentration measures will be exogenous. In addition, we also consider an augmented specification that absorbs additional variation by including an IH measure for the local concentration of *all* possible injection points (IH_k^{All}) irrespective of their treatment. Estimation results for the simple and the augmented version of equation (9) are presented in Table 7.

The estimates suggest that local treatment concentration is associated with *smaller* compliance spillovers. While the interaction effects (δ_3 and δ_4) are statistically insignificant for the IH-index (columns 1 and 2), they turn significant for the IH dummy (columns 3 and 4). Quantitatively, these interactions are quite sizable: the point estimates from the last two specifications indicate that the spillovers shrink by roughly 50 percent in networks with a positive IH index. To assess the robustness

³²Inbreeding homophily of household i is defined as the difference between the fraction of links of household i to others of the same treatment type and the fraction of i 's type in the network divided by the maximal amount of homophily possible ($1 -$ the fraction of i 's type in the network). For a closer discussion, see online Appendix C.

TABLE 7—IMPACT BY LOCAL TREATMENT CONCENTRATION

	IH index		IH ⁺ dummy	
	(1)	(2)	(3)	(4)
$Base_k$	0.2053 (0.0582)	0.2083 (0.0628)	0.2755 (0.0428)	0.2748 (0.0429)
$Threat_k$	0.2678 (0.0579)	0.2708 (0.0649)	0.3983 (0.0435)	0.3980 (0.0435)
$Base_k \times IHI_k^{Base}$	-0.0359 (0.0806)	-0.0334 (0.0823)		
$Threat_k \times IHI_k^{Threat}$	-0.1018 (0.0812)	-0.0993 (0.0830)		
IHI_k^{Base}	-0.0121 (0.0639)	-0.0106 (0.0658)		
IHI_k^{Threat}	-0.0076 (0.0574)	-0.0061 (0.0603)		
IHI_k^{All}		-0.0036 (0.0330)		
$Base_k \times IHD_k^{Base}$			-0.1527 (0.0745)	-0.1393 (0.0781)
$Threat_k \times IHD_k^{Threat}$			-0.1861 (0.0742)	-0.1758 (0.0775)
IHD_k^{Base}			0.0176 (0.0097)	0.0173 (0.0097)
IHD_k^{Threat}			0.0277 (0.0093)	0.0276 (0.0093)
IHD_k^{All}				-0.0078 (0.0116)
R^2	0.0951	0.0951	0.0958	0.0959

Notes: The table presents results from LPM estimations of equation (9). In columns 1 and 2, local treatment concentration is measured by the Inbreeding Homophily Indexes (IHI). In columns 3 and 4, we use a dummy indicating a positive inbreeding homophily ($IHD = 1$ if $IHI \geq 0$). Specifications used in columns 2 and 4 control for the local concentration of all experimental (type 1) households in a network k . All estimates are based on 14,987 observations from 3,764 networks. Standard errors, clustered at the network level, are in parentheses.

of these results, we replicated the estimations using IH measures that pool the two mailing treatments (see Table B.9 in online Appendix B). The exercise confirms the negative interaction effects and improves statistical power: with a higher treatment concentration, we find significantly smaller spillovers. The evidence therefore suggests that spreading the interventions more broadly—and thus assuring that more people “hear” about the mailings—produces larger spillovers.

Our findings are consistent with evidence highlighting the role of one’s peers’ exposure to law enforcement for the formation of one’s own risk perceptions and legal compliance (see Lochner 2007, and Section I.B in Chalfin and McCrary 2017). The fact that a high local treatment concentration reduces spillovers is also in line with theoretical results showing that similarity in beliefs hampers social learning (Golub and Jackson 2012). Facing the same treatment most likely implies similar posttreatment beliefs (e.g., about the probability of detection). One can thus interpret our findings as being consistent with the theoretical results.

VI. Conclusions

This paper studied the spread of compliance with TV license fees among neighbors in geographic networks in Austria. We exploited exogenous variation from a field experiment which randomly assigned potential license fee evaders into a control or different treatment conditions (baseline or threat mailing). Building on this variation, we tested whether untreated households, which were not covered by the experiment, respond to mailings received by their network neighbors. Our analysis identified strong and precisely estimated spillover effects. Controlling for the total number of households in a network involved in the experiment, an increase in the fraction of treated neighbors significantly raises the likelihood that an untreated cheater switches from evasion to compliance. Overall, the magnitude of the indirect treatment effects is similar to the direct treatment impact, indicating a social multiplier of roughly two (Glaeser, Sacerdote, and Scheinkman 2003). We also document heterogeneity across treatments: the threat mailings trigger a spillover that is roughly 40 percent larger than the spillover from the baseline mailings. As we focus on small municipalities, it is hard to assess if our results apply to urban areas, too. On the one hand, one might expect smaller spillovers as neighborhood networks are less relevant in larger towns and cities. On the other hand, Fellner, Sausgruber, and Traxler (2013) document slightly larger direct treatment effects in urban areas (in particular, for the threat mailing). This, in turn, could also boost spillovers, which might occur among friends or peers at work rather than spatial neighbors.

The paper further studied how information on license fee enforcement diffused among neighbors in geographic networks. By exploiting variation within and between over 3,750 networks, we derived the following results: First, relative to the threat treatment, the spillover from the baseline is more concentrated among households that are—in terms of network distance—closer to the households targeted by the experiment (the “injection points”). Second, the spillovers from the mailing treatments are increasing with the injection points’ diffusion centrality. Third, a higher local treatment concentration results in a significantly smaller spillover—a finding which can be interpreted in terms of homophily (Golub and Jackson 2012). Together with several null results from various placebo exercises, these findings indicate that the structure of our geographic networks is useful to capture patterns of information diffusion among neighbors.

The results carry important implications for the optimal targeting of policies that play on word-of-mouth diffusion among neighbors. Our findings suggest that one could, *ceteris paribus*, greatly improve the intervention’s impact by targeting central households and by spreading the treatment broadly within a network. The easy availability of geographic information facilitates the computation of neighborhood networks and allows to implement these strategies at fairly low costs.

REFERENCES

- Akerlof, George A. 1980. “A Theory of Social Custom, of Which Unemployment May Be One Consequence.” *Quarterly Journal of Economics* 94 (4): 749–75.

- Alatas, Vivi, Abhijit Banerjee, Arun G. Chandrasekhar, Rema Hanna, and Benjamin A. Olken. 2016. "Network Structure and the Aggregation of Information: Theory and Evidence from Indonesia." *American Economic Review* 106 (7): 1663–1704.
- Allcott, Hunt. 2011. "Social Norms and Energy Conservation." *Journal of Public Economics* 95 (9–10): 1082–95.
- Alstadsæter, Annette, Wojciech Kopczuk, and Kjetil Telle. 2018. "Social Networks and Tax Avoidance: Evidence from a Well-Defined Norwegian Tax Shelter." NBER Working Paper 25191.
- Aral, Sinan, and Dylan Walker. 2011. "Creating Social Contagion through Viral Product Design: A Randomized Trial of Peer Influence in Networks." *Management Science* 57 (9): 1623–39.
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2013. "The Diffusion of Microfinance." *Science* 341 (6144): 1236498.
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2019. "Using Gossips to Spread Information: Theory and Evidence from Two Randomized Controlled Trials." http://stanford.edu/~arungc/BCDJ_gossip.pdf.
- Bayer, Patrick, Stephen L. Ross, and Giorgio Topa. 2008. "Place of Work and Place of Residence: Informal Hiring Networks and Labor Market Outcomes." *Journal of Political Economy* 116 (6): 1150–96.
- Beaman, Lori, Ariel BenYishay, Jeremy Magruder, and Ahmed Mushfiq Mobarak. 2015. "Can Network Theory-Based Targeting Increase Technology Adoption?" https://cdep.sipa.columbia.edu/sites/default/files/cdep/MNW_Columbia_Magruder.pdf.
- Beaman, Lori, and Jeremy Magruder. 2012. "Who Gets the Job Referral? Evidence from a Social Networks Experiment." *American Economic Review* 102 (7): 3574–93.
- Benhabib, Jess, Alberto Bisin, and Matthew O. Jackson, ed. 2010. *The Handbook of Social Economics*. (Volumes 1A and 1B). Amsterdam: North Holland.
- BenYishay, Ariel, and A. Mushfiq Mobarak. 2019. "Social Learning and Incentives for Experimentation and Communication." *Review of Economic Studies* 86 (3): 976–1009.
- Berger, Melissa, Gerlinde Fellner-Röhling, Rupert Sausgruber, and Christian Traxler. 2016. "Higher Taxes, More Evasion? Evidence from Border Differentials in TV License Fees." *Journal of Public Economics* 135: 74–86.
- Bernheim, B. Douglas. 1994. "A Theory of Conformity." *Journal of Political Economy* 102 (5): 841–77.
- Bollinger, Bryan, and Kenneth Gillingham. 2012. "Peer Effects in the Diffusion of Solar Photovoltaic Panels." *Marketing Science* 31 (6): 900–912.
- Boning, William C., John Guyton, Ronald H. Hodge, II, Joel Slemrod, and Ugo Troiano. 2018. "Heard It through the Grapevine: Direct and Network Effects of a Tax Enforcement Field Experiment." NBER Working Paper 24305.
- Borgatti, Stephen P., and Pacey C. Foster. 2003. "The Network Paradigm in Organizational Research: A Review and Typology." *Journal of Management* 29 (6): 991–1013.
- Brollo, Fernanda, Katja Maria Kaufmann, and Eliana La Ferrara. 2017. "Learning about the Enforcement of Conditional Welfare Programs: Evidence from Brazil." IZA Discussion Paper 10654.
- Bruhin, Adrian, Lorenz Götte, Simon Haenni, and Lingqing Jiang. 2014. "Spillovers of Prosocial Motivation: Evidence from an Intervention Study on Blood Donors." IZA Discussion Paper 8738.
- Centola, Damon, and Michael Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American Journal of Sociology* 113 (3): 702–34.
- Chalfin, Aaron, and Justin McCrary. 2017. "Criminal Deterrence: A Review of the Literature." *Journal of Economic Literature* 55 (1): 5–48.
- Chetty, Raj, John N. Friedman, and Emmanuel Saez. 2013. "Using Differences in Knowledge across Neighborhoods to Uncover the Impacts of the EITC on Earnings." *American Economic Review* 103 (7): 2683–2721.
- Coleman, James S. 1958. "Relational Analysis: The Study of Social Organizations with Survey Methods." *Human Organization* 17 (4): 28–36.
- Currarini, Sergio, Matthew O. Jackson, and Paolo Pin. 2009. "An Economic Model of Friendship: Homophily, Minorities, and Segregation." *Econometrica* 77 (4): 1003–45.
- Currarini, Sergio, Matthew O. Jackson, and Paolo Pin. 2010. "Identifying the Roles of Race-Based Choice and Chance in High School Friendship Network Formation." *PNAS* 107 (11): 4857–61.
- DeGroot, Morris H. 1974. "Reaching a Consensus." *Journal of the American Statistical Association* 69 (345): 118–21.
- DeMarzo, Peter M., Dimitri Vayanos, and Jeffrey Zwiebel. 2003. "Persuasion Bias, Social Influence, and Unidimensional Opinions." *Quarterly Journal of Economics* 118 (3): 909–68.
- Dustmann, Christian, Albrecht Glitz, Uta Schönberg, and Herbert Brücker. 2016. "Referral-Based Job Search Networks." *Review of Economic Studies* 83 (2): 514–46.

- Dwenger, Nadja, Henrik Kleven, Imran Rasul, and Johannes Rincke. 2016. "Extrinsic and Intrinsic Motivations for Tax Compliance: Evidence from a Field Experiment in Germany." *American Economic Journal: Economic Policy* 8 (3): 203–32.
- Fellner, Gerlinde, Rupert Sausgruber, and Christian Traxler. 2013. "Testing Enforcement Strategies in the Field: Threat, Moral Appeal and Social Information." *Journal of the European Economic Association* 11 (3): 634–60.
- Glaeser, Edward L., Bruce I. Sacerdote, and Jose A. Scheinkman. 2003. "The Social Multiplier." *Journal of the European Economic Association* 1 (2–3): 345–53.
- Golub, Benjamin, and Matthew O. Jackson. 2010. "Naïve Learning in Social Networks and the Wisdom of Crowds." *American Economic Journal: Microeconomics* 2 (1): 112–49.
- Golub, Benjamin, and Matthew O. Jackson. 2012. "How Homophily Affects the Speed of Learning and Best-Response Dynamics." *Quarterly Journal of Economics* 127 (3): 1287–1338.
- Granovetter, Mark. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83 (6): 1420–43.
- Imbens, Guido W., and Jeffrey M. Wooldridge. 2009. "Recent Developments in the Econometrics of Program Evaluation." *Journal of Economic Literature* 47 (1): 5–86.
- Jackson, Matthew O. 2008. *Social and Economic Networks*. Princeton: Princeton University Press.
- Jackson, Matthew O. 2016. "The Past and Future of Network Analysis in Economics." In *The Oxford Handbook of the Economics of Networks*, edited by Yann Bramoullé, Andrea Galeotti, and Brian Rogers, 117–33. Oxford: Oxford University Press.
- Jackson, Matthew O., Tomas Rodriguez-Barraquer, and Xu Tan. 2012. "Social Capital and Social Quilts: Network Patterns of Favor Exchange." *American Economic Review* 102 (5): 1857–97.
- Jacobs, Jane. 1961. *The Death and Life of Great American Cities*. New York: Random House.
- Khan, Adnan Q., Asim I. Khwaja, and Benjamin A. Olken. 2016. "Tax Farming Redux: Experimental Evidence on Performance Pay for Tax Collectors." *Quarterly Journal of Economics* 131 (1): 219–71.
- Kleven, Henrik Jacobsen, Martin B. Knudsen, Claus Thustrup Kreiner, Søren Pedersen, and Emmanuel Saez. 2011. "Unwilling or Unable to Cheat? Evidence from a Tax Audit Experiment in Denmark." *Econometrica* 79 (3): 651–92.
- Kuhn, Peter, Peter Kooreman, Adriaan Soetevent, and Arie Kapteyn. 2011. "The Effects of Lottery Prizes on Winners and Their Neighbors: Evidence from the Dutch Postcode Lottery." *American Economic Review* 101 (5): 2226–47.
- Landry, Craig E., Andreas Lange, John A. List, Michael K. Price, and Nicholas G. Rupp. 2006. "Toward an Understanding of the Economics of Charity: Evidence from a Field Experiment." *Quarterly Journal of Economics* 121 (2): 747–82.
- Lochner, Lance. 2007. "Individual Perceptions of the Criminal Justice System." *American Economic Review* 97 (1): 444–60.
- Marmaros, David, and Bruce Sacerdote. 2006. "How Do Friendships Form?" *Quarterly Journal of Economics* 121 (1): 79–119.
- McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. "Birds of a Feather: Homophily in Social Networks." *Annual Review of Sociology* 27: 415–44.
- Miguel, Edward, and Michael Kremer. 2004. "Worms: Identifying Impacts on Education and Health in the Presence of Treatment Externalities." *Econometrica* 72 (1): 159–217.
- Olken, Benjamin A. 2007. "Monitoring Corruption: Evidence from a Field Experiment in Indonesia." *Journal of Political Economy* 115 (2): 200–249.
- Pomeranz, Dina. 2015. "No Taxation without Information: Deterrence and Self-Enforcement in the Value Added Tax." *American Economic Review* 105 (8): 2539–69.
- Rincke, Johannes, and Christian Traxler. 2011. "Enforcement Spillovers." *Review of Economics and Statistics* 93 (4): 1224–34.
- Sah, Raaj K. 1991. "Social Osmosis and Patterns of Crime." *Journal of Political Economy* 99 (6): 1272–95.
- Shimeles, Abebe, Daniel Zerfu Gurara, and Firew Woldeyes. 2017. "Taxman's Dilemma: Coercion or Persuasion? Evidence from a Randomized Field Experiment in Ethiopia." *American Economic Review* 107 (5): 420–24.