

Nationwide Diffusion of Technology: Experimental Evidence from Multiple Networks*

Deivy Houeix[†]

[\[Click here for the most recent version\]](#)

May 3, 2026

Abstract

When does a new technology spread between firms, and through which connections? I randomize access to a digital payment product among 2,196 taxi businesses in Senegal and trace diffusion through phone-contact directories and mobile money records covering most of the country. Three forces travel through different networks and pull in opposite directions: strong social ties carry learning across industries and cities, geographic proximity carries competition that decays sharply with distance, and shared customers carry demand complementarities that dominate within industry—even between firms with no social connection. Aggregate spillovers can mislead: average within-industry exposure looks negative even where learning is large and positive. A structural decomposition shows that targeting high-degree firms raises adoption by +61.7% over random, while targeting the most visible ones backfires (-13.5%).

*I am grateful to Isaiah Andrews, David Atkin, Abhijit Banerjee, Esther Duflo, Dean Eckles, Justine Knebelmann, Sahil Loomba, Tishara Garg, Ben Olken, Kunal Sangani, Frank Schilbach, Tavneet Suri, Rob Townsend and participants to the various MIT lunches for valuable suggestions and comments. I am also indebted to an amazing research team, particularly Koffi Adjassou, William Amedanou, Justin Chery, Martin Duflo, Cheikh Gaye, Ismael Keita, Loic Laurent, Sara Niang, Yacine Ouahioune, and Francois Soleiman for their outstanding field assistance, as well as Hannah S. Ono, Prabhav Sharma, Alexis Yi, and Lawrence Yong at MIT for their excellent research assistance. I am especially thankful to the entire team at Wave Mobile Money, the implementing partner company, for trusting me to carry out this study independently. I gratefully acknowledge financial support from JPAL-DigiFI Africa, PEDL, The Weiss Fund, and Shultz, and data collection support from SDI. I received IRB approval from MIT, protocol #2012000286. This RCT was registered as AEARCTR-0009155.

[†] Department of Economics, Harvard University

Technology diffusion is a first-order determinant of firm productivity and economic growth.¹ A large literature has now established that technology spillovers exist across a range of settings. A central question remains open: under what conditions does a new technology spread from one firm to another, and when does it fail to? Some studies find large positive spillovers within industries (Bloom et al., 2013; Myers and Lanahan, 2022) or negative ones (Cefalà et al., 2025); others find that geographic proximity is key (Marshall, 1920; Moretti, 2021; Baum-Snow et al., 2024); while others emphasize strategic complementarities among customers (Crouzet et al., 2023; Higgins, 2024) or social learning among peer firms (Chaurey et al., 2025). A natural interpretation is that different channels dominate in different contexts. But because most studies observe only one network layer at a time, it remains difficult to determine which channels are active, when, and how they interact.

A key reason relates to data limitations. Identifying technology diffusion requires three objects rarely observed together: a technology shock, the adoption decisions of connected firms across the economy, and the networks through which information and market pressure travel. Experimental studies that cleanly identify the shock are typically limited to the networks they can survey, which tend to be local and industry-specific (Cai et al., 2015; Duflo et al., 2026; Chaurey et al., 2025). Observational studies with broader geographic reach exploit policy shocks but face greater difficulty isolating the channel through which exposure operates (Crouzet et al., 2023; Higgins, 2024; Crouzet et al., 2024). Studies using administrative data—VAT networks, patent citations, trade linkages—capture formal economic ties but may miss the informal social relationships through which firm owners exchange information (Bloom et al., 2013; Alfaro-Ureña et al., 2022). This matters especially in developing economies, where firms are predominantly small, informal, and owner-operated, and where social and financial lives are deeply intertwined (Bernhardt et al., 2019). Observing one network layer while missing others does not merely reduce precision; it can reverse the sign of estimated spillovers.

I address this problem by combining a randomized experiment with two data sources that together reveal a substantial share social and economic connections among informal firms across Senegal. The experiment randomized early access to a merchant payment product among 2,196 taxi businesses in Dakar, with 60% assigned to treatment. I then trace diffusion beyond the experimental sample using encrypted phone contact directories covering roughly 7.2 million individuals (about 80% of the adult population) and the complete mobile money transaction ledger on the country’s leading platform. The technology spread naturally outside the experimental sample—to other taxi firms, to busi-

¹Arrow (1962); Jaffe (1986); Romer (1990); Grossman and Helpman (1991).

nesses in different industries, and to cities hundreds of kilometers from Dakar—creating the variation needed to study diffusion across industries, across space, and across network layers from a single experimental shock.

The phone network data constructed are, to my knowledge, novel in economics. Upon first login, the mobile money operator collects a hashed fingerprint of each user’s phone contact book (consent rate of 86%). Each saved contact number is mapped into one of $2^{16} = 65,536$ hash buckets, a compression that preserves set-intersection and set-overlap operations while making it infeasible to recover raw phone numbers. Two users are linked if each of their encrypted phone appears in each other’s contact book—a mutual-presence requirement that removes one-sided or noisy entries—and tie intensity is measured by the overlap share of their contact sets. The approach draws on the computer science literature on secure set-similarity measurement, beginning with [Broder \(1997\)](#) and extended to privacy-preserving settings by [De Cristofaro et al. \(2013\)](#) and [Swamidass et al. \(2015\)](#). Because most firms in this setting are owner-operated, these fingerprints provide unusually direct evidence on the social networks of firm owners, not just their formal business relationships. The resulting network links the experimental taxi owners to approximately 4 million individuals across Senegal, a scale that makes it possible to study diffusion across industries and space from a single experimental sample.

This combination of experimental variation with economy-wide network data allows me to identify the conditions under which technology diffuses or not, decomposing the net effect into three core channels, never separated before—social learning, within-industry competition, and customer-demand complementarities—and showing how each depends on tie type, industry, and distance. I show that different network layers carry different channels. Strong social ties transmit learning across industries and across cities; shared customers transmit demand complementarities; and local proximity primarily transmits competition and negative spillovers. Because these channels can have opposite signs, aggregate spillovers can be thus misleading.

Three findings stand out. First, technology diffuses across industries and across space, but only through meaningful social ties. Non-taxi merchants connected to treated taxis through strong phone ties ($\geq 5\%$ contact-book overlap) adopt at significantly higher rates (+10.1% effect relative to the control mean, $p_{RI} = 0.030$). The same exposure on weaker ties produces a precise zero. The pattern extends beyond Dakar: businesses in other Senegalese cities connected through strong ties also adopt at elevated rates (+22.3%), with a proportional effect roughly twice as large, consistent with a lower baseline amplifying the same force. Because these receiving firms do not compete with taxis, the cross-industry estimates come closest to isolating a social-learning channel. Survey data collected after

the experiment is consistent with spillovers operating through friends, frequent interlocutors, and bridging ties. Cross-city diffusion appears partly driven by seasonal migration: individuals connected to treated drivers in Dakar who visit secondary cities—for religious events, family, or economic activity—carry knowledge of the technology and adopt there. Seasonal and circular migration between primary and secondary cities is pervasive across developing countries but remains largely absent from the technology diffusion literature; these results suggest it may be an important vehicle through which urban innovations propagate across space.

Second, within-industry spillovers are fundamentally heterogeneous because learning and competition pull in opposite directions. The average same-industry coefficient is significantly negative (-8.4%), but survey-identified friends generate a large positive effect ($+36.1\%$). For the 3% of taxi merchants with strong ties, learning dominates; for the 97% connected mainly through weaker ties, competition dominates. Evidence from taxi garages sharpens this interpretation: the competitive effect is large and significant below 100 meters and decays steeply with distance. While [Baum-Snow et al. \(2024\)](#) find positive agglomeration spillovers from geographic proximity among innovating firms, the same spatial concentration generates *negative* spillovers in this setting because nearby informal taxi merchants compete directly for the same customers and parking spots, and adoption by a neighbor diverts business rather than creating knowledge externalities.

Third, customer-demand complementarities are the dominant within-taxi channel, operating independently of social ties. I identify this channel using shared customers—individuals who make pay-to-business mobile money transactions with both a treated driver and an untreated taxi merchant. These are regular digital-payment users (median 43 unique merchants paid in the window) for whom taxi rides are a small fraction of their transaction activity. When such a customer begins paying a treated driver digitally, the same customer’s transactions at other merchants create demand pressure for those merchants to adopt. The within-taxi customer coefficient is $+0.421$ ($p_{RI} = 0.001$, $N = 10,151$, effect size = 57.1%), and the effect persists even among merchants with no social connection to the treated driver, ruling out learning as an explanation. The cross-industry customer coefficient is weaker, consistent with customers creating demand especially within the sector where they observe digital payment in use.

A structural model ties these findings together. Using six reduced-form moments and the distance gradient of garage coefficients (six bins, four overidentifying restrictions) as identifying conditions, and four survey moments as untargeted validation, I estimate a discrete-time adoption model that decomposes each firm’s adoption probability into social learning (separately for strong and weak ties), geographic competition, and customer

demand. The model has seven parameters, all precisely estimated. Strong ties transmit substantially more learning than weak ties ($\hat{\beta}_{\text{learn}}^s = +0.097^{**}$ vs. $\hat{\beta}_{\text{learn}}^w = -0.001 \approx 0$). Geographic competition is intense and local ($\hat{\beta}_{\text{comp}}^{\text{geo}} = -6.279^{***}$, half-life = 186 m). Customer demand generates large positive complementarities ($\hat{\beta}_{\text{demand}} = +1.335^{***}$): non-garage taxi merchants sharing customers with treated drivers are 57.1% more likely to adopt ($p_{\text{RI}} = 0.001$). Within-taxi adoption is driven by the interaction of all three channels: demand dominates competition on average, but competition dominates at short distances near garages. Cross-industry diffusion operates entirely through strong-tie learning.

Policy counterfactuals show why this decomposition matters: selecting high-degree taxi businesses produces +61.7% more adoption than random assignment, while garage-only targeting—which maximizes local competition—backfires, producing -13.5% fewer adopters. The 75-percentage-point gap between the best and worst targeting rule comes entirely from understanding which channels are positive and which are negative.

This paper makes three contributions. First, the paper combines a randomized technology shock with privacy-preserving administrative network data at near-economy-wide scale, generating evidentiary scope that has previously been out of reach in the experimental literature on technology diffusion. Experimental work has largely been confined to networks that can be enumerated by survey—typically a single industry and a small set of villages or neighborhoods (see [Suri and Udry \(2024\)](#) for a review of agricultural settings, and [Chaurey et al. \(2025\)](#) for a contemporary urban-manufacturing study). Observational work using administrative data achieves broader reach but typically captures only formal economic ties, missing the informal social relationships through which firm owners exchange information ([Bloom et al., 2013](#); [Alfaro-Ureña et al., 2022](#)). The encrypted contact-book fingerprints introduced here link the experimental firms to roughly four million individuals through the mutual-contact phone network; combined with the universe of mobile-money transactions on Senegal’s leading platform, they make it possible to follow a single experimental shock as it travels across industries, across cities, and through multiple distinct network layers, without surveying recipient firms directly.

Second, the paper shows that the sign, magnitude, and economic content of technology spillovers depend systematically on which network carries the exposure—a finding that helps reconcile the heterogeneous estimates documented in prior work. The literature has produced large positive within-industry spillovers ([Bloom et al., 2013](#); [Myers and Lanahan, 2022](#)) alongside negative ones ([Cefalà et al., 2025](#)); central roles for geographic proximity ([Moretti, 2021](#); [Baum-Snow et al., 2024](#)) and face-to-face interaction ([Cai and Szeidl, 2017](#); [Atkin et al., 2024](#)); and distinct channels operating through supply chains ([Alfaro-Ureña et al., 2022](#)), demographics ([Crouzet et al., 2024](#)), and knowledge

hoarding (Cefalà et al., 2025). Observing social, customer, and geographic networks separately within one setting reveals that each transmits a distinct economic force: strong social ties carry learning, consistent with embedded ties transmitting detailed, credible, and repeated information and longer ties carrying novel information across social distance (Granovetter, 1973; Aral, 2016; Jahani et al., 2023); shared customers carry demand complementarities that dominate within industry and operate independently of any social connection; and geographic proximity carries competition that is intense and sharply distance-decaying. Because these forces have opposite signs and vary in strength across tie types, aggregate “network exposure” measures conflate them—obscuring offsetting effects that only become visible once each channel is identified separately.

Third, this paper speaks directly to the growing literature on network effects in electronic payment adoption (Björkegren, 2018; Fafchamps et al., 2023; Alvarez et al., 2023; Crouzet et al., 2023; Higgins, 2024). A persistent challenge has been disentangling social learning from strategic complementarities, since both predict higher adoption when nearby firms adopt, and the two are difficult to separate when data observe only geographic proximity or aggregate exposure. Crouzet et al. (2023) and Higgins (2024) provide compelling evidence that strategic complementarities are quantitatively important, using geographic variation from large policy shocks; Alvarez et al. (2023) develop a structural model in which complementarities drive equilibrium dynamics. These papers recognize that social learning may also contribute, but their data environments make it difficult to quantify the relative importance of each channel. Because I observe social and customer networks independently of each other and of geography, I can make significant progress on this decomposition. Within the taxi sector, customer demand generates the largest positive effect, consistent with the emphasis on strategic complementarities. But social learning is the main channel that can be detected operating across industries and cities. Within industry, the aggregate phone effect is negative—not because learning is absent, but because geographic competition offsets it on weak ties. Without separating these forces, one would attribute the entire within-industry effect to a single “network externality,” missing the positive and negative components that partially cancel. The targeting counterfactuals make the stakes concrete: a planner who targets garages to maximize local visibility produces -13.5% fewer adopters than random assignment, while one who targets high-degree drivers produces +61.7% more—a 75-percentage-point gap that is invisible without the structural decomposition.

The rest of the paper proceeds as follows. Section 1 describes the setting, experimental design, and data. Section 2 documents empirical patterns in firms’ social networks. Section 3 presents the empirical framework. Section 4 provides causal estimates of tech-

nology adoption spillovers across industry and space. Section 5 examines the mechanisms driving diffusion. Section 6 develops the structural model and estimation. Section 7 presents counterfactual targeting exercises. Section 8 concludes.

1 Setting, Experiment, and Data

1.1 The technology and market environment

The paper studies the diffusion of a new digital payment technology introduced in Senegal in 2022. The product is a peer-to-business mobile payment tool that allows customers to pay a merchant directly from an existing mobile money account by scanning a QR code displayed by the business, at a 1% fee for drivers (see Figure 2).² In the taxi sector, the technology changes payment in two economically meaningful ways. First, business payments are effectively irreversible, unlike ordinary person-to-person transfers that can often be reversed through a simple call to the provider. Second, the QR interface makes payment substantially more convenient than manually typing a driver’s phone number into the mobile money application. These features matter in a cash-intensive sector such as taxis, where transactions are frequent, small, and operationally costly.

The technology is especially well suited to studying diffusion because adoption is one-sided. Customers do not need to adopt a new consumer technology; they use existing mobile money accounts that were already widespread in Senegal. The adoption decision therefore lies with the business owner. At the same time, the technology was not specific to taxis. The same business-payment product was rolled out across many sectors—restaurants, pharmacies, boutiques, mechanics, and street vendors, among others—during the same period (see a subset of business adopters in Figure A6). This makes the setting unusually useful for studying whether the same technology diffuses beyond the sector in which access is initially randomized.

The market environment also helps isolate the forces behind adoption. The monetary cost of adoption is low. The main adoption cost is time: the merchant must locate a registration agent, provide the required paperwork, and install the business application. In practice, the relevant requirements are possession of an official ID, a personal mobile

²I deliberately do not focus on the 1% fee barrier in this paper. The technology is popular, with over two million unique Senegalese users making business payments during the experiment, though cash remains dominant. All drivers who adopted the technology were aware of the fee, and the primary focus of this study is adoption, not subsequent usage. Additionally, for all drivers in this experiment, the fee was waived for the first CFA 50,000 (approximately USD 83) collected digitally.

money account, and a smartphone.³ These features matter because they make the technology similar to many other digital business tools in low- and middle-income settings: adoption is constrained much more by information, trust, and coordination than by large fixed capital costs.

The broader rollout was rapid and followed an S-shaped curve (see [Jack and Suri \(2014\)](#); [Suri \(2017\)](#) for similar digital technologies). At the start of the experiment, essentially no taxis had adopted the technology, and adoption outside the taxi sector was still limited. Over the following months, however, the same product spread quickly across the economy: within a year, nearly 200,000 businesses in Senegal had adopted the digital payment technology (see [Figure 1](#)).⁴ By early 2024, data show that approximately 17,000 taxis and 220,000 other businesses had adopted the technology, suggesting that by the end of the experiment in October 2022, a substantial share of the potential non-taxi business population had already adopted, while taxi adoption was still growing rapidly (~7,000 of an eventual ~17,000 by early 2024). The technology adoption was not only located in Dakar, the capital city, where the experiment took place, but also across many cities in Senegal (see [Figure A7](#)). This rapid expansion is an asset rather than a complication for the paper. It allows us to ask whether a randomized intervention in one urban sector can generate diffusion of the same technology across industries and across cities.

The taxi sector in Dakar provides a particularly informative experimental environment. Taxis are numerous, socially connected, and economically important (taxis are ubiquitous in Dakar, employing about 4-6% of the urban male population, with about 21,000 active taxis in 2019 (CETUD)).⁵ The sector is predominantly informal and includes two key stakeholders: the taxi business owners and the drivers. In this industry, the owner is the legal entity that owns the car and the majority of the experimental sample consists of single taxi owners who drive their own taxis. In this study, I have grouped all drivers and owners together under the term “taxi businesses,” each considered as a business entity. Drivers interact with many customers each day, regularly encounter other drivers at garages and waiting points, and frequently exchange information. On average, a taxi driver meets about 14 customers a day, and 90% of drivers report engaging in con-

³Smartphone adoption in urban Senegal is relatively high, with the vast majority of taxi businesses in this experiment possessing one. This high rate is also documented across Sub-Saharan Africa more broadly—about 51% including rural areas ([Figure A1](#)).

⁴The company has one main competitor, less popular but that started digitizing payments around the same time. Some businesses adopted both technologies to satisfy more customers.

⁵The population under study predominantly consists of urban poor (according to the PPI Index developed by Innovations for Poverty Action), typically male workers aged 30-50. The taxi industry is often a family enterprise, with family members involved in about 50% of the cases. The vast majority of taxi drivers and owners have not completed primary education (68%) and about half of the drivers were unable to save any money over the past three months.

versation with customers at least half of the time, with 50% doing so most of the time. Additionally, a typical taxi driver reports interacting and regularly spending time with about 15 other taxi drivers. These features make all three channels studied in the paper plausible *ex ante*: social learning through owner networks, competition among nearby same-industry firms, and customer-demand complementarities when customers begin to expect digital payment.

Technological Benefits The economic impact of the technology on businesses is the focus of a separate paper ([Houeix, 2025](#)): from the point of view of this paper, the key takeaway is that this technology is beneficial for businesses. To summarize, the adoption of digital payment technology significantly reduced various costs associated with using cash transactions for drivers. Passengers usually negotiate fares at the beginning of a ride and predominantly paid in cash before the technology is introduced. While peer-to-peer (P2P) mobile money payments were available, they were rarely used due to revocability concerns related to transaction reversibility. The technology decreased the time spent searching for small change and the incidence of change-related issues, such as mistakes in giving change. In addition, the technology improved several non-monetary outcomes, including better account management, reduced anxiety about theft, and fewer incidents of digital robbery. Overall, taxi businesses expressed substantial willingness to pay for this technology. See [Houeix \(2025\)](#) for more details.

1.2 Experimental design

The experiment randomized early access to the technology among taxi businesses in Dakar. A large listing operation was first conducted from March to May 2022, covering garages, car wash stations, informal meeting points, and major traffic locations throughout the city ([Figure A3](#)). This listing exercise identified roughly 3,600 eligible owners and drivers. Eligibility required, among other things, possession of an Android smartphone and the ability to register for the merchant product.

The baseline survey followed immediately after the listing exercise and was conducted from March to June 2022 in twelve waves. Owners and drivers were invited to designated survey locations in Dakar, and treatment status was revealed only at the end of the baseline survey in order to avoid differential survey participation or attrition. The final randomized sample contains 2,196 taxi businesses, with approximately 60 percent assigned to treatment and the remainder to a waitlist control group. Randomization was computerized and stratified on baseline characteristics collected during listing and surveying.

The realized experimental sample is balanced across treatment arms.

The unit of randomization is the taxi business. In most cases this corresponds to a single owner-driver or a very small owner-operated firm. Treated businesses were provided with the QR payment card, the merchant application, and training on how to use the product. Control businesses did not receive the technology during the experimental period and were offered access only after the experiment ended, in October–December 2022. Importantly, the experimental protocol did not encourage treated businesses to market the technology to others or to share information about it. Any adoption among non-experimental firms is therefore naturally interpretable as a spillover of the randomized intervention.

The design serves two purposes. First, it identifies the private effects of the technology on adopting taxi businesses, which are studied elsewhere. Second, and central for this paper, it creates exogenous variation in the treatment status of taxi firms that are embedded in much larger social and commercial networks. The core empirical strategy then compares connected firms that differ in the share of their experimentally linked taxis who happened to be randomized into treatment.

1.3 Measuring firm networks at scale

A central object in the paper is social proximity between non-experimental firms and treated taxi businesses. This is measured using encrypted phone-contact fingerprints from the mobile money operator. Unlike the network measures more common in economics—typically survey links or realized transactions—contact-book fingerprints capture the set of contacts a user chose to save in the phone. They are therefore well suited to measuring durable social proximity among small owner-operated firms. Related objects have been studied in computer science for secure set-overlap measurement and private contact discovery, but, to my knowledge, not for economic measurement (De Cristofaro et al., 2013; Swamidass et al., 2015; Demmler et al., 2018).

The operator collected these fingerprints for roughly 7 million mobile money users in Senegal. On first login to the application, users are asked whether they consent to share their contact list; the consent rate is about 86 percent. For consenting users, the operator stores an encrypted fingerprint of the contact book, which it uses operationally to facilitate transfers to saved contacts. Because most firms in this setting are small and owner-operated, these fingerprints provide unusually direct evidence on the social networks through which business information may travel.

Each phone number saved in a user’s contact list is hashed into one of $2^{16} = 65,536$

buckets using a 16-bit FNV-1a hash (Noll et al., 2026). Let F_i denote the set of hashed contacts saved by user i , and let h_i denote the hash of user i 's own phone number. These fingerprints preserve set operations while preventing recovery of raw phone numbers.

The first step is to define an edge. Two users i and j are linked in the phone network if and only if each appears in the other's contact book:

$$G_{ij}^{\text{phone}} = \mathbf{1}\{h_j \in F_i\} \cdot \mathbf{1}\{h_i \in F_j\}.$$

This mutual-presence requirement is important since it removes one-sided or noisy links. For example, a customer may save a taxi driver's number for future rides, or a driver may receive repeated calls from a dispatcher or agent, without the number being saved in return. Such one-sided contacts do not generate an edge.⁶

Conditional on a mutual link, tie intensity is measured by the overlap in the two contact books:

$$\text{OverlapShare}_{ij} = \frac{|F_i \cap F_j|}{\min(|F_i|, |F_j|)}.$$

This is the overlap coefficient normalized by the smaller contact list, so users with larger contact books are not mechanically penalized. Economically, it is best interpreted as a measure of *social proximity*: conditional on having each other's number saved, how embedded are the two users in overlapping social circles? This interpretation is also consistent with the network-science literature, where common-neighbor and neighborhood-overlap measures are informative about latent social proximity and link formation (Adamic and Adar, 2003; Liben-Nowell and Kleinberg, 2007).

Because the fingerprints are hashed to only 16 bits, some observed overlap is mechanical. The analysis therefore uses two related network definitions. The *broad mutual-contact network* includes all reciprocal pairs whose observed overlap weakly exceeds the overlap expected from independent hash collisions. The *high-overlap network* restricts the same reciprocal pairs to those with raw overlap share of at least 5 percent.

The 5 percent cutoff is selected using raw fingerprint metadata on contact-book size for 1,907 experimental drivers. For each candidate threshold $t \in [0, 30]$, the implied number of mutual contacts for each driver—keeping only pairs with raw overlap share at least t —is compared with the driver's actual contact-book size. The threshold that minimizes

⁶The unit of observation in the phone network is a mobile money account, identified by phone number. Using the operator's identity system, 85% of phone numbers can be linked to unique individuals; about 16.5% of individuals hold multiple accounts. Aggregating to the individual level using the operator's legal-entity identifiers changes mutual-contact edge counts by less than 0.5% and strong-tie counts by 1%, because multi-account individuals are typically connected to different drivers through each account. The phone-level analysis is therefore essentially equivalent to a person-level analysis but avoids the loss of outcome linkage that individual-level identifiers would entail.

mean absolute error is $t^* = 5.0\%$. Alternative criteria yield a narrow range of 4.6% to 5.4%: at $t = 4.6\%$, the median ratio of inferred to actual contacts is closest to one, while at $t = 5.4\%$, the mean inferred count is closest to the mean actual count. By contrast, at $t = 0\%$, the average inferred degree is about 8,123, compared with a true average of about 590. The unrestricted mutual-contact graph is therefore far too dense to be interpreted as the underlying phone-contact network.⁷

At the same time, links below the 5 percent cutoff may be economically relevant. A threshold calibrated to recover the core social network is not automatically the right threshold for every mechanism. Competition may operate through weak mutual contacts—for example, nearby rivals, former customers, or peripheral industry contacts. Weak or long ties may also bridge otherwise distant circles and therefore matter for cross-industry or cross-space diffusion, even if they carry less intense information flow.

Network measurement post-experiment. Timing is an important concern. The phone-contact fingerprints are observed in February 2024, after the experiment ended, so they must be interpreted as a post-period proxy for durable pre-existing social proximity. Table B2 supports this interpretation: treated and control drivers are balanced on contact-book size, degree, overlap intensity, and the composition of their measured phone networks ($p = 0.58$ for the joint F -test). A more demanding test checks whether treatment caused selective tie strengthening—for example, treated drivers saving more contacts related to the technology. Appendix Table B3 reports the results. Treated drivers’ overlap distributions do not shift rightward (if anything, they are marginally *lower*); the composition of strong ties ($\geq 5\%$ overlap) is indistinguishable across arms for every dimension tested—number of ties, share taxi, share merchant, mean overlap (all $p > 0.20$); and a placebo on weak ties ($< 1\%$) shows no differential composition. Separately testing taxi contacts, merchant contacts, and non-merchant contacts reveals no distributional differences attributable to treatment ($p = 0.14$ for taxi contacts). These results are inconsistent with treatment-induced tie formation or strengthening.

Of the 2,196 experimental drivers, 1,918 (87.3%) have at least one phone number with fingerprint data (checking all registered numbers per driver), and 1,907 (86.8%) appear in the mutual-contact phone network. The 12.7% attrition from the experimental sample to the fingerprint data is driven by non-collection: the mobile money operator’s fingerprint extraction did not reach these users by the February 2024 observation date. Treated

⁷This calibration also has a simple statistical rationale. For a false mutual match, the expected raw overlap share is approximately $\max(|F_i|, |F_j|)/65,536$. At median driver contact-book sizes, that noise floor is around 0.7 percent. Thresholds in the 2–3 percent range would already remove most purely accidental overlap. A 5 percent cutoff is therefore conservative and lies comfortably within the calibration range.

drivers are more likely to appear (88.1% vs. 85.1%, $p = 0.017$), plausibly because treatment increased digital engagement and app usage, which triggers fingerprint collection. This difference is small and, as Table B2 shows, conditional on appearing in the network sample the two arms are balanced on all network characteristics ($p = 0.58$ for the joint F -test). To formally bound the impact of this differential attrition, Appendix Table B4 reports tightened Lee bounds (Lee, 2009) computed within contact-book-size quartiles. The bounds on the main result (Panel A, cross-industry Dakar) are $[+2.48, +3.83]$, both significant, confirming that differential attrition does not drive the estimated spillovers. After filtering and degree trimming, the phone network contains roughly 12 million mutual edges linking these drivers to approximately 4 million individuals.

Table B1 validates the overlap measure against survey-reported relationships. In the endline survey, 1,693 of the 2,196 experimental drivers completed the network module, in which each was asked about up to 16 randomly selected mobile money contacts. For each contact, the survey elicited relationship type (wife, family, friend, taxi colleague, client, and so on) and interaction frequency (daily/weekly, monthly, rarely/never), as described in Section 1.4.1. The table also reports the share of pairs above the 5% raw-overlap threshold used in the main analysis. Of the 24,136 survey pairs, 13,253 (55%) have a contact whose phone number appears in the fingerprint data and therefore enter the analysis sample.

The survey evidence shows a strong gradient in overlap by reported tie strength. At the top of the distribution, wives share 17.63% of phone contacts on average, and 69.3% of wife pairs exceed the 5% threshold. Family members follow at 11.31% (59.4% above threshold), taxi contacts at 9.93% (57.6%), contacts with whom respondents discussed the technology at 10.70% (59.5%), taxi friends at 8.57% (52.8%), and non-taxi friends at 6.69% (36.6%).

At the bottom of the distribution, taxi clients average 3.90% overlap (10.1% above threshold), serendipitous contacts—numbers the respondent could not identify when prompted, likely incidental or one-time transaction partners—average 3.46% (9.5%), and rarely contacted pairs average 4.33% (17.5%). The same pattern appears by interaction frequency: daily or weekly contacts have 9.63% overlap, versus 4.33% for contacts seen less than monthly or never—a $2.2\times$ difference.

Overall, the survey evidence supports the interpretation of contact-book overlap as social proximity. The 5% raw-overlap threshold separates stronger from weaker ties in a useful, though not sharp, way.

1.4 Survey, transaction, and location data

To interpret the mechanisms behind diffusion, the paper combines the phone network with survey evidence, transaction records, and location data.

1.4.1 Survey Data

The survey data serve two purposes. First, they provide the baseline characteristics used to describe the experimental sample and assess randomization balance. Second, they characterize the content of economic relationships. Between September and December 2023, respondents were asked about a random subset of 16 links drawn from their transaction network. For each link, the survey recorded the relationship type, communication frequency, whether the respondent exchanges advice or other important information with the contact, whether the respondent would borrow money from the contact, and whether the two discussed the payment technology. The questionnaire follows [Banerjee et al. \(2013\)](#). These data are useful for interpreting the administrative network measures. In Table B1, higher contact-book overlap is associated with ties that respondents themselves describe as closer and more socially meaningful.

1.4.2 Transaction Network Data

The transaction data provide a second network layer. Because the mobile money provider first operated peer-to-peer transfers and only later expanded into business payments, the data contain the universe of person-to-person mobile money transactions on the platform before the experiment. These records are used to construct direct transaction links and measures of interaction intensity. The transaction network is much sparser than the phone network, but it has two advantages: it is observed before treatment, and repeated transfers are a revealed form of economic or social interaction. In the reduced-form analysis, it provides an important cross-check on the phone-based results.

The transaction data are also used to construct the customer-link objects relevant for the customer-demand channel. The paper distinguishes two networks. The first is a *pre-intervention customer network*, built from person-to-person transactions between April 2021 and March 2022. Customer-like transfers are identified using the range typical of taxi payments, and repeated contact is required to exclude incidental transfers. Because this network predates the experiment, it provides the cleanest evidence that shared-customer exposure is not mechanically created by treatment. The second is a *post-period customer network*, built from merchant-payment data between November 2022 and November 2023. This later network is a cleaner measure of realized customer overlap, but because it

is observed after treatment it is used more cautiously: as supporting evidence and, later, as a calibration object rather than the sole basis for identification.

The transaction network is less exhaustive than the phone network. At baseline, the average user transacted with 35 other users, compared with 724 phone contacts on average. It nevertheless contains richer information on financial interaction. In the 12 months before the experiment, the average user sent or received 85 transfers totaling CFA 2,013,490 (USD 3,356). The average transfer size was CFA 23,686 (USD 39). These data make it possible to exploit interaction intensity—for example, the number or frequency of transactions between a taxi business and a given contact.

1.4.3 Location and Migration Data

The location data come from deposits and withdrawals made through the provider’s nationwide network of mobile money agents (see Figure A4). Because each cash-in and cash-out transaction is geolocated using the coordinates of the handling agent, these data support a user-month panel of locations. The resulting panel is unusually informative about temporary migration and geographic diffusion. It makes it possible to identify users who remain in a city throughout the experiment, users who temporarily migrate into Dakar, and treated taxi drivers who travel to other cities during the rollout.⁸ These measures are essential for distinguishing remote diffusion through social ties from diffusion through physical co-presence.

1.4.4 Technology Adoption Data

Finally, adoption of the business payment technology is observed directly in the provider’s administrative records. For each adopting business, the data record the registration date, merchant category, and registration location. The main outcome is adoption during the experimental window, defined as the period from the start of the baseline survey until the control group became eligible for treatment. These data allow adoption to be tracked separately for taxis in Dakar, non-taxi merchants in Dakar, and non-taxi merchants elsewhere in Senegal. That distinction is central to the paper because it links the same randomized intervention to diffusion within industry, across industries, and across space (see Figure 1).

⁸As a check on the construction, the data show substantial population movements into Touba during the pilgrimage month; see Figure A5.

2 Empirical Patterns in Firms' Networks

I begin by documenting two key patterns in the networks of small firms, using the comprehensive network data. These patterns motivate the investigation of technology diffusion within firms' social networks.

Stylized Fact 1 - Firms' networks not only are local and intra-industry, but also expand across industries and regions. On average, the taxi businesses in the experiment have approximately 590 phone contacts (unique entries in the encrypted contact-book fingerprint). Using the mutual-contact criterion, each driver is linked to a mean of 6,476 individuals, of which 550 are strong ties ($\geq 5\%$ overlap) and the remainder weak ties.⁹ The data reveal that the taxi businesses' network extends significantly beyond Dakar, where the experiment took place: over half of their phone contacts reside outside the capital city, and most of these are located in urban areas across the country.

As a descriptive exercise, I identify business links within the phone contact network by matching mutual contacts to all business adopters of digital payments by April 2023. On average, each driver is connected to 250 other business owners through mutual phone contacts, of which 12 are taxi merchants and 237 are in other industries (3.9% of all mutual contacts). Among strong ties specifically, each driver is connected on average to 29 business owners (1 taxi, 28 non-taxi). This extensive cross-industry reach is a critical feature of the data: it allows me to track technology diffusion across both industry and space from a single experimental sample, and it reflects the broader characteristics of firm networks in this low-income, owner-operated context.

Stylized Fact 2 - Substantial information sharing among social networks. Small taxi businesses not only share information and advice among themselves but also with firms beyond the taxi industry. As described earlier, I extracted a random subset of about sixteen mobile money recipients and senders from the experimental participants and surveyed each subject about the nature of their relationships with these individuals. As shown in Table B7, taxi businesses are familiar with 58% of the extracted contacts, the majority being family or friends. Notably, many of these family and friends also own businesses, often outside the taxi industry, providing additional support that many social connections are with other business owners. Among the people they know, businesses communicate frequently, interacting daily or weekly with 68% of them. Taxi businesses

⁹These counts are computed on the analysis sample after keeping only the most informative phone per individual and applying top-2% degree trimming. Before trimming, the mean mutual-contact count is somewhat higher (Table B2).

share important advice and information with about 74% of the people they know, particularly other businesses and taxi drivers.¹⁰ More specifically, they reported having discussed or recommended the digital payment technology to 12% of the people they know. This underscores the high level of information sharing and potential for social learning within business networks, motivating the exploration of social learning as a key mechanism for technology diffusion.

3 Empirical Framework

The experiment generates exogenous variation in firms' exposure to early adopters. I exploit this variation to estimate spillovers: how the adoption decision of a non-experimental firm responds when a larger share of the taxi businesses in its network were assigned to treatment. The estimand throughout is the causal effect of randomized exposure on connected firms' adoption, estimated separately across network layers, industries, and locations.

3.1 Exposure-based reduced-form design

The main specification regresses adoption on the share of a focal unit's experimental taxi contacts that were treated. Let $\mathcal{E}$ denote the set of experimental taxi businesses, and let $T_i \in \{0, 1\}$ indicate treatment assignment for experimental taxi $i \in \mathcal{E}$. For focal unit j and network layer n , let $\mathcal{E}_j^{(n)} \subseteq \mathcal{E}$ denote the experimental taxis connected to j in network n , with $N_j^{(n)} = |\mathcal{E}_j^{(n)}| \geq 1$. The share-treated exposure variable is

$$S_j^{(n)} = \frac{1}{N_j^{(n)}} \sum_{i \in \mathcal{E}_j^{(n)}} T_i.$$

The reduced-form specification is

$$Y_j = \alpha_{g,n} + \beta_{g,n} S_j^{(n)} + \lambda_{g,n} \bar{S}_j^{(n)} + u_j, \quad (1)$$

where Y_j indicates whether unit j adopts during the experimental window, g indexes the receiving population, and $\bar{S}_j^{(n)} = \mathbb{E}_{\mathcal{R}}[S_j^{(n)} \mid \mathcal{E}_j^{(n)}]$ is the expected exposure under the randomization protocol $\mathcal{R}$, computed from 1,000 re-randomization draws following [Borusyak and Hull \(2023\)](#). Under complete randomization with a common treatment

¹⁰The exact survey question was "More generally, do you provide important information/advice to [Name]?" in the spirit of [Banerjee et al. \(2013\)](#).

probability, this control is closely related to the number of experimental contacts; under the stratified design used here, it also absorbs differences in the composition of a focal unit’s experimental neighbors across strata.

The coefficient $\beta_{g,n}$ captures the causal effect on adoption of increasing the proportion of treated taxis among a focal unit’s experimental contacts in layer n . The share-treated parameterization is preferred because the network layers differ substantially in density: a ten-percentage-point increase in exposure has a comparable interpretation whether measured in the phone network, the transaction network, or the customer network.

A complementary number-treated specification replaces the share with the count:

$$Y_j = \tilde{\alpha}_{g,n} + \tilde{\beta}_{g,n} T_j^{(n)} + \tilde{\lambda}_{g,n} \bar{T}_j^{(n)} + \tilde{u}_j, \quad T_j^{(n)} = \sum_{i \in \mathcal{E}_j^{(n)}} T_i, \quad (2)$$

where $\bar{T}_j^{(n)} = \mathbb{E}_{\mathcal{R}}[T_j^{(n)} \mid \mathcal{E}_j^{(n)}]$ is the expected number treated under the randomization protocol. This specification asks how adoption changes when one additional experimental taxi contact is treated, controlling for expected exposure. The two parameterizations yield similar results.

Equation (1) resembles the linear-in-means formulation in the peer-effects literature (Manski, 1993; Brock and Durlauf, 2001) and is close in spirit to recent exposure-based designs that use randomized peer treatment to study learning and spillovers (Cai et al., 2015; Duflo et al., 2026; Chaurey et al., 2025). The key distinction from a classical peer-outcomes regression is that the regressor is peers’ *randomized treatment status*, not their endogenous adoption choices. Conditional on network position, variation in $S_j^{(n)}$ is generated by the experiment itself, so the usual reflection problem does not arise.

The reduced-form coefficients capture *net* spillover effects through the network used to define exposure. A positive coefficient in the cross-industry phone network can be read as evidence of social learning, because same-industry competition is absent in that population. A coefficient in the within-taxi phone network, by contrast, aggregates learning and competition. Estimating the same design across multiple network layers and samples—holding the empirical framework constant while varying the channel content of the exposure measure—is what allows the reduced-form evidence to speak to the mechanism decomposition developed later in the structural model.

3.2 Outcomes, samples, and network definitions

The outcome is an indicator for adopting the business payment technology during the experimental window, from the start of baseline surveying on March 28, 2022, until the

control group became eligible for treatment on October 5, 2022. Adoption is observed directly in the administrative records of the mobile money provider.

Three receiving populations define the main spillover samples. The first is non-experimental taxi merchants in Dakar, capturing within-industry diffusion. The second is non-taxi merchants in Dakar, capturing cross-industry diffusion within the treatment city. The third is non-taxi merchants outside Dakar, capturing diffusion across space. In all cases, the experimental taxis themselves are excluded from the spillover regressions, so the estimand is the response of untreated connected firms to randomized exposure in their network (Athey et al., 2018).

The estimating sample is defined by the network layer used to measure exposure: a focal unit enters a regression only if it has at least one experimental taxi contact in that layer. This restriction keeps estimation on the support of feasible exposure and avoids comparing connected firms to structurally isolated ones. Because the relevant connected set differs across network layers, the sample changes across specifications.

Several network layers are used, each loading on different spillover channels. The *broad mutual-contact phone network* includes all reciprocal contacts with nonnegative corrected overlap and captures generic exposure, weak ties, and competitive proximity. The *high-overlap phone network* restricts to reciprocal contacts with at least 5 percent raw contact-book overlap and is preferred for isolating social learning; as argued in Section 1.3, this threshold best recovers the true scale of drivers' contact books and sharply reduces hash-collision noise. The *pre-experiment transaction network* captures direct economic interaction and provides a cross-check on the phone-network results. The *customer network*, measured both before and after the intervention, captures shared customers between taxi merchants and identifies the customer-demand channel; the pre-intervention version provides the cleanest evidence, while the post-period version is measured more precisely but must be interpreted with greater caution. Finally, the *garage network* links merchants through local spatial proximity and isolates competitive exposure; because few experimental taxis share garages, the main regressor there is an indicator for whether any treated experimental taxi operates within a given radius.

3.3 Inference under network dependence

Overlapping network connections generate dependence across focal units: two merchants may share experimental neighbors, be linked to each other, or belong to the same local market. Rather than imposing a parametric dependence structure, I base inference on the experimental randomization itself.

For each specification, I compute randomization-inference p -values by redrawing treatment assignment according to the original stratified design, recomputing exposure, and re-estimating the coefficient. The focal units, the network, and the specification are held fixed. Because focal units are never directly treated, the relevant sharp null is that the treatment assignments of all experimental taxis have no effect on focal-unit outcomes. Under this null, the randomization distribution provides the appropriate finite-sample benchmark (Athey et al., 2018).

The tables report heteroskedasticity-robust standard errors alongside randomization-inference p -values, but the latter are the primary measure of statistical significance. In the main reduced-form tables, p -values are two-sided. In mechanism specifications where the sign is predicted ex ante—most notably the garage-competition regressions—one-sided p -values are also reported.

4 Technology Diffusion Across Industry and Space

This section documents three reduced-form facts. First, non-taxi firms in Dakar adopt at higher rates when more of their social contacts among taxi drivers were assigned to treatment. Second, this diffusion extends beyond Dakar: businesses in other Senegalese cities connected to treated taxis also adopt at elevated rates. Third, within-industry spillovers are heterogeneous—positive on strong ties, negative on weak ones—reflecting the opposing forces of social learning and competition. Together, these results establish broad causal diffusion while showing that the sign and magnitude of spillovers depend on the type of tie through which exposure operates.

4.1 Cross-industry diffusion in Dakar

Non-taxi merchants in Dakar adopt at higher rates when more of their taxi contacts were assigned to treatment. Because these firms operate outside the taxi sector and do not compete directly with the experimental taxis, the cross-industry estimate is largely free of the within-industry confounds that complicate taxi-on-taxi results.

Table 1 shows that the effect depends on tie strength. Panel A restricts to mutual contacts with at least 5 percent raw contact-book overlap—the threshold that best recovers the true scale of drivers’ contact lists and sharply reduces hash-collision noise (Section 1.3). On these stronger ties the cross-industry coefficient is positive, significant, and economically large: moving the treated share from zero to one raises adoption by about 0.3 percentage points ($\hat{\beta} \times 1,000 = +2.93$, $p_{\text{RI}} = 0.030$), a +10.1% effect relative to the

control mean. On the weaker ties, by contrast, the coefficient is indistinguishable from zero.¹¹

This pattern—null on weaker ties, significant on high-overlap ties—is consistent with social learning operating through relationships with enough shared context to transmit credible information about a technology’s reliability and profitability. Section 5 confirms this interpretation: competition and customer-demand complementarities matter primarily within the taxi sector, so the cross-industry estimates isolate a nearly pure social-learning margin.

Survey evidence reinforces this conclusion. Table 2 shows that cross-industry spillovers operate across a range of relationship categories. Estimates on stronger ties are positive, though imprecise, and exposure through weaker ties—former clients, serendipitous encounters, and other non-core contacts—is also economically important, consistent with such ties bridging otherwise separate parts of the economy. Many of these nominally weaker survey ties still exceed the 5 percent phone-overlap threshold (Table B1).

4.2 Diffusion across cities

Technology adoption also spreads beyond the capital. Panel A of Table 1 shows that businesses in other Senegalese cities connected to treated Dakar taxis adopt at significantly higher rates, with a proportional effect roughly twice as large as the within-Dakar cross-industry estimate—consistent with a lower baseline outside the capital amplifying the same diffusion force. As in Dakar, the effect concentrates on high-overlap ties; the coefficient on weaker contacts is small and insignificant. These results show that knowledge spillovers need not be local in the classic Marshallian sense (Marshall, 1920).

Table 6 shows that this pattern is not driven by a single destination. Across the major secondary cities with enough connected exposure, point estimates are generally positive. Precision varies with the density of network links into each city, but the qualitative pattern is stable. Given the low baseline adoption rate outside Dakar, these treatment-induced increases represent large proportional shifts in adoption behavior.

¹¹Because the phone network is observed after the experiment, identification relies on balance of network characteristics across treatment arms. Treated and control drivers are statistically indistinguishable on the composition and intensity of their measured phone networks, including the share of merchants, the share of high-overlap contacts, and mean overlap. If anything, control drivers have more detected mutual contacts, working against finding positive spillovers from treatment. See Table B2.

4.3 Within-industry exposure and tie heterogeneity

Within the taxi sector, same-industry exposure combines social learning with within-market competition. Learning pushes toward positive spillovers; competition pushes toward negative ones; and the average phone-exposure coefficient reflects their net effect. Table 1 confirms this tension: the within-taxi coefficient on the broad mutual-contact network is significantly negative. On the high-overlap network (Panel A), the effect size is larger in magnitude but imprecisely estimated, reflecting the much smaller sample of strong same-industry ties.

Disaggregating by tie type reveals that learning remains potent within industry on meaningful relationships. Table 2 reports that exposure through friends generates a substantial positive effect (+36.1%), while exposure through frequent interlocutors produces a similar effect (+24.3%). The contrast with the negative average phone result implies that competition dominates on the full same-industry contact set, but learning offsets—and can exceed—competitive pressure on stronger ties.

This heterogeneity aligns with the broader sociology of diffusion. Embedded ties transmit detailed, credible, and repeated information, while longer ties carry novel information across social distance (Granovetter, 1973; Aral, 2016; Jahani et al., 2023). In the present setting, strong within-industry ties help overcome knowledge-hoarding frictions (Cefalà et al., 2025) and support learning even amid rivalry. Weak same-industry contacts, by contrast, are more exposed to competitive crowding and information withholding, depressing the average spillover.

The reduced-form evidence therefore establishes broad diffusion across industry and space, but also shows that within-industry spillovers are fundamentally heterogeneous. The next section decomposes these effects into three channels: social learning, local competition, and customer demand.

5 Mechanisms: Social Learning, Competition, and Customer Demand

The reduced-form results in Section 4 point to multiple channels. Cross-industry and inter-city spillovers are consistent with social learning; heterogeneous within-industry results suggest that competition also matters. This section organizes the evidence around three channels—social learning on meaningful ties, local competitive crowding within industry, and customer-demand complementarities—and adds a fourth result on geographic diffusion through remote interactions. Each subsection identifies its channel

through a distinct empirical variation and ties it to the structural decomposition developed later.

5.1 Social learning on meaningful social ties

Spillovers are strongest where relationships are most capable of transmitting information. This pattern—not predicted by generic market complementarities—is the core evidence for social learning.

Survey-elicited relationship types sharpen the interpretation. Table 2 shows that within the taxi sector, spillovers concentrate among friends (+36.1%, $p_{RI} = 0.007$) and frequent interlocutors (+24.3%, $p_{RI} = 0.013$), while exposure through non-friends produces a smaller and insignificant effect (+2.8%). Close contacts are the people from whom a firm owner can learn how the technology works, whether customers use it, and whether payments are reliable.

Cross-industry survey evidence shows a complementary pattern. Table 2, Panel C reports that cross-industry spillovers are driven by serendipitous encounters (+47.7%, $p_{RI} = 0.003$) and former or regular taxi clients (+93.7%, $p_{RI} = 0.151$). The strongest ties (Panel B: friends, frequent contacts) show positive but imprecise cross-industry effects. This fits the idea that novel information often arrives through ties bridging separate circles (Atkin et al., 2024): inside the taxi sector, close ties overcome rivalry-induced information withholding; across industries, weaker ties are valuable precisely because they connect firms to different market environments.

The transaction network provides a final cross-check. Table B8 shows positive point estimates for within-taxi ($\hat{\beta} = 2.937$, effect size +13.7%) and cross-industry adoption in Dakar ($\hat{\beta} = 3.901$, effect size +8.2%), though neither reaches conventional significance ($p_{RI} = 0.139$ and 0.128). Despite the sparsity of this network, the sign pattern is consistent with the phone and survey evidence: diffusion occurs on ties that plausibly carry information.

5.2 Local competition and crowding out within industry

Within the taxi sector, competition partially offsets the positive learning force. This accounts for the weaker and sometimes negative average same-industry coefficients in Table 1: same-industry exposure activates a negative competitive channel that is absent across sectors.

The garage evidence cleanly isolates this channel. For the 20% of taxi drivers operating in garages, Table 3 estimates spillovers from exposure to treated taxis in the same local

garage environment. Panel A, which includes all merchants near garages (identified by the mobile money agents visited throughout the experiment), shows negative point estimates at every radius from 50 to 500 meters, with the coefficient at 300 meters reaching marginal significance ($\hat{\beta} = -0.0036$, $p_{\text{RI}} = 0.084$). Panel B restricts to merchants who are also mutual phone contacts of experimental drivers—a sample where both learning and competition are active and where the merchant population is measured with less noise. The negative effects are large and significant: at 50 meters, $\hat{\beta} = -0.181$ ($p_{\text{RI}} = 0.014$); at 100 meters, $\hat{\beta} = -0.067$ ($p_{\text{RI}} = 0.031$). The gradient is steep, consistent with competitive crowding operating through overlapping foot traffic and customer diversion in very local markets (Baum-Snow et al., 2024). Figure 4 displays the decay of the competition effect with distance.

Two qualifications are important. First, the negative effects are local and strongest where market overlap is most direct; they do not imply that the taxi sector as a whole is characterized by pure crowding out. Second, as Table 2 shows, strong social ties inside the sector still generate positive spillovers. Local rivalry depresses the average same-industry coefficient, but strong relationships transmit enough information to overcome competitive drag.

5.3 Customer-demand complementarities

When customers become accustomed to paying digitally with treated taxis, other merchants serving the same customers may find adoption more attractive—not because they learn from or compete with the treated firm, but because their customers now expect or prefer digital payments. This channel is conceptually distinct from social learning (it operates even when two merchants do not know each other) and from competition (it raises rather than lowers the private return to adoption). The operative link is the customer, not the firm.

The identification strategy exploits the customer network directly. For each non-experimental merchant j , we observe which experimental drivers share customers with j and compute the treated share among those drivers. The key regressor is therefore the fraction of j 's customer-linked experimental contacts that were randomly assigned to treatment, controlling for j 's expected exposure—a function of j 's network position that absorbs any mechanical correlation between connectivity and adoption propensity. The estimating equation asks: conditional on how connected a merchant is to the experiment through shared customers, does it matter whether those connected drivers were actually treated? A positive coefficient implies that customers who experience digital

payments with a treated driver subsequently demand them from other merchants they patronize. Importantly, this test does not require the two merchants to interact or even be aware of each other; the customer is the sole conduit. Access to the full universe of mobile-money transactions—an unusually rich data structure for studying technology diffusion—is what makes this decomposition possible.¹²

Pre-intervention customer network. Table B9 provides the cleanest test by constructing customer links from person-to-person transactions between April 2021 and March 2022, before the experiment began, ruling out treatment-induced link formation by construction.¹³ In this network, exposure to treated taxis through shared customers has a large positive effect among non-social taxi merchants ($\hat{\beta} \times 1,000 = 122.16$, $p < 0.001$, $N = 3,490$), corresponding to roughly 17.9 percent of the dependent-variable mean.¹⁴ The cross-industry coefficient is a precise zero (-7.44 , $p = 0.572$, $N = 6,613$). This pair of findings matches the customer-demand prediction exactly: the channel is real and specific to the sector in which merchants share customers.

Post-experiment customer network. Table 4 provides higher-signal evidence from the merchant-payment network observed during the rollout (shared customers with at least two business transactions post experiment, when the control group could adopt, November 2022–November 2023). The within-taxi coefficient is large and significant: $\hat{\beta} \times 1,000 = +0.421$ ($p_{RI} = 0.001$), an effect size of +57.1% (column 1, $N = 10,151$). Restricting to merchants without mutual phone ties to the treated driver—the non-social subsample in column 3—yields a similar magnitude ($p_{RI} = 0.001$). That shared-customer links predict adoption even absent any social relationship is the cleanest evidence that customer demand operates independently of learning. When the customer link coincides with a social tie (column 2), the effect roughly doubles, consistent with demand and learning being additive: overlapping customers raise the return to adoption, while a social relationship makes it easier to understand and trust the technology.

The larger post-period magnitude is expected: the merchant-payment network is a far more exhaustive measure of actual customer overlap than the pre-period P2P proxy,

¹²Most diffusion studies must infer customer overlap from geographic proximity or survey recall. Here, the platform’s transaction ledger directly reveals which customers patronize which merchants, permitting a network-based test that does not rely on spatial assumptions.

¹³Because the technology was not rolled out pre-experiment, we proxy customer overlap using P2P transfers of values similar to taxi fares (500–3,000 FCFA). This yields a sparse network, but we present it first precisely because it is free of the endogeneity concern that affects the richer post-experiment measure.

¹⁴The customer regressions exclude garage-affiliated drivers when isolating the non-social channel, to avoid confounding customer demand with the local-competition effects documented in Section 5.2. Table B11 shows that customer-demand effects tend to offset the competition effect within garages.

and customer turnover attenuates the predictive power of historical links. At the same time, because treatment increases merchant-payment activity, some measured customer overlap may itself be treatment-induced. The pre-period result therefore remains the primary evidence that the channel is genuine; Table 4 serves as higher-powered supportive evidence and, in the structural section, as a calibration target that may overstate the true magnitude.¹⁵

The cross-industry customer coefficients in Table 4 are positive but small and insignificant (column 4: $\hat{\beta} \times 1,000 = 13.06$, $p_{RI} = 0.444$, effect size +2.1%). This null is not mechanical: customers who transact with treated taxis could in principle also patronize non-taxi merchants. The most likely explanation is that cross-industry customer links in this network are sparse—few non-taxi merchants share a substantial customer base with taxi drivers. The positive sign is consistent with a small force in the right direction, but the data lack power to detect it. In the structural exercise below, the customer-demand channel is accordingly set to zero for cross-industry spillovers.

5.4 Geographic diffusion through remote interactions

The experiment takes place in Dakar, but Section 4.2 showed that exposure to treated taxis causes cross-industry adoption in other Senegalese cities. The natural explanation is physical movement: treated drivers or their contacts travel to secondary cities and transmit the technology in person. Mobile money geolocation data—which records the city of each transaction during the experiment window (March–October 2022)—allow us to test this directly.

Seasonal migration is widespread. Among the individuals in the strong-tie sample (Table 1 Panel A, $\geq 5\%$ overlap), approximately three-quarters can be matched to geolocation records. Of these, roughly two-thirds appeared in Dakar at least once during the experiment, while one-third were never observed in the capital despite maintaining strong phone ties with Dakar-based drivers. Seasonal outmigration from Dakar is common: roughly a third of Dakar-based merchants (33.7%) and drivers (34.1%) transacted outside the capital in at least one month. The rate is balanced across treatment arms (32.6% treated vs. 36.1% control), confirming that migration is not driven by the intervention. These figures establish that physical mobility between Dakar and secondary cities

¹⁵Table B10 provides a temporal-stability diagnostic by splitting the post-experiment customer network into two windows (November 2022–April 2023 and May–November 2023). The positive coefficient appears in both windows, confirming that the result is not an artifact of a particular observation period.

is substantial—precisely the kind of pattern that would sustain a contact-based diffusion channel.

Cross-industry adoption outside Dakar. Table 1 Panel A already establishes cross-city diffusion: the cross-industry coefficient outside Dakar is +1.103 ($p_{RI} = 0.034$), significant at 5%. Table 5 decomposes this by whether the individual was observed in Dakar during the experiment. The effect is substantially larger among individuals never observed in Dakar (+0.318, $p_{RI} = 0.084$) than among Dakar visitors (+0.053, $p = 0.145$). This suggests that learning from treated Dakar contacts can occur remotely—the individual need not have been physically present in Dakar to acquire information about the technology.

City-level decomposition: learning is remote, adoption requires presence. Table 6 provides a complementary decomposition, splitting by whether the individual visited the *specific secondary city* where adoption is measured. For each of the four largest secondary cities—Mbour, Thiès, Touba, and St-Louis—individuals connected to treated Dakar drivers through strong phone ties are split into seasonal migrants (who appeared in that city at least once, per mobile money geolocation data) and non-migrants.

The effects concentrate sharply among individuals who visited the destination city: Mbour (+0.60**, $p_{RI} = 0.024$), Thiès (+0.35*, $p_{RI} = 0.076$), and St-Louis (+1.14**, $p_{RI} = 0.011$). Among non-migrants—individuals who never appeared in the city—the coefficients are close to zero in three of four cities, with the exception of St-Louis (+0.004**). The effect sizes among migrants are large (70–170% of control means).

Taken together, these results point to seasonal migration as a key channel for cross-space technology diffusion. Individuals from the experimental sample who are connected to treated drivers and who seasonally visit secondary cities—for religious events (Touba), family visits (Thiès, Mbour), or economic activity (St-Louis)—carry knowledge of the technology and adopt there. Roughly a third of individuals in the strong-tie network transacted outside Dakar during the experiment, and the significant cross-city effects concentrate among these seasonal migrants. Seasonal migration between Dakar and secondary cities is a large but understudied phenomenon in West Africa; these results suggest it may be an important vehicle for technology diffusion across space.¹⁶

¹⁶The weaker result in Touba ($p_{RI} = 0.148$) may reflect that city’s distinct commercial ecosystem dominated by religious commerce; see Figure A5.

6 Structural Model

The reduced-form evidence identifies three channels of technology diffusion, each with a different sign, operating through a different network layer. Social learning transmits through phone contacts and is positive. Geographic competition transmits through physical proximity to garages and is negative. Customer demand transmits through shared transaction partners and is positive. Because these channels have opposite signs, aggregate spillovers can be misleading: the average within-industry coefficient appears negative even though learning is strongly positive on meaningful social ties. This section estimates a structural model that decomposes adoption into these three channels, quantifies each, and uses the estimates for counterfactual targeting. Appendix Table B6 consolidates the construction, timing, and use of each network object.

6.1 A model of adoption with learning, competition, and demand

Treatment was randomized on March 28, 2022; adoption is measured through October 5, 2022, yielding $T = 6$ monthly periods. Adoption is absorbing: once adopted, the decision is irreversible. At each month, non-adopter j adopts with probability:

$$p_j = \Lambda \left(\alpha_g + \underbrace{[\beta_{\text{learn}}^s \cdot \mathbf{1}[\text{strong}_j] + \beta_{\text{learn}}^w \cdot \mathbf{1}[\text{weak}_j]] \cdot S_j^{\text{phone}}}_{\text{social learning}} + \underbrace{\beta_{\text{comp}}^{\text{geo}} \cdot S_j^{\text{geo}} \cdot \mathbf{1}[\text{taxi}_j]}_{\text{geographic competition}} + \underbrace{\beta_{\text{demand}} \cdot S_j^{\text{cust}} \cdot \mathbf{1}[\text{taxi}_j]}_{\text{customer demand}} \right), \quad (3)$$

where $\Lambda(\cdot)$ is the logistic CDF and $\alpha_g \in \{\alpha_{\text{cross}}, \alpha_{\text{taxi}}\}$ is an industry-specific intercept. Cumulative adoption over six months is $\Pr(\text{adopt}_j) = 1 - (1 - p_j)^6$.

Each channel operates through a distinct network:

- *Phone network* (S^{phone}): the share of j 's mutual phone contacts who are treated. Strong ties ($\geq 5\%$ phone-contact overlap) and weak ties ($< 5\%$) have separate learning rates. This is the only channel available to cross-industry merchants.
- *Customer network* (S^{cust}): the share of j 's customer-linked drivers who are treated, constructed from 133 million pay-to-business (P2B) mobile money transaction records.¹⁷

¹⁷A customer qualifies if they made ≥ 2 P2B transactions with any non-garage experimental driver between November 2022 and November 2023. A merchant–driver pair is linked if they share ≥ 1 qualifying customer. This yields 95,114 merchants linked to 1,521 experimental drivers. Correlation with S^{phone} : $r = 0.005$.

When a customer who pays both a treated driver and taxi merchant j starts using digital payments, the customer creates demand pressure on j to adopt.

- *Geographic network* (S^{geo}): the treatment intensity at j 's nearest garage, discounted by distance:

$$S_j^{\text{geo}} = \text{ShareTreated}_{\text{nearest}(j)} \cdot \exp(-\delta \cdot d(j, \text{nearest}(j))), \quad (4)$$

where d is distance in km and $\delta > 0$ is an estimated decay rate. The exponential kernel is standard in economic geography (Ahlfeldt et al., 2015) and fits the data: Table 3 reports coefficients at six distance bins (50–500 m), and $\log |\hat{b}(d)|$ is approximately linear in d ($R^2 = 0.93$), precisely what exponential decay predicts.

The model has seven parameters: two intercepts ($\alpha_{\text{cross}}, \alpha_{\text{taxi}}$), two learning rates ($\beta_{\text{learn}}^s, \beta_{\text{learn}}^w$), geographic competition ($\beta_{\text{comp}}^{\text{geo}} \leq 0$), customer demand (β_{demand}), and distance decay (δ).

6.2 Identification and estimation

I match reduced-form moments to the model's predictions via minimum distance estimation (Chamberlain, 1984; Newey and McFadden, 1994). Six moments (m1–m6) each identify one parameter. The distance gradient (m7) uses six distance-bin coefficients to identify two parameters (δ from the slope, $\beta_{\text{comp}}^{\text{geo}}$ from the level), leaving four overidentifying restrictions tested by the log-linear fit ($R^2 = 0.93$). Because all exposure variables are functions of the exogenous treatment assignment, the model has a closed-form adoption probability—no cascade simulation is needed, in contrast to Banerjee et al. (2013).

#	Source	Target	Identifies
m1	Table 1, $\geq 5\%$ overlap, cross-ind.	OLS coef.	β_{learn}^s
m2	Table 1, $< 5\%$ overlap, cross-ind.	OLS coef.	β_{learn}^w
m3	Table 1, all overlap, within-taxi	OLS coef.	$w_s \beta_1^s + w_w \beta_1^w$
m4	Table 4, column 1	OLS coef.	β_{demand}
m5	Cross-industry adoption rate	Rate	α_{cross}
m6	Within-taxi adoption rate	Rate	α_{taxi}
m7	Table 3, 6 distance bins	Slope + level (4 overid.)	$\delta, \beta_{\text{comp}}^{\text{geo}}$

The identification logic is sequential. Cross-industry moments (m1, m2) identify learning separately for strong and weak ties, because competition and demand are absent across industries. The within-taxi phone moment (m3) loads only on learning ($w_s \beta_{\text{learn}}^s + w_w \beta_{\text{learn}}^w$), since competition operates through the separate geographic exposure S^{geo} . The within-taxi customer moment (m4) isolates demand: among non-garage

taxi merchants sharing customers (defined by ≥ 2 P2B transactions, November 2022–November 2023) with experimental drivers, the coefficient is $+0.421^{***}$ ($p_{\text{RI}} = 0.001$, $N = 10,151$). Adoption rates (m5, m6) pin down the intercepts. The gradient of garage coefficients across distance bins (m7) identifies δ from the log-linear slope and $\beta_{\text{comp}}^{\text{geo}}$ from the level (Appendix C).

6.3 Results

Table 7 presents the estimates.

Learning: strong ties matter, weak ties do not. $\hat{\beta}_{\text{learn}}^s = +0.097^{**}$ versus $\hat{\beta}_{\text{learn}}^w = -0.001 \approx 0$. Information flows through meaningful social relationships—characterized by frequent contact and substantial phone overlap—and is negligible on the 97% of ties that are weak. The quality of a driver’s contacts matters more than their number for cross-industry diffusion.

Geographic competition: intense and local. $\hat{\beta}_{\text{comp}}^{\text{geo}} = -6.279^{***}$, with a half-life of 186 m ($\hat{\delta} = 3.73^{***} \text{ km}^{-1}$). At 50 m from a garage, competitive pressure is $\sim 170\times$ stronger than through the average phone contact; at 500 m it has fallen to $\sim 26\times$. Setting amplification to unity implies an effective phone-contact distance of ~ 19 km—roughly the diameter of Dakar (Appendix F). This reflects the local nature of taxi competition: drivers at the same garage compete for the same customers and parking spots, while drivers a few hundred meters away operate in largely separate markets.

Separating competition from the phone network resolves a puzzle in the reduced-form evidence. The within-taxi phone coefficient (m3) is small and negative. In a model where competition runs through S^{phone} , this would imply that learning is positive but offset by competition. In the geographic model, m3 loads only on learning—and the population-weighted learning effect is close to zero, correctly rationalizing the small aggregate coefficient without requiring a large negative offset.

Customer demand. $\hat{\beta}_{\text{demand}} = +1.335^{***}$: non-garage taxi merchants sharing P2B customers (≥ 2 transactions, November 2022–November 2023) are 57.1% more likely to adopt (Table 4, column 1; $p_{\text{RI}} = 0.001$). What kind of customers create this demand linkage? Among the 28,609 qualifying customers, 23% are *repeat patrons* of at least one driver (≥ 2 transactions with the same driver), and 86% transact with ≥ 2 experimental drivers. These are regular mobile money taxi riders who rotate across the fleet. Crucially, 99%

of them also make P2B payments to non-taxi merchants (median 26 non-taxi merchants per customer, 78 total transactions in the window). When a treated driver onboards a customer into digital payments, that customer’s pre-existing spending at dozens of other merchants creates demand pressure for adoption.

The demand channel is essentially inactive near garages: fewer than 1% of garage-proximate taxi merchants share mobile money customers with experimental drivers. Customers who ride with drivers at a garage are passengers, not the transaction customers in the mobile money network.

Demand complementarity is specific to the taxi industry; extending it to cross-industry merchants would strengthen the counterfactual results. Among cross-industry merchants who are also social ties of drivers (Table 4, column 5), the customer coefficient is positive but imprecise (+165.1, $p_{RI} = 0.190$, $N = 1,037$), suggesting demand may operate across industries when mediated by a social relationship. This noisier estimate is not used in the structural model.

Validation. The model correctly predicts the sign of all four untargeted survey moments (Table 2): positive for taxi friends (m9) and non-friends (m10), positive for cross-industry friends (m8), and near-zero for serendipitous contacts (m11). The sign of m9 is the sharpest test: taxi friends are same-industry competitors, yet the data show a positive coefficient (+0.098). The model explains this through demand dominating competition. It underpredicts magnitudes by 3–30 $\times$, primarily because survey “friends” are the right tail of tie strength and because baseline adoption is rare (Appendix D).

7 Counterfactuals and Targeting

For the same onboarding budget of $N_T = 1,256$ drivers, how much additional adoption could alternative targeting rules generate?¹⁸

7.1 Setup

I fix all three networks at their observed values, recompute each individual’s exposure under alternative treatment assignments, and predict six-month adoption. The counter-

¹⁸The experimental sample contains 2,196 drivers ($\sim 60\%$ treated). Of these, 2,150 have a matched mobile money account (1,256 treated, 894 control), and 1,907 appear in the mutual-contact phone network (Section 1.3). The counterfactual treatment vector operates on the 2,150 matched drivers; the network exposure is computed for all individuals connected to these drivers.

factuals are evaluated on a *combined reachable network* of 168,964 individuals for whom at least one structural channel is active (Appendix E).

What changes. Changing which drivers are treated shifts all three exposure variables simultaneously. S^{phone} shifts for every connected individual. S^{cust} shifts for merchants sharing customers with different drivers. S^{geo} shifts because the treatment share at each garage changes—a garage where all drivers are treated has $\text{ShareTreated} = 1$, generating intense local competition. The comparison between degree targeting (high $\bar{S}^{\text{phone}}$, diffuse geography) and garage targeting (lower $\bar{S}^{\text{phone}}$, concentrated geography) is sharply informative about the costs of geographic concentration.

7.2 Targeting rules

All scenarios hold the budget fixed at $N_T = 1,256$:

1. *Random*: the realized experimental assignment ($\bar{S}^{\text{phone}} = 0.569$).
2. *All-garage*: what if all 1,256 treated drivers had the network characteristics of garage drivers? Phone exposure falls ($\bar{S}^{\text{phone}} = 0.499$ vs. 0.569) because garage drivers have fewer contacts; customer demand is similar ($\bar{S}^{\text{demand}} = 0.555$); and geographic competition reaches its maximum ($\text{ShareTreated}_g = 1$ at every garage). See Appendix E for the reweighting procedure.
3. *Top mobile money users*: the 1,256 heaviest platform users. Requires no network data.
4. *Highest degree*: the 1,256 drivers with the most phone contacts ($\bar{S}^{\text{phone}} = 0.898$).
5. *Geographic reach*: the 1,256 drivers with the most out-of-Dakar contacts.

7.3 Results

Table 8 reports treatment-induced adoption—predicted adoption minus the no-treatment baseline—under each rule, decomposed by margin.

Cross-industry adoption scales with phone exposure. This margin accounts for 87% of induced adoption under random onboarding (146 of 167) and is powered entirely by strong-tie learning. Degree targeting raises it by +59 ($\bar{S}^{\text{phone}} = 0.898$ vs. 0.569). All-garage reduces it by -12.6% because garage drivers have fewer phone contacts ($\bar{S}^{\text{phone}} = 0.499$).

Within-taxi: competition and learning dominate. Under random assignment, demand outweighs competition, producing 44 treatment-induced taxi adopters. The all-garage counterfactual reduces this to 9 (-25.9%). Two forces drive the decline: geographic competition reaches its maximum ($\text{ShareTreated} = 1$ at every garage) and phone learning weakens ($\bar{S}^{\text{phone}}$ falls from 0.569 to 0.499). Customer demand exposure is largely unchanged ($\bar{S}^{\text{demand}} = 0.555$ vs. 0.538), confirming that the cost of garage targeting operates through competition and network reach, not through the demand channel.

Degree targeting moves in the opposite direction, nearly doubling within-taxi adoption to 84 (+89). High-degree drivers are dispersed across Dakar, diluting S^{geo} while raising S^{phone} and S^{cust} citywide.

Total: garage targeting is the worst rule. All-garage produces 145 induced adopters (-13.5% vs. random)—the only rule that performs worse than random. Degree targeting produces 270 (+61.7%). The gap of 75 percentage points between best and worst would be impossible to anticipate without the structural decomposition.

7.4 Policy implications

Three lessons emerge for a social planner maximizing technology adoption from a fixed onboarding budget.

Network targeting is powerful. Selecting high-degree drivers produces +61.7% more adoption than random, implementable from phone fingerprint data already collected for the overlap computation. Geographic reach targeting performs comparably (+59.9%). The gains come overwhelmingly from the cross-industry learning channel, which serves the vast majority of non-taxi individuals in the combined network.

Simple heuristics work. Onboarding heavy mobile money users—requiring no network data—produces 197 induced adopters (+17.7% over random). Heavy users tend to occupy more central network positions, so the heuristic partially captures the degree-targeting gains without requiring social-network information.

Geographic concentration is counterproductive. The strategy that appears most natural from a diffusion standpoint—concentrating treatment at dense, visible garages—is the only rule that does *worse* than random. Garage drivers have fewer phone contacts (less learning) and maximum geographic concentration (more competition). Customer demand exposure is similar, but it cannot offset the competition penalty. The competition half-life of 186 m means the penalty is severe within ~200 m of a garage but decays rapidly. A planner who targets garages because they are visible and easy to reach will

find that these very features activate the competition channel. The 75-point gap between best and worst is invisible without the structural decomposition.

8 Conclusion

This paper combines a randomized experiment with exhaustive network data—phone contact directories covering 80% of Senegal’s adult population and the complete mobile money transaction ledger—to study when technology diffuses across firms and when it does not.

Three channels shape diffusion, each identified through a distinct empirical variation. Social learning operates through meaningful social ties and is the only channel that reaches across industries and across cities; it is negligible on weak ties. Within-industry competition generates negative spillovers that are large at short distances and decay steeply, partially offsetting learning gains among same-sector firms. Customer-demand complementarities—transmitted through shared customers, not social ties—are the dominant force within the taxi sector. The sign and magnitude of spillovers depend entirely on which channel is activated by which type of tie, a finding that helps explain why prior studies, each observing a different slice of the diffusion process, have reached different conclusions.

The paper makes two contributions that go beyond documenting the existence of spillovers. First, because the data and experimental design allow me to separately identify and quantify each channel—learning on strong ties, competition at short distances, demand through shared customers—I can address a question that the literature has not been able to tackle: how should a social planner target a fixed onboarding budget to maximize diffusion? The answer depends entirely on the channel structure. The counterfactuals show that selecting drivers with the most phone contacts generates +61.7% more adoption than random assignment, while the intuitively appealing strategy of concentrating treatment at garages backfires (-13.5%) because it maximizes competitive crowding. The 75-percentage-point gap between the best and worst targeting rules would be impossible to anticipate—or even to sign correctly—without the structural decomposition. Second, the paper shows that technology can diffuse over long distances without physical co-presence: the significant cross-city effects come from drivers who never visited the destination, pointing to remote interactions—WhatsApp stories, phone calls, social media—as the operative channel. This result identifies a diffusion mechanism that is typically invisible in the literature and challenges the Marshallian premise that knowledge spillovers require geographic proximity.

More broadly, the combination of random assignment of technology access, exhaustive network data, and nationwide diffusion of a single technology offers a template for studying knowledge spillovers at the scale of an entire economy. Small firms are central to output and employment in low- and middle-income countries, and their adoption decisions are shaped by social networks that existing data rarely capture. How to design interventions that leverage these networks—targeting the right firms, through the right ties, while accounting for the channels that help and those that hurt—remains an important question for both research and policy.

References

- Adamic, Lada A. and Eytan Adar**, “Friends and Neighbors on the Web,” *Social Networks*, 2003, 25 (3), 211–230.
- Ahlfeldt, Gabriel M., Stephen J. Redding, Daniel M. Sturm, and Nikolaus Wolf**, “The Economics of Density: Evidence from the Berlin Wall,” *Econometrica*, 2015, 83 (6), 2127–2189.
- Alfaro-Ureña, Alonso, Isabela Manelici, and Jose P Vasquez**, “The Effects of Joining Multinational Supply Chains: New Evidence from Firm-to-Firm Linkages*,” *The Quarterly Journal of Economics*, 01 2022, 137 (3), 1495–1552.
- Alvarez, Fernando, David Argente, Francesco Lippi, Esteban Mendez, and Diana Van Patten**, “Strategic Complementarities in a Dynamic Model of Technology Adoption: P2P Digital Payments,” *Working Paper*, 2023.
- Aral, Sinan**, “The Future of Weak Ties,” *American Journal of Sociology*, 2016, 121 (6), 1931–1939.
- Arrow, Kenneth J**, “The economic implications of learning by doing,” 1962.
- Athey, Susan, Dean Eckles, and Guido W. Imbens**, “Exact p-Values for Network Interference,” *Journal of the American Statistical Association*, 2018, 113 (521), 230–240.
- Atkin, David, Keith Chen, and Anton Popov**, “The Returns to Face-to-Face Interactions: Knowledge Spillovers in Silicon Valley,” *Working Paper*, 2024.
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson**, “The Diffusion of Microfinance,” *Science*, 2013, 341 (6144), 1236498.
- Baum-Snow, Nathaniel, Nicolas Gendron-Carrier, and Ronni Pavan**, “Local Productivity Spillovers,” *American Economic Review*, April 2024, 114 (4), 1030–69.
- Bernhardt, Arielle, Erica Field, Rohini Pande, and Natalia Rigol**, “Household Matters: Revisiting the Returns to Capital among Female Microentrepreneurs,” *American Economic Review: Insights*, September 2019, 1 (2), 141–60.
- Björkegren, Daniel**, “The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda,” *The Review of Economic Studies*, 07 2018, 86 (3), 1033–1060.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen**, “Identifying Technology Spillovers and Product Market Rivalry,” *Econometrica*, 2013, 81 (4), 1347–1393.
- Borusyak, Kirill and Peter Hull**, “Nonrandom Exposure to Exogenous Shocks,” *Econometrica*, 2023, 91 (6), 2155–2185.

- Brock, William A. and Steven N. Durlauf**, “Discrete Choice with Social Interactions,” *Review of Economic Studies*, 2001, 68 (2), 235–260.
- Broder, Andrei Z.**, “On the Resemblance and Containment of Documents,” in “Proceedings of the Compression and Complexity of Sequences” IEEE 1997, pp. 21–29.
- Cai, Jing, Alain De Janvry, and Elisabeth Sadoulet**, “Social Networks and the Decision to Insure,” *American Economic Journal: Applied Economics*, 2015, 7 (2), 81–108.
- **and Adam Szeidl**, “Interfirm Relationships and Business Performance*,” *The Quarterly Journal of Economics*, 12 2017, 133 (3), 1229–1282.
- Cefalà, Luisa, Franck Irakoze, Pedro Naso, and Nicholas Swanson**, “The Economic Consequences of Knowledge Hoarding,” *Working Paper*, 2025.
- Chamberlain, Gary**, “Panel Data,” *Handbook of Econometrics*, 1984, 2, 1247–1318.
- Chaurey, Ritam, Gaurav Nayyar, Siddharth Sharma, and Eric Verhoogen**, “Social Learning among Urban Manufacturing Firms: Energy-Efficient Motors in Bangladesh,” *Working Paper*, 2025.
- Cristofaro, Emiliano De, Mark Manulis, and Bertram Poettering**, “Private Discovery of Common Social Contacts,” *International Journal of Information Security*, 2013, 12 (1), 49–65.
- Crouzet, Nicolas, Apoorv Gupta, and Filippo Mezzanotti**, “Shocks and Technology Adoption: Evidence from Electronic Payment Systems,” *Journal of Political Economy*, 2023.
- **, Pulak Ghosh, Apoorv Gupta, and Filippo Mezzanotti**, “Demographics and Technology Diffusion: Evidence from Mobile Payments,” March 2024.
- Demmler, Daniel, Peter Rindal, Mike Rosulek, and Ni Trieu**, “PIR-PSI: Scaling Private Contact Discovery,” *Proceedings on Privacy Enhancing Technologies*, 2018, 2018 (4), 159–178.
- Duflo, Esther, Daniel Keniston, Tavneet Suri, and Celine Zipfel**, “Chat Over Coffee? Technology Diffusion Through Social and Geographic Networks in Rwanda,” Technical Report 31368 2026.
- Fafchamps, Marcel, Måns Söderbom, and Monique van den Boogart**, “Adoption with Social Learning and Network Externalities,” *Oxford Bulletin of Economics and Statistics*, 2023. Open Access.
- Gouriéroux, Christian, Alain Monfort, and Eric Renault**, “Indirect Inference,” *Journal of Applied Econometrics*, 1993, 8, S85–S118.
- Granovetter, Mark**, “The Strength of Weak Ties,” *American Journal of Sociology*, 1973, 78, 1360–1380.

- Grossman, Gene M. and Elhanan Helpman**, “Trade, knowledge spillovers, and growth,” *European Economic Review*, 1991, 35 (2-3), 517–526.
- Higgins, Sean**, “Financial Technology Adoption: Network Externalities of Cashless Payments in Mexico,” *American Economic Review*, November 2024, 114 (11), 3469–3512.
- Houeix, Deivy**, “Asymmetric Information and Digital Technology Adoption: Evidence from Senegal,” *Working Paper*, 2025.
- Jack, William and Tavneet Suri**, “Risk Sharing and Transactions Costs: Evidence From Kenya’s Mobile Money Revolution,” *American Economic Review*, 2014, 104, 183–223.
- Jaffe, Adam B**, “Technological opportunity and spillovers of R&D: Evidence from firms’ patents, profits, and market value,” *The American Economic Review*, 1986, 76 (5), 984–1001.
- Jahani, Eaman, Samuel P. Fraiberger, Michael Bailey, and Dean Eckles**, “Long ties, disruptive life events, and economic prosperity,” *Proceedings of the National Academy of Sciences*, 2023, 120 (28), e2211062120.
- Lee, David S.**, “Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects,” *Review of Economic Studies*, 2009, 76 (3), 1071–1102.
- Liben-Nowell, David and Jon Kleinberg**, “The Link-Prediction Problem for Social Networks,” *Journal of the American Society for Information Science and Technology*, 2007, 58 (7), 1019–1031.
- Manski, Charles F.**, “Identification of Endogenous Social Effects: The Reflection Problem,” *The Review of Economic Studies*, 1993, 60 (3), 531–542.
- Marshall, Alfred**, “Principles of Economics,” *MacMillan*, 1920.
- Moretti, Enrico**, “The Effect of High-Tech Clusters on the Productivity of Top Inventors,” *American Economic Review*, October 2021, 111 (10), 3328–75.
- Myers, Kyle R. and Lauren Lanahan**, “Estimating Spillovers from Publicly Funded RD: Evidence from the US Department of Energy,” *American Economic Review*, July 2022, 112 (7), 2393–2423.
- Newey, Whitney K. and Daniel McFadden**, “Large Sample Estimation and Hypothesis Testing,” *Handbook of Econometrics*, 1994, 4, 2111–2245.
- Noll, Landon Curt, Kiem-Phong Vo, Donald E. Eastlake, and Tony Hansen**, “The FNV Non-Cryptographic Hash Algorithm,” RFC 9923, RFC Editor February 2026.
- Romer, Paul M.**, “Endogenous technological change,” *Journal of Political Economy*, 1990, 98 (5, Part 2), S71–S102.
- Suri, Tavneet**, “Mobile Money,” *Annual Review of Economics*, 2017, 9, 497–520.

– **and Chris Udry**, “Agricultural Technology in Africa,” *VoxDev Vol 5, Issue 2*, March 2024.

Swamidass, S. Joshua, Matthew Matlock, and Leon Rozenblit, “Securely Measuring the Overlap between Private Datasets with Cryptosets,” *PLoS ONE*, 2015, 10 (2), e0117898.

Main Figures and Tables




Figure 1: Digital Payments' Adoption in 2022 - *Author's calculations*

Notes: This figure illustrates the cumulative adoption of digital payment technology by businesses in Senegal throughout 2022, covering both taxi businesses and other industries. The adoption rate follows an S-shaped curve for both other businesses and taxi businesses. The experiment commenced on March 28, 2022, when the first taxi businesses were equipped with the technology, with the last taxi business on-boarded by the end of April 2022. A week later, taxis outside the experiment were permitted to adopt the digital payment technology, while the control group remained on a waitlist for the time of the experiment. Throughout the year, all other businesses also had the opportunity to adopt the technology. The first end-line survey was conducted on October 5, 2022, and the experiment randomizing access to digital payment technology concluded in late December 2022. By early 2024, approximately 17,000 taxis and 220,000 other businesses had adopted the technology, suggesting that the adoption rate as of October 2022 was about 75-80% of the potential business population.



Figure 2: Digital Payments for Taxis - Sénégal

Notes: These images illustrate the implementation of the technology within taxi vehicles. The technology utilizes a QR code visible from outside the vehicle, allowing potential customers to easily identify it. Drivers are provided with a sticker to place on their windshield, enhancing the visibility of the product. Additionally, drivers were equipped with the corporate mobile application and trained on how to use it by a dedicated mobile money agent. Passengers have the option to pay using their smartphones, or drivers can scan passengers' mobile money cards when they do not have smartphones. Businesses in different sectors were provided with the same QR code, which was presented as a sticker for display in their stores.




Figure 3: Illustration of the Experimental Design

Notes: This figure illustrates and simplifies the experimental strategy. Some nodes, depicted as black dots, entered the experiment and were randomly assigned to either the treatment or control group. The study aims to quantify the likelihood of technology adoption among all users connected to an experimental node, whether treated or control, essentially comparing the adoption rates among the networks of treated versus control taxi businesses. In practice, networks often overlap, leading to network interference, an issue that is discussed in the estimation section 3.1.




Figure 4: Garage proximity spillovers: effect size by distance

Notes: From Table 3, Panel B, social ties. Effect size = $\hat{\beta}_1/\bar{y}$, where $\hat{\beta}_1$ is the coefficient on AnyTreated and $\bar{y}$ is the control mean adoption rate. 95% CI = $(\hat{\beta}_1 \pm 1.96 \times SE)/\bar{y}$, with cluster-robust SEs by garage. Stars denote RI p -values. The competition effect decays sharply with distance: $-1,291\%$ at 75 m to -72% at 750 m, consistent with within-garage competitive crowding out.

Table 1: Network Spillovers by Tie Strength

	Taxi, Dakar (within)	Cross-ind., Dakar (across, local)	Cross-ind., Other (across, distant)
<i>Panel A: Mutual ($\geq 5\%$ overlap)</i>			
Share treated	−0.322 (0.276) [0.394]	+ 2.933 ** (1.174) [0.030]	+ 1.103 ** (0.490) [0.034]
Expected exposure	−0.472 (9.011) [0.378]	−32.470 (38.291) [0.818]	+1.778 (15.979) [0.179]
Effect size	−20.5%	+10.1%	+22.3%
Control mean	1.57	29.16	4.96
<i>N</i>		126,738	
<i>Panel B: Mutual ($< 5\%$ overlap)</i>			
Share treated	− 0.143 ** (0.055) [0.020]	−0.017 (0.185) [0.935]	+0.016 (0.063) [0.774]
Expected exposure	−2.305 (1.753) [0.195]	−7.626 (5.879) [0.437]	−2.712 (2.001) [0.535]
Effect size	−8.4%	−0.1%	+0.7%
Control mean	1.71	19.57	2.23
<i>N</i>		3,900,687	

Notes: Coefficients, standard errors, and control means are multiplied by 1,000 for readability. Each cell reports the OLS coefficient on share of treated drivers among individual j 's mutual phone contacts. Specification: $y_j = \beta_1 \cdot \text{ShareTreated}_j + \beta_2 \cdot \text{ExpectedExposure}_j + \alpha$, where ExpectedExposure is the average of ShareTreated across 1,000 treatment assignment draws. Standard errors in parentheses; randomization inference p -values from 1,000 permutations in brackets. Effect size = $\beta_1/\bar{y}$, expressed as percent change. All panels use network fingerprint 240222, keeping only the phone with the most contacts per individual (using the operator's identity mapping). "Top 2% both" trimming applied: individuals and drivers in the top 2% of the degree distribution are excluded. Panel A restricts to pairs with raw overlap share $\geq 5\%$, the threshold that minimizes mean absolute error between inferred and actual contact counts (Section 1.3); Panel B restricts to pairs with raw overlap $< 5\%$. * $p_{\text{RI}} < 0.10$, ** $p_{\text{RI}} < 0.05$, *** $p_{\text{RI}} < 0.01$.

Table 2: Spillover Effects by Relationship Type (Kinship Survey)

	Taxi adoption, Dakar		Cross-ind. adoption		
Subnetwork	Coef	Eff. size	Coef	Eff. size	<i>N</i>
<i>Panel A: Taxi contacts by social proximity</i>					
Taxi × Friends	9.84*** (3.75) [0.007]	0.361			614
Taxi × Non-friends	0.88 (3.94) [0.803]	0.028			549
Taxi × Freq. daily/wk	6.95** (3.04) [0.023]	0.243			916
<i>Panel B: All contacts by social proximity</i>					
Friends	3.04** (1.41) [0.025]	0.316	0.09 (1.12) [0.920]	0.013	1,983
Freq. daily/weekly	1.53* (0.85) [0.055]	0.216	0.32 (0.68) [0.630]	0.065	4,119
<i>Panel C: Casual ties</i>					
Clients	0.46 (1.23) [0.716]	0.368	3.51 (2.27) [0.134]	0.937	395
Serendipitous encounters	0.29 (0.38) [0.457]	0.164	1.80*** (0.59) [0.004]	0.479	5,112

Notes: Each row estimates $y_j = \alpha + \beta_1 \cdot \text{NumTreated}_{s,j} + \beta_2 \cdot \text{ExpectedExposure}_{s,j} + \varepsilon_j$ on the subsample with ≥ 1 connection in subnetwork s , restricted to baseline survey contacts. Coefficients $\beta_1 \times 100$ reported; HC1 robust standard errors in parentheses. Effect size = $\beta/\bar{y}_{\text{control}}$. “Taxi adoption, Dakar” = within-industry adoption; “Cross-ind. adoption” = across-industry adoption (Dakar + outside). Stars based on randomization inference (RI) p -values from 1,000 permutations of treatment assignment; conventional t -test p -values are similar throughout. Panel A shows within-industry outcomes only, as taxi contacts have no cross-industry adoption variation. “Serendipitous encounters” are contacts classified as casual or unknown relationship. * $p_{\text{RI}} < 0.10$, ** $p_{\text{RI}} < 0.05$, *** $p_{\text{RI}} < 0.01$.

Table 3: Garage Proximity Spillovers — Within-Taxi Adoption

	50 m	75 m	100 m	150 m	300 m	500 m
<i>Panel A: All merchants near garages</i>						
Any treated	−0.0104 (0.0035)	−0.0068 (0.0055)	−0.0074 (0.0038)	−0.0049 (0.0022)	−0.0036* (0.0014)	−0.0016 (0.0010)
p_{RI}	[0.107]	[0.356]	[0.189]	[0.178]	[0.084]	[0.275]
Effect size	−136%	−104%	−141%	−120%	−111%	−58%
Mean $\bar{y}$ (%)	0.76	0.66	0.52	0.41	0.32	0.27
N	48,691	64,217	119,250	219,530	588,403	1,247,597
<i>Panel B: Social ties (mutual phone contacts)</i>						
Any treated	−0.1811** (0.0198)	−0.1819*** (0.0198)	−0.0666** (0.0377)	−0.0222* (0.0169)	−0.0138* (0.0114)	−0.0085** (0.0057)
p_{RI}	[0.014]	[0.008]	[0.031]	[0.067]	[0.061]	[0.045]
Effect size	−1089%	−1291%	−615%	−245%	−242%	−162%
Mean $\bar{y}$ (%)	1.66	1.41	1.08	0.91	0.57	0.53
N	361	426	646	1,104	2,453	4,948

Notes: Each column estimates $y_j = \beta_1 \cdot \text{AnyTreated}_j + \alpha$ for non-experimental merchants within the specified radius of experimental taxi garages. Outcome: within-industry taxi adoption in Dakar. Panel A includes all merchants; Panel B restricts to merchants who are mutual phone contacts of experimental drivers. The distance gradient across Panel B bins identifies the structural parameters δ and $\beta_{\text{comp}}^{\text{geo}}$ (moment m7 in Table 7). Effect size = $\beta_1 / \bar{y}$. SEs in parentheses; RI p -values in brackets (1,000 permutations). * $p_{RI} < 0.10$, ** $p_{RI} < 0.05$, *** $p_{RI} < 0.01$.

Table 4: Spillovers via Shared Customers (≥ 2 transactions)

	Within-taxi adoption (Dakar)			Cross-industry adoption		
	(1) All	(2) Social	(3) Non-Soc.	(4) All	(5) Social	(6) Non-Soc.
Share treated	420.76*** (39.69) [0.001]	922.55** (206.11) [0.023]	403.07*** (40.21) [0.001]	13.06 (7.88) [0.444]	165.07 (119.69) [0.190]	12.14 (7.89) [0.485]
Expected exposure	3426.86 (1423.32) [0.019]	11981.82 (10396.72) [0.652]	3258.97 (1432.55) [0.018]	45.50 (247.99) [0.362]	2262.55 (3837.39) [0.250]	46.67 (248.46) [0.366]
Effect size	+57.1%	+120.9%	+54.9%	+2.1%	+30.7%	+1.9%
Dep. var. mean	736.48	763.09	734.31	626.72	537.13	628.65
N	10,151	764	9,387	49,151	1,037	48,114

Notes: Coefficients, SEs, and dep. var. means multiplied by 1,000. Sample: non-experimental merchants linked to non-garage experimental drivers through shared customers who made ≥ 2 pay-to-business (P2B) transactions with any non-garage experimental driver between November 2022 and November 2023. Columns (1)–(3): taxi merchants (TAXI, TAXI.PILOT), dep. var. = within-taxi adoption (Dakar). Columns (4)–(6): non-taxi merchants, dep. var. = cross-industry adoption (Dakar or outside). “Social” = merchants with mutual phone contacts ($I_{\text{obs}} \geq 2$) to experimental drivers; “Non-Social” = without. Specification: $y_j = \beta_1 \cdot \text{ShareTreated}_j + \beta_2 \cdot \text{ExpectedExposure}_j + \alpha$. Effect size = $\beta/\bar{y}$. HC1 robust SEs in parentheses; RI p -values (1,000 draws) in brackets. * $p_{\text{RI}} < 0.10$, ** $p_{\text{RI}} < 0.05$, *** $p_{\text{RI}} < 0.01$.

Table 5: Cross-Industry Adoption Across Space ($\geq 5\%$ Overlap)

	Cross-Industry Adoption (Outside Dakar)		
	(1) All	(2) Visited Dakar	(3) Did Not Visit
Share treated	+0.1103** (0.0447) [0.034]	+0.0533 (0.0338) [0.145]	+0.3179* (0.1544) [0.084]
Expected exposure	+0.1778 (1.4432)	+0.1667 (1.0961)	-1.0355 (4.9109)
Control mean	0.34	0.07	1.12
Effect size	+32.2%	+72.5%	+28.4%
Avg. connections	3.39	3.91	3.02
<i>N</i>	126,738	66,186	32,211

Notes: Coefficients, SEs, and control means multiplied by 100. Network restricted to mutual phone contacts with $\geq 5\%$ overlap share (Table 1 Panel A). Top 2% both trimming applied. Column (1): all individuals in the strong-tie sample. The coefficient matches Table 1 Panel A cross-industry other in raw terms; the control mean and effect size differ because this table conditions on the geolocation-matched subsample and uses city-specific adoption as the outcome. Column (2): individuals who appeared in Dakar at least once during the experiment (March–October 2022) per mobile money geolocation data. Column (3): individuals who never appeared in Dakar. Individuals without geolocation data are included in column (1) but excluded from columns (2) and (3). Specification: $y_j = \beta_1 \cdot \text{ShareTreated}_j + \beta_2 \cdot \text{ExpExposure}_j + \alpha$. HC1 robust SEs in parentheses. RI *p*-values (1,000 draws) in brackets. * $p_{\text{RI}} < 0.10$, ** $p_{\text{RI}} < 0.05$, *** $p_{\text{RI}} < 0.01$.

Table 6: Cross-Industry Adoption by City: Migrants vs Non-Migrants ($\geq 5\%$ Overlap)

	Migrants to City				Non-Migrants to City			
	(1) Mbour	(2) Thiès	(3) Touba	(4) St-Louis	(5) Mbour	(6) Thiès	(7) Touba	(8) St-Louis
Share treated	+0.5967** (0.2584) [0.024]	+0.3546* (0.1968) [0.076]	+0.3121 (0.2036) [0.148]	+1.1381** (0.4013) [0.011]	-0.0041 (0.0057) [0.505]	+0.0003 (0.0073) [0.969]	-0.0058 (0.0060) [0.354]	+0.0044** (0.0031) [0.044]
Expected exposure	+0.3221 (8.6134)	-7.4089 (6.7311)	-4.6233 (6.7316)	-7.0307 (12.4615)	+0.0363 (0.1712)	+0.1190 (0.2295)	+0.2643 (0.1837)	+0.1760 (0.1527)
Control mean	0.62	0.51	0.61	0.67	0.01	0.01	0.01	0.00
Effect size	+96.9%	+70.2%	+51.0%	+169.3%	-42.3%	+2.8%	-58.6%	—
<i>N</i>	8,775	10,894	11,847	4,674	117,963	115,844	114,891	122,064

Notes: Coefficients, SEs, and control means multiplied by 100. Network restricted to mutual phone contacts with $\geq 5\%$ overlap share (Table 1 Panel A). Top 2% both trimming applied. “Migrants to City” = individuals who appeared in that city at least once during the experiment (March–October 2022) per mobile money geolocation data. “Non-Migrants” = all other individuals. Outcome: cross-industry adoption in the specific city (outside Dakar). Specification: $y_j = \beta_1 \cdot \text{ShareTreated}_j + \beta_2 \cdot \text{ExpExposure}_j + \alpha$. HC1 robust SEs. RI *p*-values (1,000 draws) in brackets. * $p_{\text{RI}} < 0.10$, ** $p_{\text{RI}} < 0.05$, *** $p_{\text{RI}} < 0.01$.

Table 7: Structural Parameter Estimates (7-Parameter Model)

		Estimate	SE			
<i>Panel A: Structural parameters</i>						
$\hat{\alpha}_{\text{cross}}$	(cross-industry intercept)	−5.311***	(0.012)			
$\hat{\alpha}_{\text{taxi}}$	(within-taxi intercept)	−7.651	(6.110)			
$\hat{\beta}_{\text{learn}}^s$	(strong-tie learning, $\geq 5\%$)	+0.097**	(0.040)			
$\hat{\beta}_{\text{learn}}^w$	(weak-tie learning, $< 5\%$)	−0.001	(0.006)			
$\hat{\beta}_{\text{comp}}^{\text{geo}}$	(geographic competition)	−6.279***	(0.36)			
$\hat{\beta}_{\text{demand}}$	(customer demand)	+1.335***	(0.12)			
$\hat{\delta}$	(distance decay, km^{-1})	+3.73***				
Competition half-life			186 m			
#	Source	Channel	$\hat{b}$ (data)	Model pred.	Sign	
<i>Panel B: Identifying moments (exactly matched)</i>						
m1	T1 Panel A, cross-Dkr	β_1^s	+0.00293	+0.00293	—	
m2	T1 Panel B, cross-Dkr	β_1^w	−0.00002	−0.00002	—	
m3	T1, taxi (all overlap)	$w_s\beta_1^s + w_w\beta_1^w$	−0.00014	−0.00014	—	
m4	Table 4, within-taxi cust.	β_d	+0.421	+0.421	—	
m5	Cross-ind. rate	α_{cross}	2.916%	2.916%	—	
m6	Within-taxi rate	α_{taxi}	0.171%	0.171%	—	
m7	Table 3, distance gradient	$\delta; \beta_c^{\text{geo}}$	slope = −3.73***	$R^2 = 0.93$	—	
<i>Panel C: Out-of-sample validation (untargeted predictions)</i>						
m8	T2, all friends (within-taxi)	β_1^s	+0.03044	+0.00152	✓	
m9	T2, taxi friends	$\beta_1^s + \beta_c^{\text{geo}} + \beta_d$	+0.09845	+0.00330	✓	
m10	T2, taxi non-friends	$\beta_1^w + \beta_c^{\text{geo}} + \beta_d$	+0.00875	+0.00297	✓	
m11	T2, serendipitous	β_1^w	+0.01800	+0.00002	✓	
Sign accuracy				4/4		

Notes: Minimum distance estimation with indirect inference. Seven parameters identified from six scalar moments (m1–m6) plus the distance gradient of six garage coefficients (m7), which provides four overidentifying restrictions ($R^2 = 0.93$). Learning parameters ($\beta_{\text{learn}}^s, \beta_{\text{learn}}^w$) identified from cross-industry phone-network moments via indirect inference (OLS on model-predicted adoption). Customer demand (β_{demand}) identified from within-taxi customer regression (Table 4, column 1) via average marginal effect. Distance decay (δ) and geographic competition ($\beta_{\text{comp}}^{\text{geo}}$) identified from the log-linear gradient of garage coefficients across six distance bins (Table 3). Standard errors by delta method with numerical Jacobian. The large SE on $\hat{\beta}_{\text{comp}}^{\text{geo}}$ reflects the logistic-index scale: the implied reduced-form garage coefficients are precisely estimated ($p < 0.01$ at 50 m). Panel C: within-taxi validation moments (m9, m10) load on all three channels—learning, geographic competition, and customer demand—because taxi contacts may share customers and be located near garages; the model correctly predicts positive signs for both, driven by demand dominating competition. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Counterfactual Targeting: Combined Reachable Network ($N = 168,964$)

	$\bar{S}^{\text{phone}}$	$\bar{S}^{\text{demand}}$	Cross-Ind. (Dakar)		Cross-Ind. (Other)		Within-Taxi		Total	
			Induced	Δ	Induced	Δ	Induced	Δ	Induced	Δ
No treatment (baseline)	0.000	0.000	<i>Baseline: 25 taxi + 5,375 cross-Dakar + 5,280 cross-other = 10,680</i>							
Actual random ($N_T = 1, 256$)	0.569	0.538	146	—	10	—	44	—	167	—
All-garage ($N_T = 1, 256$)	0.499	0.555	127	−12.6%	8	−12.1%	9	−25.9%	145	−13.5%
Top MM users ($N_T = 1, 256$)	0.660	0.588	172	+18	11	+10.3%	51	+14	197	+17.7%
Highest degree ($N_T = 1, 256$)	0.898	0.828	231	+59	14	+46.7%	84	+89	270	+61.7%
Geographic reach ($N_T = 1, 256$)	0.888	0.819	228	+57	14	+48.5%	83	+86	267	+59.9%

Notes: Treatment-induced adoption = predicted 6-month cumulative adoption – no-treatment baseline. Competition uses nearest-garage formulation $S_j^{\text{geo}} = \text{ShareTreated}_{\text{nearest}} \times \exp(-\hat{\delta} \times d)$ with $\hat{\delta} = 3.73^{***} \text{ km}^{-1}$ (half-life = 186 m). “All-garage” reweights the 228 observed garage drivers by $1, 256/228 \approx 5.5$ (clipped at 1.0). All three channels change: $\bar{S}^{\text{phone}}$ falls to 0.499 (fewer contacts), $\bar{S}^{\text{demand}}$ falls to 0.555 (fewer customer links), and $\text{ShareTreated}_g = 1$ at every garage. Δ = % change vs. actual random. Per-margin columns do not sum to the total because some individuals belong to multiple populations (e.g., a taxi merchant in Dakar with a customer link contributes to both cross-Dakar and within-taxi); the total counts each individual once.

Online Appendix

Nationwide Diffusion of Technology: Experimental Evidence from Multiple Networks

Deivy Houeix

A	Appendix Figures	51
A.1	Fingerprint Data Processing	58
B	Appendix Tables	59
C	Estimation Details	71
D	Validation Details	72
E	Counterfactual Details	73
F	Distance Decay of Competition	74

A Appendix Figures




Figure A1: Smartphone Adoption in Sub-Saharan Africa

Notes: This graph is sourced from the GSMA Intelligence 2023 report, “The Mobile Economy Sub-Saharan Africa 2023,” on page 15 of 46. The report gathers data and forecasts smartphone adoption across several African countries. Côte d’Ivoire, which is similar to Senegal in size and geography, reports a 60% smartphone adoption rate for the entire country, a statistic that is likely much higher in urban areas.

Figure A2: Experimental Design

Notes: This figure outlines the experimental design. Each taxi driver was initially listed at various locations including garages, car wash stations, meeting points, and on the streets of Dakar (e.g., during traffic jams). The listing was comprehensive enough to cover most parts of the city, and a large portion of the taxi drivers working in the city were approached (it is estimated that about 30% of the drivers were approached, with some refusing to be surveyed, others not interested, and some not eligible, as detailed in Section 1.4.1). Subsequently, taxi owners and drivers were invited to specific locations in Dakar to be surveyed and randomized into the treatment. The primary experimental variation involved randomizing access to the technology. The treatment was then subdivided into three different “observability” treatment arms for owners to monitor their drivers’ transactions in various ways. For this paper, these sub-treatments are pooled together as most of the sample considered here consisted of taxi businesses operated solo.




Figure A3: Listing Survey Activity in Dakar - March/April 2022

Notes: The figure displays a map of Dakar, capital city of Senegal. Each blue dot is the GPS location of the listing survey conducted before the experiment. Drivers were recruited in garages, car wash stations, meeting points, and on the streets of Dakar (e.g., during traffic jams). The extent of the listing was so broad that it covers most part of the city and a large portion of the taxi drivers working in the city were approached (among which some refused to be surveyed, some were not interested, some were not eligible, as described in Section 1.4.1). Owners not driving were primarily recruited by asking drivers about their owners' contact information during the listing survey.




Figure A4: Mobile Money Agents in Senegal in 2022

Notes: This Figure represents a map of Senegal. Each green dot is the GPS location of a mobile money agent for the main mobile money platform. There are over 15,000 unique agents in Senegal from this mobile money company, which covers the full set of urban cities and the extreme majority of communes/villages.




Figure A5: Location Data - Arrivals in Touba

Notes: This Figure represents the number of arrivers in Touba over each month from February 2021 to December 2022, as measured via the withdrawal/deposit of mobile money to agents located in the city of Touba. There are peaks of arrivals in the months of September, which is directly consistent with the pilgrimage dates happening every year in this city.




Figure A6: Technology Diffusion in Dakar Throughout the Experiment

Notes: This map shows the technology diffusion, for a randomly selected sample of businesses, across the capital city of Dakar with the firm's sectors. Each colored dot shows the location of a business who adopted the digital payment technology during the time of the experiment.




Figure A7: Technology Diffusion Inter-Industry Across Space Throughout the Experiment

Notes: This map shows the technology diffusion, for a randomly selected sample of businesses, across many different urban cities in Senegal, beyond the capital city of Dakar, where the experiment took place. Each blue dot shows the location of a business who adopted the digital payment technology during the time of the experiment.




Figure A8: Local Technology Adoption Spillovers Nearby Taxi Garages

Notes: This figure displays a map of Dakar, pinpointing the geo-locations of a subset of taxi garages where at least one taxi driver participated in the experiment. Each garage was manually coded to define a neighborhood. Mobile money agents within a 200-meter radius were identified and tagged, in this illustration in purple. For these tagged agents, I compiled a dataset of inhabitants living in proximity to these agents to delineate the residents of each neighborhood. For these residents, I then analyzed their propensity to adopt the technology depending on the number of treated taxi businesses and the total number of experimental taxi businesses within their neighborhood.

A.1 Fingerprint Data Processing

The phone contact network used throughout the paper is constructed from encrypted contact-book fingerprints collected by the mobile money operator. This appendix documents the processing steps from raw fingerprint data to the analysis-ready edge list.

Step 1: Fingerprint extraction. When a user first logs into the mobile money application, the app requests access to the phone’s contact list (consent rate $\approx 86\%$). For consenting users, each saved phone number is hashed into one of $2^{16} = 65,536$ buckets using a 16-bit FNV-1a hash (Noll et al., 2026). The resulting fingerprint is a set of unique bucket values representing the user’s contact book. The February 2024 snapshot covers approximately 7 million unique users.

Step 2: Matching experimental drivers. Each of the 2,196 experimental drivers has up to three registered phone numbers. We search the fingerprint file for any number matching any driver, keeping the fingerprint with the most buckets when a driver has multiple records. Of the 2,150 drivers with a mobile money account match (`has_match`), 1,918 (87.3%) have at least one phone number with fingerprint data (Section 1.3). The remaining 278 drivers were never fingerprinted by the application.

Step 3: Minimum contact threshold. Drivers with fewer than 10 unique hash buckets in their fingerprint are excluded, as the overlap computation is dominated by hash collision noise at such small set sizes. This removes 11 additional drivers, leaving 1,907 in the network analysis sample.

Step 4: Edge detection. For each fingerprinted user (“individual j ”), an edge to driver i is formed if and only if both parties have saved each other’s number: driver i ’s self-hash bucket appears in individual j ’s fingerprint, AND individual j ’s self-hash bucket appears in driver i ’s fingerprint. This mutual-contact requirement, described in Section 1.3, removes one-sided or noisy links.

Step 5: Overlap share computation. For each edge, the overlap share is computed as:

$$\text{OverlapShare}_{ij} = \frac{\hat{I}_{ij}}{\min(|F_i|, |F_j|)},$$

where $\hat{I}_{ij}$ is the bias-corrected intersection. This normalization ensures that users with larger contact books are not mechanically assigned higher overlap. The raw (uncorrected) overlap share is also stored for robustness.

Step 6: Degree trimming. Individuals and drivers in the top 2% of the degree distribution are excluded from the analysis to reduce the influence of hub nodes whose high

connectivity may reflect institutional roles (e.g., agents, dispatchers) rather than personal social ties. This “top 2% both” trimming is applied in all main specifications.

Final analysis sample. After all processing steps, the network contains 12.4 million mutual-contact edges linking 1,907 experimental drivers to approximately 3.95 million individuals across Senegal. Of these, 1.05 million edges (8.5%) involve strong ties ($\geq 5\%$ raw overlap). The remaining 11.3 million edges are weak ties ($< 5\%$ overlap).

B Appendix Tables

Table B1: Cross-Validating Social Networks via Encrypted Phone Contacts

	(1) Wife	(2) Family	(3) Friends	(4) Serendipitous Connections	(5) Taxi	(6) Taxi Friends	(7) Taxi Clients	(8) Daily/Weekly Interactions	(9) Rare Interactions	(10) Shared Information	(11) Discussed Technology
Overlap	17.63	11.31	6.69	3.46	9.93	8.57	3.90	9.63	4.33	9.17	10.70
Share (%)	(17.30)	(10.97)	(6.42)	(4.39)	(8.50)	(7.04)	(3.74)	(9.76)	(4.41)	(9.42)	(9.72)
Share $\geq$ 5%	69.3	59.4	36.6	9.5	57.6	52.8	10.1	50.8	17.5	48.5	59.5
Overlap (%)	(46.2)	(49.1)	(48.2)	(29.3)	(49.4)	(50.0)	(30.1)	(50.0)	(38.1)	(50.0)	(49.1)
Observations	462	3,383	2,473	4,705	1,539	777	446	5,380	587	5,480	1,098

Notes: This table cross-validates the phone contact overlap measure against survey-reported relationship types. The endline survey asked each experimental driver about up to 16 randomly selected mobile money contacts, eliciting relationship type and interaction frequency. For each (driver, contact) pair, we compute phone contact overlap directly from the encrypted fingerprint data and report the mean overlap share and the fraction of pairs exceeding the 5% corrected-overlap threshold, by relationship category. Standard deviations in parentheses.

Column definitions: (1) **Wife:** respondent's spouse. (2) **Family:** any family member (including siblings, in-laws, and extended family). (3) **Friends:** non-family, non-professional acquaintances identified as friends. (4) **Serendipitous Connections:** contacts the respondent could not recall by name when prompted—likely incidental or one-time transaction partners encountered through mobile money. (5) **Taxi:** contacts identified as taxi drivers. (6) **Taxi Friends:** contacts identified as both taxi drivers and friends. (7) **Taxi Clients:** contacts identified as clients of the respondent's taxi service. (8) **Daily/Weekly Interactions:** contacts the respondent reports interacting with at least weekly. (9) **Rare Interactions:** contacts the respondent interacts with less than monthly or never. (10) **Shared Information:** contacts with whom the respondent reports sharing general information. (11) **Discussed Technology:** contacts with whom the respondent specifically discussed mobile money or other technology.

Overlap measures: Overlap Share is the raw fraction of shared phone contacts ($|A \cap B| / \min(|A|, |B|)$, where A and B are the two individuals' contact list fingerprints). Share $\geq$ 5% Overlap reports the fraction of pairs whose bias-corrected overlap exceeds 5%, where the correction subtracts expected random hash collisions ($\mu = |A| \times |B| / 65,536$).

Table B2: Phone Contact Network Balance: Treatment vs. Control Drivers

	Treated	Control	Difference	SE	<i>p</i> -value
<i>Panel A: Contact book size</i>					
Contact book size (fingerprint)	585	614	-28.8	24.5	0.240
Mutual contacts detected	8,309	8,982	-673	352	0.056*
<i>Panel B: Contact composition</i>					
Share merchant (among mutual)	0.041	0.041	0.0000	0.0002	0.868
Share taxi (among mutual)	0.0034	0.0034	-0.0001	0.0000	0.244
<i>Panel C: Overlap distribution</i>					
Share overlap $\geq 5\%$ (among mutual)	0.0031	0.0024	0.0006	0.0008	0.438
Mean overlap (among mutual)	0.0039	0.0038	0.0000	0.0001	0.727
Median overlap (among mutual)	0.0007	0.0007	-0.0000	0.0000	0.550
Observations	1,907				
<i>Kolmogorov–Smirnov tests</i>					
Contact book size (fingerprint)	KS = 0.0632			<i>p</i> = 0.0473	
Mutual contacts detected	KS = 0.0587			<i>p</i> = 0.0791	
Mean overlap (among mutual)	KS = 0.0443			<i>p</i> = 0.3135	
Mutual contacts $\geq 5\%$ overlap	KS = 0.0256			<i>p</i> = 0.9125	
Share taxi (among mutual)	KS = 0.0544			<i>p</i> = 0.1248	
Joint <i>F</i> -test	<i>F</i> = 0.763			<i>p</i> = 0.5765	

Notes: Each row reports the mean for treated drivers, the mean for control drivers, the difference, heteroskedasticity-robust standard error, and *p*-value from a regression of the variable on treatment status. Contact book size (fingerprint) counts the number of 16-bit buckets in each driver’s phone contact fingerprint as of February 2024 (one bucket per saved contact). “Mutual pairs detected” counts all edges where both parties’ hashed numbers appear in each other’s fingerprints; this count is inflated by hash collisions at 16-bit resolution (see Section 1.3). “Mutual contacts $\geq 5\%$ overlap” restricts to pairs sharing at least 5% of contacts, filtering out collision-driven false matches. Taxi/non-taxi classification uses the merchant registry restricted to merchants registered before the fingerprint date (February 22, 2024), so that post-fingerprint merchant registrations—which could be treatment-induced—do not contaminate the classification. The phone network is observed post-experiment (February 2024) and is used as a proxy for pre-existing social proximity; see Section 1.3 for discussion. **p* < 0.10, ***p* < 0.05, ****p* < 0.01.

Table B3: Test for Treatment-Induced Tie Strengthening

	Treated	Control	Diff	<i>t</i> -stat	<i>p</i> -value
<i>Panel A: Overlap distribution (edge-level, 12.4M mutual edges)</i>					
Mean overlap	0.0242	0.0244	−0.0002		
Median overlap	0.0193	0.0194	−0.0001		
p75 overlap	0.0307	0.0305	+0.0002		
p99 overlap	0.0959	0.0997	−0.0038		
KS statistic	0.008		<i>p</i> < 0.001		
<i>Panel B: Composition of strong ties (≥ 5% overlap, driver-level)</i>					
Number of strong ties	668	749	−81	−0.51	0.608
Taxi contacts (strong)	0.79	0.85	−0.06	−0.22	0.825
Merchant contacts (strong)	32.1	35.3	−3.2	−0.57	0.571
Share taxi (strong)	0.070%	0.060%	+0.010%	+1.29	0.196
Share merchant (strong)	5.65%	5.61%	+0.05%	+0.64	0.521
Mean overlap (strong)	0.0738	0.0738	−0.0001	−0.28	0.781
<i>Panel C: Total ties (driver-level)</i>					
Mutual contacts	7,887	8,557	−671	−1.89	0.059
Strong ties (≥ 5%)	668	749	−81	−0.51	0.608
Weak ties (< 5%)	7,219	7,808	−589	−1.98	0.048
<i>Panel D: Overlap by contact type (KS tests)</i>					
Taxi contacts	KS = 0.015		<i>p</i> = 0.14		
Merchant contacts	KS = 0.007		<i>p</i> < 0.001		
Non-merchant contacts	KS = 0.008		<i>p</i> < 0.001		

Notes: Tests for whether treatment caused selective tie strengthening in the phone contact network, observed February 2024. Panel A compares the full overlap distribution across 12.4 million mutual edges; the KS statistic is mechanically significant due to sample size but the quantile differences are economically negligible (< 0.04 percentage points at every quantile) and go in the wrong direction for a strengthening story (treated overlap is weakly *lower*). Panel B compares the composition of strong ties ($\geq 5\%$ overlap) at the driver level; no dimension is significant at 10%. Panel C compares total tie counts; treated drivers have marginally *fewer* contacts, inconsistent with tie formation. Panel D tests overlap distributions separately by contact type; the taxi-contact test—the most relevant for treatment-induced strengthening—has $p = 0.14$. Welch *t*-tests with unequal variances; $N_T = 1,256$, $N_C = 894$.

Table B4: Lee Bounds for Differential Fingerprint Attrition

	Taxi, Dakar (within)	Cross-ind., Dakar (across, local)	Cross-ind., Other (across, distant)
<i>Panel A: Mutual ($\geq 5\%$ overlap)</i>			
Share treated	−0.322 [0.394]	+ 2.939 ** [0.030]	+ 1.125 ** [0.030]
Lee lower bound	−0.366	+2.963**	+0.997*
Lee upper bound	−0.297	+4.612***	+1.586***
<i>N</i>		126,799	
<i>Panel B: Mutual ($< 5\%$ overlap)</i>			
Share treated	− 0.148 ** [0.024]	−0.006 [0.971]	+0.018 [0.778]
Lee lower bound	−0.134**	+0.071	+0.033
Lee upper bound	−0.139**	+0.144	+0.011
<i>N</i>		3,671,918	

Notes: Tightened Lee bounds (Lee, 2009) for differential attrition from the experimental sample to the fingerprint network sample (treated: 88.1%, control: 85.1%). Bounds are computed within contact-book-size quartiles: within each quartile, excess treated drivers (35 total) are trimmed from the bottom (lower bound) or top (upper bound) of the within-cell degree distribution. Coefficients $\times 1,000$. RI p -values in brackets. The key result (Panel A, cross-industry Dakar) has bounds $[+2.96, +4.61]$, both significant, confirming that differential attrition does not drive the estimated spillovers. * $p_{RI} < 0.10$, ** $p_{RI} < 0.05$, *** $p_{RI} < 0.01$.

Table B5: Robustness: Alternative Overlap Thresholds

	Taxi, Dakar (within)	Cross-ind., Dakar (across, local)	Cross-ind., Other (across, distant)
<i>Panel A: $\geq 4.6\%$ overlap (median-ratio criterion)</i>			
Share treated	-0.112 [0.748]	+0.481 [0.695]	+0.358 [0.394]
<i>Panel B: $\geq 5.0\%$ overlap (min-MAE criterion, main specification)</i>			
Share treated	-0.322 [0.394]	+2.933** [0.030]	+1.103** [0.034]
<i>Panel C: $\geq 5.4\%$ overlap (mean-count criterion)</i>			
Share treated	-0.507 [0.200]	+1.166 [0.455]	+1.098* [0.070]

Notes: Coefficients $\times 1,000$. RI p -values in brackets. Each panel replicates Table 1 Panel A using a different overlap threshold from the calibrated range (Section 1.3): 4.6% minimizes the median ratio of inferred to actual contacts; 5.0% minimizes mean absolute error (main specification); 5.4% minimizes the difference in mean inferred vs. mean actual counts. Signs are stable across the range: within-taxi is negative (competition), cross-industry is positive (learning). The cross-industry effect outside Dakar is positive at all three thresholds and significant or marginally significant at 5.0% and 5.4%.

* $p_{\text{RI}} < 0.10$, ** $p_{\text{RI}} < 0.05$, *** $p_{\text{RI}} < 0.01$.

Table B6: Network Objects: Construction, Timing, and Use

Network	Construction	Observation period	Scale	Used in
<i>Panel A: Phone network (social ties)</i>				
Mutual phone contacts	User i saved j 's number AND j saved i 's number (encrypted fingerprints)	Feb 2024	12.4M edges, 3.95M individuals	Tables 1, 2; structural model (S^{phone}); counterfactuals
Strong ties ($\geq 5\%$)	Subset: overlap share $\geq 5\%$ of smaller contact book	Feb 2024	1.05M edges, 89,809 individuals	Table 1 (cross-industry); structural β_{learn}^s
Weak ties ($< 5\%$)	Subset: overlap share $< 5\%$	Feb 2024	11.3M edges, 3.8M individuals	Table 1 (all overlap); structural β_{learn}^w
<i>Panel B: Customer network (demand channel)</i>				
Pre-intervention customers (P2P)	Person-to-person transfers in taxi-fare range (500–3,000 FCFA), ≥ 2 transactions with same driver	Apr 2021– Mar 2022	—	Reduced-form Table 3 (pre-intervention specification)
Post-experiment customers (P2B)	Pay-to-business transactions, ≥ 2 with any non-garage XP driver	Nov 2022– Nov 2023	28,609 customers, 95K merchants, 3.4M edges	Reduced-form Table 3 (main); structural β_{demand} ; counterfactual S^{cust}
<i>Panel C: Geographic network (competition channel)</i>				
Nearest-garage proximity	Haversine distance from individual (imputed via driver-contact GPS average) to nearest garage	Cross-section	42 garages (33 with XP drivers), 168,964 individuals	Table 4; structural $\beta_{\text{comp}}^{\text{geo}}$, δ ; counterfactual S^{geo}
<i>Panel D: Transaction network (robustness)</i>				
P2P transaction ties	Person-to-person mobile money transfers, ≥ 1 transfer in either direction	Apr 2021– Mar 2022	870K individuals with tx	Robustness checks; pre-treatment placebo

Notes: This table summarizes the network objects used in the analysis. The phone network (Panel A) is the primary network for measuring social proximity. It is observed post-experiment (February 2024) and interpreted as a proxy for pre-existing social ties; Table B2 and Table B3 support this interpretation. The customer network (Panel B) links merchants and drivers through shared customers; the post-experiment P2B version is used in the structural model because it directly observes business payments (which did not exist pre-experiment for taxis). The geographic network (Panel C) captures competitive proximity to taxi garages. The transaction network (Panel D) is observed pre-experiment and provides robustness checks.

Table B7: Kinship Survey in the Transaction Network - Summary Statistics

	Share
<i>Panel A. Nature of the Relationships</i>	
Family	0.23 (0.16)
Friends	0.16 (0.15)
Any Business	0.34 (0.20)
Taxi Drivers	0.10 (0.11)
Taxi Drivers Same Garage	0.01 (0.04)
Business Owners	0.19 (0.15)
Both Personal and Professional	0.23 (0.17)
Casual Interactions - Do Not Remember	0.42 (0.23)
<i>Panel B. Frequency of Conversations in Past 3 Months</i>	
Every Day	0.37 (0.26)
Every Week	0.31 (0.23)
At Least Once a Month	0.21 (0.21)
Less Than Once a Month	0.06 (0.12)
Never	0.02 (0.07)
Would Borrow Money To If Needed	0.51 (0.29)
<i>Panel C. Information Sharing</i>	
Generally Share Important Advice/Information	0.74 (0.28)
Discussed or Recommended the Digital Payment Technology	0.12 (0.18)
Observations	1952

Notes: This table displays summary statistics about the kinship survey data collected from September to December 2023 among RCT-taxi businesses. The survey involved businesses successfully followed up with and who consented to participate in this kinship survey. Sixteen links were identified for each respondent. Eight links were randomly selected from the pre-experiment transaction network, while the other eight were chosen from transactions occurring during or after the experiment. Respondents were asked to define their relationship and several survey questions with each of these links.

Table B8: Transaction Network Spillovers

	Taxi, Dakar (within)	Cross-ind., Dakar (across, local)	Cross-ind., Other (across, distant)
Share treated	0.003677*** (0.001132)	0.004270** (0.001771)	−0.000357 (0.000511)
Expected exposure	0.065780** (0.032807)	0.015035 (0.049220)	0.025621* (0.013662)
Effect size	+0.142	+0.071	−0.079
Control mean	0.02594	0.05997	0.00451
N		73,102	

Notes: OLS coefficient on share of treated drivers among individual j 's pre-treatment financial transaction partners. Specification: $y_j = \beta_1 \cdot \text{ShareTreated}_j + \beta_2 \cdot \text{ExpectedExposure}_j + \alpha$. HC1 robust standard errors in parentheses. Effect size = $\beta/\bar{y}$. Transaction network constructed from pre-treatment mobile money transfer records (2020–2022Q1) between experimental drivers and all individuals. Edge filter: “Exclude incoming-only” retains edges where the driver sent at least one transfer to individual j . * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B9: Pre-Intervention P2P Customer Network (≥ 2 transactions, Apr 2021–Mar 2022)

	Within-taxi adoption (Dakar)		Cross-industry adoption	
	(1) All	(2) Non-Soc.	(3) All	(4) Non-Soc.
Share treated	121.34*** (23.51)	122.16*** (23.57)	−7.46 (13.14)	−7.44 (13.17)
Expected exposure	−196.24 (761.38)	−187.82 (761.85)	641.61 (421.51)	633.78 (422.60)
Effect size	+17.9%	+17.9%	−1.0%	−1.0%
Dep. var. mean	679.49	680.80	761.12	760.62
<i>N</i>	3,510	3,490	6,635	6,613

Notes: Coefficients, SEs, and dep. var. means multiplied by 1,000. Sample: non-experimental merchants linked to non-garage experimental drivers through shared customers identified from pre-intervention P2P transaction data (April 2021–March 2022, 500–3,000 FCFA per transaction, ≥ 2 transactions per customer). Merchants who adopted before March 2022 are excluded ($N = 1,785$). Matched (experimental) merchants are also excluded. Columns (1)–(2): taxi merchants, dep. var. = within-taxi adoption (Dakar). Columns (3)–(4): non-taxi merchants, dep. var. = cross-industry adoption (Dakar or outside). “Non-Social” excludes merchants with mutual phone contacts to experimental drivers. Only 20 taxi and 22 non-taxi merchants have mutual phone ties in this pre-intervention sample, so Social subsamples are omitted. HC1 robust SEs in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B10: Customer Network Temporal Stability Diagnostic

	W1 Nov 2022–Apr 2023	W2 May–Nov 2023
<i>Panel A: Non-Social subsample (no mutual phone ties)</i>		
ShareTreated ($\times 100$)	+27.28*** (2.33)	+6.59*** (2.51)
<i>N</i>	4,450	7,757
$\bar{y}$	0.784	0.744
<i>Panel B: All customer-linked merchants</i>		
ShareTreated ($\times 100$)	+28.18*** (2.33)	+6.47*** (2.50)
<i>N</i>	4,495	7,906

Notes: Each column constructs the shared-customer network G^{cust} from transactions within that window only (≥ 2 transactions per customer). Garage-affiliated drivers excluded. Outcome: within-taxi DKR mobile payment adoption by endline. Share-Treated = $\sum T_i / N_j^{\text{cust}}$ is the share of customer-linked experimental drivers assigned to treatment. Robust SEs in parentheses. The Non-Social subsample restricts to merchants with no mutual phone ties ($I_{\text{obs}} < 2$) with any experimental driver. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table B11: Spillovers via Shared Customers — Garage Drivers (≥ 2 transactions)

	Within-taxi adoption (Dakar)			Cross-industry adoption		
	(1) All	(2) Social	(3) Non-Soc.	(4) All	(5) Social	(6) Non-Soc.
Share treated	288.83 *** (31.46) [0.004]	822.93 * (178.02) [0.082]	272.07*** (31.91) [0.006]	-26.85* (8.36) [0.066]	-162.25 (124.73) [0.258]	-26.24* (8.37) [0.073]
Expected exposure	-1769.78 (1228.62) [0.049]	11500.64 (7187.00) [0.976]	-2049.02 (1239.14) [0.056]	467.95 (269.58) [0.599]	2012.04 (4589.93) [0.095]	464.16 (270.01) [0.615]
Effect size	+39.9%	+112.0%	+37.6%	-4.4%	-30.9%	-4.3%
Dep. var. mean	724.48	735.00	724.04	615.54	525.23	616.80
<i>N</i>	10,108	400	9,708	38,787	535	38,252

Notes: Coefficients, SEs, and dep. var. means multiplied by 1,000. Sample: non-experimental merchants linked to **garage** experimental drivers through shared customers (≥ 2 transactions, November 2022–November 2023). Columns (1)–(3): taxi merchants (TAXI, TAXI_PILOT), dep. var. = within-taxi adoption (Dakar). Columns (4)–(6): non-taxi merchants, dep. var. = cross-industry adoption (Dakar or outside). “Social” = merchants with mutual phone contacts to experimental drivers; “Non-Social” = without. Specification: $y_j = \beta_1 \cdot \text{ShareTreated}_j + \beta_2 \cdot \text{ExpectedExposure}_j + \alpha$. Effect size = $\beta/\bar{y}$. HC1 robust SEs in parentheses; RI p -values (1,000 draws) in brackets. * $p_{\text{RI}} < 0.10$, ** $p_{\text{RI}} < 0.05$, *** $p_{\text{RI}} < 0.01$.

Table B12: Garage Proximity Spillovers — Within-Taxi Adoption (Share-Treated)

	50 m	75 m	100 m	150 m	300 m	500 m
<i>Panel A: All merchants near garages</i>						
Share treated	−0.0037 (0.0056)	−0.0015 (0.0054)	−0.0035 (0.0044)	−0.0025 (0.0027)	−0.0020 (0.0019)	−0.0005 (0.0014)
p_{RI}	[0.663]	[0.845]	[0.499]	[0.463]	[0.348]	[0.740]
Effect size	−49%	−23%	−67%	−61%	−64%	−20%
Mean $\bar{y}$ (%)	0.76	0.66	0.52	0.41	0.32	0.27
N	48,691	64,217	119,250	219,530	588,403	1,247,597
<i>Panel B: Social ties (mutual phone contacts)</i>						
Share treated	−0.0877** (0.0486)	−0.0742** (0.0483)	−0.0275 (0.0234)	−0.0021 (0.0130)	−0.0033 (0.0083)	−0.0027 (0.0052)
p_{RI}	[0.011]	[0.014]	[0.119]	[0.856]	[0.649]	[0.552]
Effect size	−528%	−527%	−253%	−24%	−58%	−51%
Mean $\bar{y}$ (%)	1.66	1.41	1.08	0.91	0.57	0.53
N	361	426	646	1,104	2,453	4,948

Notes: Same specification as Table 3 but using ShareTreated (continuous: fraction of treated drivers at merchant j 's nearest garage) instead of AnyTreated (binary). Specification: $y_j = \beta_1 \cdot \text{ShareTreated}_j + \beta_2 \cdot \text{ExpectedShareTreated}_j + \alpha$, where ExpectedShareTreated is the mean of ShareTreated across 1,000 treatment assignment draws. Cluster-robust SEs (at garage level) in parentheses; RI p -values in brackets. * $p_{RI} < 0.10$, ** $p_{RI} < 0.05$, *** $p_{RI} < 0.01$.

C Estimation Details

Overview. The seven parameters are solved sequentially, exploiting the block structure of the moment conditions. All exposure variables are functions of the exogenous treatment assignment, so the model has a closed-form adoption probability—no cascade simulation is needed.

Learning ($\beta_{\text{learn}}^s, \beta_{\text{learn}}^w$). For each candidate β , I compute model-implied adoption probabilities $P_j(\theta) = 1 - (1 - \Lambda(\alpha + \beta \cdot S_j^{\text{phone}}))^6$ for all cross-industry individuals, then run the same OLS regression on P_j that produced the reduced-form Table 1 coefficients—an indirect inference step (Gouriéroux et al., 1993). The value of β that equates the model-predicted OLS coefficient to the data coefficient is found separately for the $\geq 5\%$ overlap subsample ($m1 \rightarrow \beta_{\text{learn}}^s$) and the $< 5\%$ subsample ($m2 \rightarrow \beta_{\text{learn}}^w$). Each is a one-dimensional root-finding problem (Brent’s method).

Customer demand (β_{demand}). Identified via the average marginal effect: $\hat{\beta}_{\text{demand}} = \hat{b}_{m4}/[\bar{p}(1 - \bar{p})]$, where $\bar{p} = 0.736$ is the mean adoption rate among customer-linked non-garage taxi merchants (Table 4, column 1).

Intercepts ($\alpha_{\text{cross}}, \alpha_{\text{taxi}}$). Each solved by bisection from the rate moments ($m5, m6$): find α such that the sample average of $P_j(\alpha, \hat{\beta})$ matches the observed six-month adoption rate.

Distance decay (δ) and **geographic competition** ($\beta_{\text{comp}}^{\text{geo}}$). The reduced-form coefficient at distance d is approximately:

$$\hat{b}(d) \approx \beta_{\text{comp}}^{\text{geo}} \cdot \exp(-\delta \cdot d) \times T \cdot \bar{p}(1 - \bar{p}), \quad (\text{A1})$$

where the second factor is the average marginal effect of the logistic-binomial model.¹⁹ Taking logs gives $\log |\hat{b}(d)| \approx \text{const} - \delta \cdot d$. Regressing across six bins yields $\hat{\delta} = 3.73^{***}$ km⁻¹ (half-life = 186 m). The level pins down $\hat{\beta}_{\text{comp}}^{\text{geo}} = \hat{b}(d_0)/[\exp(-\hat{\delta} \cdot d_0) \cdot T \bar{p}(1 - \bar{p})]$. Using the 100 m coefficient ($\hat{b}(0.1) = -0.0065$) and $\bar{p} = 0.048\%$, this gives $\hat{\beta}_{\text{comp}}^{\text{geo}} = -6.279^{***}$.

Fit of the distance gradient. The log-linear specification fits well ($R^2 = 0.93$). Predicted-to-actual ratios range from 0.76 (150 m) to 1.35 (50 m), with 3 of six bins within 15%.

¹⁹Using $\partial \Lambda / \partial \eta = \bar{p}(1 - \bar{p})$ and $\partial P / \partial p = T(1 - p)^{T-1} \approx T$ when $p \ll 1$. Exact in the limit $\bar{p} \rightarrow 0$.

Standard errors. Standard errors for $\hat{\alpha}_{\text{cross}}$, $\hat{\alpha}_{\text{taxi}}$, $\hat{\beta}_{\text{learn}}^s$, and $\hat{\beta}_{\text{learn}}^w$ use the delta method with a numerical Jacobian of the moment conditions.

The remaining three parameters ($\hat{\delta}$, $\hat{\beta}_{\text{comp}}^{\text{geo}}$, $\hat{\beta}_{\text{demand}}$) require different approaches because the AME inversion that identifies them amplifies noise on the logit-index scale. A naive delta-method SE for $\hat{\beta}_{\text{comp}}^{\text{geo}}$ would divide by $T\bar{p}(1-\bar{p}) \approx 0.003$, inflating the SE $\sim 300\times$ even though the underlying reduced-form garage coefficients are significant at $p < 0.01$.

$\hat{\delta}$ and $\hat{\beta}_{\text{comp}}^{\text{geo}}$: Both are identified from the log-linear regression $\log |\hat{b}(d)| = a - \delta \cdot d$ across six distance bins, which has four overidentifying restrictions. The OLS standard errors of this regression directly give $\text{SE}(\hat{\delta})$ from the slope and $\text{SE}(\hat{\beta}_{\text{comp}}^{\text{geo}})$ from the intercept via $|\hat{\beta}_{\text{comp}}^{\text{geo}}| \cdot T\bar{p}(1-\bar{p}) = \exp(\hat{a})$, so $\text{SE}(\hat{\beta}_{\text{comp}}^{\text{geo}}) = \exp(\hat{a}) \cdot \text{SE}(\hat{a}) / [T\bar{p}(1-\bar{p})]$.

$\hat{\beta}_{\text{demand}}$: Identified via $\hat{\beta}_{\text{demand}} = \hat{b}_{\text{m4}} / [\bar{p}(1-\bar{p})]$, where $\bar{p}$ is the control-group adoption rate. The SE is computed by moment bootstrap: drawing $\hat{b}_{\text{m4}}$ from $\mathcal{N}(\hat{b}, \hat{\sigma}_b^2)$ and perturbing $\bar{p}$, re-solving for $\hat{\beta}_{\text{demand}}$ in each of 500 draws, and taking the empirical standard deviation. This correctly propagates uncertainty through the nonlinear AME inversion.

ShareTreated vs. AnyTreated. The reduced-form Table 3 uses binary AnyTreated for RI power. The structural model requires continuous ShareTreated because counterfactuals change treatment intensity. Of the 33 garages with drivers (mean 56 per garage), 29 have AnyTreated = 1 and only 4 have AnyTreated = 0—each with exactly one control driver. The binary variable has minimal cross-sectional variation; RI provides the power. Appendix Table B12 reports ShareTreated regressions with RI p -values for comparison.

D Validation Details

Sign predictions. The model correctly predicts the sign of all four untargeted survey moments and the correct ordering ($m9 > m10$; $m8 > m11$):

- $m8$ (friends, cross-industry): data +0.030, model +0.002. Positive, driven by $\hat{\beta}_{\text{learn}}^s$.
- $m9$ (taxi friends): data +0.098, model +0.003. Positive despite same-industry competition, because demand (+0.92 in the linear index) dominates competition (−0.22).
- $m10$ (taxi non-friends): data +0.009, model +0.003. Positive, demand dominates.
- $m11$ (serendipitous): data +0.018, model ≈ 0 . Near-zero, consistent with $\hat{\beta}_{\text{learn}}^w \approx 0$.

Magnitude attenuation. The model underpredicts magnitudes by $3\times$ ($m10$) to $30\times$ ($m9$). Three factors contribute. First, *logistic compression*: baseline adoption is rare (0.196%

for taxi, 2.735% for cross-industry), placing the model in the left tail where large index effects produce small probability differences. Second, *subsample selection*: survey “friends” are the right tail of tie strength, not the population average that identifies the parameters. Third, *unmodeled channels*: taxi friends likely share customers at above-average rates; serendipitous contacts may transmit through in-person interactions.

Because Panel C predictions use representative population-mean exposures (not individual-level data from the survey sample), they are indicative of qualitative implications rather than exact forecasts. The sign predictions are the appropriate validation criterion.

E Counterfactual Details

Relevant population. The full analysis sample ($N \approx 3.9\text{M}$) is not the right basis: the vast majority have only weak ties ($\hat{\beta}_{\text{learn}}^w \approx 0$), making them structurally unresponsive to targeting. The combined reachable network includes: (A) 89,809 with $\geq 5\%$ overlap phone ties (learning); (B) 7,917 taxi merchants (all three channels); (C) 33,508 non-taxi within 500 m of a garage; (D) 59,055 with customer links. Populations C–D enter only through phone learning—competition and demand require $1[\text{taxi}_j]$ per equation (3). Total: 168,964.

Customer exposure from transactions. S_j^{cust} is constructed from P2B mobile money transactions between November 2022 and November 2023. A customer qualifies if they made ≥ 2 P2B transactions with any non-garage experimental driver during this window; a merchant–driver pair is linked if they share at least one qualifying customer. This yields 95,114 merchants linked to 1,521 drivers. Correlation with S^{phone} : $r = 0.005$, ensuring demand and learning respond independently to targeting.

All-garage reweighting. The counterfactual asks: what if all 1,256 treated drivers had the network characteristics of garage drivers? The same 1,256 drivers remain treated (preserving the budget and network positions), but each driver’s contribution to exposure is scaled by the ratio of average garage-driver to average treated-driver characteristics.

Formally, define the driver-level treatment vector $\mathbf{t} \in \mathbb{R}^{n_d}$. Under random assignment, $t_i = 1$ for treated drivers. Under all-garage, $t_i^{\text{phone}} = \bar{d}_{\text{garage}}/\bar{d}_{\text{treated}}$ for all 1,256 treated drivers, where $\bar{d}$ denotes mean phone degree. The phone exposure for individual j is:

$$S_j^{\text{phone}} = \min\left(1, \frac{\sum_i B_{ji} t_i^{\text{phone}}}{\sum_i B_{ji}}\right), \quad (\text{A2})$$

where $B_{ji} = 1$ if driver i is a mutual phone contact of j . Customer demand exposure S_j^{cust} is computed analogously, scaling by the ratio of average customer links (from the incidence matrix B^{cust}). Geographic exposure is set to its maximum: $\text{ShareTreated}_g = 1$ at every garage.

The ratios are computed from the incidence matrices directly: among treated drivers, garage drivers have fewer mutual phone contacts but a similar number of customer links as the average treated driver. Phone exposure falls from $\bar{S}^{\text{phone}} = 0.569$ to 0.499; customer demand is largely unchanged ($\bar{S}^{\text{demand}} = 0.555$ vs. 0.538). The main penalty comes from geographic competition ($\text{ShareTreated} = 1$ at all garages) and reduced phone learning: total induced adoption falls -13.5%.

F Distance Decay of Competition

Amplification factors. I recover $w(r)$ by converting each Panel B garage coefficient to the latent scale ($\hat{b}(r)/[\bar{p}(r)(1 - \bar{p}(r))]$), subtracting learning ($\hat{\beta}_{\text{learn}}^s$), and dividing by the phone-network baseline (≈ -0.066):

Radius (m)	$\hat{b}(r)$	$\bar{p}(r)$	$w(r)$
50	-0.181	0.017	170×
75	-0.182	0.014	201×
100	-0.067	0.011	96×
150	-0.022	0.009	39×
300	-0.014	0.006	39×
500	-0.009	0.005	26×

Parametric fits. A power law $w(d) = A \cdot d^{-\gamma}$ yields $\hat{A} = 6,421$, $\hat{\gamma} = 0.891$. Setting $w(d^*) = 1$ gives $d^* \approx 19$ km—approximately the diameter of Dakar, consistent with the average phone contact being dispersed across the city. An exponential fit ($w(d) = 322 \cdot e^{-d/97\text{m}}$) implies a half-life of 67 m, shorter than the 186 m structural half-life. The difference arises because the amplification factors $w(r)$ incorporate both the exponential kernel *and* variation in baseline adoption $\bar{p}(r)$ across distance bins: merchants closer to garages have higher baseline adoption, so the AME conversion $[\bar{p}(1 - \bar{p})]^{-1}$ amplifies the latent coefficient more at short distances, making $w(r)$ decline faster than the pure structural decay $\exp(-\hat{\delta} \cdot d)$.