Nationwide Diffusion of Technology: Experimental Evidence from Multiple Networks*

Deivy Houeix†

[Click here for the most recent version]

March 27, 2026

Abstract

When does a new technology spread from one firm to another, and through which network connections? I study this question using a randomized rollout of a digital payment technology among 2,196 taxi businesses in Dakar, Senegal, combined with encrypted phone contact directories covering roughly 80% of the adult population and the complete transaction ledger of the country's leading mobile money platform. I show that diffusion depends sharply on the network through which exposure arrives. Strong social ties transmit learning across industries and across cities. Geographic proximity within the taxi sector primarily transmits competition, generating negative spillovers concentrated at very short distances. Shared customers generate demand complementarities, and this is the dominant source of within-industry diffusion, operating even between taxi businesses with no social connection. Because these channels have opposite signs, aggregate spillovers can be misleading: average within-industry exposure appears negative even though learning is strongly positive on meaningful social ties. Cross-city diffusion is driven primarily by remote social connections rather than by treated taxis physically visiting destination cities. A structural model that matches these patterns shows that targeting high-degree taxi businesses generates 59% more adoption than random assignment, whereas targeting garages backfires by intensifying local competition, substantially reducing effectiveness that is only visible once the channels are separately identified.

*I am grateful to Isaiah Andrews, David Atkin, Abhijit Banerjee, Esther Duflo, Dean Eckles, Justine Knebelmann, Sahil Loomba, 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 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 (Banerjee et al., 2013; Duflo et al., 2023; 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 is data. 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., 2023; 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 small 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; in this setting, it can reverse the sign of estimated spillovers.

I address this problem by combining a randomized experiment with two data sources that together reveal the near-universe of social and economic connections among firms in Senegal. The experiment randomized early access to a QR-code-based 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,

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

to businesses 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 $\approx$ 85%). Each saved contact number is mapped into one of $2^{16} = 65536$ 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 appears in the 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 roughly 1,900 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 and the conditions under which it does not, decomposing the net effect into three channels—social learning, within-industry competition, and customer-demand complementarities—and showing how each depends on tie type, industry, and distance. The central result is simple: 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. Because these channels can have opposite signs, aggregate spillovers can be 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: a ten-percentage-point increase in the treated share raises adoption by about one percentage point, a $+9.8\%$ effect. 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 ($+21.4\%$), 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 evidence is consistent with spillovers operating through friends, frequent interlocutors, and bridging ties (Granovetter, 1973; Aral, 2016; Jahani et al., 2023). The cross-city evidence further suggests that diffusion operates through remote interactions rather than physical co-presence (e.g., via seasonal migration) in the Marshallian sense (Marshall, 1920): connected treated taxis who actually visited a secondary city have no detectable effect there, while the significant effects come from treated taxis who never visited.

Second, within-industry spillovers are fundamentally heterogeneous because learning and competition pull in opposite directions. The average same-industry coefficient is significantly negative (-7.5%), but survey-identified friends generate a large positive effect ($+37.4\%$). The structural model rationalizes this pattern: strong ties transmit substantially more learning than weak ties ($\hat{\beta}^{s}_{\text{learn}} = +0.100$ vs. $\hat{\beta}^{w}_{\text{learn}} = +0.001 \approx 0$), while within-industry exposure also activates a negative competition channel ($\hat{\beta}_{\text{comp}} = -0.102$). For the 3.7% of taxi merchants with strong ties, learning dominates; for the 96.3% 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 (Baum-Snow et al., 2024).

Third, customer-demand complementarities are the dominant within-taxi channel, operating independently of social ties. When customers grow accustomed to paying digitally with a treated taxi, other merchants serving the same customers adopt at higher rates even absent any social connection between the firms. This channel produces effects of $+17.9\%$ on pre-intervention customer links and $+41.8\%$ on the post-experiment network restricted to non-social merchants. The cross-industry customer coefficient is a precise zero. The structural decomposition confirms that customer demand is the largest contributor to within-taxi adoption ($+0.147$ pp over six months, vs. -0.005 pp from the net phone channel). Social learning, by contrast, is the only channel that operates across industries and across cities ($+0.008$ pp).

A structural model ties these findings together. Using six reduced-form moments as identifying conditions and five 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), competition, and customer demand. Policy counterfactuals show why this decomposition matters: selecting high-degree taxi businesses produces 59% more adoption than random assignment, while garage-only targeting—which maximizes local exposure—backfires, producing 84% fewer adopters. The $10\times$ gap between the worst and best rule comes entirely from understanding which channels are positive and which are negative.

This paper makes three contributions. First, it provides, to my knowledge, the first experimental evidence of technology diffusion operating simultaneously across industries, cities, and network layers within a single economy. The existing experimental literature has focused primarily on agricultural technologies in small rural communities (see Suri and Udry (2024) for a review), where costs and heterogeneous returns make diffusion context-dependent; a notable contemporary exception is Chaurey et al. (2025), who study social learning among urban manufacturing firms. By leveraging administrative network data covering the near-universe of economic actors, this paper reveals pathways—cross-industry spillovers through strong social ties and inter-city propagation through remote interactions—that are difficult to detect with survey- or neighborhood-based data. The 116,613 individuals linked to 2,196 experimental firms through the strong-tie phone network alone represent a scale of observation not previously available in an experimental diffusion study.

Second, the paper advances the micro-foundations of knowledge spillovers by showing that the sign, magnitude, and channel of spillovers depend on the network connection through which exposure operates. The literature has focused on intra-industry spillovers in local clusters (Moretti, 2021; Baum-Snow et al., 2024), typically mediated by face-to-face interaction (Cai and Szeidl, 2017; Atkin et al., 2024). Recent work delineates specific channels—supply chains (Alfaro-Ureña et al., 2022), demographics (Crouzet et al., 2024), knowledge hoarding (Cefalà et al., 2025)—but simultaneous observation of social, customer, and geographic networks within a single setting has remained out of reach. This paper observes all three and shows that each carries a distinct channel: social ties transmit learning, customer overlap transmits demand, and geographic proximity transmits competition. That these channels can have opposite signs—and that their relative strength depends on tie type—offers one way to reconcile the heterogeneous results found across the literature.

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 this decomposition. The answer is nuanced. Within the taxi sector, customer demand is the dominant force ($+0.147$ pp, vs. -0.005 pp from the net phone channel), consistent with the emphasis on strategic complementarities. But social learning is the only channel that operates across industries and cities ($+0.008$ pp). Within industry, the aggregate phone effect is negative—not because learning is absent, but because 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 structural decomposition quantifies these contributions directly. The targeting counterfactuals make the stakes concrete: a planner who treats all exposure as complementarity and targets garages produces $10\times$ fewer adopters than one who targets high-degree taxi businesses to maximize learning through the broader network.

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 technology 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

[2] I 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.

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 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, about 75-80% of the potential business population had adopted the technology. 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.

[3]Smartphone 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).

[4]The 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.

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 conversation 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, 2024): 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 (2024) for more details.

[5]The 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.

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 85 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.

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 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.[6]

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. If treatment had systematically created new ties, such balance would be difficult to reconcile with the data.

After filtering and degree trimming, the phone network contains roughly 16 million mutual edges linking about 1,900 experimental drivers to approximately 4 million individuals. This scale makes it possible to study whether the randomized technology diffuses not only within the treated sector, but also across industries and across space

[6]This 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.

throughout Senegal.

Table B1 validates the overlap measure against survey-reported relationships. In the endline survey, each of the 1,693 experimental drivers 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, communica-

tion 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.[7] 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).

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 587 phone contacts. While I cannot retrieve the exact phone contacts due to the encryption process, I can determine the number of contacts for each user and utilize the overlap share of contacts to demonstrate the geographical reach of the network. The data reveal that

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

the taxi businesses' network extends significantly beyond Dakar, where the experiment took place. Notably, over half of their phone contacts reside outside the capital city, as indicated in the third row. Most of these contacts are located in urban areas across the country, as shown in the second row.

As a descriptive exercise, I identify business links within the phone contact network of the taxi businesses. To achieve this, I leverage data on all business adopters of digital payments by April 2023 and their registered industry. I find that, on average, each business is connected to about 48 other business owners, which represents a non-trivial share of their phone contacts. As in most business sectors, taxi businesses' connections span various industries. Notably, the taxi industry tends to be more concentrated within its own sector, which might downward bias the spillover effects observed across different industries compared to other sectors. This extensive network spread is a critical feature of the data, which not only allows me to track technology diffusion across both space and industry, but also reflects the broader characteristics of firm networks in this low-income 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 B3, 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 share important advice and information with about 74% of the people they know, particularly other businesses and taxi drivers.[8] 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.

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

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

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

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 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:

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

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., 2023; 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: a ten-percentage-point increase in the treated share raises adoption likelihood by about one percentage point, a 9 percent effect relative to the mean. On the weaker ties, by contrast, the coefficient is indistinguishable from zero.[9]

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

[9]Because 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.

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.

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 and frequent interlocutors generates a substantial positive effect (approximately $+37\%$), while exposure

through casual ties is smaller but still positive (roughly $+10\%$). 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 ($+37.4\%$, $p_{\text{RI}} = 0.010$) and frequent interlocutors ($+32.1\%$, $p_{\text{RI}} = 0.004$), while exposure through non-friends produces a smaller and insignificant effect ($+10.4\%$, $p_{\text{RI}} = 0.435$). 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 ($+47.7\%$, $p_{\text{RI}} = 0.003$) and former or regular taxi clients ($+93.7\%$, $p_{\text{RI}} = 0.151$). 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 B4 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 ($p_{\text{RI}} = 0.139$ 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.[10]

Pre-intervention customer network. Table B5 provides the cleanest test by constructing customer links from person-to-person transactions between April 2021 and March

[10]Most 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.

2022, before the experiment began, ruling out treatment-induced link formation by construction.[11] 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.[12] 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 transactions, November 2022–November 2023). The within-taxi coefficient is large and significant: $\hat{\beta} \times 1{,}000 = 338.49$ ($p_{\text{RI}} = 0.003$), an effect size of $+45.3\%$ (column 1). Restricting to merchants without mutual phone ties to the treated driver—the non-social subsample in column 3—yields $\hat{\beta} \times 1{,}000 = 311.55$ ($p_{\text{RI}} = 0.008$, effect size $+41.8\%$). 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: $\hat{\beta} \times 1{,}000 = 716.57$, $p_{\text{RI}} = 0.080$, effect size $+93.7\%$), 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, 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.[13]

The cross-industry customer coefficients in Table 4 are positive but small and insignificant (column 4: $\hat{\beta} \times 1{,}000 = 13.38$, $p_{\text{RI}} = 0.461$, effect size $+2.2\%$). This null is not mechan-

[11] Because 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.

[12] The 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 B7 shows that customer-demand effects tend to offset the competition effect within garages.

[13] Table B6 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.

ical: 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 163,390 individuals in the strong-tie sample ($\geq$ 5% overlap), 120,223 (73.6%) can be matched to geolocation records. Of these, 80,031 (66.6%) appeared in Dakar at least once during the experiment, while 40,192 (33.4%) 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 is substantial—precisely the kind of pattern that would sustain a contact-based diffusion channel.

Cross-industry adoption outside Dakar. Table 5 estimates the same specification as Table 1 Panel A with cross-industry adoption *outside Dakar* as the outcome. The full sample ($N = 163{,}390$) yields a positive effect ($+0.0835$, $p_{\text{RI}} = 0.056$). Splitting by whether the individual was observed in Dakar during the experiment reveals positive point estimates in both subsamples: $+0.0472$ among Dakar visitors ($p_{\text{RI}} = 0.159$, effect size $+60.7\%$) and $+0.2437$ among non-visitors ($p_{\text{RI}} = 0.127$, effect size $+22.0\%$). Neither subsample achieves conventional significance individually, but the consistently positive coefficients—and the fact that individuals never observed in Dakar also respond—already hint that physical co-presence may not be necessary for diffusion.

City-level decomposition: diffusion without physical contact. Table 6 provides the sharpest test. For each of the four largest secondary cities—Touba, Thiès, Mbour, and St-Louis—we decompose the treatment-exposure variable by whether the connected driver actually visited that city during the experiment. If diffusion operates through physical co-presence, the effect should load on drivers who traveled to the destination; drivers who never set foot in a city should have no influence on adoption there.

The data reject this prediction cleanly. Among drivers who visited each city, all four coefficients are insignificant with large p-values. The significant effects come instead from drivers who *never visited* the destination city across three of four main cities.[14] This is the opposite of what a standard contact-diffusion model would predict: the channel that *should* matter (physical co-presence) does not, and the channel that *should not* matter (remote ties to drivers who never visited) does.

The most parsimonious interpretation is that connected individuals acquire general information about the usefulness of mobile payments—not city-specific merchant recommendations—and act on it when they encounter adoption opportunities in secondary cities. The technology's value proposition is portable: once a person learns that digital payments work reliably from a trusted Dakar contact, that knowledge applies wherever the platform operates. Anecdotal evidence suggests a concrete mechanism: during fieldwork, treated drivers were observed posting WhatsApp stories and Instagram content showcasing their use of digital payments as a mark of professional modernization. Such broadcasts reach a driver's entire contact list instantaneously and irrespective of geographic distance, providing exactly the kind of remote, non-targeted information transmission that could explain why drivers who never visited a city nonetheless influence adoption there. Effect sizes are large in proportional terms (55–150%) but operate on small base rates (control adoption of 0.02–0.07%), so the absolute number of additional adopters is modest. Nevertheless, the result identifies a channel that is typically invisible in the diffusion literature: technology adoption propagating over long distances through remote interactions alone, without requiring that either party travel to the destination.

[14] The null result in Touba—across both driver migration groups—may reflect that city's distinct commercial ecosystem; see Figure A5.

6 Structural Model and Estimation

6.1 A parsimonious model of adoption and spillovers

The reduced-form evidence identifies three distinct spillover channels, each active in a different subpopulation. This section estimates a structural model that decomposes adoption into social learning, within-industry competition, and customer demand, using six reduced-form moments from Tables 1 and 4 as identifying conditions and five additional moments from Tables 2 and 3 as untargeted validation tests. The model is a parsimonious discrete-time adoption framework informed by finite-horizon diffusion models with network transmission (Banerjee et al., 2013) and network-interaction models with positive and negative peer incentives (Jackson and Yariv, 2007). Social learning is allowed to vary with tie strength: strong ties ($\geq 5\%$ phone-contact overlap) and weak ties ($< 5\%$ overlap) transmit information at different rates.

Timing and adoption process. Treatment was randomized on March 28, 2022; adoption is measured through October 5, 2022, yielding $T = 6$ monthly periods. Adoption is an absorbing process: at each period t, each non-adopter j draws an independent adoption decision; once adopted, the decision is irreversible.

Adoption probability. Because all exposure variables are functions of the time-invariant treatment assignment $T_i \in \{0, 1\}$, the per-period adoption probability is constant across periods. At each $t = 1, \ldots, 6$, non-adopter j adopts with probability:

$$p_j = \Lambda\Big(\alpha_g + \big[\beta^s_{\text{learn}} \cdot \mathbf{1}[\text{strong}_j] + \beta^w_{\text{learn}} \cdot \mathbf{1}[\text{weak}_j]\big] \cdot S_j^{\text{phone}} + \beta_{\text{comp}} \cdot S_j^{\text{phone}} \cdot \mathbf{1}[\text{taxi}_j] + \beta_{\text{demand}} \cdot S_j^{\text{cust}} \cdot \mathbf{1}[\text{taxi}_j]\Big), \tag{3}$$

where $\Lambda(\cdot)$ is the logistic CDF, $\alpha_g \in \{\alpha_{\text{cross}}, \alpha_{\text{taxi}}\}$ is an industry-specific intercept, $\mathbf{1}[\text{strong}_j]$ indicates $\geq 5\%$ phone-contact overlap (with $\mathbf{1}[\text{weak}_j] = 1 - \mathbf{1}[\text{strong}_j]$), S_j^{phone} is the share of j's phone contacts who are experimentally treated, S_j^{cust} is the share of customer-linked drivers who are treated, and $\mathbf{1}[\text{taxi}_j]$ indicates whether j operates in the taxi industry.

The cumulative adoption probability over the six-month window is:

$$\Pr(\text{adopt}_j) = 1 - (1 - p_j)^6 = 1 - \big(1 - \Lambda(\eta_j)\big)^6, \tag{4}$$

where η_j is the linear index from equation (3).

Parameters. The model has six parameters:

1. α_{cross}, α_{taxi}: industry-specific intercepts capturing baseline adoption propensity absent network exposure. Separate intercepts are necessary because baseline adoption rates differ sharply (2.735% for cross-industry vs. 0.196% for within-taxi over six months).
2. β^{s}_{learn}: the effect of phone-contact exposure for individuals with $\geq 5\%$ overlap—strong-tie social learning.
3. β^{w}_{learn}: the corresponding parameter for contacts with $< 5\%$ overlap—weak-tie social learning.
4. β_{comp}: the additional effect of phone-contact exposure for taxi merchants, capturing within-industry competition. The net phone effect for a taxi merchant with strong ties is $\beta^{s}_{\text{learn}} + \beta_{\text{comp}}$; with weak ties, $\beta^{w}_{\text{learn}} + \beta_{\text{comp}}$.
5. β_{demand}: the effect of shared-customer exposure on taxi merchant adoption, capturing customer-demand complementarities.

6.2 Identification

Estimation approach. I use minimum distance estimation (Chamberlain, 1984; Newey and McFadden, 1994), matching six reduced-form OLS coefficients to the model's predictions. With six moments and six parameters, the system is exactly identified.

The estimator solves:

$$\hat{\theta} = \arg\min_{\theta} \ \left[\hat{b} - h(\theta)\right]' W \ \left[\hat{b} - h(\theta)\right], \tag{5}$$

where $\hat{b} = (\hat{b}_1, \ldots, \hat{b}_6)'$ is the vector of reduced-form moments and $h(\theta)$ maps parameters to model-predicted moments. With exact identification, point estimates are invariant to W.

Mapping structural parameters to reduced-form moments. The structural model is nonlinear (logistic), but the reduced-form targets are OLS coefficients. The bridge between the two is indirect inference (Gouriéroux et al., 1993): I compute model-implied cumulative adoption probabilities $P_j(\theta) = 1 - (1 - \Lambda(\eta_j))^6$ at each individual's observed exposure, then run the same OLS regression on the simulated outcomes that was run on

the data:[15]

$$P_j(\theta) = a_m + h_m(\theta) \cdot S_j + c_m \cdot \text{ExpectedExposure}_j + \varepsilon_j, \tag{6}$$

where $h_m(\theta)$ is the OLS coefficient on ShareTreated. Using OLS as the auxiliary model—rather than logit—is required for consistency: the data targets $\hat{b}$ are OLS coefficients, and indirect inference requires the same estimator on both sides. The OLS projection accounts for the full distribution of exposures and maps the nonlinear structural model to the linear estimand.

This approach is used for the three phone-network moments (m1–m3). The customer-demand moment (m4) instead uses the average marginal effect $h_m = \beta_{\text{demand}} \cdot \bar{p}(1-\bar{p})$, because the customer network is observed separately from the phone network and the relevant OLS coefficient maps directly to this expression. The two rate moments (m5, m6) use exact sample averages of $P_j(\theta)$.

Identifying restrictions. Each channel is identified from a subsample where only that channel is active:

#	Source	Target	Channel loading
Identifying moments			
m1	Table 1, $\geq 5\%$ overlap, cross-ind.	ShareTreated coef.	β^s_{learn}
m2	Table 1, $< 5\%$ overlap, cross-ind.	ShareTreated coef.	β^w_{learn}
m3	Table 1, all overlap, within-taxi	ShareTreated coef.	$w_s\beta^s_{\text{l}} + w_w\beta^w_{\text{l}} + \beta_{\text{c}}$
m4	Table 4, non-social customers	ShareTreated coef.	β_{demand}
m5	Cross-industry adoption rate	Endline rate	α_{cross}
m6	Within-taxi adoption rate	Endline rate	α_{taxi}
Validation moments (untargeted)			
m7	Table 3, garage 100 m (social ties)	AnyTreated coef.	$\beta^s_{\text{l}} + \beta_{\text{c}}$
m8	Table 2, friends (all-industry)	ShareTreated coef.	β^s_{learn}
m9	Table 2, taxi friends	ShareTreated coef.	$\beta^s_{\text{learn}} + \beta_{\text{comp}}$
m10	Table 2, taxi non-friends	ShareTreated coef.	$\beta^w_{\text{learn}} + \beta_{\text{comp}}$
m11	Table 2, serendipitous (cross-ind.)	ShareTreated coef.	β^w_{learn}

The identification logic is sequential. Cross-industry moments (m1, m2) identify learning separately for strong and weak ties, because competition and customer demand are absent across industries. The within-taxi phone moment (m3) loads on $w_s\beta^s_{\text{learn}} + w_w\beta^w_{\text{learn}} +$

[15]The indirect inference principle estimates a structural model's parameters by matching the output of a simpler *auxiliary model*—here OLS—applied to simulated data with the same auxiliary model applied to real data. The auxiliary model need not be correctly specified; it serves as a summary statistic that both data and model must match.

β_{comp}, where $w_s = 3.7\%$ and $w_w = 96.3\%$ are the population shares of strong and weak ties; with learning pinned down from m1 and m2, this moment identifies competition. The non-social customer moment (m4) isolates demand complementarities among taxi merchants who share customers but have no phone ties to treated drivers. The two adoption rates (m5, m6) pin down the intercepts.

The validation moments carry different loadings depending on tie strength. Survey-identified friends (m8, m9) are classified as strong ties, since friendship implies meaningful social interaction consistent with $\geq 5\%$ overlap. Non-friends (m10) and serendipitous contacts (m11) are classified as weak ties. The garage moment (m7) loads on $\beta^s_{\text{learn}} + \beta_{\text{comp}}$ because customer demand is empirically negligible among garage-proximate merchants: only 32% share any customers with same-garage drivers, and among those who do, the customer-overlap coefficient is negative (-0.150, $p_{\text{RI}} = 0.098$), indicating competitive diversion rather than demand complementarity.

Standard errors. Standard errors use the sandwich formula with a block-diagonal moment covariance matrix that accounts for overlap across moments estimated on related subsamples: block 1 groups (m1, m5), block 2 is m2, block 3 groups (m3, m6), and m4 is a separate scalar. Counterfactual confidence intervals use parametric bootstrap (500 draws from $\mathcal{N}(\hat{\theta}, \hat{V}(\hat{\theta}))$).

6.3 Results

Table 7 presents the estimated parameters (Panel A), the exact fit to identifying moments (Panel B), the out-of-sample validation (Panel C), and the counterfactual decomposition (Panel D).

Heterogeneous learning. Strong ties transmit substantially more learning than weak ties: $\hat{\beta}^s_{\text{learn}} = +0.100$ versus $\hat{\beta}^w_{\text{learn}} = +0.001 \approx 0$. Information about technology adoption flows through meaningful social relationships characterized by frequent contact and substantial phone-network overlap. Weak ties, which comprise 96.3% of the phone network, contribute negligibly to diffusion of a complex technology like mobile payments.

Competition and the net phone effect. The model estimates $\hat{\beta}_{\text{comp}} = -0.102$: within-industry phone exposure carries a negative competitive component. The net phone effect depends on tie strength. For strongly-tied taxi merchants (3.7% of the sample), learning dominates. For weakly-tied taxi merchants (96.3%), competition dominates:

$\hat{\beta}^{w}_{\text{learn}} + \hat{\beta}_{\text{comp}} = -0.100 < 0$. Because weak ties are the vast majority, the population-weighted net phone effect is negative—rationalizing the negative average same-industry coefficient in Table 1.

Customer demand. Customer demand is the dominant within-taxi channel ($\hat{\beta}_{\text{demand}} = +1.643$). The decomposition in Panel D confirms this: customer demand contributes $+0.147$ percentage points to within-taxi adoption over six months, versus a net phone contribution of -0.005 percentage points. For cross-industry merchants, social learning on strong ties is the only active channel.

Validation. Panel C compares model predictions to five untargeted moments (RMSE $= 0.061$, correct sign for 4 of 5). The key test is m9 (taxi friends, Table 2): the data show a positive coefficient ($+0.102$), and the model correctly predicts a positive sign ($+0.001$) because friends are strong ties for whom learning ($\hat{\beta}^{s}_{\text{learn}}$) exceeds competition ($\hat{\beta}_{\text{comp}}$). This prediction is genuinely out of sample—$\hat{\beta}^{s}_{\text{learn}}$ was identified from cross-industry contacts (m1), yet it correctly predicts the sign of a within-industry moment. The remaining sign mismatch (m10: taxi non-friends) likely reflects customer-demand contamination in the survey subsample, since non-friends who share customers with experimental drivers face positive demand pressure that the phone-network model does not capture.

Distance decay of competition. The garage coefficients in Table 3 are dramatically largerDrivers Did Not Visit City than the phone-network competition parameter would predict at any single-contact level. This reflects competitive pressure—discouragement, information withholding, customer diversion—evaluated at physical co-location ($\sim$100 m) rather than at the diffuse distance of an average phone contact across metropolitan Dakar. Inverting the reduced-form coefficients using the structural estimates yields competition amplification factors ranging from $139\times$ at 50 m to $21\times$ at 500 m. Setting amplification to unity and solving for the implied distance yields approximately 32 km—close to the diameter of metropolitan Dakar, a sensible effective distance for the average phone contact (Appendix B.1).

Despite the large per-contact intensity at garages, the garage channel is narrow: only 43% of taxi merchants are located near an experimental garage. The phone channel reaches all taxi merchants at much lower per-contact intensity but dominates total competitive exposure in aggregate.

7 Counterfactuals and Targeting

7.1 Setup

The structural estimates from Table 7 allow us to ask a question relevant to any social planner or platform seeking to maximize technology adoption: for the same onboarding budget of $N_T = 1{,}256$ drivers, how much additional adoption could be generated by choosing *which* drivers to onboard? We fix the phone and customer networks at their observed values, recompute each individual's exposure S_j under alternative targeting rules, and predict cumulative six-month adoption using the estimated logistic model. The exercise holds structural parameters constant across rules—ignoring potential general-equilibrium responses and treating the network as exogenous. Despite these caveats, the counterfactual is informative because the parameters are identified from exogenous experimental variation and the targeting rules change only *who* is onboarded, not the technology or market environment.

We report three adoption margins: cross-industry in Dakar, cross-industry outside Dakar (among the 4,720 individuals with merchant activity in non-Dakar regions, using a separate intercept $\alpha_{\text{cross-other}}$ calibrated to the observed non-Dakar adoption rate with the same learning coefficient), and within-taxi.

Demand exposure. The structural model loads all three within-taxi channels—learning, competition, and customer demand—on the same phone-network exposure S_j^{phone}. While learning plausibly requires direct communication, customer demand and competition operate through the broader market: a customer who rides with both driver A and driver B generates demand pressure on B regardless of whether A and B are phone contacts. To verify that this proxy does not distort counterfactual rankings, we construct a separate demand measure S_j^{demand} that weights each driver by their observable customer reach. For each driver d, we compute the number of their phone contacts who also appear as merchants sharing customers with d in the transaction data (Table 4), and normalize by the driver-level mean ($\bar{n}_{\text{cust}} = 2.74$) to obtain a customer weight w_d. The demand exposure is then:

$$S_j^{\text{demand}} = \frac{\sum_{d \in \mathcal{T} \cap \mathcal{N}_j} w_d}{\sum_{d \in \mathcal{N}_j} w_d},$$

where $\mathcal{T}$ is the treated set and $\mathcal{N}_j$ is j's set of driver contacts. The within-taxi prediction becomes $\eta_j^{\text{taxi}} = \alpha_{\text{taxi}} + (\beta_{\text{learn}} + \beta_{\text{comp}}) \times S_j^{\text{phone}} + \beta_{\text{demand}}^{\text{adj}} \times S_j^{\text{demand}}$, with $\beta_{\text{demand}}^{\text{adj}}$ rescaled to match aggregate within-taxi adoption under actual treatment. The customer weight w_d is only weakly correlated with phone degree ($r = 0.04$), so reweighting by customer

reach genuinely shifts the exposure measure rather than replicating phone degree.[16] This customer-weighted specification produces identical targeting rankings and only modestly compresses magnitudes relative to the phone-only model; we use it throughout as the more conservative specification.

The relevant population. The full analysis sample ($N \approx 3.9$M) is not the right basis for counterfactual evaluation. Over 98% of those individuals have only weak phone ties ($< 5\%$ overlap), through which social learning is negligible ($\hat{\beta}^{w}_{\text{learn}} = +0.001$). Targeting cannot create strong ties where none exist—tie strength is determined by network topology, not treatment assignment—so the vast majority of the sample is structurally unresponsive to any reallocation of treatment. We therefore construct a *combined reachable network*: the union of all individuals for whom at least one structural channel is active. This includes (A) 116,613 individuals with $\geq 5\%$ overlap phone ties (social learning); (B) 8,583 taxi merchants with any mutual phone tie (learning, competition, and customer demand); (C) 37,543 non-taxi merchants within 500m of an experimental garage (geographic exposure); and (D) 66,844 non-taxi merchants with shared-customer links to experimental drivers (demand proximity). In this sample, the majority of ties are strong and the remainder are reachable through other channels.

7.2 Network-based targeting

We evaluate five targeting scenarios against the no-treatment baseline (10,376 predicted adopters over six months). All scenarios hold the budget fixed at $N_T = 1{,}256$ except garage-only ($N = 228$, reflecting the smaller pool of garage-based drivers):

1. *Actual random onboarding*: the realized experimental assignment.
2. *Garage-only*: onboard only drivers operating from taxi garages—the rule a field team with no network data would most naturally implement.
3. *Top mobile money users*: select the 1,256 heaviest platform users.
4. *Highest degree*: select drivers with the most mutual phone contacts.
5. *Geographic reach*: select drivers with the most contacts registered as merchants outside Dakar.

[16] The raw count $n_{\text{cust},d}$ correlates more strongly with phone degree ($r = 0.39$): drivers with more phone contacts also have more shared-customer links in levels. But because w_d normalizes by $\bar{n}_{\text{cust}}$ and S_j^{demand} divides by the sum of weights over all contacts, the *relative* customer weight does not increase with phone-network size. The customer and phone channels scale together in levels but not in proportions.

Results. Table 9 reports predicted adoption under each rule. Highest-degree targeting produces 327 treatment-induced adopters—$+59\%$ more than the 205 induced by random onboarding—by raising average phone exposure from $\bar{S}^{\text{phone}} = 0.568$ (random) to 0.899 (degree). Geographic reach targeting performs nearly identically (324 induced, $+58\%$), because drivers with many out-of-Dakar contacts also tend to have high overall degree.

All three adoption margins respond to targeting and move in the same direction. Cross-industry Dakar adoption—the dominant margin, accounting for 86% of induced adoption—rises from 176 (random) to 280 (degree, $+59\%$). Within-taxi adoption rises from 15 to 25 induced ($+69\%$). Cross-industry adoption outside Dakar rises from 15 to 22 induced ($+49\%$; geographic reach achieves $+51\%$ on this margin, slightly outperforming degree targeting). This alignment arises because every individual in the combined reachable network has at least one active channel, so shifting treatment toward higher-degree drivers increases effective exposure throughout.

Top mobile money users represent an intermediate case (241 induced, $+17\%$ over random). This rule requires no network data—any platform observes its own transaction records—yet outperforms random because heavy users tend to occupy more central network positions.

Garage targeting: the intuitive mistake. Garage-only targeting is a rule a social planner could naturally consider. Garages are visible, spatially concentrated, and easy to reach—they are where a field team would go first. Onboarding drivers in a shared location also creates maximal local exposure: every other driver in the garage observes the technology firsthand, a seemingly ideal condition for diffusion. Yet this rule produces only 32 induced adopters (-84% vs. random), the worst outcome by a wide margin. The failure has two reinforcing causes. First, garages contain only 228 eligible drivers, mechanically reducing the treatment group and collapsing network-wide exposure ($\bar{S}^{\text{phone}} = 0.092$). Second—and more subtly—the very feature that makes garages appealing (dense local interaction) activates the competition channel most intensely. As Section 5.2 documented, same-garage exposure generates strongly negative spillovers at short distances: treated drivers divert customers from untreated neighbors, offsetting the learning gains that spatial proximity would otherwise produce. Garage targeting thus concentrates treatment precisely where the negative channel is strongest, while starving the rest of the network of the positive learning and demand channels that drive aggregate adoption. It is, in short, a strategy that looks optimal through the lens of a pure-learning model but is actively harmful once competition is accounted for.

Policy implications. From the perspective of a social planner seeking to maximize technology adoption from a fixed budget, the counterfactuals deliver three lessons. First, network-based targeting is powerful: selecting drivers with the most phone contacts produces $+59\%$ more adoption than random assignment, implementable from phone fingerprint data already collected for overlap computation. Geographic reach targeting, which leverages merchant registration data, performs comparably ($+58\%$). Second, even a simple heuristic—onboarding heavy mobile money users, requiring no network data at all—outperforms random by $+17\%$, offering a low-cost improvement for platforms without access to social-network information. Third, and most importantly, the planner must account for negative spillovers when designing targeting rules. The garage result is a cautionary tale: the strategy that appears most natural from a diffusion standpoint—concentrating treatment where agents interact most densely—backfires because it maximizes competitive crowding while minimizing network reach. A $10\times$ gap in effectiveness separates the worst rule (garage-only, 32 induced) from the best (highest degree, 327 induced), and the entire difference comes from understanding which channels are positive and which are negative. More broadly, the results show that a social planner with access to the structural decomposition—knowing that learning operates on strong ties, competition on local proximity, and demand through shared customers—can do substantially better than one who treats all spillovers as homogeneous.

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 59% more adoption than random assignment, while the intuitively appealing strategy of concentrating treatment at garages backfires because it maximizes competitive crowding. The $10\times$ 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.

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. © 2016 by The University of Chicago.

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).

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," NBER Working Paper 31368, National Bureau of Economic Research 2023.

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* (S1), 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*, 2024.

Jack, William and Tavneet Suri, "Risk Sharing and Transactions Costs: Evidence From Kenya's Mobile Money Revolution," *American Economic Review*, 2014, *104*, 183–223.

Jackson, Matthew O. and Leeat Yariv, "Diffusion of Behavior and Equilibrium Properties in Network Games," *American Economic Review*, 2007, *97* (2), 92–98.

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.

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

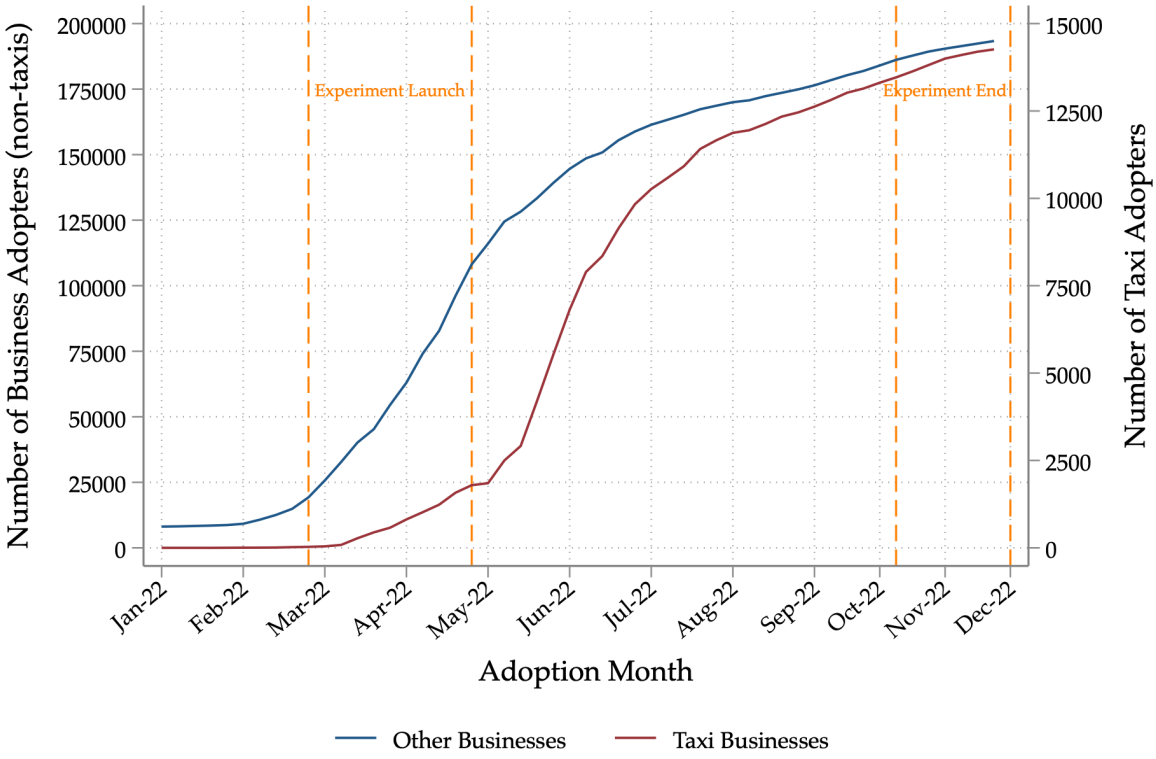


Figure 1: Digital Payments' Adoption in 2022 - *Author's calculations*

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 onboarded 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 endline 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.

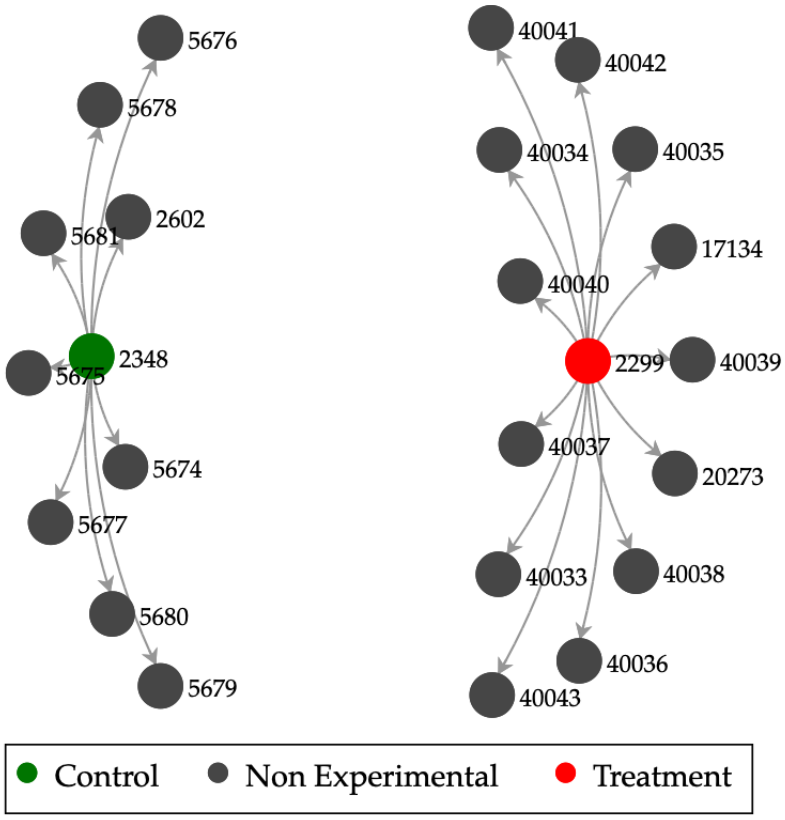


Figure 3: Illustration of the Experimental Design

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.

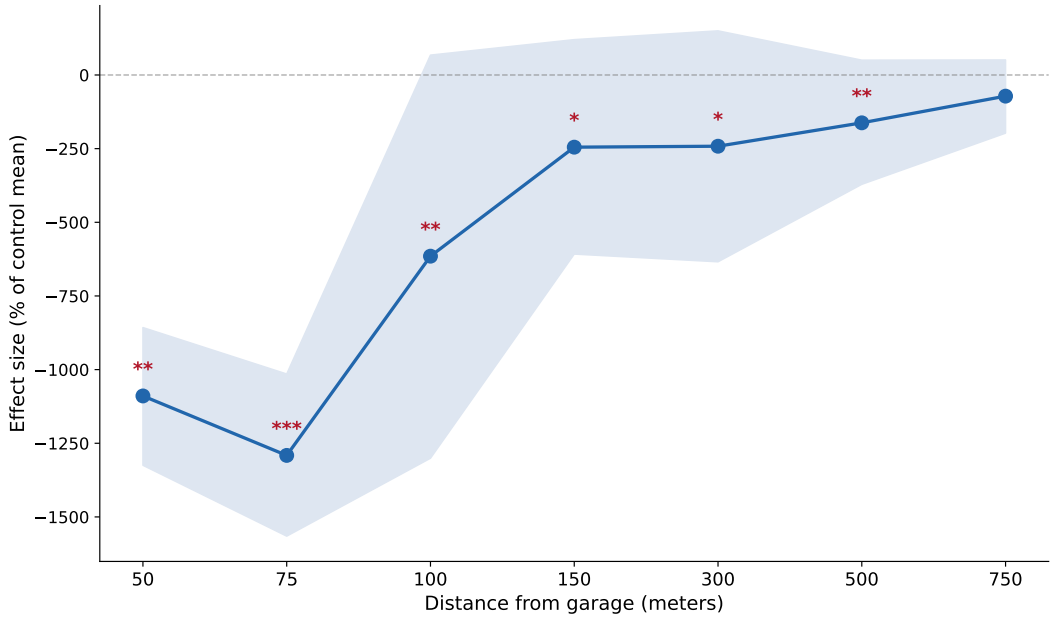


Figure 4: Garage proximity spillovers: effect size by distance

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 $= (\hat{\beta}_1 \pm 1.96 \times \text{SE})/\bar{y}$, 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)
Panel A: Presumed mutual contacts ($\geq$ 5% overlap)			
Share treated	-0.179	$+2.686^{**}$	$+0.829^{*}$
	(0.344)	(1.024)	(0.379)
	$[0.631]$	$[0.045]$	$[0.055]$
Effect size	-7.5%	$+9.8\%$	$+21.4\%$
Control mean	2.38	27.35	3.88
N		116,613	
Panel B: Weaker ties ($<$ 5% overlap)			
Share treated	-0.195^{***}	$+0.038$	-0.003
	(0.056)	(0.159)	(0.023)
	$[0.003]$	$[0.814]$	$[0.910]$
Effect size	-9.9%	$+0.2\%$	-0.5%
Control mean	1.98	18.22	0.53
N		3,898,583	

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. Robust standard errors in parentheses; randomization inference p-values from 1,000 permutations in brackets. Effect size $= \beta/\bar{y}$, expressed as percent change. All panels use network fingerprint 240222. Edges with $I_{\text{obs}} = 0$ (mutual contacts sharing zero encrypted contact buckets) and $I_{\text{obs}} > 500$ are excluded. "Top 2% both" trimming applied: individuals and drivers in the top 2% of the degree distribution are excluded to reduce the influence of hub nodes. Panel A restricts to pairs with raw overlap share $\geq 5\%$, the threshold that minimizes mean absolute error between true (fingerprint-verified) and inferred contact counts; Panel B restricts to pairs with raw overlap $< 5\%$, isolating weak ties where social learning is minimal. $^{*}p_{\text{RI}} < 0.10$, $^{**}p_{\text{RI}} < 0.05$, $^{***}p_{\text{RI}} < 0.01$.

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

	Taxi adoption, Dakar		Cross-ind. adoption		
Subnetwork	Coef	Eff. size	Coef	Eff. size	N
Panel A: Taxi contacts by social proximity					
Taxi × Friends	10.19^{***}	$+37.4\%$			614
	(3.73)				
	$[0.010]$				
Taxi × Non-friends	3.29	$+10.4\%$			549
	(4.11)				
	$[0.435]$				
Taxi × Freq. daily/wk	9.20^{***}	$+32.1\%$			916
	(3.11)				
	$[0.004]$				
Panel B: All contacts by social proximity					
Friends	2.98^{**}	$+30.9\%$	0.12	$+1.9\%$	1,983
	(1.41)		(1.12)		
	$[0.049]$		$[0.917]$		
Freq. daily/weekly	1.63^{*}	$+23.0\%$	0.33	$+6.7\%$	4,119
	(0.85)		(0.70)		
	$[0.100]$		$[0.653]$		
Panel C: Casual ties					
Taxi clients	0.46	$+36.8\%$	3.51	$+93.7\%$	395
	(1.23)		(2.27)		
	$[0.792]$		$[0.151]$		
Serendipitous encounters	0.30	$+17.3\%$	1.79^{***}	$+47.7\%$	5,112
	(0.38)		(0.59)		
	$[0.465]$		$[0.003]$		

Notes: Each row estimates $y_j = \alpha + \beta_1 \cdot \text{ShareTreated}_{s,j} + \beta_2 \cdot \text{ExpectedExposure}_{s,j} + \varepsilon_j$ on the subsample with ≥ 1 connection in subnetwork s, restricted to baseline survey contacts. ShareTreated = num. treated / num. connections; ExpectedExposure = mean $P(\text{treated})$ across 1,000 simulated treatment draws. Coefficients $\beta_1 \times 100$ reported; HC1 robust SEs in parentheses. RI p-values (1,000 permutations) in brackets. Effect size = $\beta / \bar{y}_{\text{control}}$, in percent. "Taxi adoption, Dakar" = within-industry; "Cross-ind. adoption" = across-industry (Dakar + outside). Panel A: within-industry outcomes only (taxi contacts lack cross-industry variation). "Serendipitous encounters": 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
Panel A: All merchants near garages						
Any treated	-0.0104	-0.0068	-0.0074	-0.0049	-0.0036^{*}	-0.0016
	(0.0035)	(0.0055)	(0.0038)	(0.0022)	(0.0014)	(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
Panel B: Social ties (mutual phone contacts)						
Any treated	-0.1811^{**}	-0.1819^{***}	-0.0666^{**}	-0.0222^{*}	-0.0138^{*}	-0.0085^{**}
	(0.0198)	(0.0198)	(0.0377)	(0.0169)	(0.0114)	(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 **100 m** column in Panel B is the structural validation moment (m6), mapping to $\beta_{\text{learn}} + \beta_{\text{comp}}$. Effect size $= \beta_1/\bar{y}$. SEs in parentheses; RI p-values in brackets (1,000 permutations). $^{*}p_{\text{RI}} < 0.10$, $^{**}p_{\text{RI}} < 0.05$, $^{***}p_{\text{RI}} < 0.01$.

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	338.49***	716.57*	311.55***	13.38	25.14	13.15
	(42.40)	(165.37)	(43.69)	(7.99)	(98.81)	(8.01)
	[0.003]	[0.080]	[0.008]	[0.461]	[0.834]	[0.460]
Expected exposure	1335.09	4164.68	1219.04	−392.68	−1842.96	−366.85
	(1613.67)	(8343.70)	(1638.40)	(254.85)	(3412.26)	(255.43)
	[0.967]	[0.996]	[0.948]	[0.328]	[0.469]	[0.323]
Effect size	+45.3%	+93.7%	+41.8%	+2.2%	+5.0%	+2.2%
Dep. var. mean	747.93	764.71	745.72	606.05	498.67	609.74
N	10,045	1,173	8,872	56,670	1,879	54,791

Notes: Coefficients, SEs, and dep. var. means multiplied by 1,000. Sample: non-experimental merchants linked to non-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 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	+0.0835*	+0.0472	+0.2437
	(0.0377)	(0.0307)	(0.1366)
	[0.056]	[0.159]	[0.127]
Expected exposure	+0.5470	+0.2089	+0.8465
	(1.2099)	(1.0019)	(4.3406)
	[0.651]	[0.835]	[0.845]
Control mean	0.32	0.08	1.11
Effect size	+25.8%	+60.7%	+22.0%
Avg. connections	3.36	3.95	3.07
N	163,390	80,031	40,192

Notes: Coefficients, SEs, and control means multiplied by 100. Network restricted to mutual phone contacts with $\geq 5\%$ overlap share (Table 1 Panel B). Top 2% both trimming applied. Column (1): all individuals (replicates Table 1 Panel B cross-industry other). Column (2): individuals who appeared in Dakar at least once during the experiment (March–October 2022) per geolocation data. Column (3): individuals who never appeared in Dakar. 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: Split by Driver Migration ($\geq 5\%$ Overlap)

	Drivers Visited City				Drivers Did Not Visit City			
	(1) Mbour	(2) Thiès	(3) Touba	(4) St-Louis	(5) Mbour	(6) Thiès	(7) Touba	(8) St-Louis
Share treated	$+0.041$	-0.062	$+0.024$	-0.123	$+0.003$	$+0.037^{**}$	$+0.016$	$+0.031^{**}$
	(0.049)	(0.067)	(0.080)	(0.118)	(0.015)	(0.015)	(0.018)	(0.013)
	[0.462]	[0.323]	[0.776]	[0.790]	[0.852]	[0.024]	[0.420]	[0.032]
Expected exposure	-0.033	$+0.092$	-0.002	$+0.100$	-0.013	-0.070	-0.040	-0.054
	(0.057)	(0.102)	(0.111)	(0.196)	(0.051)	(0.052)	(0.066)	(0.126)
Control mean	0.06	0.06	0.08	0.05	0.04	0.04	0.06	0.03
Effect size	$+70.3\%$	-99.1%	$+28.8\%$	-244.8%	$+6.6\%$	$+96.3\%$	$+29.4\%$	$+115.8\%$
N				116,613				
Drivers visited	96 (52 T)	166 (97 T)	116 (74 T)	22 (14 T)	—	—	—	—

Notes: Coefficients, SEs, and control means multiplied by 100. Network restricted to mutual phone contacts with $\geq 5\%$ overlap share (Table 1 Panel B). Top 2% both trimming applied. "Drivers Visited City" columns use a city-specific treatment variable: $\text{ShareTreated}_j^{\text{visited}}$ = (number of j's treated driver contacts who visited city c at least once during the experiment) / (total driver contacts of j). "Drivers Did Not Visit City" uses the complementary set of treated drivers who never appeared in city c. Driver visits identified from mobile money geolocation data (AG_CITY), March–October 2022. Outcome: cross-industry adoption (non-taxi, non-Dakar) in the specific city during the experiment. Specification: $y_j = \beta_1 \cdot \text{ShareTreated}_j^c + \beta_2 \cdot \text{ExpExposure}_j^c + \alpha$. 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 7: Structural Model: Heterogeneous Learning Estimates and Out-of-Sample Validation

	Estimate	SE
Panel A: Structural parameters		
$\hat{\alpha}_{\text{cross}}$ (cross-industry intercept)	-5.433^{***}	(0.026)
$\hat{\alpha}_{\text{taxi}}$ (within-taxi intercept)	-8.770^{**}	(3.424)
$\hat{\beta}^{s}_{\text{learn}}$ (strong-tie learning, $\geq 5\%$)	$+0.100^{***}$	(0.037)
$\hat{\beta}^{w}_{\text{learn}}$ (weak-tie learning, $< 5\%$)	$+0.001$	(0.006)
$\hat{\beta}_{\text{comp}}$ (within-industry competition)	-0.102	(0.343)
$\hat{\beta}_{\text{demand}}$ (customer demand)	$+1.643^{***}$	(0.230)
Net phone (strong ties): $\hat{\beta}^{s}_{\text{l}} + \hat{\beta}_{\text{c}}$	-0.002	
Net phone (weak ties): $\hat{\beta}^{w}_{\text{l}} + \hat{\beta}_{\text{c}}$	-0.100	

#	Source	Channel	$\hat{b}$ (data)	Model pred.	Sign
Panel B: Identifying moments (exactly matched)					
m1	T1, cross-ind. $\geq 5\%$	β^{s}_{l}	$+0.00269$	$+0.00269$	—
m2	T1, cross-ind. $< 5\%$	β^{w}_{l}	-0.00019	-0.00019	—
m3	T1, taxi $\geq 0\%$	$w_s\beta^{s}_{\text{l}}+w_w\beta^{w}_{\text{l}}+\beta_{\text{c}}$	-0.00019	-0.00019	—
m4	T3 col. (3) Non-Soc. cust.	β_{d}	$+0.31155$	$+0.31155$	—
m5	Cross-ind. rate	α_{cross}	2.735%	2.735%	—
m6	Within-taxi rate	α_{taxi}	0.196%	0.196%	—
Panel C: Out-of-sample validation (untargeted predictions)					
m7	T4, garage 100m (social)	$\beta_{\text{l}}+\beta_{\text{c}}$	-0.06665	-0.00104	✓
m8	T2, friends	β^{s}_{l}	$+0.02978$	$+0.00872$	✓
m9	T2, taxi friends	$\beta^{s}_{\text{l}}+\beta_{\text{c}}$	$+0.10190$	$+0.00035$	✓
m10	T2, taxi non-fr.	$\beta^{w}_{\text{l}}+\beta_{\text{c}}$	$+0.03285$	-0.02173	×
m11	T2, serendip.	β^{w}_{l}	$+0.01706$	$+0.00005$	✓
Validation RMSE				0.061	
Sign accuracy				$3/5$	

	Cross-industry		Within-taxi	
	Rate	Δ	Rate	Δ
Panel D: Decomposition (6 months, $\bar{S} = 0.58$)				
Observed (data)	2.735%		0.196%	
Model baseline (no treatment)	2.581%	—	0.093%	—
+ social learning	2.589%	$+0.008$	0.093%	$+0.000$
+ competition only	—	—	0.088%	-0.005
Learning + competition	2.589%	$+0.008$	0.088%	-0.005
+ customer demand only	—	—	0.240%	$+0.147$
All channels	2.589%	$+0.008$	0.227%	$+0.134$

Notes: Minimum distance estimation. The model has 6 parameters matched to 6 identifying moments (m1–m6); the system is exactly identified (zero residual on identifying moments). Standard errors: delta method for α_{cross}, β^{s}_{learn}, β^{w}_{learn}, β_{demand}; numerical Jacobian for β_{comp}. For phone-network moments (m1–m3), model predictions use indirect inference: compute model-implied $P_j(\theta)$, run the same OLS as the reduced-form, extract the partial coefficient on S_j^{phone}. Customer moment (m4) uses AME at subsample mean. Rate moments (m5, m6) use exact sample averages. Panel C reports untargeted out-of-sample predictions. Strong ties ($\geq 5\%$ overlap) comprise 3.7% of the full sample; weak ties ($< 5\%$) the remaining 96.3%. The net phone effect is approximately zero for strong ties ($\hat{\beta}^{s}_{\text{l}} + \hat{\beta}_{\text{c}} = -0.002$) and negative for weak ties (-0.100). Panel D decomposes the treatment effect by channel. $^{*}p<0.10$, $^{**}p<0.05$, $^{***}p<0.01$.

Table 8: Counterfactual Targeting Simulations (6-Parameter Model)

	$\bar{S}^{\text{phone}}$	Cross-Industry Adopters			Within-Taxi Adopters			Total	
		Count	[95% CI]	Δ	Count	[95% CI]	Δ	Count	Δ
No treatment (baseline)	0.000	37,974	[37,617, 38,342]	—	719	[565, 946]	—	38,693	—
Actual random ($N_T = 1256$)	0.572	37,979	[37,450, 38,509]	—	1,975	[1,352, 3,034]	—	39,954	—
Strong-tie cross-industry ($N_T = 1256$)	0.818	37,985	[37,369, 38,625]	$+0.0\%$	2,729	[1,760, 4,477]	$+38.2\%$	40,715	$+1.9\%$
Garage-first ($N_T = 1256$)	0.575	37,979	[37,449, 38,511]	-0.0%	1,983	[1,356, 3,048]	$+0.4\%$	39,962	$+0.0\%$
Highest degree ($N_T = 1256$)	0.891	37,982	[37,322, 38,649]	$+0.0\%$	2,998	[1,905, 5,028]	$+51.8\%$	40,979	$+2.6\%$
Top MM users ($N_T = 1256$)	0.654	37,980	[37,430, 38,542]	$+0.0\%$	2,211	[1,477, 3,457]	$+11.9\%$	40,190	$+0.6\%$

Notes: Predicted adoption counts from the 6-parameter structural model (Table 7) under alternative treatment targeting rules. Budget fixed at $N_T = 1{,}256$ treated drivers. Garage-first selects all 228 garage drivers plus 1,028 random non-garage drivers. $\bar{S}^{\text{phone}}$ = mean share of treated phone contacts across the analysis sample. 95% CIs from parametric bootstrap (500 draws from $N(\hat{\theta}, \hat{V}(\hat{\theta}))$). Δ = percentage change vs. actual random allocation. Cross-industry adopters predicted using $\hat{\beta}^{s}_{\text{learn}}$ for strong ties ($\geq 5\%$ overlap) and $\hat{\beta}^{w}_{\text{learn}}$ for weak ties. Within-taxi adopters include learning, competition, and customer demand channels. Cumulative adoption over $T = 6$ months.

Table 9: Counterfactual Targeting: Combined Reachable Network ($N = 205{,}669$)

	$\bar{S}^{\text{phone}}$	$\bar{S}^{\text{demand}}$	Within-Taxi		Cross-Ind. (Dakar)		Cross-Ind. (Other)		Total	
			Induced	Δ	Induced	Δ	Induced	Δ	Induced	Δ
No treatment (baseline)	0.000	0.000	*Baseline: 8 taxi + 5,088 cross-Dakar + 5,280 cross-other = 10,376*							
Actual random ($N_T = 1{,}256$)	0.568	0.534	15	—	176	—	15	—	205	—
Garage-only ($N = 228$)	0.092	0.081	2	−84%	28	−84%	2	−83%	33	−84%
Top MM users ($N_T = 1{,}256$)	0.662	0.590	17	+11%	208	+18%	16	+11%	241	+17%
Highest degree ($N_T = 1{,}256$)	0.899	0.828	25	+69%	280	+59%	22	+49%	327	+59%
Geographic reach ($N_T = 1{,}256$)	0.891	0.821	25	+68%	277	+58%	22	+51%	324	+58%

Notes: Treatment-induced adoption = predicted 6-month cumulative adoption − no-treatment baseline. Sample is the union of four populations with at least one active structural channel: (A) 116,613 individuals with $\geq 5\%$ overlap mutual phone ties, after top-2% trimming (social learning), (B) 8,583 taxi with any mutual phone tie (learning + competition + demand), (C) 37,543 non-taxi within 500m of a garage (geographic proximity), (D) 66,844 non-taxi with shared-customer links (demand proximity). 56.7% have strong ties ($\hat{\beta}^{s}_{\text{learn}} = +0.100$); the remainder have weak ties ($\hat{\beta}^{w}_{\text{learn}} = +0.001$) but are reachable through other channels. Customer-weighted demand model: learning/competition use S^{phone}; demand uses S^{demand}. Overlap filter consistent with Table 1 regression sample: $I_{\text{obs}} > 0$, $I_{\text{obs}} \leq 500$, top-2% degree trimming. Δ = % change in induced adoption vs. actual random. 95% CIs (bootstrap, 500 draws) for total: random [10,350, 12,360]; highest degree [10,458, 12,916].

Online Appendix

Nationwide Diffusion of Technology Within Firms' Social Networks

Deivy Houeix

A Appendix Figures **52**

B Appendix Tables **60**
B.1 Distance Decay of Competition *67*

A Appendix Figures

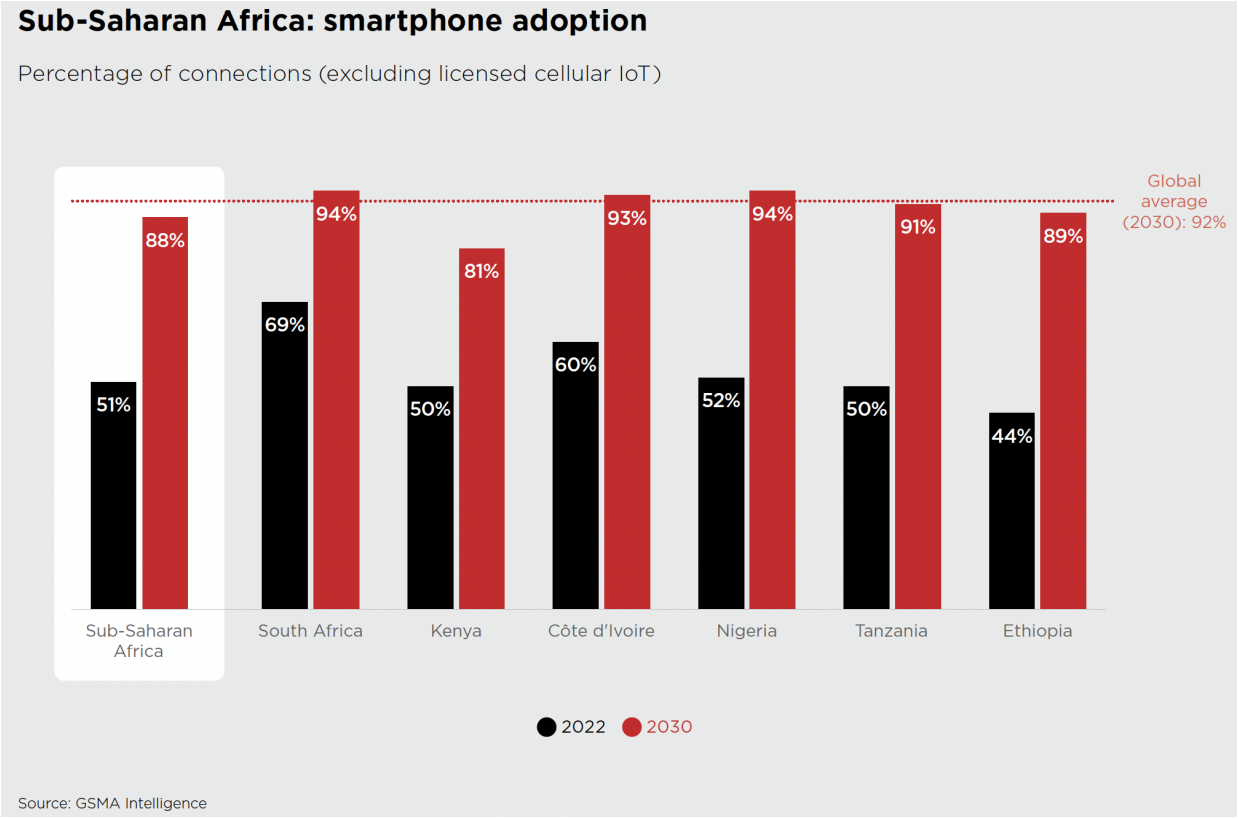


Figure A1: Smartphone Adoption in Sub-Saharan Africa

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. Cote 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.

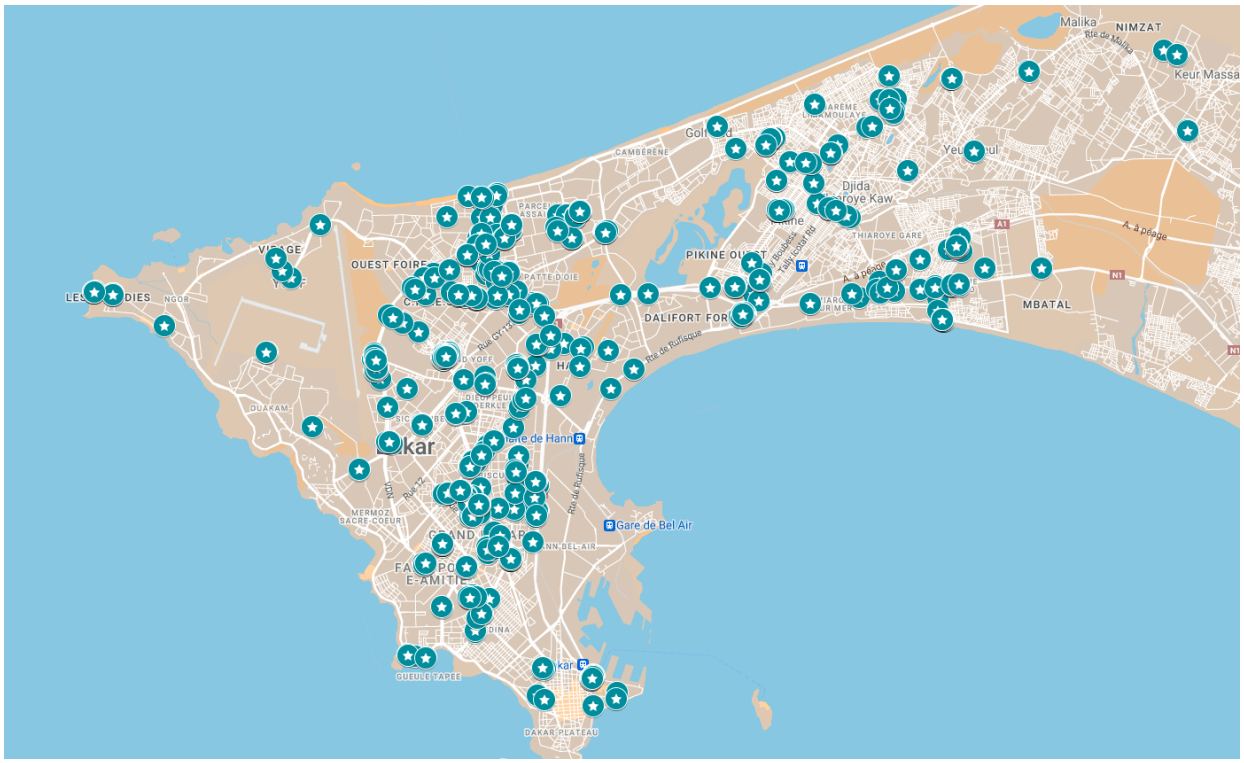


Figure A3: Listing Survey Activity in Dakar - March/April 2022

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.

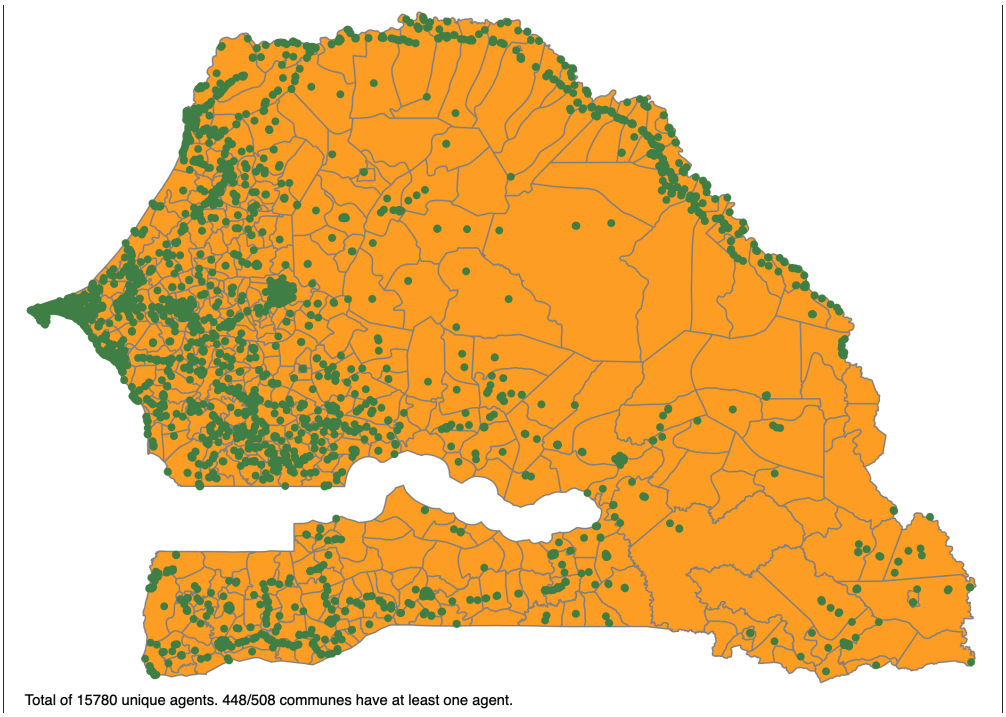


Figure A4: Mobile Money Agents in Senegal in 2022

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.

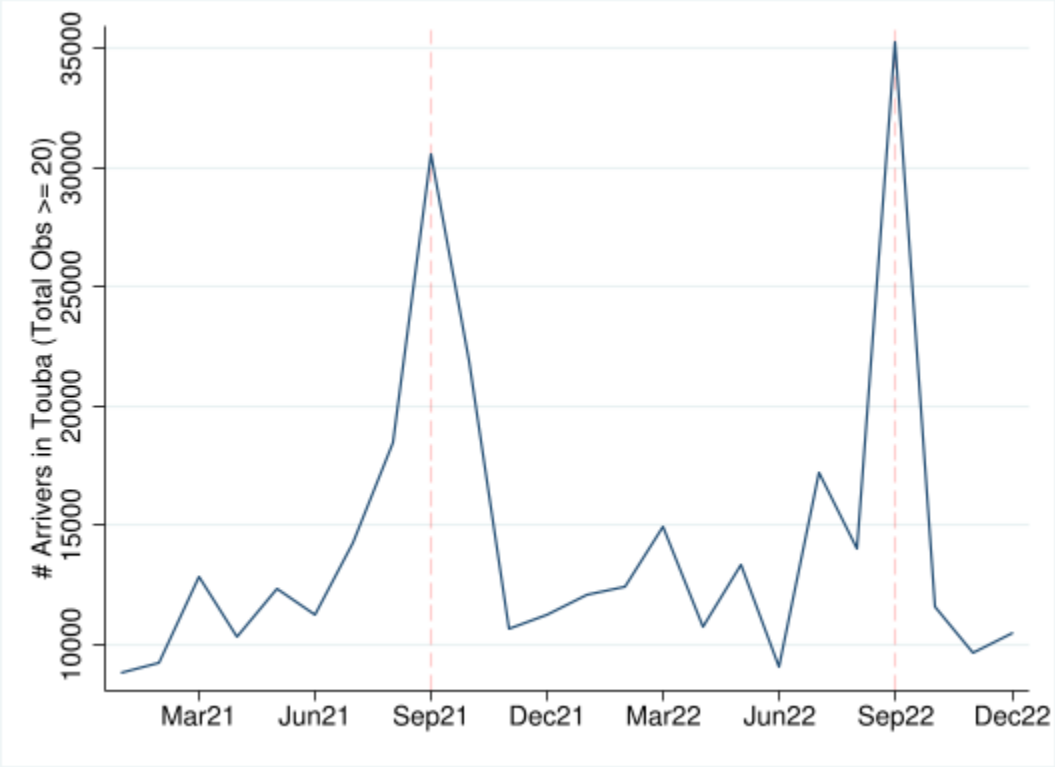


Figure A5: Location Data - Arrivals in Touba

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.

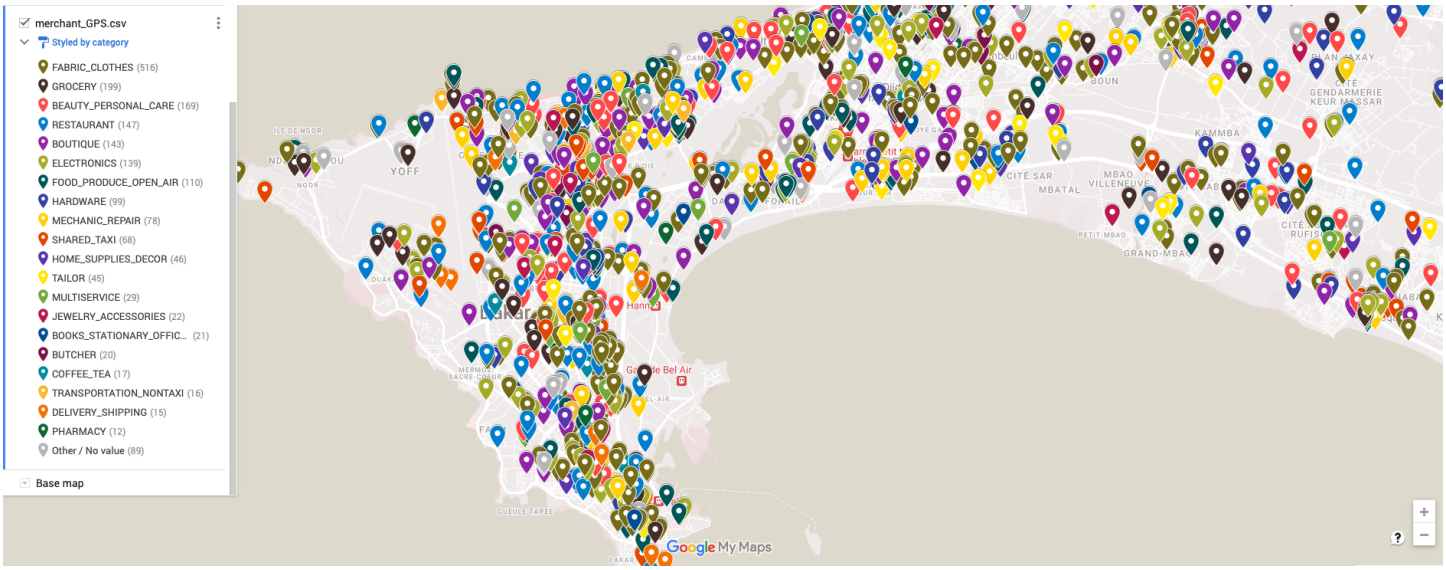


Figure A6: Technology Diffusion in Dakar Throughout the Experiment

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.

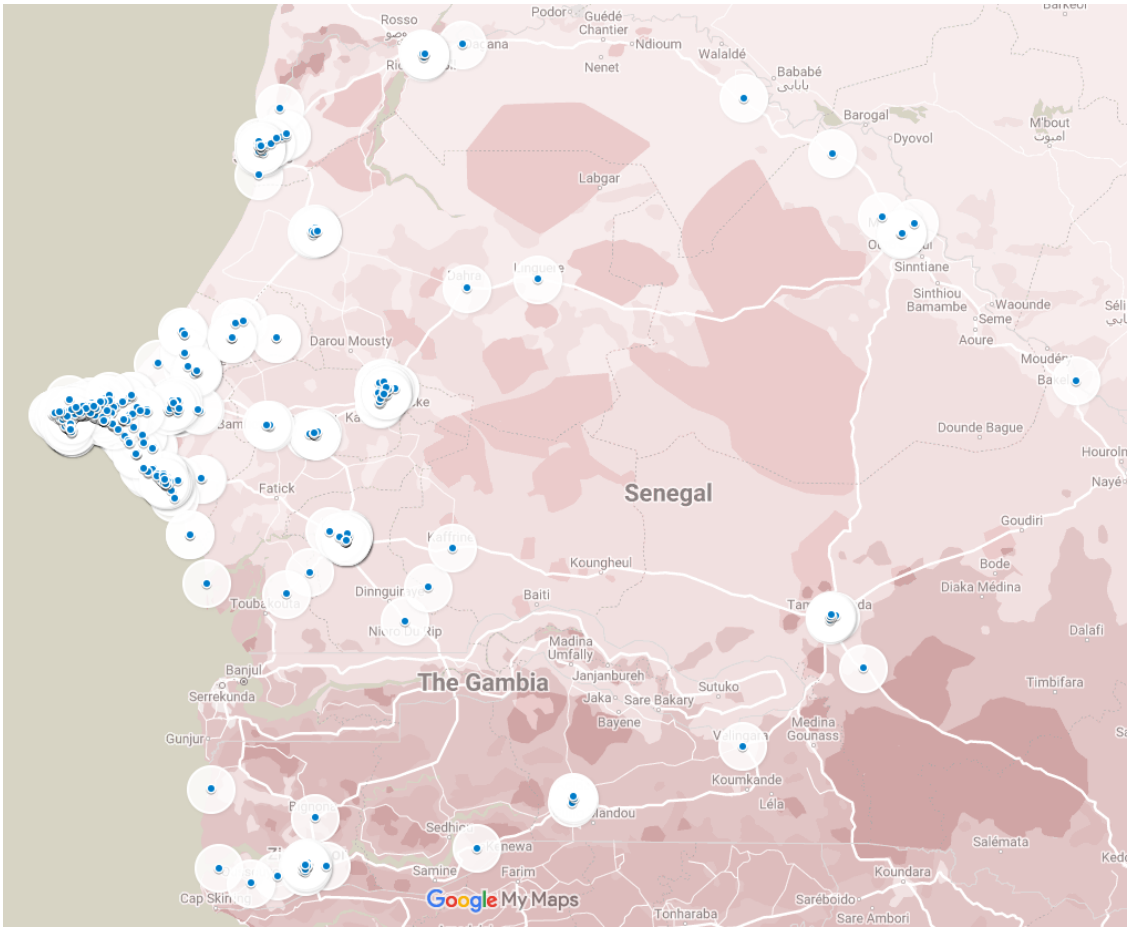


Figure A7: Technology Diffusion Inter-Industry Across Space Throughout the Experiment

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.

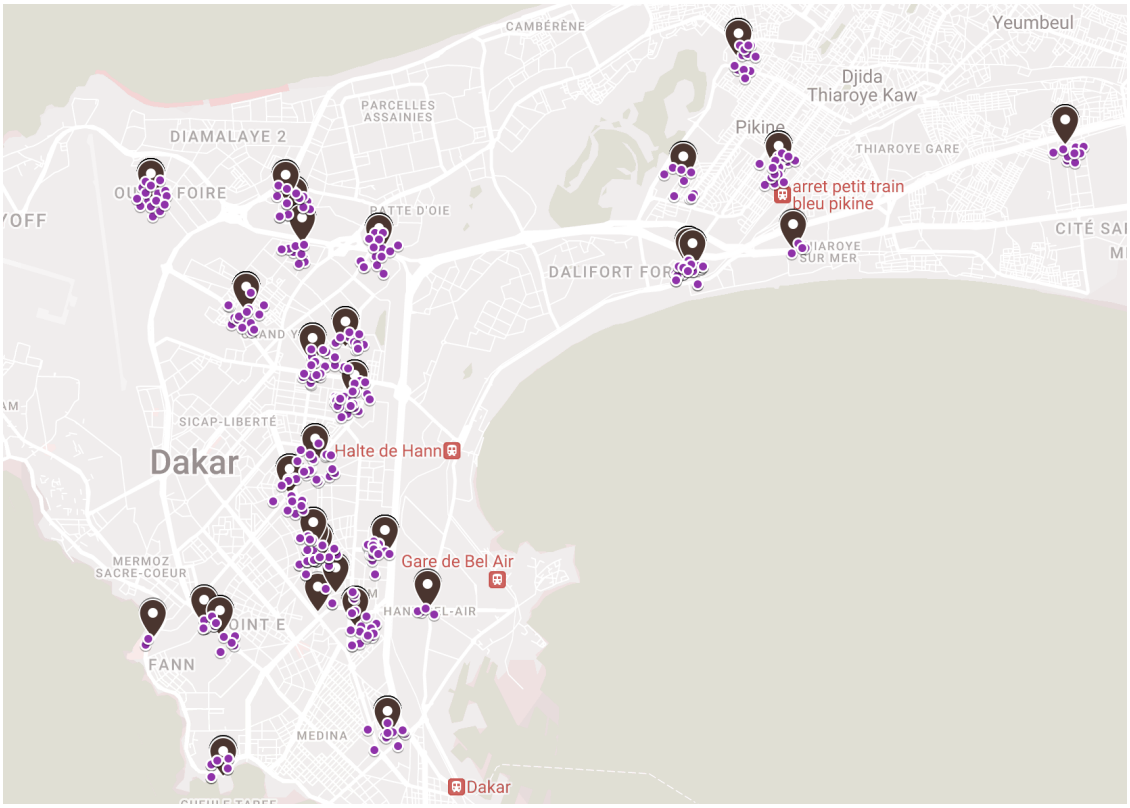


Figure A8: Local Technology Adoption Spillovers Nearby Taxi Garages

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.

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	p-value
Panel A: Contact book size					
Contact book size (fingerprint)	585	614	-28.8	24.5	0.240
Mutual contacts detected	8,309	8,982	-673	352	0.056*
Panel B: Contact composition					
N merchant (among mutual)	342	370	-27.9	14.6	0.056*
N taxi merchant (among mutual)	28.0	31.0	-2.99	1.25	0.017**
N non-taxi merchant (among mutual)	314	339	-24.9	13.4	0.063*
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
Panel C: Overlap distribution					
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		
Kolmogorov–Smirnov tests					
Contact book size (fingerprint)		KS $= 0.0632$		$p = 0.0473$	
Mutual contacts detected		KS $= 0.0587$		$p = 0.0791$	
Mean overlap (among mutual)		KS $= 0.0443$		$p = 0.3135$	
Mutual contacts ¿=5% overlap		KS $= 0.0256$		$p = 0.9125$	
Share taxi (among mutual)		KS $= 0.0544$		$p = 0.1248$	
Joint F-test		$F = 0.763$		$p = 0.5765$	

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: Kinship Survey in the Transaction Network - Summary Statistics

	Share
Panel A. Nature of the Relationships	
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)
Panel B. Frequency of Conversations in Past 3 Months	
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)
Panel C. Information Sharing	
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 B4: Transaction Network Spillovers

	Taxi, Dakar (within)	Cross-ind., Dakar (across, local)	Cross-ind., Other (across, distant)
Share treated	$+2.937$	$+3.901$	-0.663
	(1.029)	(1.577)	(0.431)
	$[0.139]$	$[0.128]$	$[0.206]$
Expected exposure	$+65.780$	$+15.035$	$+25.621^{*}$
	(32.807)	(49.220)	(13.662)
	$[0.918]$	$[0.938]$	$[0.091]$
Effect size	$+13.7\%$	$+8.2\%$	-20.8%
Control mean	21.51	47.48	3.19
N		74,199	

Notes: Coefficients, SEs, and control means multiplied by 1,000. OLS 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 SEs in parentheses; RI p-values (1,000 pre-computed treatment draws) in brackets. Effect size $= \beta/\bar{y}$. Transaction network from pre-treatment mobile money transfers (2020–2022Q1). "Exclude incoming-only": retains edges where the driver sent ≥ 1 transfer to individual j. $^{*}p_{\text{RI}} < 0.10$, $^{**}p_{\text{RI}} < 0.05$, $^{***}p_{\text{RI}} < 0.01$.

Table B5: 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*****	122.16***	−7.46	−7.44
	(23.51)	(23.57)	(13.14)	(13.17)
Expected exposure	−196.24	−187.82	641.61	633.78
	(761.38)	(761.85)	(421.51)	(422.60)
Effect size	+17.9%	+17.9%	−1.0%	−1.0%
Dep. var. mean	679.49	680.80	761.12	760.62
N	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, $\geq$ 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 B6: Customer Network Temporal Stability Diagnostic

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

Notes: Each column constructs the shared-customer network G^{cust} from transactions within that window only ($\geq$ 2 transactions per customer). Garage-affiliated drivers excluded. Outcome: within-taxi DKR mobile payment adoption by endline. ShareTreated $= \sum T_i / N_j^{\text{cust}}$ 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 B7: 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	$\mathbf{288.83}^{***}$	$\mathbf{822.93}^{*}$	272.07^{***}	-26.85^{*}	-162.25	-26.24^{*}
	(31.46)	(178.02)	(31.91)	(8.36)	(124.73)	(8.37)
	[0.004]	[0.082]	[0.006]	[0.066]	[0.258]	[0.073]
Expected exposure	-1769.78	11500.64	-2049.02	467.95	2012.04	464.16
	(1228.62)	(7187.00)	(1239.14)	(269.58)	(4589.93)	(270.01)
	[0.049]	[0.976]	[0.056]	[0.599]	[0.095]	[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
N	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. $^{*}p_{\text{RI}} < 0.10$, $^{**}p_{\text{RI}} < 0.05$, $^{***}p_{\text{RI}} < 0.01$.

B.1 Distance Decay of Competition

This appendix details the estimation of competition amplification factors from the garage proximity coefficients in Table 3, Panel B (social ties).

Procedure. The structural model predicts that the garage coefficient at radius r loads on $\beta^s_{\text{learn}} + \beta_{\text{comp}} \times w(r)$, where $w(r)$ is the competition amplification at distance r and $w = 1$ at the phone-network baseline by normalization. I recover $w(r)$ in three steps.

First, for each radius $r \in \{50, 75, 100, 150, 300, 500\}$ m, convert the reduced-form coefficient $\hat{b}(r)$ from Table 3 Panel B to the latent-index scale using the subsample adoption rate $\bar{p}(r)$:

$$\widehat{\text{net}}(r) = \frac{\hat{b}(r)}{\bar{p}(r)\,(1 - \bar{p}(r))}.$$

Second, extract $w(r)$ using the already-identified learning and competition parameters:

$$w(r) = \frac{\widehat{\text{net}}(r) - \hat{\beta}^s_{\text{learn}}}{\hat{\beta}_{\text{comp}}}.$$

Standard errors on $w(r)$ are propagated from the garage coefficient standard errors via the delta method. Third, fit parametric decay curves to the six $(r, w(r))$ points by nonlinear least squares.

Amplification factors. The estimated amplification and fitted values under three functional forms are:

Radius (m)	$\hat{b}(r)$	$\bar{p}(r)$	$w(r)$	w^{power}	w^{exp}	w^{inv}
50	-0.181	0.017	139	157	152	154
75	-0.182	0.014	164	114	127	119
100	-0.067	0.011	79	91	107	97
150	-0.022	0.009	32	66	75	71
300	-0.014	0.006	32	39	26	39
500	-0.009	0.005	21	26	6	25

Parametric decay curves. Three functional forms are fitted:

- Power law: $w(d) = A \cdot d^{-\gamma}$ ($\hat{A} = 3{,}367.5$, $\hat{\gamma} = 0.783$)
- Exponential: $w(d) = A \cdot e^{-d/\lambda}$ ($\hat{A} = 216$, $\hat{\lambda} = 142$ m)
- Inverse distance: $w(d) = A/(1 + d/\lambda)$ ($\hat{A} = 371.6$, $\hat{\lambda} = 35$ m)

The exponential provides the best overall fit. Competition amplification ranges from $139\times$ at 50 m to $21\times$ at 500 m, confirming steep decay with distance.

Effective phone-contact distance. Setting $w(d^*) = 1$ under the power law yields $d^* = A^{1/\gamma} = 3{,}367.5^{1/0.783} \approx 32$ km. This is approximately the diameter of metropolitan Dakar, consistent with the average phone contact being dispersed across the city rather than co-located at a garage.