September 15, 2026

NBER Workshop on Network and Development Economics
National Bureau of Economic Research
1050 Massachusetts Avenue, Cambridge, MA 02138

Dear Professors Breza, Chandrasekhar, and Viviano,

My name is Deivy Houeix. I am a Prize Fellow at Harvard University. I received my PhD in Economics from MIT in 2025. My advisors were Ben Olken (bolken@mit.edu), Esther Duflo (eduflo@mit.edu), David Atkin (atkin@mit.edu), and Tavneet Suri (tavneet@mit.edu). Feel free to contact them, or let me know if you would like them to provide a recommendation.

I study how firms in low-income countries grow and adopt new technologies. Most of my papers reflect my interest in relationships within firms, between owners and their employees, and across firms, whether horizontally, among firms linked by a shared market or a social tie, or vertically, between a firm and its customers and suppliers. The enclosed paper is a new draft and an example of this line of work. Using novel network data, I investigate how a technology spreads between firms across industries and across space, and through which connections.

I would welcome feedback on the paper, as I am still working through the framing and some of the results. One example: exposure to treatment is defined on overlapping networks, and I would love to brainstorm with others about inference and exposure design under interference. I also hope to learn about empirical network research and to meet other researchers working at the intersection of networks and development.

My CV and the paper follow.

Thanks so much for your consideration.

Deivy Houeix

# Deivy Houeix

deivyhoueix@fas.harvard.edu ◇ deivyhoueix.com

**Current Positions**

Assistant Professor of Economics, Columbia University, Graduate School of Business 2027 -

Prize Fellow in Economics, History and Politics, Harvard University 2025 - 2027

**Education**

Ph.D. in Economics, Massachusetts Institute of Technology (MIT) 2025
*Advisors*: Benjamin Olken, Esther Duflo, David Atkin, Tavneet Suri

M.Res. in Economics, Ecole Polytechnique and ENSAE Paris 2017

B.A. in Economics, Ecole Normale Superieure (ENS) Paris-Saclay 2015

**Previous Relevant Positions**

Research Manager, Stanford University 2017–2019
Prof. Pascaline Dupas and Prof. Marcel Fafchamps

Research Assistant, World Bank Development Research Group 2016–2017
Dr. Quy-Toan Do

**Research Interests**

Development Economics, Organizational Economics, Trade

**Working Papers**

1. Asymmetric Information and Digital Technology Adoption: Evidence from Senegal
2. Nationwide Diffusion of Technology Within Firms' Social Networks
3. Relational Frictions Along the Supply Chain: Evidence from Senegalese Traders (with Edward Wiles)
4. Financial Inclusion and Rural Electrification: Evidence from Togo (with Paul Brimble, Axel Larrinaga, and Toni Oki)
5. Eliciting Poverty Rankings from Urban or Rural Neighbors: Methodology and Empirical Evidence (with Pascaline Dupas and Marcel Fafchamps). R&R, *Quantitative Economics*

**Non-Refereed Publications**

"Technology and Development" (with Julieta Caunedo, Tommaso Porzio, David Argente, Jaedo Choi, Yulu Tang, Danial Lashkari, Jacob Moscona, Karthik Sastry, Federico Rossi, Luisa Cefala, Erin Kelley, Conor Walsh), *VoxDevLit 23(1)*, June 2026.

"Asylum seekers in the European Union: building evidence to inform policy making" (with Mohamed Abdel Jelil, Paul Andres Corral, Anais Dahmani, Maria Davalos, Giorgia Demarchi, Neslihan Demirel, Quy-Toan Do, Rema Hanna, Sara Lenehan, and Harriet Mugera), *World Bank Flagship Report*, June 2018.

"Urban Development in Africa: Preliminary Report on the Addis Ababa SEDRI Study" (with Girum Abebe, Daniel Agness, Pascaline Dupas, Marcel Fafchamps, Tigabu Getahun), *Report*, April 2018.

**Teaching**

JPAL-ENSEA Policy Impact Evaluation, Abidjan 2025

Development Economics: Macroeconomics, PhD, TA for Prof. Robert Townsend (MIT 14.772) 2024

Applied Economics for Managers, MIT Sloan Executive MBA, TA for Prof. Tavneet Suri (MIT 15.024) 2024

JPAL Development Methodologies Summer School, Abidjan and Rabat 2022, 2023

*Guest Lectures*: Harvard (The Economics of Human Mobility, 2025); Brown (Development Economics, 2024)

**Awards, Honors, and Fellowships**

Ronald H. Coase Best Dissertation Award Honorable Mention, SIOE 2026

Daniel Cohen Award 2025

J-PAL Fellowship 2023-2024

MIT Presidential Graduate Fellowship 2019-2020

Best Master's Thesis Award 2017

Full-Ride Academic Scholarship - Elève Normalien 2014-2018

Valedictorian - B.A., and M.Res. 2015 and 2017

**Research Grants**

"Adoption and Impacts of Digital Payment Technologies"
Private Enterprise Development in Low-Income Countries (PEDL), The Shultz Fund, The Weiss Fund, Digital Identification and Finance Initiative in Africa (JPAL DigiFI Africa)

"Relational Frictions Along the Supply Chain: Evidence from a Randomised Experiment Among Senegalese Traders"
Private Enterprise Development in Low-Income Countries (PEDL), The Shultz Fund, SurveyCTO Research Grant

"Hidden Barriers to Rural Electrification: Evidence from Togo"
Private Enterprise Development in Low-Income Countries (PEDL) - Climate Change

"Eliciting Poverty Rankings from Urban or Rural Neighbors: Methodology and Empirical Evidence"
Innovations for Poverty Action (IPA) - Research Methods Initiative

"Digitalization of Local Tax Collection in Côte d'Ivoire"
USAID - Development Innovation Ventures (DIV), JPAL Governance Initiative

**Conferences and Seminars**

Central European University CEU Vienna; Cornell Economics; CSAE Conference Oxford; PEDL 15th Anniversary Conference; NOVAfrica Nova School of Business and Economics; Yale PEDL Early-Career Researchers Workshop 2026

Harvard Business School TOM; Harvard Economics; World Bank DECRG; Oxford Economics; Northwestern Economics; UCLA Economics; Columbia Business School; Wharton BEPP; GPRL Rookiefest 2025 (deferred); IGC 9th Urbanization and Development Conference - Cape Town South Africa; PEDL 2025 Workshop on AI, Finance and Entrepreneurship - Nairobi Kenya; NBER International Trade SI; NBER Development SI; Daniel Cohen Award PSE; USC Economics; Dartmouth College; Yale Firms, Trade, and Development Conference; Stanford - A Tribute to Marcel Fafchamps; Brown Economics 2025

Harvard Cities & Development Workshop; Yale Firms, Trade, and Development Conference 2024

CEPR WEFIDEV Finance and Development - Northwestern 2023

IPA Methods and Measurement Conference - Northwestern 2021

**Refereeing**

American Economic Journal: Applied Economics, American Economic Review, Economic Development and Cultural Change, Journal of Development Economics, Journal of the European Economic Association, Journal of Public Economics, Journal of Political Economy, The Quarterly Journal of Economics,

# Nationwide Diffusion of Technology: Cross-Industry Learning Through Networks*

Deivy Houeix†

*PRELIMINARY DRAFT, DO NOT CITE*

September 16, 2026

**Abstract**

How does a new technology spread among informal firms, and how far? I randomize access to a digital payment technology among 2,196 taxi businesses in Dakar, Senegal, and trace its diffusion through new network data covering most informal firms in the country: social ties (phone contacts), location (GPS pings), and shared customers (mobile money ledgers). I show how the technology spread across the Senegalese economy through knowledge spillovers. Adoption rises across industries and cities, partly through weak ties and remote interactions, and most where knowledge is scarce and least among rivals within the same industry. Knowledge travels along social ties that span the country, so distance may be a poor proxy for how far a spillover reaches.

*I am grateful to Isaiah Andrews, David Atkin, Abhijit Banerjee, Luisa Cefalà, Esther Duflo, Dean Eckles, Tishara Garg, Justine Knebelmann, Sahil Loomba, Vishan Nigam, Ben Olken, Kunal Sangani, Frank Schilbach, Tavneet Suri, Nicholas Swanson, Rob Townsend and participants in 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 the 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

How does a new technology spread among informal service-sector firms, and how far does it reach? Since Marshall (1920), the leading paradigm has been that proximity speeds the flow of ideas between firms: the "mysteries of the trade . . . are as it were in the air." A large literature establishes that geography matters: firms in denser environments are more productive and innovate more (Moretti, 2021; Baum-Snow et al., 2024), and diffusion between firms is a first-order determinant of productivity and growth (Arrow, 1962; Romer, 1990; Grossman and Helpman, 1991). Yet the forces governing spillovers are typically documented within a single, bounded geographic area, which leaves open how local each force actually is and whether distance is a good proxy. The distinction matters for reach: a local spillover may require a few seeds in every place, whereas a network spillover may target nodes to reach a country.

The difficulty is that proximity bundles two relationships that are rarely observed separately. Firms that are close to each other tend to transmit useful knowledge from one to the other (*learning*). Nearby firms also often share customers, so that one firm's adoption raises another's return (*demand*), especially for two-sided technologies like digital payments.

Unpacking these channels requires three ingredients that are rarely available together: an exogenous technology shock, the adoption data of connected firms, and, importantly, the distinct networks along which information and market incentives operate. This paper brings the three pieces together. I randomize early access to a new digital payment technology among 2,196 taxi owners and drivers, spanning 1,454 businesses, in Dakar, Senegal, with 58% assigned to treatment. I trace the technology's diffusion well beyond the experimental sample using three economy-wide data sources from the country's leading mobile money company: hashed phone-contact fingerprints covering roughly 7 million people, about 70–80% of the adult population; the complete ledger of mobile money transactions; and agents' GPS pings covering the locations of most informal businesses and mobile money users. The technology spreads on its own to more than 230,000 businesses: other taxis, firms in other industries, and firms in cities hundreds of kilometers from Dakar.

This breadth of diffusion allows me to identify mechanisms often difficult to pin down: standard spillover designs measure indirect effects on a held-out control group inside the experiment that is often local and defined by the study itself (Crépon et al., 2013; Baird et al., 2018). Here, I measure the organic diffusion of the technology across Senegal's entire service sector from a single experimental seed.

The mechanisms I highlight are general, and the taxi business can be seen as a typical small informal firm in a lower-income country. It usually has a single owner, it serves

many customers each day, and its owner has many social ties. A new general-purpose technology such as digital payments reaches such a firm the same two ways it reaches any other: through what the owner hears and through whom the firm serves.

I measure the social and customer networks separately, and I locate every firm with agent GPS pings, so that learning can be traced through social ties and demand through shared customers. The *social network* is constructed from contact-book fingerprints, which are, to my knowledge, novel in economics. On first login, users are asked to share with the app a hashed fingerprint of their phone contact book, which the operator uses to facilitate transfers to saved contacts; 86% consent. Each saved number is hashed into one of 65,536 buckets, so I observe the set of buckets a user saved without recovering any number. I link two users when each appears in the other's contact book and measure tie strength by their overlap share: how much of their contact lists they hold in common, net of expected hash collisions. The approach draws on the computer science literature on secure set-similarity measurement, which starts with Broder (1997), and was extended to privacy-preserving settings by De Cristofaro et al. (2013) and Swamidass et al. (2015). Since most firms in this setting are owner-operated, the fingerprints provide an unusually direct measure of the social networks of firm owners, covering both mutual contacts and contacts with high overlap. The resulting network links all experimental taxi business owners to other mobile money users across Senegal, at varying degrees of tie strength. To validate this novel measure of social network, I conduct an independent relationship-listing survey. I draw a random sample of each respondent's past transaction partners and ask them to describe their relationship. This survey confirms that corrected overlap tracks social closeness—spouses share 20.8% of contacts, taxi colleagues 13.6%, family members 11.4%, and friends 6–9%, while serendipitous contacts (phone numbers a respondent could not place when asked, typically one-time transaction partners) average only 3.2% and customers average 1.9% corrected overlap, just below the 2% mark. I define a *community tie* as a corrected overlap of at least 2%, so that the typical customer relationship falls outside the social network (results are not sensitive to this definition, and thresholds of 2, 5, and 10% are analyzed throughout the paper).

The other two data layers come from location and transaction data. First, the location data place each firm and each contact in space, which is what the distance profiles of the two networks and the migration tests use. Geolocated cash-in and cash-out transactions at mobile money agents not only locate users geographically but also record who moves between cities throughout the experiment. Second, the *customer network* is built from the mobile money transaction ledger: two firms are linked if they share regular customers—people who make repeated payments to both firms, including at least one experimental

driver. This ledger is a rare window onto customer-sharing links, which are typically unobservable, especially in informal economies where cash transactions and the absence of value-added-tax records leave no trail of who transacts with whom. The network is measured both before the experiment (peer-to-peer, when business payments did not exist), and during/after (peer-to-business, when treated and then control taxis gain access to digital payments and appear as payees).

Three empirical patterns in these networks motivate the paper. First, social ties decay with distance but are far from local: 66% of a taxi owner's community ties, and the same share of the strong ones, lie more than 30 km from home, most of them in urban centers across the country. Second, customer links are local: the median link joins a driver to a business 7.6 km away, 62% of links lie within 10 km of the driver's home, and only 10% reach beyond 30 km. The two networks therefore stretch over different distances. Third, contact books track a firm's commercial life: merchants with a 10% larger contact book receive about 2% more payments and about 7% more value. The relationship is descriptive, but it ties network size to revenue. Together, these patterns suggest that learning can travel wherever social ties reach, while demand stops where shared customers stop. They also raise a distributional concern: if diffusion runs through contact books, and contact books grow with revenue, the technology may reach profitable firms first and the rest late or never.

My empirical strategy leverages the random assignment: treated drivers gain early access to the technology while control drivers are excluded for the first six months, generating clean exogenous variation in firms' exposure to early adopters. I measure each firm's exposure as the share of its experimental taxi connections assigned to treatment, and identify spillovers from variation in that share conditional on *expected* exposure under the randomization (Borusyak and Hull, 2023); a count parameterization yields similar results. Expected exposure depends on how connected a firm is and on which strata its experimental contacts fall into, so conditioning on it holds network position fixed and removes the confound that more central firms, those with many contacts or many nearby rivals, may differ from less central ones for reasons unrelated to treatment.

I write a simple framework to organize the reduced-form evidence around the learning vs. demand effect. Learning can be seen as a gate, with a firm adopting only once informed by a social tie, and demand is the payoff, since an informed firm adopts when enough of its customers are willing to pay digitally. The two multiply, and can be separated by reach. Demand requires a shared customer whereas learning only requires a tie. The framework yields three main predictions. (1) Adoption should rise with treated ties across industries and cities, especially where baseline information is scarce. (2) These ties

should transmit less within the taxi industry due to rivalry, except through friends. (3) If ties carried demand, their effect would run through riders who paid a treated driver, whereas learning would raise usage whether or not customers are shared.

I have three main findings. First, I identify *social learning* by documenting spillover effects on firms in other industries that are connected to treated taxi drivers. A non-taxi firm neither competes with a taxi nor typically shares its customers, so demand is muted; a positive adoption response to a treated taxi contact therefore primarily reflects learning. In the reduced form, the effect is significant on strong ties: non-taxi firms connected to treated taxis through high corrected-overlap contacts ($\geq 5\%$) adopt at a rate $+39.0\%$ above the control mean ($p = 0.017$), while the weak-tie estimate is positive but imprecise. Per-tie transmission is approximately uniform across the strength distribution. Survey evidence corroborates this channel: taxi businesses report sharing advice or information with about 74% of the people they know, and having discussed the technology with, or recommended it to 12% of them. Interestingly, the spillovers appear in the community network, but not among direct mutual phone contacts. Because control drivers were also told about the technology, mutual contacts are more likely to be infra-marginal. Their mutual contacts learn as much as they would learn from a treated tie. Community contacts instead learn by observation, either directly or through second-order links (treated drivers anecdotally report posting photos of the digital technology on social media). Transmission comes from well-connected treated drivers, including in Dakar for drivers whose reach extends across the country, and it is strongest where knowledge is scarce.

Spillovers expand to firms in cities beyond Dakar, where effect sizes are roughly twice as large. Such reach contrasts with the spirit of Marshall's local spillovers and with what has been documented in the literature so far. A natural candidate for how diffusion can cross cities is seasonal migration, which is pervasive in this context—roughly a third of Dakar-based business owners are observed moving outside the capital during the experiment. Yet I find that physical travel is not necessary for transmission: the effect is precisely estimated through treated drivers who never visited the destination city ($p = 0.018$), and it is significant among receiving businesses never observed in Dakar ($p = 0.027$). The effect of exposure through drivers who did travel is larger in magnitude but imprecisely estimated, consistent with migration reinforcing the remote channel. This shows the extent to which information travels across space in this economy, within the community network, without requiring face-to-face interaction.

Second, within the taxi industry the same ties transmit little: an additional treated tie does not raise adoption at any threshold, in and outside Dakar. I can be read as rivalry within the industry (Cefalà et al., 2025; Luo et al., 2025). Friendship offsets this attenua-

tion: a driver with a treated taxi *friend*, as identified in the survey, adopts at a higher rate than one with a control friend ($p = 0.007$).

Third, the tie effect does not run through customers. Digital payments are two-sided, so a business's return to adoption depends on whether its customers pay digitally, and the tie effect could in principle reflect riders who first paid a treated driver and then asked other drivers to accept the same. Two tests show otherwise. First, among within-taxi adopters, riders who had paid a treated taxi account for a quarter of payments, and social exposure does not raise that share under any of several definitions of such a rider. Second, among adopters, usage rises with social exposure, but it rises most outside Dakar, where customers cannot be shared with the experimental drivers. What travel along the social tie is thus more consistent with know-how.

This paper contributes to three literatures. First, it clarifies what distance measures in the diffusion of technology between firms. Development economics has typically focused on local networks: neighbors, villages, clusters, narrow radii (Cai et al., 2015; Suri and Udry, 2024; Duflo et al., 2026; Chaurey et al., 2025), and the broader literature treats proximity and face-to-face contact as the leading conduits (Marshall, 1920; Moretti, 2021; Baum-Snow et al., 2024; Cai and Szeidl, 2017; Atkin et al., 2024). That focus may well have been a good approximation: when exchanging information requires co-presence, the local network is the diffusion network. However, there is good reason to think diffusion is not only local in today's economy. Mobile phones hold ties together across the country, and mobile money moves value along them. Consistent with this, I find that the forces spreading the technology do not require closeness: learning runs through social ties at any distance, whereas the force that is local is demand, which needs a shared customer, and customers are shared within a city. The distinction matters most for non-tradable services, a growing share of employment in developing economies (Rodrik, 2016), whose customers are local.

Second, the paper makes progress on a long-standing identification problem in the literature on network effects and the adoption of digital payments (Björkegren, 2018; Fafchamps et al., 2023; Alvarez et al., 2025; Crouzet et al., 2023; Higgins, 2024; Crouzet et al., 2024): separating social learning from strategic complementarities. Both predict higher adoption when connected firms adopt. Because payments are two-sided and a firm's incentive to adopt rises with its customers' digital use, the two mechanisms are easily conflated when the data capture only proximity or aggregate exposure. Recent work identifies complementarities from policy-induced geographic variation (Crouzet et al., 2023; Higgins, 2024) and structural models (Alvarez et al., 2025), while noting that learning is hard to quantify in the same setting. Observing the two conduits separately shows

that the forces differ less in existence than in shape. Demand is bounded by the customer network itself: demand effects mostly operate between firms that share customers, and the cross-industry customer effect is thus limited. Social learning, however, may be weaker per tie, but spread across each firm's thousands of community ties. The network literature documents the same pattern for weak ties: each transmits little, but in aggregate they carry most novel information (Granovetter, 1973; Aral and Van Alstyne, 2011; Bakshy et al., 2012). It is the primary channel that crosses industries and cities in this setting. Network-based targeting has so far been evaluated by seeding on centrality (Banerjee et al., 2013) or on threshold models of reinforcement (Beaman et al., 2021), whereas here the transmission is carried by exposure through well-connected treated drivers. This breadth also helps explain why the reach of learning has been hard to quantify: a design confined to one sector or one neighborhood retains only the small fraction of ties that stay inside its boundary.

Third, the paper expands the measurement frontier for firm networks in developing economies, where most firms are informal, owner-operated, and absent from administrative records. In these settings, the owner's personal and business networks are deeply intertwined (Bernhardt et al., 2019). Diffusion studies have often been bounded by what a survey can enumerate, typically one industry in a handful of villages or neighborhoods (Cai et al., 2015; Chuang and Schechter, 2015; Suri and Udry, 2024; Duflo et al., 2026; Chaurey et al., 2025), or by administrative records on formal firms in richer countries (Bloom et al., 2013). I leverage hashed contact-book fingerprints, to my knowledge new to economics, which map the social ties of roughly 7 million people (70–80% of Senegal's adult population) without directly revealing phone numbers. I also use the mobile money ledger to measure the customer links that cash typically leaves untraced. Relationships that earlier work inferred from physical proximity or reconstructed from surveys thus become observable at national scale. The networks I measure show the limits of the distance proxy: ties can be sparse and stretch far beyond the neighborhood, and 66% of taxi owners' social ties lie more than 30 km from their home. The ties that carry the diffusion are community ties, rarely appearing in survey rosters. Such data now exist across lower-income countries, so the same measurement is feasible in many other settings. The questions asked of formal production networks (how firms find partners, how they upgrade, how shocks and technologies propagate) can then be further addressed in informal economies. (Bernard and Moxnes, 2018; Breza et al., 2019; Verhoogen, 2023).

The rest of the paper proceeds as follows. Section 1 describes the setting, the experimental design, and the network data. Section 2 documents empirical patterns in firms' networks. Section 3 presents the empirical framework. Section 4 lays out a simple model

of learning and demand. Section 5 estimates social learning across industries and cities. Section 6 tests the within-industry and customer-demand predictions. Section 7 concludes.

# 1 Setting, Experiment, and Data

## 1.1 The Technology and Market Environment

**The Digital Payment Technology.** 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 business directly from an existing mobile money account by scanning a QR code displayed by the business, at a 1% fee for drivers (see Figure 1).[1] In the taxi sector, the technology changes payment in two economically meaningful ways. First, business payments are effectively irreversible, unlike ordinary person-to-person (p2p) transfers that can often be reversed through a simple call to the provider. Second, the QR interface makes payment substantially more convenient than manually typing a driver's phone number into the mobile money application. These features matter in a cash-intensive sector such as taxis, where transactions are frequent, small, and operationally costly (see Houeix (2026)).

The technology is especially well suited to studying diffusion. Here, customers use existing mobile money accounts that were already widespread in Senegal. The adoption decision therefore lies with the taxi 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 1). This makes the setting particularly useful for studying whether the same technology diffuses beyond the sector in which access is initially randomized.

The market environment also helps isolate the mechanisms behind adoption beyond monetary cost. The latter is low in this setting as the main adoption cost is time: the business must locate a registration agent, provide the required paperwork, and install the business application for free. In practice, the relevant requirements are possession of an

[1] 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.

official ID, a personal mobile money account, and a smartphone.[2] These features matter because they make the technology similar to many other digital business tools in low- and middle-income settings, where adoption is constrained by information, trust, and coordination, and less so 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 payment 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, over 230,000 businesses in Senegal had adopted the digital payment technology (see Figure 2).[3] In particular, by early 2024, data show that approximately 17,000 taxis and 220,000 other businesses had adopted the technology, suggesting that by the end of the experiment in October 2022, a substantial share of the potential non-taxi business population had already adopted (taxi adoption was still growing rapidly: ∼7,000 of an eventual ∼17,000 by early 2024). The technology adoption was not only located in Dakar, the capital city, where the experiment took place, but also across many cities in Senegal (see Figure A7). This rapid expansion represents an opportunity to study whether a randomized intervention in one urban service sector can generate diffusion of the same technology across industries and across cities.

**The Taxi Sector.** 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)).[4] 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 individually considered as a business entity.

Drivers interact with many customers each day, regularly encounter other drivers at

[2]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).

[3]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.

[4]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.

garages and waiting points, and frequently exchange information; about 14% of drivers operate from a garage. 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. A typical taxi driver reports interacting and regularly spending time with about 15 other taxi drivers. These features make both channels studied in the paper plausible ex ante: social learning through social networks, and customer-demand complementarities when customers begin to expect digital payment.

Overall, the taxi industry is well suited to isolate different mechanisms. For a firm rooted in one location, the two networks largely coincide: its neighbors are at once its customers and its acquaintances, so exposure measured on either network proxies for the other. Taxi drivers move through the city, which pulls the networks apart—social ties possibly stretch across the country and customers span Dakar—so each mechanism can be measured on a network where the other is muted.

**Technological Benefits.** The economic impact of the technology on businesses is the focus of a separate paper (Houeix, 2026): from the point of view of this paper, the key takeaway is that this technology is beneficial for business adopters. 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 was introduced. While 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. See Houeix (2026) for more details. This paper abstracts from the owner–driver relationship: the observability sub-treatments are pooled, and diffusion is studied at the level of the business. Many individuals in the experiment are owners driving their own cab, and most adopters outside the experiment are owner-operated firms with no employees, so the internal monitoring margin plays limited role in the spillovers studied here.

## 1.2 Experimental Design

The experiment randomized early access to the technology among taxi businesses in Dakar.

**Listing Survey.** Between March and May 2022, drivers and owners were listed at garages, car washes, meeting points, and taxi stands across Dakar (Figure A3). Non-driving owners were contacted through their drivers, as owner contact details were required for adoption.

**Baseline Survey.** Starting from the listing, I contacted 2,072 taxi owners and asked each to update their current roster of drivers. Baseline data collection was conducted in twelve rolling batches across three Dakar sites, from March to June 2022. In total, 3,011 individuals (owners + drivers) were invited to the baseline survey, with owners and drivers interviewed separately in 45-minute questionnaires. The baseline dataset eventually includes 2,196 unique taxi stakeholders: 1,821 are drivers (including owner-drivers who operate their own taxi) and the remainder are non-driving owners who employ at least one driver. Overall attrition from invitation to participation was 26%, with no significant differences across treatment arms (treatment assignments were disclosed only after baseline surveys to prevent differential attrition). Roughly half of attrition reflects baseline ineligibility (e.g., the driver lacked an Android phone, a license, or an ID), rather than refusal or no-shows.

The unit of randomization is the taxi business. The sample comprises three types of stakeholders: of the 2,196 surveyed, 1,821 are drivers—owner-drivers who operate their own vehicle or employed drivers working for an owner—and the remaining 375 are non-driving owners who employ at least one driver (608 owner–driver pairs are linked at baseline). All types are included as a single unit for the following analysis—the taxi business—regardless of whether the respondent is the owner or the driver. Approximately 58 percent were assigned to treatment and the remainder to a waitlist control group.[5] Randomization was computerized and stratified on baseline characteristics collected during listing and surveying. The realized experimental sample is balanced across treatment arms. Treated businesses were provided with the QR payment card, the business application, and training on how to use the product. For the sake of this study, control businesses were prevented from receiving the technology during the experimental period and were offered access only after the experiment ended. Importantly, the experimental protocol did not encourage treated businesses to market the technology to

[5]Sample accounting: 2,196 stakeholder records randomized at the business level → 2,150 matched to a mobile money account (1,256 treated, 894 control; 46 could not be matched), spanning 1,454 taxi businesses with treatment constant within business → 1,912 with phone fingerprint data → 1,907 in the phone network after minimum-contact filtering. Throughout, "experimental drivers" refers to these matched records—including non-driving owners' accounts—whose phones define the experimental side of the network. All reduced-form analyses use the 2,150 matched records for exposure computation; the phone network analyses condition on the 1,907 appearing in the fingerprint data.

others or to share information about it. The control businesses were still made aware of the technology and could possibly share information about it. Any adoption among non-experimental firms therefore reflects transmission from the experimental sample rather than direct contact with the study.

**Timeline.** The key dates are: (i) first treated onboarding: March 28, 2022 (start of outcome window); (ii) last treated onboarding: June 2022 (final baseline wave); (iii) adoption-window cutoff: October 5, 2022 (when control-group rollout began and could freely adopt); (iv) $74.1\%$ of control vehicles ($79.4\%$ by November 2023). Drivers onboarded in later waves had fewer than six months of exposure, but since adoption is measured as a binary indicator by October 5, the six-month window is an upper bound on treated exposure for all drivers.

The experimental design serves two purposes. First, it identifies the private effects of the technology on adopting taxi businesses, which are studied in a separate paper (Houeix, 2026). Second, the focus of this paper, it creates exogenous variation in the treatment status of taxi firms that are embedded in much larger social, local, and customer 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

### 1.3.1 Social Networks from Hashed Phone Contact Fingerprints

A central object in the paper is the social proximity between non-experimental firms and treated taxi businesses, which I measure using encrypted phone-contact fingerprints collected by the mobile money company at first login, with about 86% of users consenting. Unlike the survey links more common in economics, contact-book fingerprints capture the set of contacts a user chose to save in the phone and the contact overlap, making them well suited to measuring social proximity among informal firms. The operator collected fingerprints for roughly 7 million users: on first login to the application, users are asked whether they consent to share their contact list (consent rate about 86 percent), and for those who consent the operator stores an encrypted fingerprint that it uses operationally to facilitate transfers to saved contacts. This is a routine request among consumer apps. 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. Related objects have been studied in computer science to recover set-overlap measurement and private contact securely, but, to my knowledge, not for

economic measurement (De Cristofaro et al., 2013; Swamidass et al., 2015; Demmler et al., 2018).

**Constructing the Network.** Each phone number saved in a user's contact list is hashed into one of $2^{16} = 65{,}536$ buckets using a 16-bit system FNV-1a hash (Noll et al., 2026). Let $F_i$ denote the set of hashed contacts saved by user $i$, and $h_i$ the hash of user $i$'s own phone number; these fingerprints preserve set operations while preventing recovery of raw phone numbers. The fingerprints yield two complementary objects for any pair of users $i$ and $j$: directed *save* indicators, $\mathbf{1}\{h_j \in F_i\}$ ($i$ saved $j$'s number) and $\mathbf{1}\{h_i \in F_j\}$, whose conjunction defines a *mutual contact*; and the contact-book *overlap* $|F_i \cap F_j|$, which is defined whether or not either party saved the other. The analysis network is built from the overlap measure—two people who share a meaningful fraction of their saved contacts are socially proximate even if neither stored the other's number—with the *save* indicators retained as tie attributes and analyzed separately.[6]

**Measuring Tie Strength.** Tie intensity is measured by the *bias-corrected overlap share* of the two contact books:

$$\text{OverlapShare}_{ij} = \frac{\max(I_{ij} - \mu_{ij},\, 0)}{\min(|F_i|, |F_j|)}, \qquad \mu_{ij} = \frac{|F_i|\,|F_j|}{65{,}536},$$

where $I_{ij}$ is the observed intersection of the two hashed fingerprints and $\mu_{ij}$ is the intersection expected from independent hash collisions, so the numerator retains only overlap in excess of chance. The correction matters because hashing maps every saved number into one of 65,536 buckets: two users not belonging to the same community network will still share some buckets by accident. For two typical contact books of about 600 numbers each, $\mu_{ij} = 600 \times 600/65{,}536 \approx 5.5$, so roughly five to six of their buckets coincide by chance alone—about 0.9 percent of either book.[7] An observed intersection of that size carries no information about a relationship, and the corrected share nets it out. Normalizing by the smaller contact list ensures that users with larger contact books are

[6]The unit of observation in the phone network is a mobile money account, identified by phone number. Using the fingerprint data, 85% of phone numbers can be linked to unique individuals; about 16.5% of individuals can be linked to multiple mobile money accounts. Aggregating to the individual level using the operator's legal-entity identifiers changes mutual-contact edge counts by less than 0.5% and strong-tie counts by 1%, because few individuals appear in the network through more than one account. The phone-level analysis is therefore essentially equivalent to a person-level analysis.

[7]With $B = 65{,}536$ buckets, if the two fingerprints were independent uniform subsets of their observed sizes, each bucket occupied by $i$ would be occupied by $j$ with probability $|F_j|/B$, so $\mathbb{E}[I_{ij}] = \sum_{b \in F_i} \Pr(b \in F_j) = |F_i||F_j|/B = \mu_{ij}$. The corrected share thus subtracts the intersection that pure chance would generate between two unrelated users.

not mechanically penalized.

The correction follows the chemoinformatics literature, which subtracts the overlap expected under random hashing when comparing molecular fingerprints folded into a small bit space (Swamidass and Baldi, 2007). In economics, this is best interpreted as a measure of *social proximity*: how embedded are the two users in overlapping social circles? The interpretation is also consistent with the network-science literature, where neighborhood-overlap measures are informative about latent social proximity and link formation (Adamic and Adar, 2003; Liben-Nowell and Kleinberg, 2007).

**Tie Strength and the Community Network.** I define a *community tie* as a pair with corrected overlap of at least 2 percent—two people who demonstrably share a meaningful fraction of their social environment, whether or not either saved the other's number. The 2 percent floor sits well above the collision-noise scale[8] and is validated directly against survey ground truth (clients and serendipitous contacts average only 1.9 and 3.2 percent corrected overlap and clear it just 22.8 to 27.7 percent of the time, see Appendix Table B1 and below).

The community graph is deliberately broader than a contact book: it measures *shared social environment*—after hub trimming[9], the median driver is connected to 7,462 individuals (mean 12,721). *Strong ties* are defined as community ties with corrected overlap of at least 5 percent: a median of 556 per driver (mean 1,521). The main specifications are reported at the 2, 5, and 10 percent thresholds.[10] Reciprocity is measured separately: a *mutual contact* is a pair in which each party saved the other's hashed phone number. Mutual ties are a smaller subset of the community network with higher baseline adoption, and the analysis later shows they carry within-industry adoption across cities, but do not amplify diffusion across industry (Appendix Table B5). Non-mutual contact pairs above the floor are reported as close relationships 37 percent of the time versus 18 percent below it (see further discussion below and details in Appendix Table B2).

[8] After the $\mu$ correction, the overlap of two unrelated users is centered on zero, with standard deviation $\sqrt{\mu_{ij}}/\min(|F_i|, |F_j|) \approx 1/\sqrt{65{,}536} \approx 0.4$ percent for comparable contact books; the 2 percent floor is roughly five standard deviations above this null, so almost no purely accidental overlap clears it.

[9] The community graph includes ties with widely-held numbers (agents, popular businesses, commercial phone numbers, etc.) that likely reflect no personal relationship. I remove these by dropping the 2% most-connected contacts and, in a second step, every contact linked to a top-2-percent-connectivity driver, which leaves each driver's idiosyncratic ties (median 172,700 $\rightarrow$ 7,462). The point estimates are essentially unchanged by this procedure; trimming primarily sharpens inference by removing contacts whose treatment exposure is near-collinear across the sample.

[10] Results are insensitive to the exact strength cutoff: exclusive strength bands from 2 to 17+ percent show approximately uniform per-band transmission for cross-industry adoption (Appendix Figure A8).

**Validating Overlap as Social Proximity.** Three pieces of survey evidence support the interpretation of overlap as social closeness. In the endline survey, 1,693 of the 2,196 experimental drivers completed a network module, each reporting on up to 16 randomly selected mobile money contacts—relationship type (wife, family, friend, taxi colleague, client, and so on) and interaction frequency—as described in Section 1.4. Of the 23,520 survey pairs, 16,470 had fingerprint data on both sides (Appendix C).

*First, corrected overlap rises monotonically with reported closeness.* Spouses share 20.8% of contacts on average (86.5% of pairs above the 2% community floor; 76.8% above the 5% strong-tie threshold), taxi colleagues 13.6% (81.4%; 62.8%), family members 11.4% (80.8%; 64.1%), and friends 6–9%. At the bottom of the distribution, serendipitous contacts—numbers the respondent could not place when prompted, typically incidental or one-time transaction partners—average 3.2% with only 27.7% above the community floor, and taxi clients 1.9% (22.8%): the 2% floor aims to separate durable relationships from incidental contact. The same gradient appears by interaction frequency: contacts seen weekly or more average 10.6% corrected overlap, versus 5.5% for contacts seen less than monthly, a $1.9\times$ difference (Table B1).

*Second, reciprocity marks relationship depth.* Mutual contacts are reported as close 61.1% of the time. In contrast, one-sided contacts (36.8%) and community ties where neither party saved the other (37.1%) are indistinguishable—a one-way save adds nothing beyond community membership—while pairs *below* the community floor are "close" (spouse, family, friends) only 18.2% of the time (Table B2). Reciprocal saving is thus a meaningful marker of tie depth—which is why the mutual layer is analyzed separately (Appendix Table B5)—but overlap without reciprocity still identifies real relationships, and it is these community ties that carry the diffusion effects in Table 1.

*Third, overlap carries social-proximity content at any distance.* Regressing survey-reported closeness on corrected overlap, geographic distance, reciprocal-save (mutual) status, and their interactions (Table B3) shows that overlap predicts closeness at every distance—and is, if anything, *more* informative for distant pairs, with the overlap slope roughly five times steeper for contacts outside Dakar: nearby, shared environment generates overlap even among casual acquaintances, while at distance only genuine relationships sustain it. Conditional on overlap, reciprocal saving carries substantial *additional* closeness content (a 31-percentage-point higher probability of a close relationship, stable across distance): overlap measures the shared social environment, while reciprocity marks the depth of the relationship within it.

**Network Observed After the Experiment.** The phone-contact fingerprints are observed in February 2024, after the experiment ended, and are interpreted as a post-period proxy for durable *pre-existing* social proximity. Table B4 supports this interpretation in two steps. First, treated and control drivers are balanced on network levels: no count or composition measure differs across arms, either on the raw fingerprint network (Panel A) or on the community-tie network used in the analysis (Panel B; joint $F$-test, $p = 0.69$), and the two marginal raw-network comparisons—control drivers have slightly *more* mutual contacts ($p = 0.056$) and marginally larger contact books—work *against* finding positive spillovers from treatment. Second, a more demanding test asks whether treatment *strengthened* ties, e.g., through treated drivers saving more technology-related contacts (Panel C): treated drivers' tie-strength distributions do not shift rightward—pooled overlap is, if anything, weakly *lower* for treated, on both the raw edge-level and the corrected community measures—the composition of strong ties ($\geq 5\%$), the sample of interest below, is indistinguishable across arms on every dimension tested (all $p > 0.19$), and contact-type splits show the same (taxi contacts, the most relevant for the mechanism, $p = 0.14$). The February-2024 network can therefore be treated as pre-determined with respect to treatment; any residual bias from differential appearance is bounded below by the Lee bounds, see Table B8 and further discussion below.

**Test for Differential Appearance Across Treatment.** Of the 2,150 experimental drivers with a matched mobile money account, 1,912 (88.9%) have at least one phone number with fingerprint data, and 1,907 appear in the phone network; the 11.1% gap reflects non-collection, as the operator's fingerprint extraction did not reach these users by the February 2024 snapshot. Treated drivers are somewhat more likely to appear (89.9% vs. 87.6%, $p = 0.093$), plausibly because treatment raised app usage, which triggers collection—but the difference is small and, conditional on appearing, the economically relevant variables (overlap composition, share business, share taxi) show no meaningful imbalance (Table B4). To formally bound any resulting bias, I also compute tightened Lee bounds (Lee, 2009) within contact-book-size quartiles (see Table B8 and results below).

### 1.3.2 Customer Network from Transaction Data

The transaction data are primarily used to construct the customer-link objects relevant for the customer-demand channel. The *customer network* is built from business-payment data between November 2022 and November 2023, when both treated and control taxis gain access to digital payments and appear as payees. Two firms are linked if they share regular customers—people who make repeated payments to both firms (at least two transac-

tions with each party), including at least one experimental driver. This network gives the distance profile of customer links (Section 2) and identifies the riders who paid a treated driver in the payer decomposition of Section 6.2.

In addition, the transaction data provide a second social 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 network is less exhaustive than the phone network. It nevertheless contains richer information on financial interaction. In the 12 months before the experiment, the median experimental driver sent or received 119 transfers totaling CFA 1,492,655 (USD 2,488)—predominantly as a sender, consistent with fares not yet being payable digitally—with a median transfer size of CFA 12,397 (USD 21).

### 1.3.3 Geographic Network from 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 particularly 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.[11] These measures are essential for distinguishing remote diffusion through social ties from diffusion through physical co-presence. The same data locate firms for the distance profiles in Section 2 and the density split in Section 5.4.

## 1.4 Survey and Technology Adoption Data

**Survey Data.** To interpret the mechanisms behind diffusion, the paper combines the phone network with survey evidence. The survey data serve two purposes. First, they

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

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

**Technology Adoption Data.** Adoption of the business payment technology is observed directly in the provider's administrative records. For each adopting business, the data record the exact registration date, business 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 businesses in Dakar, and non-taxi businesses 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 2). The records also give usage among adopters, measured as the number of business payments a firm receives in the first six months after adoption, which Section 6.2 uses as an outcome.

# 2 Empirical Patterns in Firms' Networks

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

**Stylized Fact 1 - Firms' social ties decay with distance but are not local.** On average, the taxi businesses in the experiment have approximately 590 phone contacts (unique entries in the encrypted contact-book fingerprint). The community-tie graph is far broader than the contact book by construction—it links people who share social environment, not only saved numbers. After hub trimming, the median driver has 7,462 community ties

(mean 12,721), of which 556 are strong ($\geq 5\%$); the reciprocal-save (mutual) layer is far smaller still ($\sim$100 per driver).[12] The data reveal that the taxi businesses' network extends significantly beyond Dakar, where the experiment took place: 66% of their community ties and 66% of their strong ties lie more than 30 km from the driver's home, and most of these are located in urban areas across the country (Figure 4).

As a descriptive exercise, I identify business links within the phone contact network by matching contacts to all business adopters of digital payments as of April 2023. Through mutual phone contacts alone (before hub trimming), the average driver is connected to 59 other business owners—15 taxi businesses and 44 in other industries (5.6% of all mutual contacts). Among strong mutual ties ($\geq 5\%$ corrected overlap), the composition reverses: of 10 business contacts, 8 are taxi businesses—the closest professional relationships concentrate *within* the industry, foreshadowing the within-industry, cross-market role the mutual layer plays below. The community-tie network extends this reach further (Table B4). Together, these layers allow me to track technology diffusion across both industry and space from a single experimental sample, and they reflect how deeply firm networks are interwoven with social networks in this low-income, owner-operated context.

**Stylized Fact 2 - Customer links are local.** Customer links, by contrast, are local. Figure 5 plots the distance profile of the shared-customer network, where each observation is a business–driver customer link plotted at the distance between the driver's home and the linked business, so that a business linked to several drivers appears once per link. The median distance between a driver and a linked business is 7.6 km, 62% of customer links lie within 10 km of the driver's home, and only 10% lie beyond 30 km, against 66% of community ties. Taxi businesses account for 39% of these links and businesses in other sectors for 61%.

The social network and the customer network connect different pairs of people. In the validation survey, taxi clients share on average 1.9% of contacts with the driver—below the community-tie floor—and only 22.8% of client relationships clear it (Appendix Table B1). Matching each surveyed taxi contact to the customer network, taxi *friends* share customers with the respondent less than half as often as non-friend taxi colleagues (9% versus 24% of pairs, with 4.46 versus 23.55 shared customers on average). A driver's friends typically serve other markets. Because social closeness and customer overlap sort

[12] These counts are computed on the analysis sample after two filtering steps: (i) keeping only the most informative phone per individual (best-phone-per-person, Appendix C), and (ii) applying top-2% degree trimming (removing hub nodes). Before these steps, the mean mutual-contact count is $\sim$8,200, of which $\sim$700 clear the strong-tie threshold (Table B4, Panels A and C, which also reports the intermediate pre-trimming values).

into largely different relationships, each layer can be measured separately and can carry a different economic force, which is the premise of the empirical design that follows.

**Stylized Fact 3 - Bigger contact books go with more business.** Merchants with larger contact books do more business. Figure 6 plots the number and the value of business payments a merchant received in the six months before January 2024 against the size of its contact-book fingerprint, for all merchants trading in that window, in twenty equal-count bins after absorbing category, adoption-month, 2 km grid-cell, and onboarding-agent fixed effects: a 10% larger contact book goes with about 2% more payments and about 7% more value received. The relationship is descriptive, since the fingerprint is contemporaneous with the outcome window, and where it predates the business the elasticity falls to $+0.013$ ($t = 0.58$), so the pattern says that active merchants keep larger contact books, without showing that contacts generate business.

**Stylized Fact 4 - 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 B10, 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.[13] More specifically, they reported having discussed or recommended this digital payment technology to 12% of the people they know. This underscores the significant 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.

[13]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. To make the distinction concrete, consider a clothing shop in Dakar whose owner has five experi-

mental taxi drivers in her mutual phone network, three of them randomly assigned early access. Her exposure in the phone layer, $S_j^{(\text{phone})}$, equals $3/5$. Equation (1) asks whether her probability of adopting during the experimental window rises with this treated share, holding fixed her expected exposure $\bar{S}_j^{(\text{phone})}$—which is pinned down by how many experimental taxi contacts she has, and so absorbs her network position. A positive $\beta_{g,n}$ means that exposure to early adopters through her phone contacts raises her own adoption. Because she is in a different industry, she does not share their passengers, so the estimate isolates social learning.

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. The within-industry results are reported in this parameterization because the model is stated in counts.

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

The reduced-form coefficients capture *net* spillover effects through the network used to define exposure. A positive coefficient in the cross-industry phone network can be read as evidence of social learning, because same-industry rivals are absent in that population. A coefficient in the within-taxi phone network, by contrast, aggregates learning and whatever rivals withhold. 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 test the predictions of Section 4.

## 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 populations define the main spillover samples. The first is non-experimental taxi businesses in Dakar, capturing within-industry diffusion. The second is non-taxi businesses in Dakar, capturing cross-industry diffusion within the treatment city. The third is non-taxi businesses 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. 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 *community-tie phone network* includes all pairs with bias-corrected contact-book overlap of at least 2 percent—the calibrated boundary of a person's measured social world (Section 1.3)—and captures social exposure whether or not the parties saved each other's numbers; the *strong-tie* specification restricts to corrected overlap of at least 5 percent. The *mutual-contact layer* (reciprocal saves) is analyzed separately. The *pre-experiment transaction network* captures direct economic interaction and provides a cross-check on the phone-network results. The *customer network* measures shared customers between taxi businesses and identifies the riders who paid a treated driver (Section 6.2).

## 3.3 Inference Under Network Dependence

Overlapping network connections generate dependence across focal units: two businesses may share experimental neighbors, be linked to each other, or belong to the same local market. 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 randomization-inference $p$-values as the primary measure of statistical significance. In the main reduced-form tables, $p$-values are two-sided.

# 4 A Simple Model of Learning and Demand

This section lays out a two-stage model of adoption in which learning is a gate and demand is the payoff, and uses it to organize the evidence in Sections 5 and 6.

**Setup.** Business $j$ operates in location $\ell$, has $k_j$ treated ties, and serves a share $c_j$ of customers willing to pay digitally; in the count specification of Section 3.1, $k_j$ is $T_j^{(n)}$, the number of treated experimental ties in network $n$.

**Information.** Business $j$ is informed about the technology with probability

$$P_I(k_j) = 1 - (1 - \lambda_\ell)\,(1 - q)^{k_j}. \tag{3}$$

The background rate $\lambda_\ell$ captures marketing and agents in location $\ell$, and each treated tie transmits with probability $q$, which depends on the pair and is the same at any distance.

**Adoption.** An informed business adopts if and only if $x_j\, d_j \geq \kappa$. The private value $x_j$ is drawn from $U[0,1]$, demand is $d_j = d_\ell^0 + c_j$, and $\kappa$ is the fixed cost of adoption. The adoption probability is

$$\Pr(A_j) = P_I(k_j) \cdot G(d_j), \qquad G(d) = 1 - \kappa/d. \tag{4}$$

**Usage.** Among adopters, payments per customer are $y_j = d_j\,\big(1{+}\eta{\cdot}\mathbf{1}[\text{learned from a tie}]\big)$. The parameter $\eta$ is know-how: a business that learned the technology from a tie uses it more intensively.

**One more treated tie.** Consider one more treated cross-industry tie, with $c_j$ unchanged. The adoption probability rises by

$$\Delta \Pr(A_j) = \underbrace{(1 - \lambda_\ell)}_{\text{not informed yet}} \quad \underbrace{(1 - q)^{k_j}}_{\text{no other tie got through}} \quad \underbrace{q}_{\text{this one does}} \quad \underbrace{G(d_j)}_{\text{willing}}. \tag{5}$$

Learning is the gate and demand is the payoff, and the two multiply.

**Rivalry.** A treated driver may transmit less to a rival than to a friend: $q_{\text{rival}} \leq q$ and $q_{\text{friend}} = q$.

**Predictions.** The model yields three predictions.

*P1. Learning across industry and space.* If cross-industry ties carry learning and $q$ does not fall with distance, adoption rises with treated strong ties across industries and cities, without anyone traveling, and the effect is largest where $\lambda_\ell$ is low.

*P2. Within-industry transmission.* If rivals hoard, so that $q_{\text{rival}} \leq q$ while $q_{\text{friend}} = q$, the same ties transmit less in the taxi industry, except through friends.

*P3. Demand and know-how.* Demand needs shared customers ($c_j > 0$), whereas know-how needs only a tie, so if ties carried demand their effect would run through riders who paid a treated driver, and if they carry learning, usage rises whether or not customers are shared.

Reach therefore separates the two forces: demand is bounded by shared customers, whereas learning is bounded only by the ties a business has.

# 5 Social Learning Across Industry and Space

This section documents two main 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 through remote interactions along social ties: businesses in other Senegalese cities connected to treated taxis also adopt at higher rates, whether or not they or the drivers physically move between cities. In both populations, customer demand is muted, so the estimates identify social learning; survey evidence on relationship types corroborates this reading. Section 6 turns to the taxi sector, where the same ties meet rivals and shared customers.

## 5.1 Cross-Industry Diffusion in Dakar

Non-taxi businesses 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 share few customers with the experimental taxis (Stylized Fact 2, Section 2), the cross-industry estimate isolates learning.

Table 1 shows that the effect operates throughout the community network, reporting the specification at the 2, 5, and 10 percent corrected-overlap thresholds (Section 1.3); Appendix Figure A6 maps the resulting adoption across Dakar for a selected sample. At the strong-tie threshold the cross-industry coefficient is positive, significant, and economically large: moving the treated share from zero to one raises adoption by $\hat{\beta} \times 1{,}000 = +4.16$ ($p_{\text{RI}} = 0.017$), a $+39\%$ effect relative to predicted adoption at zero exposure, rising to $+5.16$ ($p_{\text{RI}} = 0.015$, $+69\%$) at the 10 percent threshold. Decomposing the effect into exclusive strength bands (Appendix Figure A8), per-band transmission operates throughout the strength distribution—effect sizes of $+15\%$ to $+58\%$ of the zero-exposure baseline from the 2–5% band through 17%+—consistent with transmission throughout the community-tie strength distribution; stronger ties yield more precise estimates, without evidence of larger per-tie effects.[14]

Reciprocity helps distinguish which ties carry each margin (Appendix Table B5). Cross-industry diffusion—local and distant—operates entirely through *non-mutual* community ties: pairs who share a social environment without having saved each other's numbers. Mutual contacts, despite baseline adoption rates several times higher, show no incremental response to treatment; their exposure to the technology is saturated irrespective of the experiment (inframarginality). Two features of the setting explain why the community layer, and not the mutual layer, carries the marginal information. First, control drivers were offered and told about the technology, so the close contacts of a control driver could have nearly the same direct information as those of a treated driver (no direct change in *awareness*). The treatment contrast on mutual ties is potentially small by design. Second, community contacts, who share the driver's environment, could learn instead through observation and second-order links, which may respond to whether the driver actually carries the technology. E.g., treated drivers anecdotally report posting photos of their QR code on social media, a channel that reaches the broad community further than close contacts alone.

**Count Specification.** Table 2 replaces the treated *share* with the *number* of treated ties. Cross-industry learning survives: an additional treated strong tie raises adoption by $+1.67$ per thousand at the $\geq 5\%$ threshold ($p_{\text{RI}} = 0.015$; $+2.01$, $p_{\text{RI}} = 0.008$ at $\geq 10\%$), in Dakar and beyond ($+0.30$, $p_{\text{RI}} = 0.007$).

[14] 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 businesses, 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 B4.

**Diffusion by Business Category.** Cross-industry diffusion reaches many kinds of firms. Figure 7 estimates the count specification separately for the ten largest non-taxi business categories, with an indicator for registering a business in that category as the outcome, on strong ties and pooling Dakar and outside Dakar. Seven of the ten categories respond to one additional treated strong tie, with effects of $+7\%$ to $+19\%$ of the category's baseline adoption rate, and they span retail, food, services, and durables: fabric and clothes, beauty and personal care, restaurants, boutiques, open-air food produce, electronics, and hardware. The three nulls are mechanics, tailors, and grocers, and Appendix Figure A10 extends the exercise to all 29 categories.

**Lee Bounds for Differential Appearance.** Because drivers enter the network sample only if the February-2024 snapshot contains their fingerprint, and treated drivers are slightly more likely to appear (89.9% vs. 87.6%, $p = 0.093$, see Section 1.3), the estimation sample is potentially selected by treatment. Following Lee (2009), I bound the resulting bias by trimming the outcome distribution of the over-represented arm by the appearance-rate differential—comparing equally selected samples and bracketing the true effect under worst-case selection, under the assumption that treatment only (weakly) raises appearance—with trimming within contact-book-size quartiles to tighten the bounds. The bounds are tight around every main estimate (Table B8): $[+4.079^{**}, +4.198^{**}]$ for cross-industry Dakar and $[+0.680^{***}, +0.685^{***}]$ outside Dakar; on the community network, the distant cross-industry effect is bounded at $[+0.554^{**}, +0.686^{***}]$. Even worst-case trimming preserves sign, magnitude, and significance throughout, so differential appearance does not drive the estimated spillovers.

**Removing Shared-Customer Pairs and Garage-Based Drivers.** As an additional check, Appendix Table B7 re-estimates the cross-industry specification of Table 1 after two exclusions: every driver–contact pair sharing a regular customer, and every pair of a driver who operates from a garage, where drivers share a stream of passengers. The estimates are essentially unchanged—$+3.63^{**}$ [0.037] in Dakar and $+0.61^{**}$ [0.027] outside Dakar at the strong-tie threshold, against $+4.162 * *$ and $+0.681 * * *$ in the full sample—so the cross-industry diffusion documented here is carried by social learning alone, with shared customers contributing nothing to the estimates.

## 5.2 Widespread Diffusion via Remote Interactions

Technology adoption also spreads beyond the capital. Table 1 shows that businesses in other cities connected to treated Dakar taxis adopt at significantly higher rates at every threshold, with proportional effects of $+81\%$ to $+138\%$ ($p_{\text{RI}} = 0.011$, $0.009$, and $0.030$ at the 2, 5, and 10 percent cutoffs), roughly twice the within-Dakar cross-industry effect sizes, consistent with a lower baseline outside the capital amplifying the same diffusion force. These results show that knowledge spillovers need not be local in the classic Marshallian sense (Marshall, 1920).

One possible mechanism is physical movement: treated drivers or their contacts temporarily 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 me to test this directly.

**Seasonal Migration is Widespread.** Among the individuals in the community-tie sample (Table 1 Panel A, $\geq 2\%$ corrected overlap), approximately half can be matched to geolocation records. Of these, roughly two-thirds appeared in Dakar at least once during the experiment, while one-third were never observed in the capital despite maintaining strong phone ties with Dakar-based drivers. Seasonal outmigration from Dakar is common: roughly a third of Dakar-based businesses (33.7%) and drivers (34.1%) cashed in/out outside the capital in at least one month during the experiment. The rate is similar across treatment arms (32.6% treated vs. 36.1% control, $p = 0.103$ from a two-proportion $z$-test), so migration does not seem driven by the technology. 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.

**Business Owner's Physical Travel is not Required for Diffusion.** Table 4 shows the overall effect on the community-tie sample ($+0.68$ per 1,000, $p_{\text{RI}} = 0.009$, $+138\%$ of the zero-exposure baseline) and decomposes it by physical mobility, measured from mobile money geolocation. Panel A splits receiving businesses by their *own* movement: the effect is statistically indistinguishable from zero among those observed in Dakar during the experiment ($p_{\text{RI}} = 0.118$) and is concentrated among businesses *never* observed in the capital, whose coefficient is five times larger ($+1.82$ per 1,000, $p_{\text{RI}} = 0.027$), although a similar effect size. Physical presence in Dakar therefore does not appear required for adoption to respond to treated Dakar contacts.

Panel B of Table 4 tests the complementary channel: whether diffusion to a secondary city runs through treated *drivers* who physically travel there. For each of the eight largest

secondary cities $c$, I code each experimental driver as having *visited* $c$ if any of their mobile-money wallets transacted at an agent located in $c$ at least once during the experiment, and I split individual $j$'s treated share into two city-specific exposures: the share of $j$'s driver contacts that are treated *and* visited $c$, and the share treated but never observed there. The outcome is adoption *in city* $c$, and the two exposures enter a single regression stacked across the eight cities with city fixed effects, an expected-exposure control for each share (its mean across the 1,000 placebo assignments, which absorbs the fact that individuals connected to $c$-visiting drivers disproportionately live near $c$), standard errors clustered by individual, and randomization inference over the driver assignment. I find that the visiting-driver coefficient is three times larger ($+0.021$ vs. $+0.007$ per 100) but imprecise ($p_{\text{RI}} = 0.133$), because few treated drivers visit any given city—so I cannot rule out a substantial role for drivers physically carrying the technology, and seasonal migration may well be part of the story. What the design does establish is that physical travel is not *necessary*: exposure to treated drivers who never visited the destination raises adoption there by $+125\%$ of the control mean ($p_{\text{RI}} = 0.018$). It indicates that remote phone interactions move the technology across cities.

**Tie Strength and Cross-Industry Diffusion Across Space.** Across the exclusive strength bands (Appendix Figure A8), the distant cross-industry effect is large wherever the data can measure it: effect sizes of +83% to +144% in the 2–5, 5–8, and 8–11% bands and +126% in the open tail (all $p_{\text{RI}} \leq 0.064$), with two statistically null cells in the thin 11–17% range (252 and 153 adopters). As within Dakar, transmission does not rise systematically with tie strength—but at every band the distant effect size exceeds the local one: any community tie carries the technology across cities, to people the driver knows who live and work elsewhere.

Because these are cross-industry businesses (the outcome is non-taxi adoption), customer demand is a relatively muted channel; the effect is attributable to social learning transmitted across space through the phone network.

## 5.3 Who Transmits: Sender Connectedness

Transmission comes from well-connected treated drivers. To show this, I split each firm's exposure by the connectedness of the treated driver on the other side of the tie, dividing drivers at the median of their strong-tie degree and entering the two exposures jointly (Figure 8). Exposure to high-degree treated drivers raises cross-industry adoption in Dakar and outside Dakar, whereas exposure to low-degree treated drivers is a null in

both populations. Since degree is balanced across treatment arms, the split does not reintroduce selection.

Drivers with countrywide reach transmit more, in Dakar too. Repeating the split with the driver's number of strong ties located outside Dakar, on the firms with ties to both groups (Figure 9), exposure to drivers with many outside-Dakar strong ties raises adoption in Dakar as well as outside it, whereas exposure to drivers with few outside ties does not. Because every driver's network has a similar Dakar share, outside reach measures total reach, and the split is again balanced across treatment arms.

Direct reach is the operative statistic, and network position adds little: splitting drivers by eigenvector centrality on the driver–driver shared-contact network does not sharpen the contrast, since the point estimates are larger for high-centrality drivers but imprecise, and transmission is significant through the low-centrality group (Appendix Figure A9).

## 5.4 Where Learning Matters Most

The effect is strongest where knowledge is scarce. Outside Dakar, the proportional effects are roughly twice the within-Dakar cross-industry effect sizes, consistent with a lower baseline outside the capital amplifying the same diffusion force (Section 5.2), and the same pattern appears within each geography. Splitting zones at the median of pre-experiment merchant density, that is, prior adoption in the zone, with 2 km neighbourhoods in Dakar and 10 km towns elsewhere, the effect of a fully treated tie portfolio on cross-industry adoption, as a share of predicted adoption at zero exposure, is $+28\%$ in Dakar zones with low prior adoption against $+14\%$ in zones with high prior adoption, and $+86\%$ against $+65\%$ outside Dakar (Figure 10). In the model, this is the prediction that learning matters most where $\lambda_\ell$ is low.

## 5.5 Survey Evidence on the Learning Channel

Survey-elicited relationship types corroborate the learning interpretation. Table 3, Panel C reports that cross-industry spillovers are significant through serendipitous encounters ($+47.5\%$, $p_{\text{RI}} = 0.004$) and positive but imprecise through former or regular taxi clients ($+94.4\%$, $p_{\text{RI}} = 0.134$). The strongest ties (Panel B: friends, frequent contacts) show positive but highly imprecise cross-industry effects. This is consistent with the idea that novel information often arrives through ties bridging separate circles (Atkin et al., 2024): across industries, weaker ties are valuable precisely because they connect firms to different market environments (note that many of these nominally weaker survey ties still exceed the 5 percent corrected phone-overlap threshold but need not be mutual contacts,

see Table B1). The marginal ties are community ties, weak ones included, rather than mutual contacts: the survey pattern matches the administrative evidence that cross-industry diffusion operates entirely through non-mutual community ties, with approximately uniform per-band transmission (Section 5.1; Appendix Table B5 and Figure A8).

# 6 Within-Industry Transmission and Customer Demand

Within the taxi sector, the same ties meet rivals and shared customers, which is where this section tests predictions P2 and P3 of Section 4.

## 6.1 Within-Industry Transmission

Within the taxi industry, an additional treated tie does not raise adoption. Table 2 reports the count specification of Section 3.1 for within-taxi adoption in Dakar and in Thiès, the only market outside Dakar where the taxi product was rolled out, at the 2, 5, and 10 percent thresholds, and every within-taxi coefficient is a statistical null, whereas the cross-industry rows of the same table show that learning survives in counts, with an additional treated strong tie raising non-taxi adoption in Dakar and outside Dakar. Prediction P2 therefore holds on the average tie: the ties that carry learning across industries transmit little within the taxi industry.

The share specification points the same way with a different sign: in Appendix Table B6, the within-taxi coefficient on the treated share is negative and significant at the stronger-tie threshold ($\hat{\beta} \times 1{,}000 = -0.98$, $p_{\text{RI}} = 0.007$, a $-30\%$ effect), and null at the community threshold ($+0.08$, $p_{\text{RI}} = 0.674$). One possible explanation is that rivals' withholding responds further to the fraction of one's circle that is treated than the count.

Disaggregating by tie type reveals that learning remains important within industry on friendship and regular relationships. Table 3 shows that within the taxi sector, spillovers concentrate among friends ($+36.2\%$, $p_{\text{RI}} = 0.007$) and frequent interlocutors ($+24.0\%$, $p_{\text{RI}} = 0.023$), while exposure through non-friends produces a null effect. Close contacts are the people from whom another taxi businesses can share information on how the technology works, whether customers use it, and whether payments are reliable. In the model, friends transmit at the full rate $q$ and rivals at the lower rate $q_{\text{rival}}$.

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 within industry. Weaker same-industry contacts, by contrast, transmit little. The cross-market evidence refines the mechanism: the pattern is closer to a strength of *long* ties than of weak ones (Granovetter, 1973; Aral, 2016) — the Thiès channel runs through dyadically strong ties that span distant local contexts, with transmission *rising* in contact-book overlap (Appendix Table B5), while in the saturated Dakar market even strong ties carry little marginal information to transmit.

**Lee Bounds.** The tightened Lee bounds of Section 5.1 bracket the within-taxi estimate at $[-1.103^{***}, -0.976^{***}]$ (Table B8): even worst-case trimming for differential appearance preserves sign, magnitude, and significance.

**Adoption Within-Industry in Thiès.** Within-industry adoption in the Thiès market—the only market outside Dakar where the taxi product was rolled out—responds positively through mutual ties at every threshold ($+1.10$ to $+5.16$, $p_{\text{RI}} \leq 0.003$)—and on the mutual layer with no overlap floor at all: exposure through a treated mutual tie raises adoption by 18% of the base rate ($p_{\text{RI}} = 0.033$), see Table B5. This channel is also confirmed in the pre-treatment transaction network (Appendix Table B11). Thiès drivers do not compete with Dakar drivers for passengers, so the withholding in P2 has no reason to operate on these ties.

## 6.2 Customer Demand

When customers become accustomed to paying digitally with treated taxis, other businesses serving the same customers may find adoption more attractive because their customers now expect or prefer digital payments. This channel is conceptually distinct from social learning (it operates even when two businesses do not know each other). The channel operates through shared customers alone.

**Riders Who Paid a Treated Driver.** The first test asks whether the tie effect runs through the drivers' own customers. Taxi merchants share a customer base with the experimental drivers: among within-taxi adopters, a quarter of payments already come from riders who paid a treated driver during the experiment. If exposure redirected demand, that share would rise with the number of treated ties. Table 5 regresses the share of an adopter's payments in its first six months that come from treated-induced customers on the adopter's number of treated ties, under four definitions of a treated-induced payer and for within-taxi and cross-industry adopters in Dakar. A payer counts as treated-

induced if they paid a treated taxi during the experiment, before their first payment to the firm, before the firm adopted, or within 30 days before paying. Within taxis, virtually every firm has some such payers, and an extra treated tie moves the mix by 1 to 2.4 percentage points, with randomization-inference $p$-values from 0.05 to 0.44. Among cross-industry adopters in Dakar, treated-induced payers are a small share of payments, with control means of 6 percent or less and between 40 and 81 percent of firms having none, and no definition shows a shift. Exposure therefore does not shift the customer mix in either group: where the drivers' customers demonstrably are, exposure does not bring more of them.

**Usage Among Adopters.** The second test uses intensity. Among adopters, I regress usage in the first six months after adoption, measured as $\log(1 + \text{payments})$ and as an indicator for at least ten payments, on social exposure through strong ties ($\geq 5\%$ corrected overlap) to treated drivers, separately for within-taxi adopters in Dakar, cross-industry adopters in Dakar, and cross-industry adopters outside Dakar (Table 6). The usage response rises with distance from the shared customer market: it is null among within-taxi adopters, whose market is contested, marginal among cross-industry adopters in Dakar, and largest among cross-industry adopters outside Dakar, where customers cannot be shared with the experimental drivers. Demand cannot produce this pattern, since it needs a shared customer, whereas know-how can: in the model, a business that learned from a tie uses the technology more, with $\eta > 0$, whether or not it shares customers.

# 7 Conclusion

This paper combines a randomized experiment with exhaustive network data—phone contact directories covering 70–80% of Senegal's adult population and the complete mobile money transaction ledger—to study when technology diffuses across firms and when it does not.

Two forces shape diffusion, and reach distinguishes them. Social learning operates throughout the community network—with responses that the data do not distinguish across the strength distribution, and not incrementally through mutual contacts—and is the only channel that reaches across industries and across cities, whereas demand needs a shared customer, and shared customers are local. Within the taxi industry, the same ties transmit little, except among friends, and the tie effect, which does not run through the drivers' own customers, raises usage most where customers cannot be shared.

The paper makes two contributions that go beyond documenting the existence of

spillovers. First, the data and experimental design separate learning from demand by their reach: learning travels through community ties at any distance, and demand stops at shared customers, which are local. I show that since exposure through well-connected drivers carries the transmission, seeding well-connected firms would reach more of the economy.

Second, the paper shows that technology can diffuse over long distances: information travels remotely through phone-contact ties, and physical movement is not required. Adoption in secondary cities does not depend on the adopting business itself visiting Dakar, and exposure to treated drivers who never traveled to the destination raises adoption there on its own (Table 4); the channel through drivers who did travel is larger in magnitude but too imprecise to certify, so migration may add to the remote channel. The result refines the Marshallian view of knowledge spillovers: neither the transmission of information nor its conversion into adoption requires co-presence. Seasonal migration is pervasive in this setting—roughly a third of Dakar-based businesses and drivers move seasonally—which makes the sufficiency of the remote channel all the more notable: social networks carry urban innovations across space beyond direct interactions do.

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.** 2003. "Friends and Neighbors on the Web." *Social Networks*, 25(3): 211–230.

**Alvarez, Fernando, et al.** 2025. "Strategic Complementarities in a Dynamic Model of Technology Adoption: P2P Digital Payments." *Working Paper*.

**Aral, Sinan.** 2016. "The Future of Weak Ties." *American Journal of Sociology*, 121(6): 1931–1939.

**Aral, Sinan, and Marshall Van Alstyne.** 2011. "The Diversity-Bandwidth Trade-off." *American Journal of Sociology*, 117(1): 90–171.

**Arrow, Kenneth J.** 1962. "The economic implications of learning by doing."

**Athey, Susan, et al.** 2018. "Exact p-Values for Network Interference." *Journal of the American Statistical Association*, 113(521): 230–240.

**Atkin, David, et al.** 2024. "The Returns to Face-to-Face Interactions: Knowledge Spillovers in Silicon Valley." *Working Paper*.

**Baird, Sarah, et al.** 2018. "Optimal Design of Experiments in the Presence of Interference." *Review of Economics and Statistics*, 100(5): 844–860.

**Bakshy, Eytan, et al.** 2012. "The Role of Social Networks in Information Diffusion." 519–528.

**Banerjee, Abhijit, et al.** 2013. "The Diffusion of Microfinance." *Science*, 341(6144): 1236498.

**Baum-Snow, Nathaniel, et al.** 2024. "Local Productivity Spillovers." *American Economic Review*, 114(4): 1030–69.

**Beaman, Lori, et al.** 2021. "Can Network Theory-Based Targeting Increase Technology Adoption?" *American Economic Review*, 111(6): 1918–1943.

**Bernard, Andrew B., and Andreas Moxnes.** 2018. "Networks and Trade." *Annual Review of Economics*, 10: 65–85.

**Bernhardt, Arielle, et al.** 2019. "Household Matters: Revisiting the Returns to Capital among Female Microentrepreneurs." *American Economic Review: Insights*, 1(2): 141–60.

**Björkegren, Daniel.** 2018. "The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda." *The Review of Economic Studies*, 86(3): 1033–1060.

**Bloom, Nicholas, et al.** 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica*, 81(4): 1347–1393.

**Borusyak, Kirill, and Peter Hull.** 2023. "Nonrandom Exposure to Exogenous Shocks." *Econometrica*, 91(6): 2155–2185.

**Breza, Emily, et al.** 2019. "Networks in Economic Development." *Oxford Review of Economic Policy*, 35(4): 678–721.

**Brock, William A., and Steven N. Durlauf.** 2001. "Discrete Choice with Social Interactions." *Review of Economic Studies*, 68(2): 235–260.

**Broder, Andrei Z.** 1997. "On the Resemblance and Containment of Documents." 21–29. IEEE.

**Cai, Jing, and Adam Szeidl.** 2017. "Interfirm Relationships and Business Performance*." *The Quarterly Journal of Economics*, 133(3): 1229–1282.

**Cai, Jing, et al.** 2015. "Social Networks and the Decision to Insure." *American Economic Journal: Applied Economics*, 7(2): 81–108.

**Cefalà, Luisa, et al.** 2025. "The Economic Consequences of Knowledge Hoarding." *Working Paper*.

**Chaurey, Ritam, et al.** 2025. "Social Learning among Urban Manufacturing Firms: Energy-Efficient Motors in Bangladesh." National Bureau of Economic Research Working Paper 34296.

**Chuang, Yating, and Laura Schechter.** 2015. "Social Networks in Developing Countries." *Annual Review of Resource Economics*, 7: 451–472.

**Crépon, Bruno, et al.** 2013. "Do Labor Market Policies Have Displacement Effects? Evidence from a Clustered Randomized Experiment." *Quarterly Journal of Economics*, 128(2): 531–580.

**Crouzet, Nicolas, et al.** 2023. "Shocks and Technology Adoption: Evidence from Electronic Payment Systems." *Journal of Political Economy*.

**Crouzet, Nicolas, et al.** 2024. "Demographics and Technology Diffusion: Evidence from Mobile Payments."

**De Cristofaro, Emiliano, et al.** 2013. "Private Discovery of Common Social Contacts." *International Journal of Information Security*, 12(1): 49–65.

**Demmler, Daniel, et al.** 2018. "PIR-PSI: Scaling Private Contact Discovery." *Proceedings on Privacy Enhancing Technologies*, 2018(4): 159–178.

**Duflo, Esther, et al.** 2026. "Chat Over Coffee? Technology Diffusion Through Social and Geographic Networks in Rwanda." 31368.

**Fafchamps, Marcel, et al.** 2023. "Adoption with Social Learning and Network Externalities." *Oxford Bulletin of Economics and Statistics*. Open Access.

**Granovetter, Mark.** 1973. "The Strength of Weak Ties." *American Journal of Sociology*, 78: 1360–1380.

**Grossman, Gene M., and Elhanan Helpman.** 1991. "Trade, knowledge spillovers, and growth." *European Economic Review*, 35(2-3): 517–526.

**Higgins, Sean.** 2024. "Financial Technology Adoption: Network Externalities of Cashless Payments in Mexico." *American Economic Review*, 114(11): 3469–3512.

**Houeix, Deivy.** 2026. "Asymmetric Information and Digital Technology Adoption: Evidence from Senegal." *Working Paper*.

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

**Jahani, Eaman, et al.** 2023. "Long ties, disruptive life events, and economic prosperity." *Proceedings of the National Academy of Sciences*, 120(28): e2211062120.

**Lee, David S.** 2009. "Training, Wages, and Sample Selection: Estimating Sharp Bounds on Treatment Effects." *Review of Economic Studies*, 76(3): 1071–1102.

**Liben-Nowell, David, and Jon Kleinberg.** 2007. "The Link-Prediction Problem for Social Networks." *Journal of the American Society for Information Science and Technology*, 58(7): 1019–1031.

**Luo, Kaicheng, et al.** 2025. "Laboratories of Autocracy: Landscape of Central–Local Dynamics in China's Policy Universe." Working paper, September 2025.

**Manski, Charles F.** 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *The Review of Economic Studies*, 60(3): 531–542.

**Marshall, Alfred.** 1920. "Principles of Economics." *MacMillan*.

**Moretti, Enrico.** 2021. "The Effect of High-Tech Clusters on the Productivity of Top Inventors." *American Economic Review*, 111(10): 3328–75.

**Noll, Landon Curt, et al.** 2026. "The FNV Non-Cryptographic Hash Algorithm." RFC Editor RFC 9923.

**Rodrik, Dani.** 2016. "Premature Deindustrialization." *Journal of Economic Growth*, 21(1): 1–33.

**Romer, Paul M.** 1990. "Endogenous technological change." *Journal of Political Economy*, 98(5, Part 2): S71–S102.

**Suri, Tavneet.** 2017. "Mobile Money." *Annual Review of Economics*, 9: 497–520.

**Suri, Tavneet, and Chris Udry.** 2024. "Agricultural Technology in Africa." *VoxDev Vol 5, Issue 2*.

**Swamidass, S. Joshua, and Pierre Baldi.** 2007. "Mathematical Correction for Fingerprint Similarity Measures to Improve Chemical Retrieval." *Journal of Chemical Information and Modeling*, 47(3): 952–964.

**Swamidass, S. Joshua, et al.** 2015. "Securely Measuring the Overlap between Private Datasets with Cryptosets." *PLoS ONE*, 10(2): e0117898.

**Verhoogen, Eric.** 2023. "Firm-Level Upgrading in Developing Countries." *Journal of Economic Literature*, 61(4): 1410–1464.

# Main Figures and Tables




Figure 1: Digital Payments for Businesses — Senegal

*Notes:* These images illustrate the implementation of the technology across different business settings. The top panels show the deployment within taxi vehicles: the QR code is visible from outside the vehicle through a windshield sticker, allowing potential customers to easily identify it. 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. The bottom panels show the same QR code technology deployed in other sectors, including open-air market vendors (left) and retail stores (right), where businesses display the QR code as a sticker at their point of sale for customers to scan and pay. Businesses in different sectors were provided with the same QR code technology, with displays adjusted to the sector.

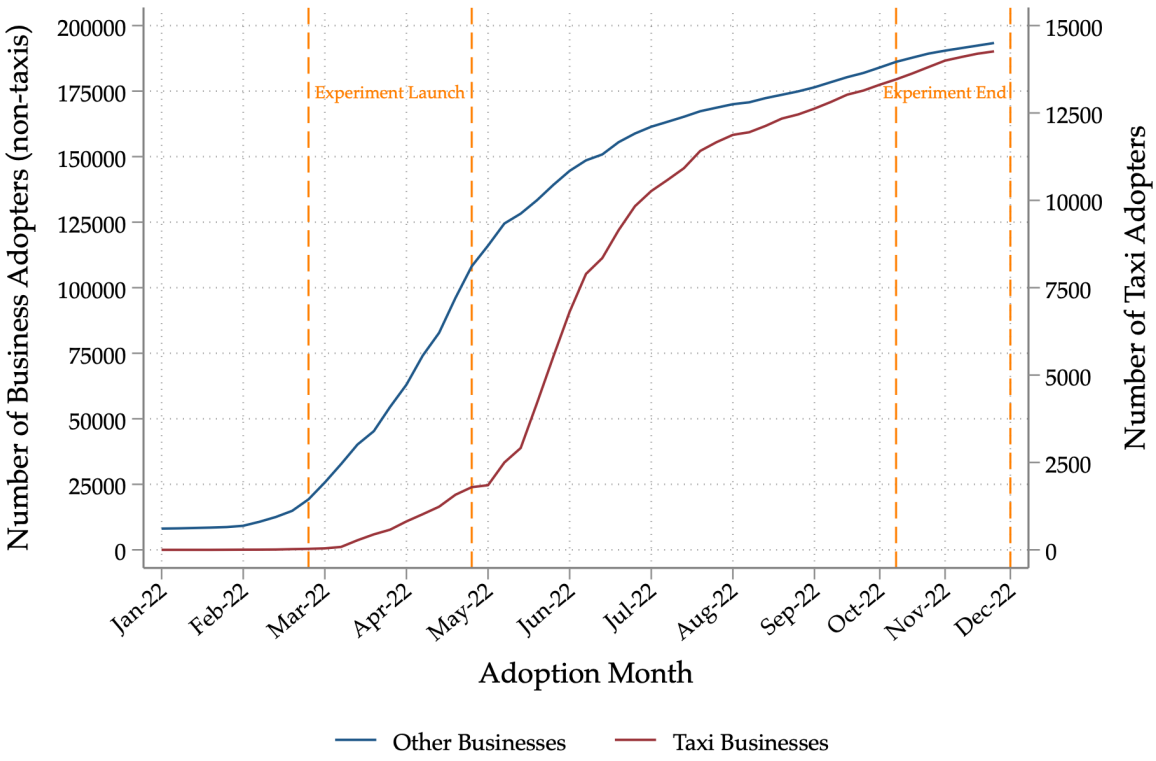



Figure 2: Digital Payments' Adoption in 2022 - *Author's calculations*

*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. A week after the end of the baseline survey, 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. By October 2022, a substantial share of non-taxi businesses had already adopted, while taxi adoption (∼7,000 of an eventual ∼17,000) was still growing.

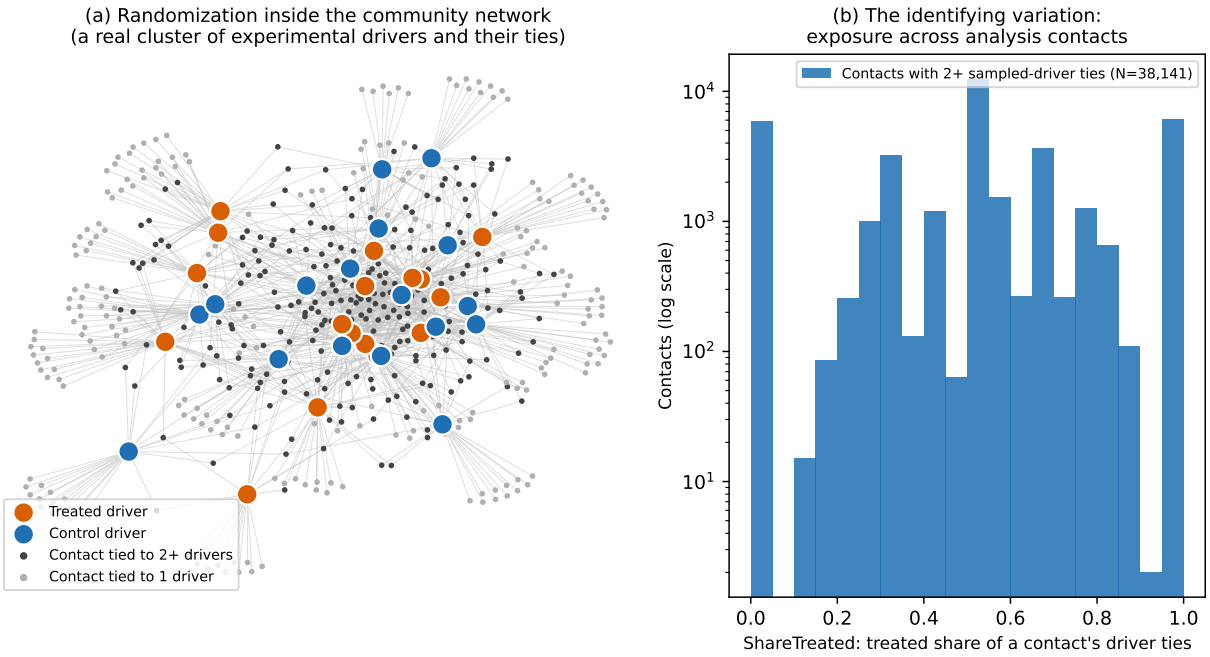



Figure 3: The Experimental Design in the Community Network

*Notes:* Both panels are drawn from the actual analysis data. *Panel (a)*: a connected cluster of 34 experimental drivers (selected by shared-contact intensity from a random sample of the experimental population, balancing arms) with a sample of their community-tied contacts (bias-corrected contact-book overlap $\geq 2\%$; Section 1.3). Orange and blue circles are drivers as randomized into treatment and control; dark gray dots are contacts tied to two or more experimental drivers—the individuals for whom randomization generates variation in exposure—and light gray dots are single-driver contacts. Treatment was assigned to drivers embedded in one connected social fabric; spillovers are identified from contacts' positions within it. *Panel (b)*: the realized distribution of $\text{ShareTreated}_j$ (the treated share of contact $j$'s experimental driver ties) among multi-tie contacts of the sampled drivers—the identifying variation used throughout the analysis (log scale; spikes at coarse fractions reflect contacts with few driver ties). A stylized schematic of the design is in Appendix Figure A2.

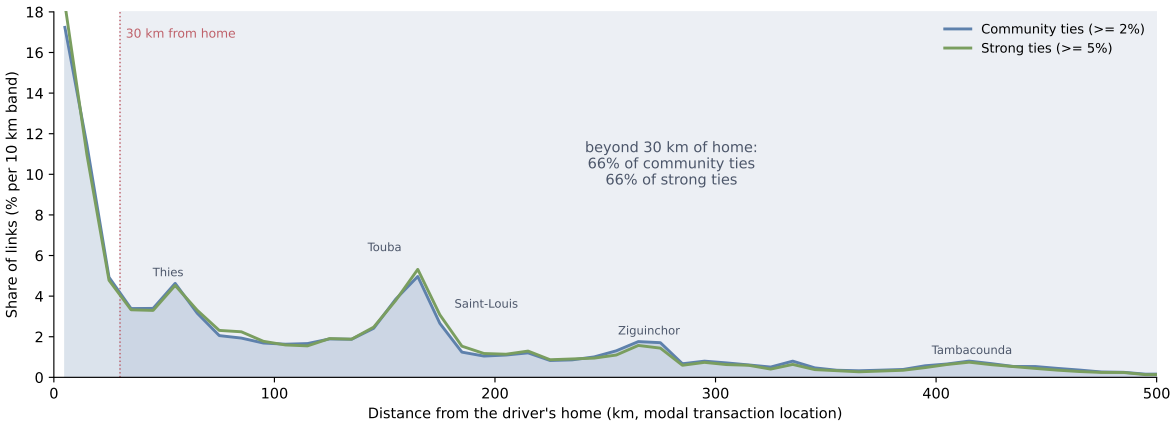



Figure 4: Social Ties by Distance: Decay and Non-Local

*Notes:* This figure plots the share of a driver's located ties per 10 km band of distance from the driver's home, for community ties (corrected overlap $\geq 2\%$) and strong ties ($\geq 5\%$). The driver's home is his modal transaction agent over March–October 2022, the experiment window; each contact is located at the modal location of its agent transactions (Section 1.3). Ties decay with distance and spike at the secondary cities. Beyond 30 km of home, 66% of community ties and 66% of strong ties remain.

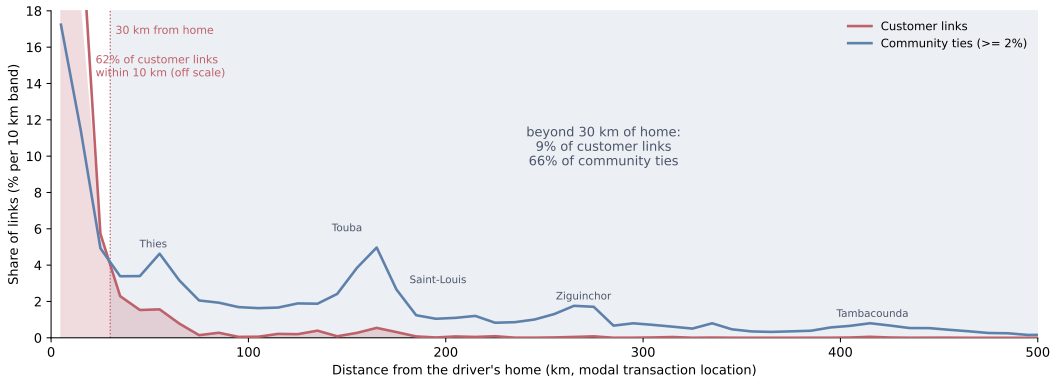



Figure 5: Customer Links Are Local

*Notes:* This figure plots the distance profile of the shared-customer network against the community-tie profile of Figure 4. Each observation is a business–driver customer link, plotted at the distance between the driver's home and the linked business and counted once per link; the vertical axis is the share of links per 10 km band. Customer links are measured over November 2022–November 2023, when both treated and control drivers have merchant accounts: a link joins an experimental driver and a business that share a regular customer, with at least two transactions with each party in that window (Section 1.3). The driver is located at his modal transaction agent over March–October 2022, and businesses at their registered coordinates. Customer links concentrate close to home: 62% lie within 10 km and 10% beyond 30 km, against 66% of community ties.

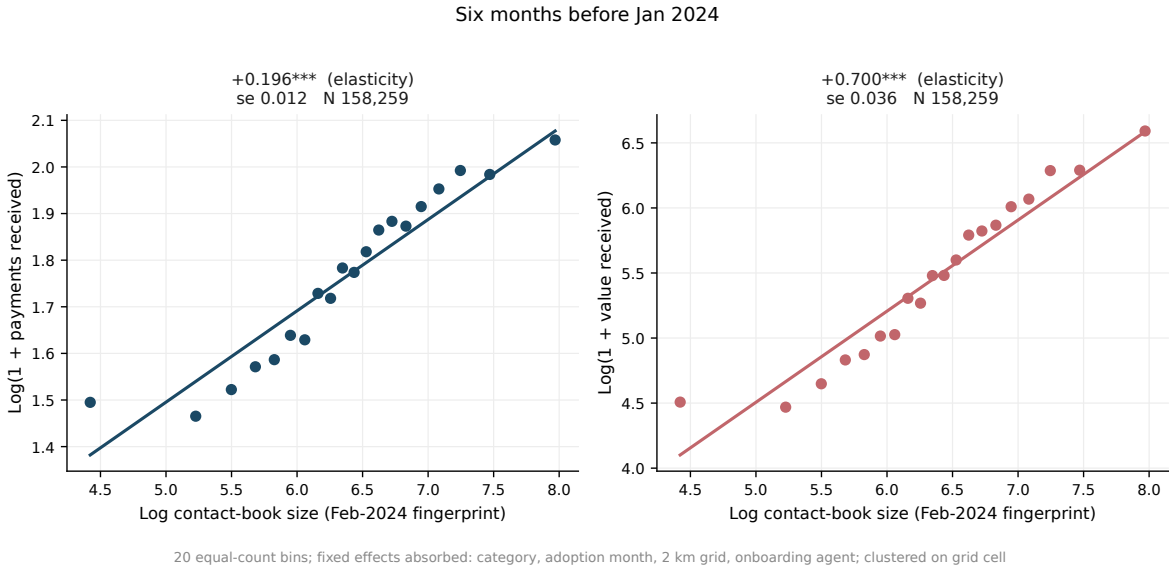
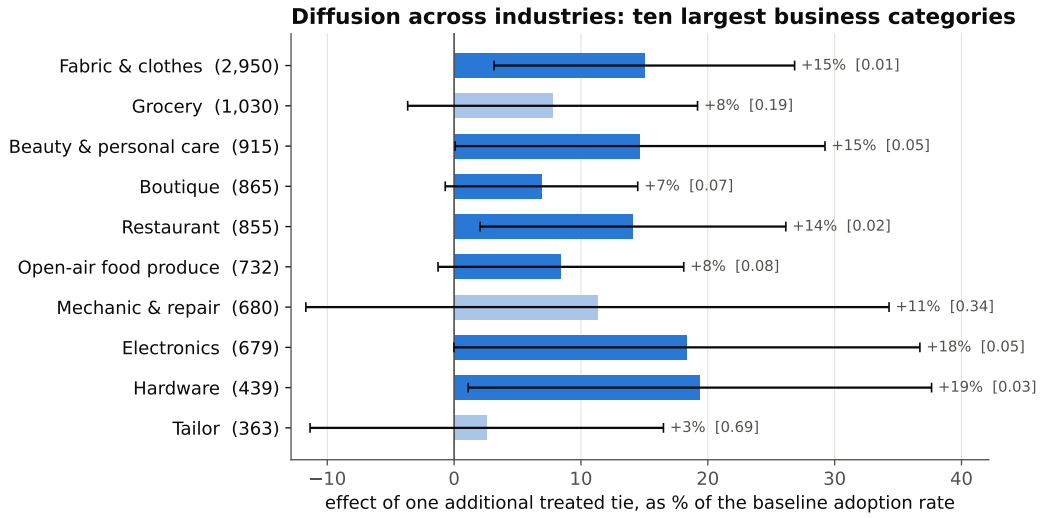



Figure 6: Contact-Book Size and Business Volume

*Notes:* Binned scatter of the log of one plus the number of business payments received (left) and the log of one plus their value (right) in the six months before January 2024, against the log size of the merchant's contact-book fingerprint (February 2024 snapshot), for all merchants trading in that window ($N = 158{,}259$). Twenty equal-count bins after absorbing category, adoption-month, 2 km grid-cell, and onboarding-agent fixed effects; standard errors clustered on grid cell. The fitted elasticities are $+0.196$ (s.e. $0.012$) for payments and $+0.700$ (s.e. $0.036$) for value. The relationship is descriptive: the fingerprint is contemporaneous with the outcome window.




Count specification (number of treated ties, controlling for its expected value under the assignment) on strong ties (corrected overlap ≥ 5%), Dakar and outside pooled; N = 1,020,391 contacts. In parentheses: adopters of that category inside the estimation sample (nationwide in-window adopters are roughly ten times larger, e.g. fabric & clothes 26,944). 95% confidence intervals from the design distribution of the estimator over 1,000 placebo assignments; brackets: RI p-value; dark bars p < 0.10.

Figure 7: Cross-Industry Diffusion by Business Category: Ten Largest Categories

*Notes:* Each bar is the effect of one additional treated strong tie (corrected overlap $\geq 5\%$) on registering a business in the named category, as a percent of the category's baseline adoption rate, from the count specification of Section 3.1 (number of treated ties, controlling for its expected value under the assignment), pooling Dakar and outside Dakar; $N = 1{,}020{,}391$ contacts. In parentheses: adopters of that category inside the estimation sample (nationwide in-window adopters are roughly ten times larger, e.g., fabric and clothes 26,944). Bars show 95% confidence intervals from the design distribution of the estimator over 1,000 placebo assignments; brackets give RI $p$-values; dark bars have $p < 0.10$.

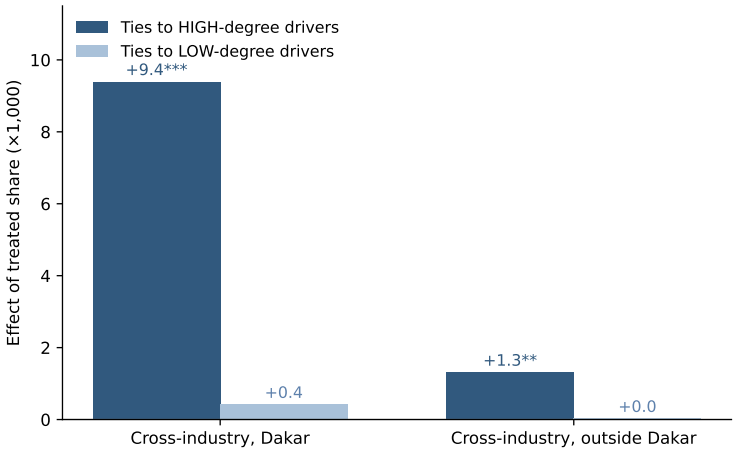



Figure 8: Cross-Industry Transmission by the Treated Driver's Connectedness

*Notes:* Each firm's treated share is split by the connectedness of the treated driver on the other side of the tie, at the median of the drivers' strong-tie degree (median strong ties 1,516 in the high group and 265 in the low group), and the two exposures enter the specification of Table 1 jointly, with their expected-exposure controls. Bars are the coefficients on the treated share ($\times 1{,}000$) for cross-industry adoption in Dakar and outside Dakar; stars denote RI $p$-values ($^{*} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$). Degree is uncorrelated with treatment (59.0% versus 59.1% treated).

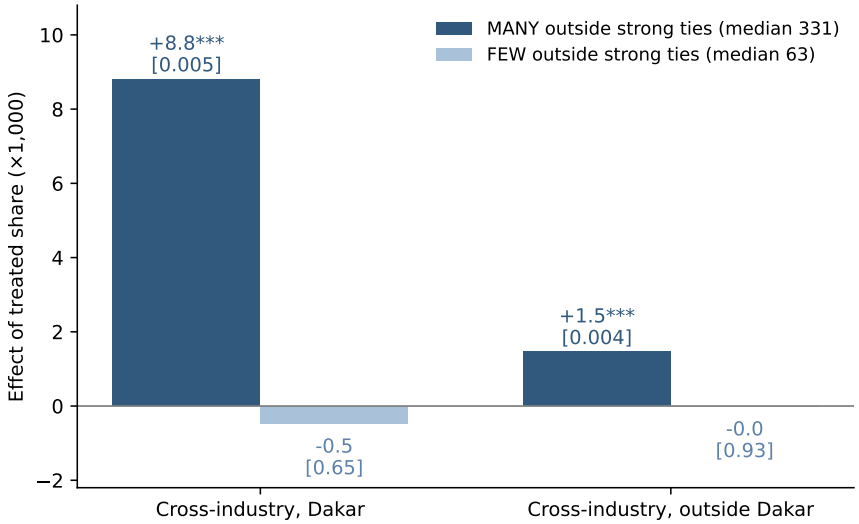



Figure 9: Cross-Industry Transmission by the Treated Driver's Countrywide Reach

*Notes:* Each firm's treated share is split by the treated driver's number of strong ties ($\geq 5\%$) located outside Dakar, at the median (331 versus 63 located ties), and the two exposures enter the specification of Table 1 jointly. The sample is the $N = 143{,}854$ firms with ties to both groups. Bars are the coefficients on the treated share ($\times 1{,}000$) for cross-industry adoption in Dakar and outside Dakar, with RI $p$-values in brackets. Every driver's network is about 40% Dakar and 60% country, so outside reach is total reach; the split is balanced across arms (59.2% versus 59.1% treated).

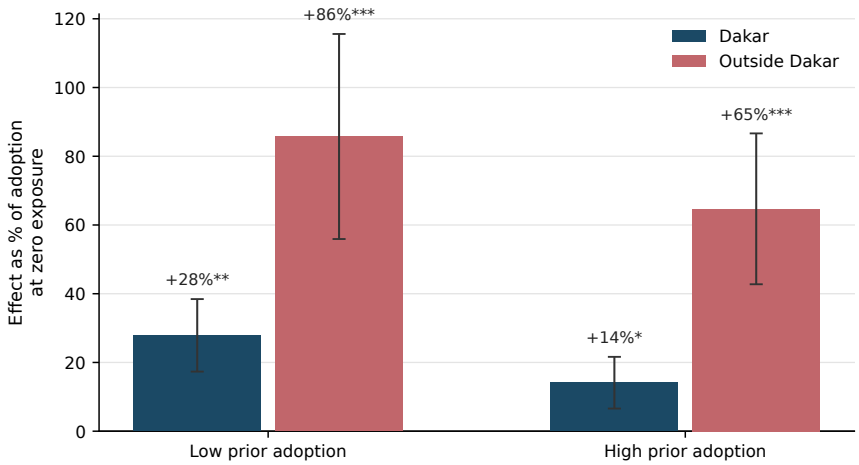



Figure 10: Cross-Industry Transmission Is Strongest Where Knowledge Is Scarce

*Notes:* Effect of a fully treated tie portfolio on cross-industry adoption, as a share of predicted adoption at zero exposure, by zone type. Zones are split at the median of pre-experiment merchant density; Dakar is divided into 2 km neighbourhoods and the rest of the country into 10 km towns. Whiskers are $\pm 1$ design standard deviation from 1,000 placebo assignments. $N = 533{,}541$.

Table 1: Cross-Industry Network Spillovers by Tie Strength

| | $\geq 2\%$ (community) | $\geq 5\%$ (strong) | $\geq 10\%$ |
|---|---|---|---|
| **Cross-industry, Dakar** | | | |
| Share treated $\beta$ | $+2.91$ {+41%} | $+4.16^{**}$ {+39%} | $+5.16^{**}$ {+69%} |
| | (0.50) | (0.29) | (0.42) |
| | [0.123] | [0.017] | [0.015] |
| Expected exposure $\lambda$ | $-1.78$ | $-22.75$ | $-32.48$ |
| | (15.92) | (9.60) | (13.96) |
| | [0.911] | [0.018] | [0.020] |
| **Cross-industry, outside Dakar** | | | |
| Share treated $\beta$ | $+0.68^{**}$ {+138%} | $+0.68^{***}$ {+81%} | $+0.73^{**}$ {+137%} |
| | (0.16) | (0.09) | (0.13) |
| | [0.011] | [0.009] | [0.030] |
| Expected exposure $\lambda$ | $+8.35$ | $-2.85$ | $-1.00$ |
| | (5.52) | (3.02) | (4.13) |
| | [0.130] | [0.345] | [0.810] |
| Baseline $\hat{y}(S{=}0)$, Dakar / outside | 7.07 / 0.49 | 10.66 / 0.84 | 7.53 / 0.54 |
| $N$ | 1,145,899 | 1,020,391 | 376,871 |

*Notes:* All estimates $\times$1,000; effect sizes in braces, computed as $\beta_1/\hat{y}(S{=}0)$, where $\hat{y}(S{=}0) = \bar{y} - \hat{\beta}_1\bar{S}$ is predicted adoption at zero treated exposure with controls at their means ($\bar{S} = 0.590$, 0.666, 0.628 at the three thresholds); HC1 standard errors in parentheses; $p$-values in brackets—randomization inference (1,000 permutations) for Share treated, HC1 for Expected exposure (a control, not re-randomized). Each cell estimates $y_j = \beta_1\,\text{ShareTreated}_j + \beta_2\,\text{ExpExposure}_j + \alpha$: $y_j$ indicates registration as a non-taxi merchant (in Dakar, or outside Dakar) during the experiment window, $\text{ShareTreated}_j$ is the treated share of contact $j$'s experimental driver contacts, and $\text{ExpExposure}_j$ is its mean across the 1,000 placebo re-randomizations, absorbing $j$'s network position. Ties are all links with bias-corrected contact-book overlap above the indicated threshold, saved or not (2% floor set above the collision-noise scale and survey-validated; Section 1.3); contacts connected to any top-2% degree driver are excluded entirely. Within-industry effects are in Appendix Table B6; per-band transmission is in Appendix Figure A8; the reciprocity split is in Appendix Table B5; results excluding shared-customer pairs and garage-based drivers are in Appendix Table B7. $^{*}p_{\text{RI}} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$.

Table 2: Count Specification: Number of Treated Ties

| | $\geq 2\%$ | $\geq 5\%$ | $\geq 10\%$ |
|---|---|---|---|
| Cross-industry, Dakar | $+0.25$ [.114] {+4%} | $+1.67^{**}$ [.015] {+16%} | $+2.01^{***}$ [.008] {+25%} |
| Cross-industry, outside | $+0.05^{**}$ [.021] {+17%} | $+0.30^{***}$ [.007] {+39%} | $+0.25^{*}$ [.052] {+37%} |
| Within-taxi, Dakar | $+0.035$ [.214] {+22%} | $-0.063$ [.840] {$-$2%} | $+0.392$ [.608] {+8%} |
| Within-taxi, Thiès | $+0.001$ [.922] {+1%} | $+0.018$ [.786] {+5%} | $+0.293$ [.154] {+53%} |
| Within-taxi baseline (Dakar / Thiès) | 0.158 / 0.062 | 2.730 / 0.338 | 4.742 / 0.552 |
| $N$ | 1,145,899 | 1,020,391 | 376,871 |

*Notes:* Coefficients $\times 1{,}000$ on the *number* of treated ties (rather than the treated share), from the specification of Table 1 with NumTreated replacing ShareTreated and the expected-count control replacing expected exposure; community-tie network at the indicated corrected-overlap thresholds, strict top-2% trimming, RI $p$-values (1,000 permutations) in brackets, effect sizes per additional treated tie in braces, computed as $\beta_1/\hat{y}(\text{NT}{=}0)$, where $\hat{y}(\text{NT}{=}0) = \bar{y} - \hat{\beta}_1\overline{\text{NT}}$ is predicted adoption at zero treated ties with controls at their means (the within-taxi baseline row reports it $\times 1{,}000$) ($\overline{\text{NT}} = 11.99, 1.74, 1.31$ at the three thresholds). Cross-industry learning survives the count specification (certified at $\geq 5\%$ and $\geq 10\%$). The within-taxi coefficients—negative and significant in the share specification (Appendix Table B6)—are null in counts at every threshold, whereas learning operates tie by tie. $^{*}p_{\text{RI}} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$.

Table 3: 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 | 98.45** | 0.362 | | | 614 |
| | (37.45) | | | | |
| | [0.011] | | | | |
| Taxi × Non-friends | 8.75 | 0.028 | | | 549 |
| | (39.39) | | | | |
| | [0.835] | | | | |
| Taxi × Freq. daily/weekly | 69.54** | 0.240 | | | 916 |
| | (30.42) | | | | |
| | [0.039] | | | | |
| *Panel B: All contacts by social proximity* | | | | | |
| Friends | 30.44** | 0.316 | 0.86 | 0.013 | 1,983 |
| | (14.12) | | (11.18) | | |
| | [0.043] | | [0.933] | | |
| Freq. daily/weekly | 15.27 | 0.213 | 3.23 | 0.065 | 4,119 |
| | (8.47) | | (6.80) | | |
| | [0.116] | | [0.641] | | |
| *Panel C: Casual ties* | | | | | |
| Clients | 4.59 | 0.369 | 35.15 | 0.944 | 395 |
| | (12.25) | | (22.75) | | |
| | [0.791] | | [0.150] | | |
| Serendipitous encounters | 2.87 | 0.165 | 18.00*** | 0.475 | 5,112 |
| | (3.84) | | (5.90) | | |
| | [0.491] | | [0.003] | | |

*Notes:* Each row is a *separate* regression: $y_j = \alpha + \beta_1 \cdot \text{NumTreated}_{s,j} + \beta_2 \cdot \text{NumContacts}_{s,j} + \varepsilon_j$, where $\text{NumContacts}_{s,j}$ is the count of $j$'s experimental contacts in $s$ (absorbing network position within the subnetwork), where exposure is computed over the indicated relationship-defined subnetwork $s$ only, on the subsample with $\geq 1$ connection in $s$, restricted to baseline survey contacts. "Friends" and "Non-friends" partition the surveyed taxi ties, so the two rows compare transmission through the two halves of an exclusive split; an individual connected through both tie types (to different drivers) may contribute to both rows, and the frequency rows re-slice the same ties along a different dimension. Coefficients $\beta_1 \times 1{,}000$ reported; HC1 robust standard errors in parentheses. Effect size $= \beta_1 / \hat{y}(0)$ with $\hat{y}(0)$ predicted adoption at zero treated contacts (controls at their means). "Taxi adoption, Dakar" = within-industry adoption; "Cross-ind. adoption" = across-industry adoption (Dakar + outside). Stars based on randomization inference (RI) $p$-values that re-randomize the business-level treatment assignment using the 1,000 stratified draws of the original protocol, rebuilding each exposure under every draw; conventional $t$-test $p$-values are similar throughout. Panel A shows within-industry outcomes only, as taxi contacts have no cross-industry adoption variation. "Serendipitous encounters" are contacts classified as casual or unknown relationship. $^{*}p_{\text{RI}} < 0.10$, $^{**}p_{\text{RI}} < 0.05$, $^{***}p_{\text{RI}} < 0.01$.

Table 4: Diffusion Across Space: Widespread, and Not Only Through Migration

| | Share treated | $p_{\text{RI}}$ | Effect size | $N$ |
|---|---|---|---|---|
| **Overall effect** | $+0.68^{***}$ | [0.009] | $+138\%$ | 1,145,816 |
| *Panel A: By the user's own mobility* | | | | |
| User **visited** Dakar | $+0.39$ | [0.118] | $+182\%$ | 281,207 |
| User **never** visited Dakar | $+1.82^{**}$ | [0.027] | $+88\%$ | 275,683 |
| *Panel B: By the driver's mobility to the destination (pooled across cities)* | | | | |
| Treated driver **visited** the city (migrant) | $+0.21$ | [0.133] | $+400\%$ | 1,145,899×8 |
| Treated driver **did not** visit the city | $+0.07^{**}$ | [0.018] | $+125\%$ | 1,145,899×8 |

*Notes:* Coefficients are multiplied by 1,000. The sample is the 1,145,816 individuals in the community-tie network (bias-corrected contact-book overlap $\geq 2\%$, the survey-validated community-tie floor; Section 1.3), under the same trimming as Table 1. Each cell reports $\beta_1$ from $y_j = \beta_1 \,\text{ShareTreated}_j + \beta_2 \,\text{ExpExposure}_j + \alpha$, where $\text{ShareTreated}_j$ is the share of individual $j$'s experimental driver contacts assigned to treatment and $\text{ExpExposure}_j$ is its mean across 1,000 placebo re-randomizations (included to absorb network position); $\beta_2$ is insignificant throughout. The overall effect is positive and significant: treatment-induced cross-industry adoption reaches secondary cities across thirteen regions of Senegal—Thiès and Mbour account for only 41% of in-network destination adopters, with Touba, Saint-Louis, Kaolack, Ziguinchor, and Tambacounda among the contributors. *Panel A* (outcome: cross-industry adoption outside Dakar) splits individuals by their *own* mobility—whether they were observed in Dakar at least once during the experiment (mobile-money geolocation): both groups respond, but the never-visited coefficient is five times larger and only it is statistically significant—the effect does not require the user to travel to the capital. *Panel B* (outcome: cross-industry adoption in the destination city $c$) splits the treatment by the *drivers'* mobility. Both rows come from the single stacked regression $y_{jc} = \beta_1 \,\text{ShareTreated}_{jc}^{\text{visit}} + \beta_2 \,\text{ShareTreated}_{jc}^{\text{no-visit}} + \gamma_1 \,\overline{S}_{jc}^{\text{visit}} + \gamma_2 \,\overline{S}_{jc}^{\text{no-visit}} + \alpha_c + \varepsilon_{jc}$, where the two shares partition $j$'s treated driver contacts by whether the driver transacted at an agent in city $c$ at least once during the experiment (agent geolocation), so they sum to $\text{ShareTreated}_j$ for every city. The specification controls for *both* an expected-exposure term for each share ($\overline{S}$, its mean across the 1,000 placebo re-randomizations, absorbing the fact that individuals connected to $c$-visiting drivers disproportionately live near $c$) *and* city fixed effects $\alpha_c$; the sample stacks the eight largest secondary cities ($1{,}145{,}899 \times 8 = 9{,}167{,}192$ individual–city observations, differing from the overall row's 1,145,816 by 83 individuals due to independent pipeline construction), with SEs clustered by individual and RI re-randomizing the driver assignment and rebuilding both shares in each draw; see Section 5.2 for interpretation. RI $p$-values (1,000 permutations) in brackets; effect size $= \beta_1/\hat{y}(S{=}0)$, predicted adoption at zero treated exposure (controls at their means). $^{*}p_{\text{RI}} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$.

Table 5: Exposure Does Not Shift the Customer Mix

| Payer counted as treated-induced if they paid a treated taxi … | Control mean | Firms with no such payer | Effect per treated tie | [$p_{\text{RI}}$] |
|---|---|---|---|---|
| **Within-taxi** (2,105 firms, 135,663 payments) | | | | |
| … during the experiment | 24.6% | 0.6% | $+1.8\%$ | [.104] |
| … before their first payment to the firm | 24.6% | 2.5% | $+1.2\%$ | [.387] |
| … before the firm adopted | 9.8% | 6.9% | $+2.4\%$ | [.051] |
| … within 30 days before paying | 6.8% | 15.2% | $+1.0\%$ | [.444] |
| **Cross-industry, Dakar** (5,810 firms, 604,562 payments) | | | | |
| … during the experiment | 6.2% | 39.7% | $+3.2\%$ | [.463] |
| … before their first payment to the firm | 3.1% | 56.1% | $+5.5\%$ | [.288] |
| … before the firm adopted | $<$0.1% | 80.6% | $+5.5\%$ | [.562] |
| … within 30 days before paying | 1.2% | 74.5% | $-0.2\%$ | [.963] |

*Notes:* Binomial logit of a firm's months 1–6 payments from treated-induced customers out of its total payments, on the number of treated ties; controls for expected treated count, degree, and adoption-month fixed effects. Payment-weighted, so firm size drops out exactly. Each row uses a different definition of a treated-induced payer: a customer who paid a treated taxi during the experiment, before their first payment to the firm, before the firm adopted, or within 30 days before paying. The control mean is the predicted share at zero treated ties; the effect is the change in that share per additional treated tie, in percentage points. $p_{\text{RI}}$ from 1,000 placebo assignments.

Table 6: Usage Among Adopters and Social Exposure, by Population

| | Within-taxi Dakar | Cross-industry Dakar | Cross-industry outside Dakar |
|---|---|---|---|
| $\log(1 + \text{payments})$ | $+0.09$ [.387] | $+0.10^{*}$ [.051] | $+0.54^{***}$ [.002] |
| $\geq 10$ payments | $+0.02$ [.435] | $+0.02^{*}$ [.076] | $+0.13^{***}$ [.001] |
| $N$ (adopters) | 2,674 | 13,706 | 1,316 |

*Notes:* Each cell is a separate regression among adopters of the indicated population. The outcome is usage in the first six months after adoption: $\log(1 + \text{payments})$, the log of one plus the number of business payments received, or an indicator for receiving at least ten payments. The regressor is social exposure through strong ties (corrected overlap $\geq 5\%$) to treated experimental drivers. RI $p$-values (1,000 permutations) in brackets. $^{*}p_{\text{RI}} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$.

Table 7: Learning Relative to All Other Forces: A Lower Bound

| Cross-industry adoption | $r = \beta/\hat{y}(S{=}0)$ |
|---|---|
| Dakar, strong ties ($\geq 5\%$) | 0.39 |
| Outside Dakar, strong ties ($\geq 5\%$) | 0.81 |
| Outside Dakar, community ties ($\geq 2\%$) | 1.38 |

*Notes:* $r = \beta/\hat{y}(S{=}0)$ is the ratio of the Share treated coefficient to predicted adoption at zero treated exposure, from Table 1 (the effect sizes in braces there; the community-tie row is the estimate of Table 4). The baseline $\hat{y}(S{=}0)$ contains every force other than learning through treated ties, including customer demand, so $r$ is a lower bound on learning relative to demand. The comparison is at full exposure; the realized mean treated share is $0.666$ on strong ties and $0.590$ on community ties (Table 1).

# Online Appendix

# Nationwide Diffusion of Technology: Experimental Evidence on Cross-Industry Learning Through Networks

Deivy Houeix

**A Appendix Figures** **53**

**B Appendix Tables** **62**

**C Fingerprint Data Processing** **74**

# A Appendix Figures

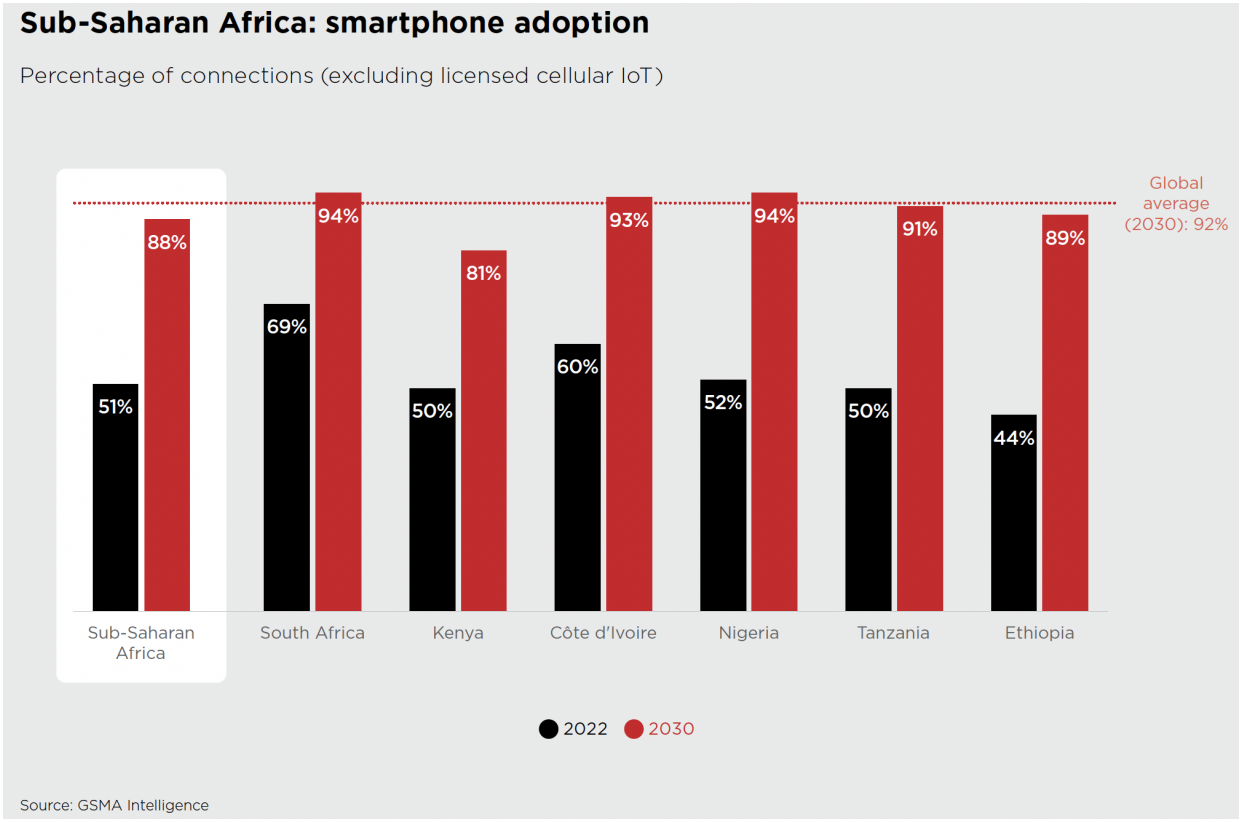



Figure A1: Smartphone Adoption in Sub-Saharan Africa

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

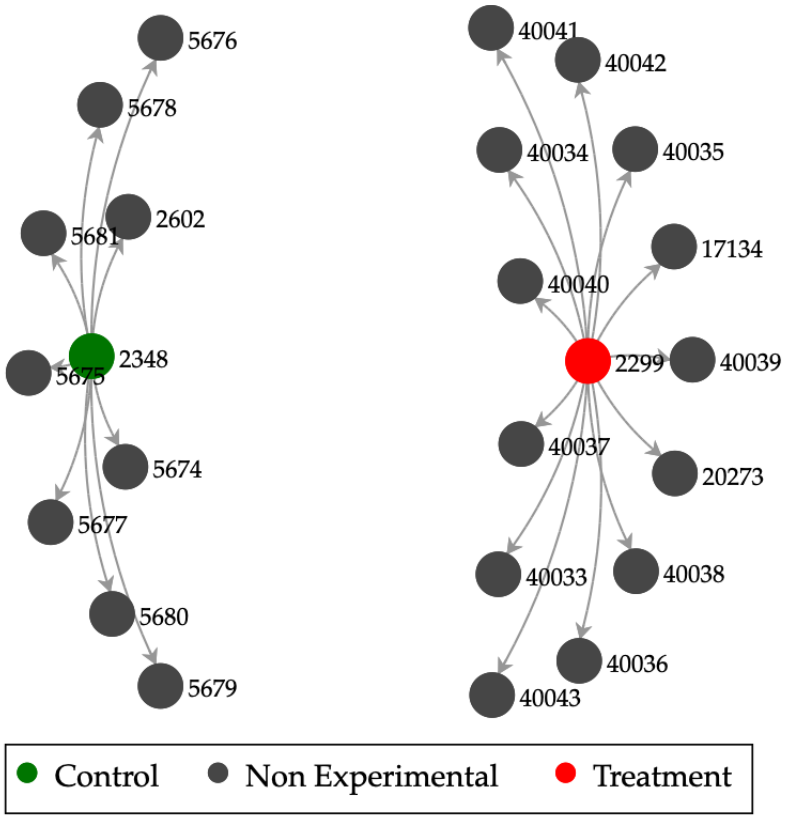



Figure A2: Illustration of the Experimental Design (Schematic)

*Notes:* This figure illustrates and simplifies the experimental strategy. Colored nodes (red/green) are experimental taxi businesses randomly assigned to treatment or control. Black nodes are non-experimental individuals connected to the experimental sample through the phone network. The study quantifies the likelihood of technology adoption among these non-experimental contacts, comparing adoption rates across the networks of treated versus control taxi businesses. In practice, networks often overlap, leading to network interference, an issue discussed in Section 3.3. Figure 3 shows the design in the actual community network.

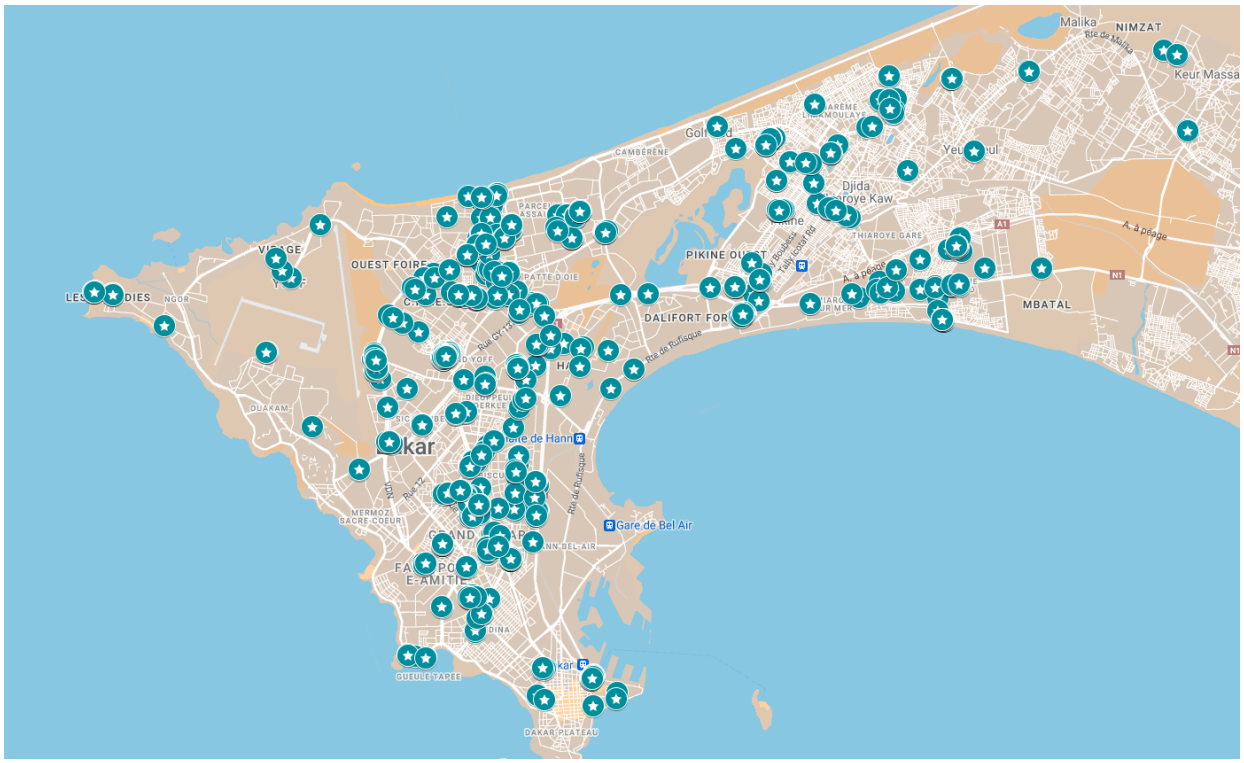



Figure A3: Listing Survey Activity in Dakar - March/April 2022

*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 parts 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, and some were not eligible). Owners not driving were primarily recruited by asking drivers about their owners' contact information during the listing survey.

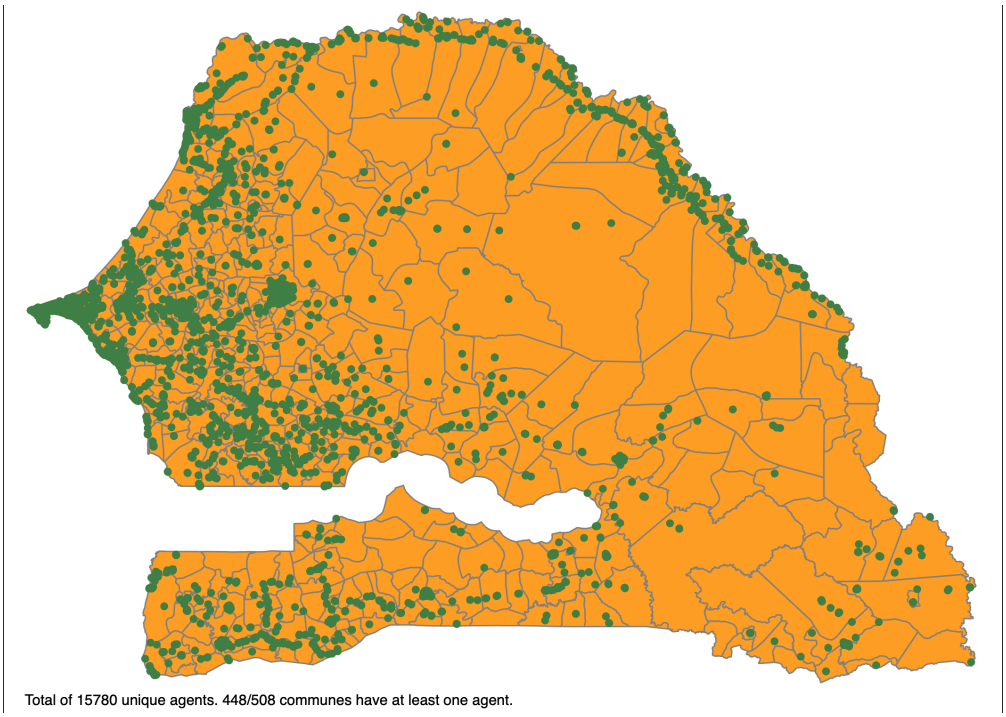



Figure A4: Mobile Money Agents in Senegal in 2022

*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 company. There are over 15,000 unique agents in Senegal from this mobile money company, which covers the full set of urban cities and the vast majority of communes/villages.

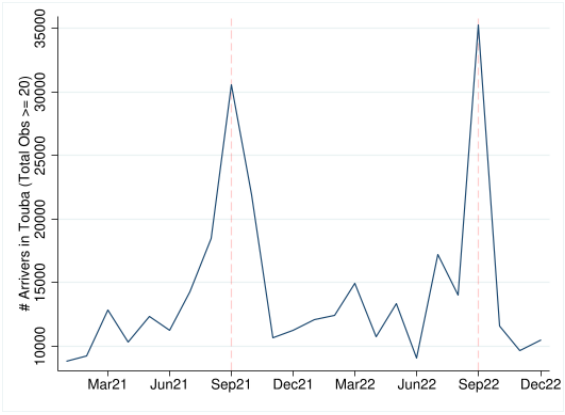



Figure A5: Location Data - Arrivals in Touba

*Notes:* This Figure represents the number of arrivals 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.

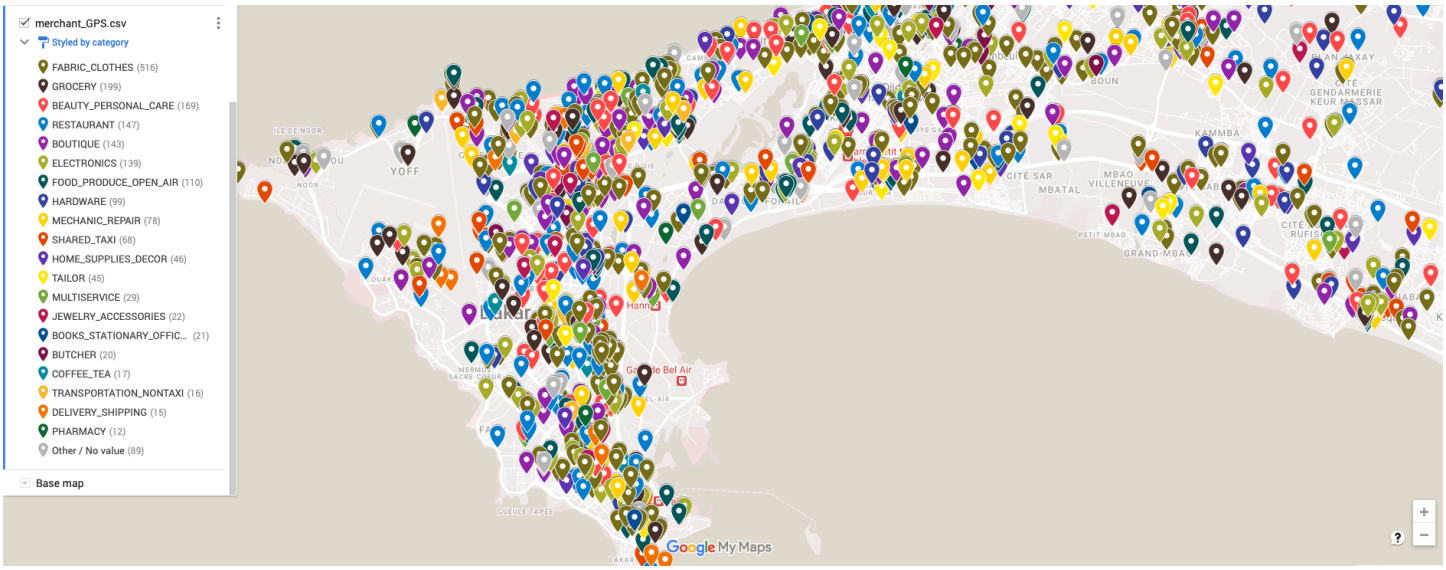



Figure A6: Technology Diffusion in Dakar Throughout the Experiment

*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 that adopted the digital payment technology during the time of the experiment.

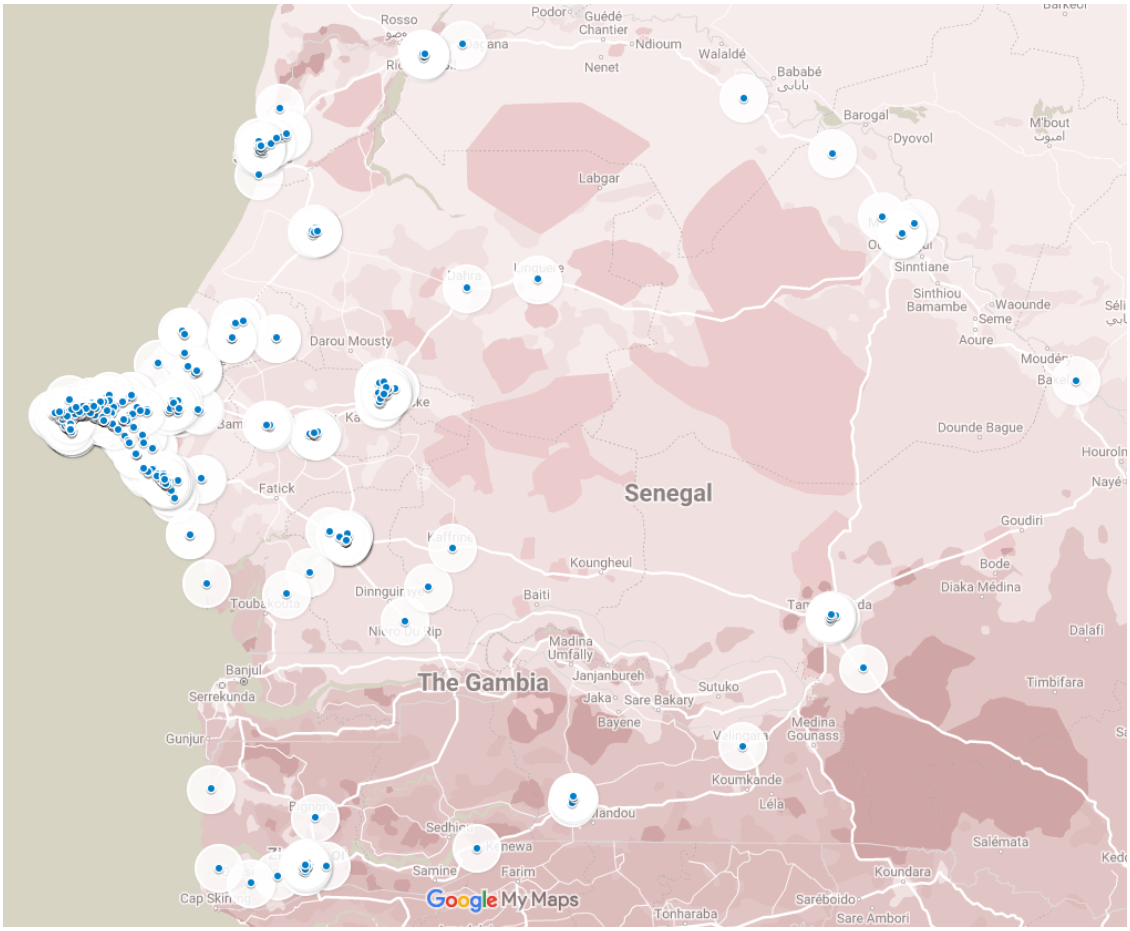



Figure A7: Technology Diffusion Inter-Industry Across Space Throughout the Experiment

*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 that adopted the digital payment technology during the time of the experiment.

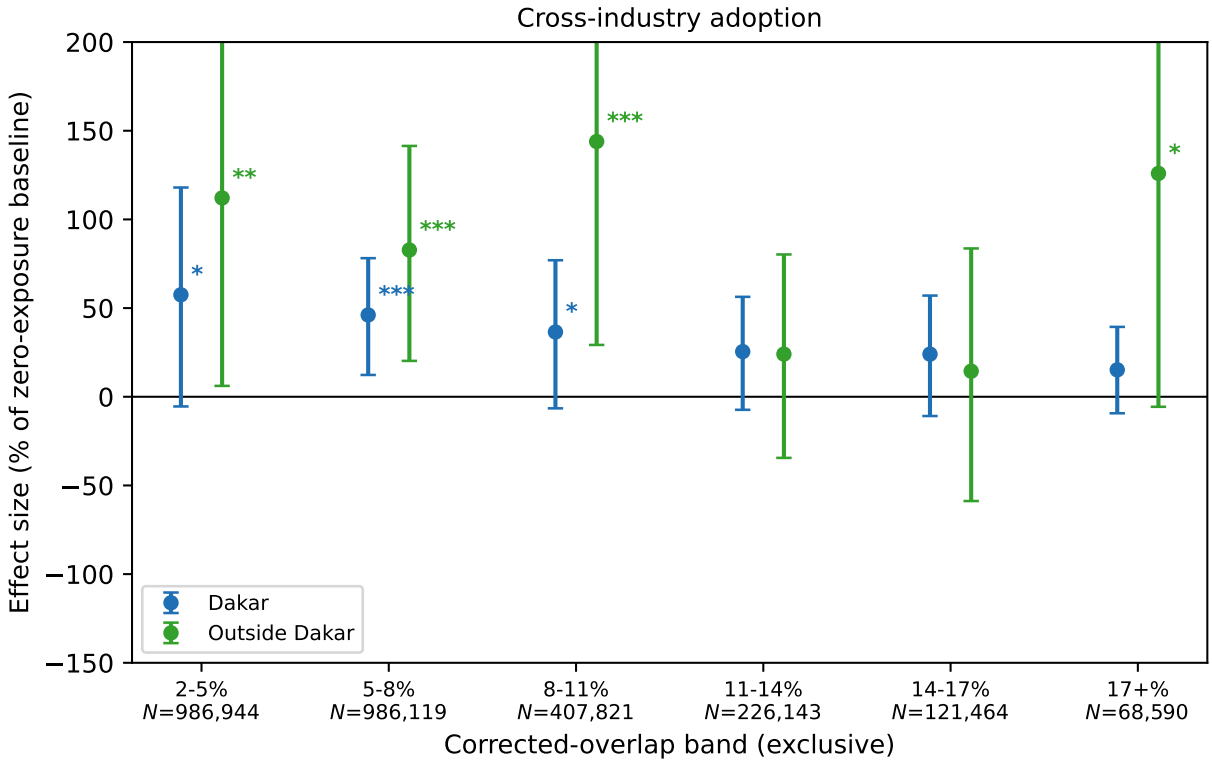



Figure A8: Per-band transmission by tie strength

*Notes:* Effect sizes ($\hat{\beta}_1/\hat{y}(S{=}0)$, in percent of predicted adoption at zero exposure, with $\hat{y}(S{=}0) = \bar{y} - \hat{\beta}_1\bar{S}$) of ShareTreated on adoption, estimated separately within exclusive bands of bias-corrected contact-book overlap (specification, sample construction, and strict trimming as in Table 1; band $N$s from 986,944 at 2–5% to 68,590 at 17%+). Error bars are 95% randomization-inference confidence intervals obtained by test inversion: the set of constant effects $\tau_0$ whose sharp null is not rejected at 5% across the 1,000 placebo re-randomizations. Stars denote RI $p$-values at the zero null ($^{*} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$). Panel (a): cross-industry transmission operates throughout the strength distribution—certified in the 5–8% band in Dakar and in the 2–5, 5–8, and 8–11% bands outside Dakar—showing that diffusion operates throughout the community network rather than only at the closest ties. Panel (b): the within-taxi effect is negative in the 5–8% band (Dakar $-29\%$, CI $[-50, -6]$) and is absent at both the weakest and strongest ties.

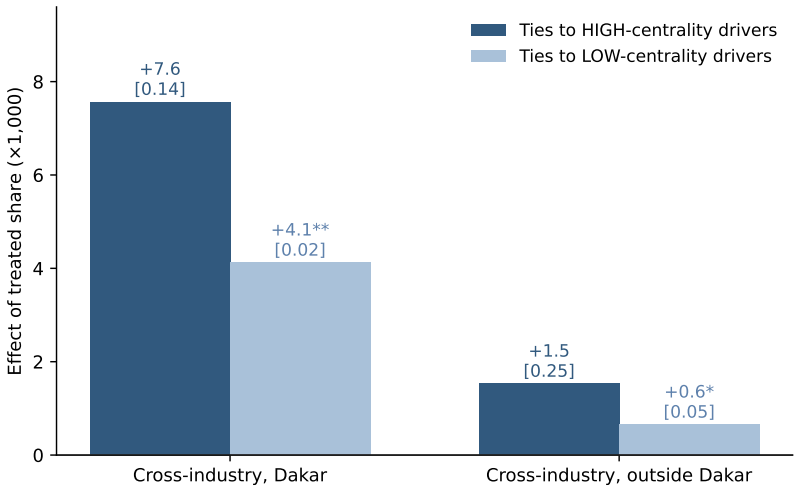
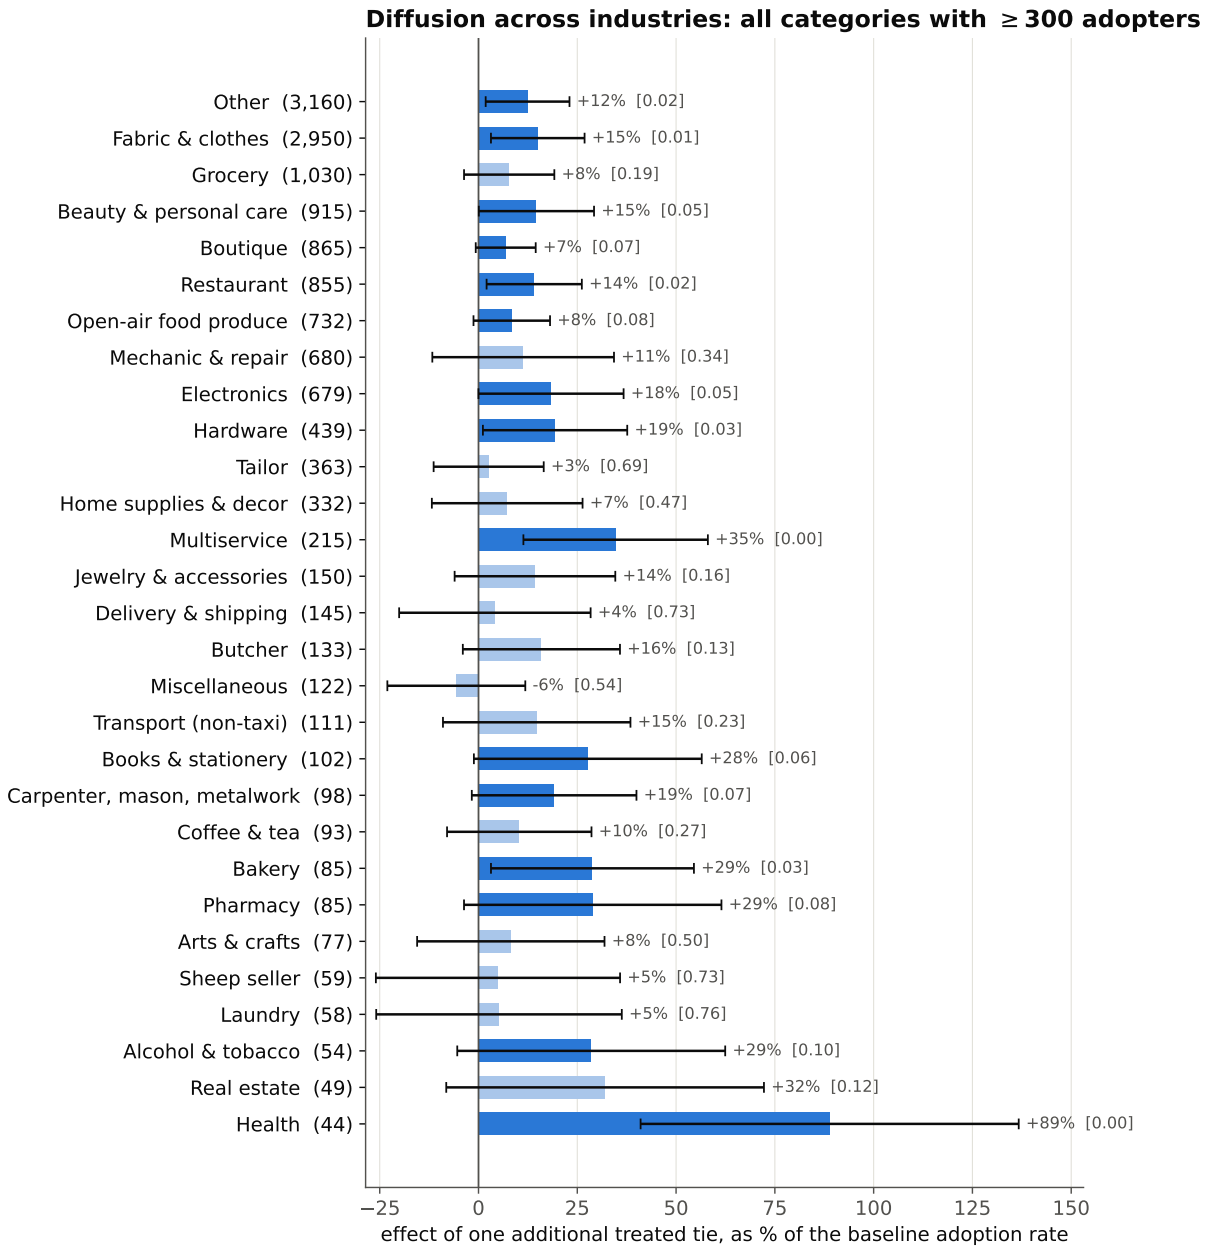



Figure A9: Robustness: Centrality Instead of Degree

*Notes:* Same joint specification as Figure 8, splitting treated drivers at the median of their eigenvector centrality on the driver–driver shared-contact network (correlation with degree 0.74). Bars are the coefficients on the treated share ($\times 1{,}000$) for cross-industry adoption in Dakar and outside Dakar, with RI $p$-values in brackets. The low-centrality group retains a mean degree of 614; the treated share is balanced across the two groups (58.1% versus 60.0%).




Count specification (number of treated ties, controlling for its expected value under the assignment) on strong ties (corrected overlap ≥ 5%), Dakar and outside pooled; N = 1,020,391 contacts. In parentheses: adopters of that category inside the estimation sample (nationwide in-window adopters are roughly ten times larger, e.g. fabric & clothes 26,944). 95% confidence intervals from the design distribution of the estimator over 1,000 placebo assignments; brackets: RI p-value; dark bars p < 0.10.

Figure A10: Cross-Industry Diffusion by Business Category: All 29 Categories

*Notes:* Same specification as Figure 7, for all 29 non-taxi business categories, ordered by the number of adopters of the category inside the estimation sample (in parentheses). Bars show 95% confidence intervals from the design distribution of the estimator over 1,000 placebo assignments; brackets give RI $p$-values; dark bars have $p < 0.10$.

# B Appendix Tables

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

| Relationship category | $n$ | Mean corrected overlap (%) | Share $\geq 2\%$ (community) | Share $\geq 5\%$ (strong) | Share $\geq 10\%$ | Share mutual |
|---|---|---|---|---|---|---|
| Wife | 495 | 20.81 | 86.5% | 76.8% | 64.4% | 72.7% |
| Taxi contact | 1,882 | 13.64 | 81.4% | 62.8% | 38.6% | 62.3% |
| Family | 3,814 | 11.44 | 80.8% | 64.1% | 41.8% | 63.8% |
| Friend (family friend) | 6,712 | 9.09 | 72.5% | 53.3% | 31.6% | 59.3% |
| Professional | 6,077 | 8.94 | 65.7% | 46.0% | 26.4% | 55.6% |
| Friend (non-family) | 2,922 | 6.00 | 61.9% | 39.1% | 18.2% | 53.5% |
| Serendipitous | 7,739 | 3.20 | 27.7% | 13.7% | 6.3% | 20.1% |
| Client | 526 | 1.94 | 22.8% | 8.7% | 3.0% | 36.7% |
| *By interaction frequency* | | | | | | |
| Weekly or more | 6,028 | 10.56 | 73.3% | 54.4% | 33.7% | 61.5% |
| Less than monthly | 2,770 | 5.54 | 52.3% | 31.4% | 14.6% | 46.2% |

*Notes:* This table cross-validates the bias-corrected contact-book 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 surveyed (driver, contact) pair matched to fingerprint data under the corrected E.164 hashing scheme (Appendix C; 22,537 of 23,520 pairs classified, of which 16,470 have fingerprints on both sides), we compute the bias-corrected overlap share $\max(I-\mu, 0)/\min(|F_i|, |F_j|)$, where $\mu$ is the expected hash-collision count, and report the category mean and the share of pairs above the 2% community-tie floor, the 5% strong-tie threshold, and 10%. Corrected overlap rises steeply with relationship closeness: wives average 20.8% (86.5% above the community floor; 76.8% strong), taxi colleagues and family 11–14% (about 81% above the floor), friends 6–9%, while serendipitous contacts (numbers the respondent could not place when prompted—typically incidental transaction partners) average 3.2% with only 27.7% above the floor, and taxi clients 1.9% (22.8%)—the calibrated 2% floor cleanly separates durable relationships from incidental contact. Contacts seen weekly or more share 1.9× the overlap of contacts seen less than monthly. The last column reports the share of each category in which both parties saved each other's number (mutual): reciprocity follows the same gradient—72.7% of wives versus 36.7% of clients and 20.1% of serendipitous contacts—consistent with reciprocal saving marking relationship depth (Table B3).

Table B2: Closeness by Phone Link Type

| | Mutual (both saved) | One-sided (either saved) | Comm. tie, not saved ($\geq 2\%$) | No comm. tie ($< 2\%$) | No fingerprint |
|---|---|---|---|---|---|
| Share close (wife/family/friend) | 61.1% | 36.8% | 37.1% | 18.2% | 34.9% |
| $N$ | 6,517 | 4,173 | 768 | 5,012 | 6,067 |

*Notes:* Survey-validated (driver, contact) pairs from the endline kinship module, matched to fingerprint data under the corrected E.164 hashing scheme (Appendix C); 22,537 classified pairs. Contacts were sampled from the pre-experiment mobile money transaction network. "Close" $= 1$ if the driver reports the contact as wife, family, or friend. Categories: *mutual* = both parties' phone hashes appear in each other's contact-book fingerprint; *one-sided* = exactly one does; *community tie, not saved* = neither saved the other, but bias-corrected contact-book overlap is at least 2%—the community-tie floor used throughout the analysis; *no community tie* = both have fingerprints, no save in either direction, and corrected overlap below 2%; *no fingerprint* = the contact lacks fingerprint data on all their numbers. Two facts stand out. First, the community floor is the sharp divide: unsaved pairs above it are reported close 37.1% of the time, versus 18.2% below it—half the rate even of pairs with no fingerprint data at all (34.9%, whose composition mirrors the population). Second, a *one-sided* save adds nothing beyond community membership (36.8% vs. 37.1%): only *reciprocal* saving marks deeper relationships (61.1%). The corrected-overlap floor, not the save indicator, is what separates genuine relationships from incidental transaction partners; reciprocity then marks depth within the community network, which is why the mutual layer is analyzed separately (Appendix Table B5).

Table B3: Overlap, Reciprocity, and Closeness Across Distance

| | (1) Close | (2) Close | (3) Close | (4) Close |
|---|---|---|---|---|
| Corrected overlap share | | $+0.555^{***}$ | $+0.221^{***}$ | $+0.083^{**}$ |
| | | (0.044) | (0.040) | (0.042) |
| Distance (100 km) | | $+0.0006$ | $+0.0009$ | |
| | | (0.0008) | (0.0010) | |
| Overlap × Distance | | | $-0.0052$ | |
| | | | (0.0049) | |
| Far (modal city ≠ Dakar) | | | | $+0.112^{***}$ |
| | | | | (0.011) |
| Overlap × Far | | | | $+0.391^{***}$ |
| | | | | (0.082) |
| Mutual (both saved) | $+0.331^{***}$ | | $+0.310^{***}$ | $+0.332^{***}$ |
| | (0.008) | | (0.009) | (0.011) |
| Mutual × Distance | | | $+0.0004$ | |
| | | | (0.0016) | |
| Mutual × Far | | | | $-0.073^{***}$ |
| | | | | (0.019) |
| $R^2$ | 0.109 | 0.025 | 0.112 | 0.126 |
| $N$ | 13,355 | 12,991 | 12,991 | 13,355 |

*Notes:* OLS regressions of survey-reported closeness (contact is wife, family, or friend) on the bias-corrected contact-book overlap share, reciprocal-save (mutual) status, geographic distance, and their interactions, for all surveyed (driver, contact) pairs with fingerprint data on both sides—overlap is computed directly from the fingerprints, so the sample is not restricted to pairs appearing in the network edge lists. Locations are transaction-weighted mean agent coordinates from the 2020–2024 individual–agent panels, supplemented by merchant-registry coordinates; distance is the haversine distance between the pair's locations (100 km units); column (4) uses an indicator for the contact's modal city lying outside Dakar. HC1 standard errors in parentheses. Three patterns. First, reciprocity marks relationship depth: mutual pairs are reported close far more often both unconditionally (column 1; 62.8% vs. 29.7%) and conditional on overlap (columns 3–4), and the premium persists at distance ($0.332 - 0.073 = 0.259$, column 4). Second, overlap predicts closeness at all distances, and is *more* informative for distant pairs (far slope $0.083 + 0.391 = 0.474$): nearby, overlap partly reflects incidental shared environment, while at distance only genuine relationships sustain it—licensing the interpretation of cross-city community ties as real relationships. Third, the continuous distance level effect is a precise zero; the positive far indicator reflects the kin composition of long-distance transaction partners (wives and family are over-represented among distant contacts). $^{*}p < 0.10$, $^{**} < 0.05$, $^{***}p < 0.01$.

Table B4: Phone Network: Balance and Tests for Treatment-Induced Rewiring

| | Treated | Control | Difference | SE / $t$ | $p$-value |
|---|---|---|---|---|---|
| *Panel A: Raw network levels (fingerprinted drivers, $N = 1{,}907$)* | | | | | |
| Contact book size (fingerprint) | 585 | 614 | $-28.8$ | 24.5 | 0.240 |
| Mutual contacts detected | 8,309 | 8,982 | $-673$ | 352 | 0.056* |
| 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 |
| Mean overlap (among mutual) | 0.0039 | 0.0038 | 0.0000 | 0.0001 | 0.727 |
| *Panel B: Community-tie network (corrected overlap $\geq 2\%$, pre-trim; $N = 1{,}907$)* | | | | | |
| Community ties per driver ($\times$1,000) | 295.3 | 305.0 | $-9.8$ | 14.0 | 0.486 |
| Strong ties per driver ($\times$1,000) | 46.6 | 46.9 | $-0.4$ | 4.7 | 0.935 |
| Mean corrected overlap (driver-level) | 0.0337 | 0.0335 | 0.0002 | 0.0005 | 0.715 |
| Merchant community ties ($\times$1,000) | 6.50 | 6.57 | $-0.07$ | 0.35 | 0.839 |
| Taxi community ties | 717 | 742 | $-25$ | 35 | 0.469 |
| Joint $F$-test (Panel B) | $F(5, 1901) = 0.615$ | | | $p = 0.688$ | |
| *Panel C: Tie-strengthening tests (full matched sample, $N_T = 1{,}256$, $N_C = 894$)* | | | | | |
| Strong ties per driver ($\geq 5\%$ raw) | 668 | 749 | $-81$ | $-0.51$ | 0.608 |
| Mean overlap among strong ties | 0.0738 | 0.0738 | $-0.0001$ | $-0.28$ | 0.781 |
| Share taxi among strong ties | 0.070% | 0.060% | $+0.010\%$ | $+1.29$ | 0.196 |
| Taxi-contact overlap distribution | KS $= 0.015$ | | | $p = 0.14$ | |
| Pooled corrected overlap (community net) | 0.0389 vs. 0.0397, treated *lower* | | | | |

*Notes:* Fingerprints observed February 2024; the network is used as a proxy for pre-existing social proximity (Section 1.3). Panels A–B report treated and control means, the difference, robust SE, and $p$-value from a regression on treatment status. Panel A uses raw (uncorrected) quantities, inflated by hash collisions. Panel B reports the corrected community-tie network used in the analysis (strong $=\geq 5\%$), before the symmetric top-2% hub trim—hence the large counts; arms are jointly indistinguishable ($p = 0.688$). Panel C tests for treatment-induced tie strengthening on the full matched sample (Welch $t$-tests; level differences from Panel A reflect the sample, not new findings). Because selective saving of technology-related contacts would fatten the upper tail of the tie-strength distribution without necessarily moving its mean, the Kolmogorov–Smirnov test compares the *full* overlap distribution of treated and control drivers' taxi contacts—the margin where treatment-induced saving would appear—and detects no shift ($p = 0.14$); pooled tie strength is, if anything, *lower* for treated drivers, the wrong direction for a strengthening story. Merchant and taxi classifications use only pre-fingerprint registrations. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table B5: Reciprocity and the Cross-Market Channel: Mutual Phone Contacts

| | $\geq 2\%$ | $\geq 5\%$ | $\geq 10\%$ |
|---|---|---|---|
| Cross-industry, Dakar | $+0.96$ [.558] {+2.9%} | $+0.87$ [.701] {+2.6%} | $+0.88$ [.790] {+2.9%} |
| Cross-industry, outside | $-0.22$ [.432] {$-10\%$} | $-0.45$ [.247] {$-23\%$} | $-0.57$ [.214] {$-33\%$} |
| Within-taxi, Dakar | $+0.74$ [.656] {+2.7%} | $+0.07$ [.990] {+0.1%} | $-0.32$ [.948] {$-0.4\%$} |
| Within-taxi, Thiès | $+\mathbf{1.10}^{***}$ [.003] {+26%} | $+\mathbf{2.84}^{***}$ [.001] {+35%} | $+\mathbf{5.16}^{***}$ [.001] {+43%} |
| $N$ | 188,948 | 73,393 | 31,880 |

*Notes:* Coefficients $\times 1{,}000$; RI $p$-values (1,000 permutations) in brackets; effect sizes in braces, computed as $\beta_1/\hat{y}(S{=}0)$ with $\hat{y}(S{=}0) = \bar{y} - \hat{\beta}_1\bar{S}$, predicted adoption at zero treated exposure (controls at their means). Each column re-estimates the specification of Table 1 with exposure computed over *mutual* phone contacts only—pairs in which each party saved the other's number—at the indicated corrected-overlap floor. The mutual layer is estimated *untrimmed*: because mutual status requires explicit reciprocal saves, the hub and collision contamination that the top-2% trimming targets is largely absent. Results are insensitive to this choice: with the strict trim reapplied, the within-taxi Thiès estimates are $+0.90$ to $+4.39$ ($p_{\text{RI}} \leq .006$), and on the full mutual layer with no overlap floor (1.30 million contacts, strict trim) the effect remains $+18\%$ of the base rate ($p_{\text{RI}} = .033$). Two patterns stand out. First, cross-industry: mutual contacts—despite baseline adoption rates 2.5–10$\times$ higher than the community network's—show no incremental response to treatment at any threshold. The cross-industry diffusion of Table 1 is therefore carried by non-mutual community ties, consistent with mutual contacts' exposure to the technology being saturated irrespective of the experiment (inframarginality). Second, within industry: adoption in the Thiès market (the only market outside Dakar where the taxi product was rolled out; 98% of outside-Dakar taxi adopters) responds through mutual ties at every threshold: profession-specific adoption crosses markets through reciprocal relationships, consistent with experiential, trust-intensive content requiring high-bandwidth ties, while awareness—abundant in a launched market—does not bind. The same channel appears in the pre-treatment transaction network (Table B11). $^{*}p_{\text{RI}} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$.

Table B6: Within-Taxi Network Spillovers by Tie Strength

| | $\geq 2\%$ (community) | $\geq 5\%$ (strong) | $\geq 10\%$ |
|---|---|---|---|
| **Within-taxi, Dakar** | | | |
| Share treated | $+0.08$ {+14%} | $-0.98^{***}$ {$-30\%$} | $-0.56$ {$-10\%$} |
| | (0.09) | (0.12) | (0.28) |
| | [0.674] | [0.007] | [0.542] |
| Expected exposure | $+4.09$ | $-1.09$ | $-0.32$ |
| | (2.79) | (3.71) | (9.41) |
| | [0.143] | [0.770] | [0.973] |
| Baseline $\hat{y}(S{=}0)$ | 0.54 | 3.28 | 5.61 |
| $N$ | 1,145,899 | 1,020,391 | 376,871 |

*Notes:* Same specification, samples, community-tie definitions, trimming, and baseline convention ($\hat{y}(S{=}0) = \bar{y} - \hat{\beta}_1 \bar{S}$) as Table 1; the outcome is registration as a *taxi* merchant in Dakar. The within-industry pattern is more nuanced than the cross-industry one. At the community threshold ($\geq 2\%$) the Dakar effect is a statistical null ($+0.08$, $p_{\text{RI}} = 0.674$; the $+14\%$ effect size reflects the very low baseline, $\hat{y}(S{=}0) = 0.54$ per 1,000, not a detectable effect); at the strong-tie threshold the effect is *negative* and significant ($-0.98$, $-30\%$, $p_{\text{RI}} = 0.007$); at $\geq 10\%$ the point estimate remains negative but loses precision as the sample shrinks and closest ties (family, intimates) dilute the colleague share. Appendix Figure A8 shows the band decomposition: the negative effect concentrates at 5–8% overlap, and within degree strata it is driven by contacts with one to five strong driver ties. In Thiès (the only market outside Dakar where the taxi product was rolled out) the community-layer effects are null throughout (not shown); the positive cross-market within-taxi channel operates through the reciprocal-save (mutual) layer and the transaction network (Appendix Tables B5 and B11), which this table's community-tie exposure does not isolate. $^{*}p_{\text{RI}} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$.

Table B7: Cross-Industry Spillovers Excluding Shared-Customer Pairs and Garage-Based Drivers

| | $\geq 2\%$ (community) | $\geq 5\%$ (strong) | $\geq 10\%$ |
|---|---|---|---|
| **Cross-industry, Dakar** | | | |
| Share treated | $+1.76$ {+23%} | $+3.63^{**}$ {+33%} | $+4.59^{**}$ {+59%} |
| | (0.47) | (0.30) | (0.43) |
| | [0.333] | [0.037] | [0.045] |
| Expected exposure | $+13.01$ | $-9.10$ | $-47.64$ |
| | (15.00) | (10.06) | (14.24) |
| | [0.386] | [0.366] | [0.001] |
| **Cross-industry, outside Dakar** | | | |
| Share treated | $+0.48^{*}$ {+78%} | $+0.61^{**}$ {+69%} | $+0.75^{**}$ {+134%} |
| | (0.15) | (0.09) | (0.13) |
| | [0.057] | [0.027] | [0.044] |
| Expected exposure | $+8.62$ | $-2.63$ | $-2.00$ |
| | (5.06) | (3.20) | (4.34) |
| | [0.089] | [0.410] | [0.645] |
| Dep. var. mean (Dkr / outside) | 8.79 / 0.90 | 13.41 / 1.29 | 10.64 / 1.02 |
| $N$ | 1,142,867 | 944,780 | 347,063 |

*Notes:* Specification, sample construction, and cell format exactly as Table 1 (all estimates $\times 1{,}000$; effect sizes in braces; HC1 SEs in parentheses; RI $p$-values in brackets for Share treated, HC1 for Expected exposure), after two purges applied to the trimmed Table 1 universe: (i) every driver–contact pair sharing a regular customer ($\geq 2$ transactions with each side), removing the demand channel; and (ii) every pair of the 285 garage-based drivers (self-reported garage-based in the endline survey or matched to a GPS garage roster), where drivers share a stream of passengers. The cross-industry learning effects survive with magnitudes close to Table 1 at the strong-tie thresholds ($+4.16 \rightarrow +3.63$ in Dakar and $+0.68 \rightarrow +0.61$ outside at $\geq 5\%$; $+5.16 \rightarrow +4.59$ and $+0.73 \rightarrow +0.75$ at $\geq 10\%$), indicating that cross-industry diffusion operates through social learning rather than shared demand. Under the same purge, the within-taxi Dakar coefficient at $\geq 5\%$ remains negative and significant ($-0.88$, $p_{\text{RI}} = 0.025$). $^{*}p_{\text{RI}} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$.

Table B8: Lee Bounds for Differential Fingerprint Attrition

| | Taxi, Dakar (within) | Cross-ind., Dakar (across, local) | Cross-ind., outside (across, distant) |
|---|---|---|---|
| *Community ties ($\geq 2\%$ corrected overlap)* | | | |
| Share treated | $+0.075$ | $+2.907$ | $\mathbf{+0.680}^{**}$ |
| Lee lower bound | $+0.027$ | $+1.994$ | $+0.554^{**}$ |
| Lee upper bound | $+0.079$ | $+2.996$ | $+0.686^{***}$ |
| $N$ | | 1,145,899 | |
| *Strong ties ($\geq 5\%$ corrected overlap)* | | | |
| Share treated | $\mathbf{-0.984}^{***}$ | $\mathbf{+4.162}^{**}$ | $\mathbf{+0.681}^{***}$ |
| Lee lower bound | $-1.103^{***}$ | $+4.079^{**}$ | $+0.680^{***}$ |
| Lee upper bound | $-0.976^{***}$ | $+4.198^{**}$ | $+0.685^{***}$ |
| $N$ | | 1,020,391 | |

*Notes:* Tightened Lee bounds (Lee, 2009) for differential attrition from the experimental sample to the fingerprint network sample, computed on the community-tie network (bias-corrected overlap; same universe, first-row dedupe, and strict top-2% trimming as Table 1—the Share treated rows replicate Table 1 exactly). Bounds are computed within contact-book-size quartiles: within each quartile, excess treated drivers are trimmed from the top (lower bound) or bottom (upper bound) of the connected-outcome distribution, and the network and estimates are rebuilt. The bounds are tight throughout—at the strong-tie threshold every bounded estimate retains its sign and significance, including the within-taxi effect ($[-1.103, -0.976]$, both $p_{\text{RI}} < 0.01$)—indicating that differential fingerprint attrition does not drive the results. Coefficients $\times 1{,}000$. RI $p$-values from 1,000 permutations. $^{*}p_{\text{RI}} < 0.10$, $^{**}p_{\text{RI}} < 0.05$, $^{***}p_{\text{RI}} < 0.01$.

Table B9: 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 From 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 | 1,693 |

*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 answer several survey questions about each of these links. $N = 1,693$ is the number of driver respondents; each observation is the driver-level share across their 16 links. Panel B statistics are computed over the 1,605 respondents with frequency answers; the advice/information question in Panel C was added after the survey began and is available for 1,371 respondents (shares among known contacts in both panels). The 27,120 survey pairs referenced in Section 1.3 come from the drivers who completed the network validation module.

Table B10: 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 From 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 | 1,693 |

*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 answer several survey questions about each of these links. $N = 1,693$ is the number of driver respondents; each observation is the driver-level share across their 16 links. Panel B statistics are computed over the 1,605 respondents with frequency answers; the advice/information question in Panel C was added after the survey began and is available for 1,371 respondents (shares among known contacts in both panels). The 23,520 survey pairs referenced in Section 1.3 come from the drivers who completed the network validation module.

Table B11: Spillovers in the Pre-Experiment Transaction Network ($\geq$ 2 P2P transfers)

| | Taxi, Dakar | Cross-ind., Dakar | Cross-ind., other | Taxi, Thiès |
|---|---|---|---|---|
| Share treated | $+1.53$ | $+0.33$ | $+0.14$ | $+\mathbf{11.19}^{***}$ |
| $p_{\text{RI}}$ | [0.605] | [0.936] | [0.845] | [0.001] |
| Effect size | $+4.8\%$ | $+0.6\%$ | $+3.4\%$ | $+83\%$ |
| Dep. var. mean | 0.0319 | 0.0542 | 0.0041 | 0.0134 |
| Adopters | 739 | 1,254 | 95 | 311 |
| $N$ | | 23,146 | | |

*Notes:* Coefficients $\times 1{,}000$. The network is constructed from pre-treatment mobile-money transfer records (2020–2022Q1): an individual $j$ is linked to an experimental driver if the driver *sent* at least two P2P transfers to $j$ before the experiment (repeated-transfer ties). Specification and trimming as in Table 1. Money-sending ties carry within-industry adoption into the Thiès market ($+83\%$ of baseline, $p_{\text{RI}} = 0.001$)—mirroring the mutual-contact channel of Appendix Table B5—and carry no significant local effects. The pattern is identical with a $\geq 1$-transfer definition ($N = 60{,}381$; Thiès $+5.23$, $+60\%$, $p_{\text{RI}} = 0.001$; all other margins null under randomization inference). RI $p$-values (1,000 permutations) in brackets. $^{*}p_{\text{RI}} < 0.10$, $^{**} < 0.05$, $^{***} < 0.01$.

# C Fingerprint Data Processing

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

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

**Step 2: Matching experimental drivers.** Each of the 2,196 experimental drivers has up to three registered phone numbers. I search the fingerprint file for any number matching any driver, keeping the fingerprint with the most buckets when a driver has multiple records (best-phone-per-person). Of the 2,150 drivers with a mobile money account match (`has_match`), 1,912 (88.9%) have at least one phone number with fingerprint data (Section 1.3). The remaining 238 drivers were never fingerprinted by the application.

**Step 3: Minimum contact threshold.** Drivers with fewer than 10 unique hash buckets in their fingerprint are excluded, as the overlap computation is dominated by hash collision noise at such small set sizes. This removes 5 additional drivers, leaving 1,907 in the network analysis sample.

**Step 4: Candidate pairs and overlap measurement.** For each fingerprinted user ("individual $j$") and each driver $i$, the raw bucket intersection $I_{ij} = |F_i \cap F_j|$ is computed. At 16-bit resolution, part of any observed intersection is mechanical: two independent contact books of sizes $|F_i|$ and $|F_j|$ collide in

$$\mu_{ij} = \frac{|F_i|\,|F_j|}{65{,}536}$$

buckets in expectation. The analysis therefore measures ties by the *bias-corrected overlap share*

$$\text{OverlapShare}_{ij}^{\text{corr}} = \frac{\max(I_{ij} - \mu_{ij},\, 0)}{\min(|F_i|, |F_j|)},$$

which subtracts the expected collision count before normalizing by the smaller contact book. For a truly unrelated pair, the corrected share is centered on zero with noise of order $1/\min(|F_i|, |F_j|)$; a 2 percent floor therefore removes nearly all purely accidental overlap while retaining genuine shared-environment pairs. The 2 percent *community-tie* floor is set an order of magnitude above this noise scale and is validated against survey-reported relationships (Section 1.3; Appendix Tables B1 and B2). Because the share is normalized by the smaller contact book, the resulting graph is deliberately broader than a contact book—it measures shared social environment—and is subsequently reduced by strict hub trimming (Step 7).

**Step 5: Save indicators as tie attributes.** Directed *save* indicators are computed separately: driver $i$'s self-hash bucket appearing in individual $j$'s fingerprint ("$j$ saved $i$"), and vice versa. The conjunction defines a *mutual contact*. These indicators do *not* define the network: the analysis network is the community-tie graph (corrected overlap $\geq 2\%$), and the save indicators are retained only as tie *attributes*—to identify the small reciprocal-save (mutual) layer, which plays a distinct role (carrying within-industry adoption across markets; Appendix Table B5), and to know the direction from which a link originates. Because a single self-hash bucket match can arise from collision, save indicators for any individual pair are noisy; they are informative in aggregate (mutual pairs are reported as substantially closer in survey data; Table B2) but are never used as a stand-alone existence criterion for a tie.

**Step 6: Pair deduplication.** Individuals can appear under multiple phone numbers. Driver–contact pairs are deduplicated keeping the first record per pair (best-phone convention) *before* any threshold is applied, so that thresholding operates on a unique-pair edge list.

**Step 7: Degree trimming.** Contacts in the top 2% of the degree distribution are dropped, and every contact connected to any top-2% degree driver is then excluded entirely (strict trimming). This removes hub nodes whose high connectivity may reflect institutional roles (e.g., agents, dispatchers) rather than personal social ties, and is applied in all main specifications.

**Final analysis sample.** After all processing steps, the community-tie network (corrected overlap $\geq 2\%$) links the experimental drivers to 1,145,899 individuals through 23.7 million

edges; 1,020,391 individuals have at least one strong tie ($\geq 5\%$), and about 164,000 are additionally connected through the reciprocal-save (mutual) layer.