Strategic Complementarities in a Dynamic Model of Technology Adoption: P2P Digital Payments*

Fernando Alvarez
University of Chicago

David Argente
Yale University

Francesco Lippi
LUISS and EIEF

Esteban Méndez
Central Bank of Costa Rica

Diana Van Patten
Yale University

December 2025

Abstract

We develop a dynamic model of technology adoption featuring strategic complementarities: the benefits of the technology increase with the number of adopters. We show that complementarities give rise to gradual adoption, multiple equilibria, multiple steady states, and suboptimal allocations. We study the planner's problem and its implementation through adoption subsidies. We apply the theory to SINPE Móvil, a peer-to-peer payment app developed by the Central Bank of Costa Rica. Using transaction-level data and user-specific networks that we construct from administrative records, we causally estimate sizable complementarities. In our calibrated model, the optimal subsidy pushes the economy to universal adoption.

JEL Classification Numbers: E4, E5, O1, O2
Keywords: Technology Adoption, Strategic Complementarities, Means of Payment

*We thank Isaac Baley, Andres Blanco, Diego Comin, Nicolas Crouzet, Joe Hazell, Rishabh Kirpalani, John Leahy, Claudio Michelacci, Ben Moll, Guillermo Noguera, Ezra Oberfield, Franck Portier, Rob Shimer, Nancy Stokey, Mathieu Taschereau-Dumouchel, Daniel Xu, Yu Zhu, and seminar participants at EIEF, Oxford, Minnesota, Chicago Booth, NYU, Princeton, Dartmouth, Columbia, Richmond Fed, Philadelphia Fed, UC Berkeley, LSE, UCL, Duke-NCSU-UNC, Rochester, Toronto, Northwestern, IIES, Chicago Fed, UCSB LAEF, Central Bank of Costa Rica, Bank of Mexico, Fed Board, the Cowles Trade Summer Conference, the BSE Summer Forum, the Hydra conference, the Bank of Canada, the Bank of Portugal, the Rome Junior Conference (Pizzanomics), the SED Wisconsin, NBER SI, NBER EFG, and Warwick-CFM-Vienna Global Macro, the Cowles Macro Summer Conference. This project received STEG grant No. 608. Lippi acknowledges financial support from the ERC grant: 101054421-DCS. The views expressed herein are those of the authors and do not necessarily represent the views of the Central Bank of Costa Rica (BCCR). All results have been reviewed by the BCCR to ensure that no confidential information is disclosed.

1 Introduction

We study the diffusion of a new technology in an economy characterized by strategic complementarities. These complementarities occur because the benefits that agents derive from the technology increase with the number of users –a phenomenon long recognized in the applied literature on technology diffusion (see Griliches (1957); Mansfield (1961)). Progress in this research area is hindered by the challenges that arise when modeling adoption dynamically –a large state space, non-linear decisions, multiple equilibria–, and by the lack of detailed data on technology diffusion. We develop a new tractable model of technology adoption and apply it to the diffusion of a new payment app, SINPE Móvil, a digital application created by the Central Bank of Costa Rica that allows instantaneous P2P transfers between bank account holders in the network.[1] By its nature, the usefulness on this app depends on others joining the network. We aim to quantify the value of this complementarity using granular data from SINPE and other sources. We use the model to discuss equilibrium existence, multiplicity of equilibrium paths, multiplicity of stationary equilibria, and the local stability of stationary equilibria (see e.g., Matsuyama (1991)). We characterize the planner's problem and its implementation through subsidies, and use a calibrated version of the model to analyze the optimal policy.

The model assumes that the benefits of the technology at time t depend on the number of agents who have adopted it, $N(t)$, and on an idiosyncratic persistent random component, $x(t)$. In particular, we assume that the flow benefit of the app is proportional to the product between these variables, $x(t)N(t)$, so that an agent is more likely to adopt if her private needs for it are high (a high x) and/or when more agents use the app (higher N). A single parameter, controlling the intensity of this interaction effect, measures the strength of the strategic complementarities. A high value of x also implies that an agent will use the technology more intensively, a feature that we leverage when calibrating the model to the data where we observe both adoption as well as the intensity of usage. Adoption entails a fixed (once and for all) cost and agents choose when to adopt taking the aggregate path of adoption as given. We show that when the idiosyncratic benefits are random the equilibrium features gradual adoption through a simple mechanism: agents wait for others to adopt. The optimal adoption rule is given by a time-dependent threshold value, denoted by $\bar{x}(t)$, such that adoption is optimal if $x(t) > \bar{x}(t)$. We assume that the economy starts with an (arbitrary) measure of agents endowed with the technology, which serves as the initial

[1] Although the app is called "SINPE Móvil," throughout we will be referring to it only as "SINPE," which stands for Costa Rica's National Electronic Payment System (by its initials in Spanish). The app was launched in May 2015 and over 60% of the adult population used it in 2021, with about 10% of the country's GDP transacted via SINPE. See Björkegren (2018) for a related network-goods analysis using data on mobile phones adoption in Rwanda.

condition of the equilibrium. Aggregation of the adoption decisions across agents yields a path for the fraction of agents that use the technology at each time t, $N(t)$. Given the initial and terminal conditions, the equilibrium has a classic fixed point structure: the optimal decision path ($\bar{x}$) depends on the aggregate path (N), and vice-versa.

The model yields three main results, each summarized by a theorem. We show that the optimal adoption rule for each agent, summarized by the threshold path $\bar{x}$, is a decreasing functional of the path of adoption N. The strength of this effect depends on the parameter that controls the strategic complementarity. Likewise, we show that the adoption path N is a decreasing functional of the path $\bar{x}$, for any initial distribution of adopters. An equilibrium is a fixed point given by the composition of these two functionals. The first theorem establishes the existence, and possibly the multiplicity, of dynamic equilibria. These equilibria form a non empty lattice, i.e., they are ordered so that there is a "largest one," N^H, and a "smallest" one, N^L. The adoption path of the largest equilibrium is above the smallest one at every point in time, $N^H(t) \geq N^L(t)$, for all t. More equilibria may exist and are bracketed between these ones (the paths of different equilibria do not cross). We establish these results using the monotone comparative statics logic by Milgrom and Shannon (1994), and Tarski's fixed point theorem. We show that there is a critical mass of adopters $\underline{N}_0$ such that, if the initial measure of adopters is below $\underline{N}_0$, there exists an equilibrium where no one will adopt eventually. We also study stationary equilibria, i.e., equilibria where N is constant through time, and show that, besides the stationary equilibrium with no adoption, the model may feature two additional interior stationary equilibria, with low- and high-adoption.

We find two types of multiple dynamic equilibria for a fixed initial condition. A first kind, which we denote as "delayed adoption", is a family of equilibria where the path of the endogenous variables is time-shifted. These equilibria differ in the length of the initial period with no adoption. The second kind is one where the equilibrium either converges to the high adoption or to the no adoption. We discuss the initial conditions and parameter conditions under which each case occurs.

The second theorem characterizes the stability of the stationary equilibria by means of a perturbation analysis with respect to the initial condition, assumed to be one of the two interior equilibria. The analysis is non-trivial because it involves the linearization of an infinite dimensional system: the distribution of adopters. We handle the problem by leveraging techniques from the Mean Field Game (MFG) literature (e.g., Alvarez et al., 2023a; Auclert et al., 2022; Bilal, 2023), which determines the local stability by inspecting the eigenvalues of a linear operator. One novelty compared to the MFG problem studied in Alvarez et al. (2023a) is the possibility of *multiple* stationary equilibria. The stability condition then depends on the particular equilibrium that is chosen. We find that the high-adoption equilibrium is locally stable, while the low-adoption is unstable, a feature that leads us to discard it from the analysis.

Equilibria are socially inefficient because agents do not internalize the fact that when they adopt they benefit all agents who already have the technology. We show how to characterize efficient allocations by solving a planner's problem which takes into account the dynamics across the entire network. The third theorem shows how to decentralize the planner's solution using a simple tool: a time-varying subsidy paid to those that use the technology.

We then leverage comprehensive data collected since SINPE was created to analyze the dynamics of adoption and usage, to document the presence of strategic complementarities, and to discipline the parametrization of the model. Our baseline analysis links data on users—both receivers and senders—within their employer-employee network.[2] We identify the presence of strategic complementarities using arguably exogenous variation in the network size due to mass layoffs. We document a causal relation between the share of agents who have adopted (N) and usage of the app, both at the extensive margin as well as at the intensive margin: a sudden decrease of the network size lowers the probability of adoption and lowers the intensity of use.[3] This effect persists across a battery of ways to define usage and networks. It also emerges after using a leave-one-out instrument and following a balanced panel of adopters to address concerns regarding selection.

We match the theory with the data in a quantitative analysis where we calibrate the model using key moments from the data with the objective to compute the optimal adoption subsidy. To capture the initial gradual diffusion of the technology, observed in each network, we supplement the model with a layer of slow-information diffusion following the seminal work of Bass (1969). The strength of the strategic complementarities is chosen using the information retrieved from the mass layoffs described above. The calibrated model shows that the optimal subsidy speeds up adoption by the agents and ultimately pushes the economy towards universal adoption of the payment app.

Related Literature. Several recent studies are related to our paper. Benhabib et al. (2021) model firms that can endogenously innovate and adopt a technology. They analyze the effect of these choices on productivity and balanced growth, but without conducting an analysis of the transition between stationary distributions; likewise, Buera et al. (2021) study

[2]Individual-to-individual transactions account for over 95% of all transactions, regardless of the time period considered. We find that 44% of all SINPE transactions occur between coworkers. Family networks and spatial "neighborhood" networks are also considered for robustness.

[3]Namely, we focus on networks of coworkers and examine the effect of network changes on the intensity of the app's usage and its adoption for workers displaced by a mass layoff. By analyzing the usage intensity of workers who had already adopted the app prior to being displaced, we are able to isolate the influence of strategic complementarities rather than the effects of learning.

policies that can coordinate technology adoption across firms. A closely related contribution is Crouzet et al. (2023), who develop a model with a unique equilibrium where the rate of adoption of electronic payment by *retailers* increases following an aggregate shock. Their analysis is motivated by 2016 Indian Demonetization, and exploits the variation in the intensity with which *firms* in Indian districts were exposed to the shock to examine the adoption of retailers. Unlike our model, which has heterogeneous agents and generates dynamics and gradual adoption *endogenously* (as agents wait for others to adopt before doing so), their model features homogeneous agents and a sluggish adjustment à la Calvo (1983), generating gradual adoption through this imposed friction. Moreover, the heterogeneity in our model allows us to accommodate, not only aggregate shocks when we analyze transition dynamics in closed-form, but also dynamics after shocks that target particular types of agents; for instance, we compare the propagation after "giving the app" to people with high vs. low idiosyncratic benefits, which in turn can be mapped to observables like wages and skills.

The paper also deals with technical issues of multiplicity and stability that have plagued the economic geography and trade literatures. Recent papers have developed algorithms that exploit the super- or sub-modularity of the objective function based on Tarski's theorem (Jia, 2008; Arkolakis et al., 2023; Alfaro-Urena et al., 2023). Our approach also leverages the monotonicity of our problem, but does so for an analysis of dynamic stability as a criterion to select an equilibrium and develops the planning problem to study efficiency.

The paper is organized as follows. The next section presents the model, Section 3 discusses different types of equilibria that may arise. Section 4 uses a perturbation method to inspect the stability of the stationary equilibria. Section 5 discusses the planning problem. Section 6 adds an information diffusion mechanism to the baseline model. Section 7 presents the data and documents the non-negligible role of strategic complementarities in the adoption and use of SINPE. A calibrated version of the model is used in Section 8 to discuss the optimal subsidy for the efficient adoption of SINPE.

2 The Model

This section presents a tractable model of technology adoption within a "network" of agents. The model fits alternative notions of network, later discussed in the empirical analysis, such as a group of co-workers, households living in the same neighborhood, or a (broad) notion of family members. The network is populated by a continuum of agents who differ in the potential benefits from adopting the technology. Let $N(t) \in [0, 1]$ be the fraction of agents who have adopted at time $t \in [0, T]$. The flow benefit at time t for an agent who has already adopted the technology is

$$x(\theta_0 + \theta_n N(t)) \tag{1}$$

where $\theta_0, \theta_n > 0$ are parameters, x is a stochastic process, independent across agents, with variance σ^2 per unit of time, no drift, and reflecting barriers at $x = 0$ and $x = U$, so that $dx = \sigma dW$ where W is a standardized Brownian motion. Later we provide a derivation of this equation, as the indirect utility benefit arising from the optimal intensity of technology use in each period, see equation (29). We let $c > 0$ be the fixed cost of adopting the technology and $r > 0$ be the time discount rate. With probability ν per unit of time agents die, so that agents discount time at rate $\rho \equiv r + \nu$. Dead agents are replaced by newborns without the technology and an x drawn from the invariant density $f(x) = 1/U$ for $x \in [0, U]$, where f is uniform because of the reflecting barriers.

2.1 Individual Decisions, Aggregation, Equilibrium

We next describe the agent's optimal decision as a function of the whole path of aggregate adoption $N : [0, T] \to [0, 1]$, discuss how to aggregate individual decision to compute the aggregate path of adoption, and define the equilibrium.

Let $a(x, t)$ be the value function of an agent who has adopted the technology and has state x at time t:

$$a(x,t) = \mathbb{E}\Big[\int_t^T e^{-\rho(s-t)}\left(\theta_0 + \theta_n N(s)\right) x(s) ds + e^{-\rho(T-t)} a_T(x) \;\Big|\; x(t) = x\Big] \tag{2}$$

for all $t \geq 0$ and $x \in [0, U]$. The agent takes the whole path N as given. For finite T we assume that $a_T(x) = (\theta_0 + \theta_n \bar{n})\,\mathbb{E}\Big[\int_T^\infty e^{-\rho(s-T)} x(s) ds \;\Big|\; x(T) = x\Big]$, where $\bar{n} \in [0, 1]$. The interpretation of a_T is the value of adopting when the fraction of adopters is a given constant $\bar{n}$.

An agent with state x, who has not yet adopted at time t, has a value function $v(x, t)$ that solves the stopping-time problem

$$v(x,t) = \max_{t \leq \tau} \mathbb{E}\left[e^{-\rho(\tau - t)}\left(a\left(x\left(\tau\right), \tau\right) - c\right) \Big| x(t) = x\right], \tag{3}$$

where τ denotes the time of the adoption and depends only on the information generated by the process for x and on calendar time t (the latter because of the dynamics of $N(t)$).

We will use the convention that for $T = \infty$ then the set $[0, \infty]$ over which the functions of interest are defined shall be interpreted as $[0, \infty)$.

Discretized Model. We consider a discretized version of the model where time is discrete, at intervals of length Δ_t, and the state x is discrete in intervals of length Δ_x. The reflecting Brownian Motion, Poisson processes, and discounting are changed accordingly, following the scheme used in finite difference approximations, see Definition 3 in Appendix B for a detailed definition. For small Δ_t, Δ_x, the discretized model converges to the decision problem in continuous time. The advantage of the discretized model is that we can compute numerical solutions for the equilibrium path for the case of T finite. Instead, the advantage of the continuous time model is that it is easier to characterize stationary solutions as well as perturbations.

Next we state a preliminary result to establish that we can represent the optimal adoption rule at time t as a threshold rule, $\bar{x}(t)$.

LEMMA 1. Fix a path N and a time $t \in [0, T]$. If it is optimal to adopt at (x_1, t), then it is also optimal to adopt at (x_2, t) where $x_2 > x_1$. This holds for the continuous time as well as for the discretized model.

For finite T define

$$D_T(x) \equiv a_T(x) - v(x, T)$$

further discussed in Section 2.2. The function $\mathcal{X}$ denotes the path for the optimal threshold as: $\bar{x} = \mathcal{X}(N; D_T)$, so that $\bar{x} : [0, T] \to [0, U]$. It is immediate from equation (2) and equation (3) that what matters for the optimal adoption decision is $D_T(x)$, which is the reason why we include D_T as an argument of $\mathcal{X}$.

Aggregation. Given the individual decision rule we can compute the implied path for the fraction of adopters, N. We start by defining the probability that an agent at s with state $x(s) = x$ survives until time t, while the value of her state remains below $\bar{x}$ during this period:

$$P(x, s, t; \bar{x}) = Pr\Big[x(\iota) \leq \bar{x}(\iota), \text{ for all } \iota \in [s, t] \;\Big|\; x(s) = x\Big]\, e^{-\nu(t-s)}. \tag{4}$$

For an agent who at time s has $x \leq \bar{x}(s)$, the value of $P(x, s, t; \bar{x})$ gives the probability that the agent will survive up to t without adopting. Let $m_0(x)$ be the density of the agents at time $t = 0$ without the technology. Given the assumption about x, we require $0 \leq m_0(x) \leq 1/U$ for all $x \in [0, U]$. The fraction of agents who have adopted the technology at time t is

$$N(t) = 1 - \int_0^U P(x, 0, t; \bar{x}) m_0(x) dx - \int_0^t \nu \left[\int_0^U P(x, s, t; \bar{x}) \frac{1}{U} dx\right] ds. \tag{5}$$

The second term on the right hand side is the fraction of agents who did not have the technology at time 0 and survived until time t without adopting. The third term considers the cohorts of agents that are born between 0 and t, and for each of these cohorts computes the fraction that survived without adopting up to t. We note that an equivalent version of equation (5) holds in the discretized version of the model, which for a given $\bar{x}$ is simply a matrix manipulation. We let $\mathcal{N}(\bar{x}; m_0)$ be the path of N as a function of $\bar{x}$ (the path of the adoption threshold) and of the initial condition m_0.

Equilibrium. The equilibrium is given by the fixed point between the forward looking optimal adoption decision, encoded in $\mathcal{X}$, and the backward looking aggregation, encoded in $\mathcal{N}$. To emphasize the forward looking nature of $\mathcal{X}$, note that it depends on the terminal value function D_T. To emphasize the backward looking nature of $\mathcal{N}$, note that it propagates the initial condition m_0. We then have the following definition.

Definition 1. Fix an initial condition m_0 and a terminal value function D_T. An equilibrium $\{N^*, \bar{x}^*\}$ solves the fixed point:

$$N^* = \mathcal{F}(N^*; m_0, D_T) \text{ where } \mathcal{F}(N; m_0, D_T) \equiv \mathcal{N}(\mathcal{X}(N; D_T); m_0) \tag{6}$$

and where $\bar{x}^* = \mathcal{X}(N^*; D_T)$.

Note that this is a canonical definition of equilibrium, where the operator $\mathcal{F}$ combines the two operators $\mathcal{N}$ and $\mathcal{X}$ defined before. This definition holds for both the continuous time and the discretized version of the model.

2.2 A Recursive Formulation of the Equilibrium

This section derives a recursive representation of the equilibrium that will be useful to study the local stability of the equilibrium and to study the planning problem. To derive the recursive representation we first consider a simple stopping time problem that combines $a(x,t)$ and $v(x,t)$ into a single equation. We consider the value function

$$D(x,t) = \min_{\tau \geq t} \mathbb{E}\Big[\int_t^\tau e^{-\rho(s-t)} \left(\theta_0 + \theta_n N(s)\right) x(s) ds + e^{-\rho(\tau - t)} c \;\Big|\; x(t) = x \Big] \tag{7}$$

with terminal condition $D(x,T) = D_T(x)$. The interpretation is that $D(x,t)$ is the optimal cost of adoption which is made of the flow-opportunity cost until adoption takes place (at τ), plus the actual discounted value of the adoption cost c. The function $D(x,t)$ is related to a and v by $D(x,t) \equiv a(x,t) - v(x,t)$. Note that $a(x,t) - c$ is the net value of adopting immediately while $v(x,t)$ is the net optimal value, that may entail adopting in the future (see equation (2) and equation (3)).

Under differentiability assumptions on $D(x,t)$ we can rewrite equation (7) as a Hamilton-Jacobi-Bellman (HJB) partial differential equation with boundaries, derived in Appendix F, which satisfies:

$$\rho D(x,t) = \min \left\{ \rho c \ , \ x(\theta_0 + \theta_n N(t)) + \frac{\sigma^2}{2} D_{xx}(x,t) + D_t(x,t) \right\} \tag{8}$$

for all $x \in [0,U]$, $t \in [0,T]$ and terminal condition $D(x,T) = D_T(x)$. Optimality requires that $D(x,t) \leq c$, which yields the value matching condition at the barrier. We are looking for a classical solution that satisfies:

$$\rho D(x,t) = x(\theta_0 + \theta_n N(t)) + \frac{\sigma^2}{2} D_{xx}(x,t) + D_t(x,t) \tag{9}$$

for all $x \in [0, \bar{x}(t)]$ and $t \in [0,T]$ with boundary conditions:

$$\begin{aligned} D(\bar{x}(t),t) &= c && \text{Value Matching} \\ D_x(\bar{x}(t),t) &= 0 && \text{Smooth Pasting} \\ D_x(0,t) &= 0 && \text{Reflecting} \end{aligned} \tag{10}$$

If the solution is regular it also features smooth pasting. Finally, since $x = 0$ is a reflecting barrier, the value function has a zero derivative at that point.

Let $m(x,t)$ denote the density of the agents with x that have not adopted at t. The law of motion of m for all $t \geq 0$ is:

$$\begin{aligned} & m_t(x,t) = \nu \left(\frac{1}{U} - m(x,t) \right) + \frac{\sigma^2}{2} m_{xx}(x,t) \text{ if } 0 \leq x \leq \bar{x}(t) \\ & m(x,t) = 0 \quad \text{ for } x \in [\bar{x}(t), U] \\ & m_x(0,t) = 0 \end{aligned} \tag{11}$$

and initial condition $m_0(x) = m(x,0)$ for all $x \in (0,U)$. The p.d.e. is the standard Kolmogorov forward equation (KFE). The density of non-adopters is zero to the right of $\bar{x}(t)$, since this is an exit point. The last boundary condition is obtained from our assumption that x reflects at $x = 0$. The fraction of agents that have adopted the technology is thus given by

$$N(t) = 1 - \int_0^{\bar{x}(t)} m(x,t)dx. \tag{12}$$

We use these equations to provide an equilibrium definition, equivalent to Definition 1, which emphasizes the dynamic nature of the equilibrium.

DEFINITION 2. An equilibrium is given by the functions $\{D, m, \bar{x}, N\}$ satisfying the coupled p.d.e.'s for D and m in (9) and (11), and the boundary conditions in (10), (11), and (12).

We note that this system of p.d.e.'s is involved for two reasons. First, the equations are coupled through $\bar{x}$ and N. Second, the equations feature a time-varying free boundary, which is known to be non-trivial.

3 Equilibria

In this section we establish equilibrium existence, the possibility of multiple equilibria and equilibria with "delayed adoption" (relative to an efficient benchmark). We also characterize equilibria with no adoption, i.e., situations in which given an initial condition m_0, no one will use the technology eventually and conclude by discussing stationary equilibria.

3.1 Monotonicity and Existence of Equilibrium

The next lemma shows that the function $\mathcal{X}$, giving the path of the optimal threshold $\bar{x}$ as a function of the path N, is monotone decreasing. Thus, an agent facing a higher path of adoption will choose to adopt earlier. Moreover, the lemma shows that an agent facing larger values of θ_0 and/or θ_n, will also adopt earlier.

LEMMA 2. If $T < \infty$ let the terminal value function be $D_T(x)$ and $\theta_n \geq 0$. Let $\bar{x}$ be the optimal threshold path for an agent facing the path N. Consider two paths such that $N'(t) \geq N(t)$ for all $t \in [0, T]$, then $\bar{x}'(t) \leq \bar{x}(t)$ for all $t \in [0, T]$. Moreover, let $\theta \equiv (\theta_0, \theta_n)$ with the corresponding optimal threshold path $\bar{x}$. If $\theta' \geq \theta$ then $\bar{x}'(t) \leq \bar{x}(t)$ for all $t \in [0, T]$.

Lemma 2 also holds in the discretized version of the model.[4] The proof holds as we verify the conditions to use Topkis (1978). Thus, once we reformulate the problem in terms of stopping times, we can apply the monotone comparative statics logic developed by Milgrom and Shannon (1994) to characterize the policy function.

Next, we show that given the initial condition $m_0(x)$, if the path $\bar{x}(t) \leq \bar{x}'(t)$ then $N'(t) \leq N(t)$ for all t. We need to show that the fraction of non-adopters is decreasing in $\bar{x}(t)$. This implies that $\mathcal{N}$ is monotone decreasing.

[4]For instance, it holds for a finite difference approximation, which we use for some computations, and which converges to the continuous-time version.

LEMMA 3. Fix m_0 and consider two paths for the thresholds $\bar{x}$ and $\bar{x}'$, satisfying $\bar{x}'(t) \geq \bar{x}(t)$ for all $t \in [0, T]$. Let $N' = \mathcal{N}(\bar{x}'; m_0)$ and $N = \mathcal{N}(\bar{x}; m_0)$. Then $N'(t) \leq N(t)$ for all $t \in [0, T]$. Moreover, fix a threshold path $\bar{x}$, and consider two initial measures with $m_0'(x) \geq m_0(x)$ for all $x \in [0, U]$, then $N' = \mathcal{N}(\bar{x}; m_0')$ and $N = \mathcal{N}(\bar{x}; m_0)$. Then $N'(t) \leq N(t)$ for all $t \in [0, T]$.

The next theorem uses the monotonicity of $\mathcal{X}$ and $\mathcal{N}$, proven in Lemma 2 and Lemma 3, to establish through equation (6) that $\mathcal{F}$ is monotone. This allows us to use Tarski's theorem and establish the existence, and possibly the multiplicity, of equilibria.

THEOREM 1. Consider either the discretized model or its continuous time limit, let T be the terminal horizon and $\theta_n \geq 0$. Fix an initial condition m_0 and a terminal value function D_T.
(i) The equilibria of this model are a non-empty lattice. Hence the model has a smallest equilibrium, $\{\bar{x}^L, N^L\}$, and a largest one, $\{\bar{x}^H, N^H\}$, and any other equilibrium path $\{\bar{x}, N\}$ satisfies $N^L \leq N \leq N^H$ and $\bar{x}^L \geq \bar{x} \geq \bar{x}^H$ for all $t \in [0, T]$.
(ii) Let $\theta' \geq \theta$, and $m_0' \leq m_0$ for all $x \in [0, U]$. Consider the equilibrium $\{\bar{x}', N'\}$ with the largest N' corresponding to $\{\theta', m_0'\}$ and the equilibrium $\{\bar{x}, N\}$ with largest N corresponding to $\{\theta, m_0\}$. Then $\bar{x}' \leq \bar{x}$ and $N' \geq N$ for all $t \in [0, T]$.

An important consequence of part (i) of the theorem is that the equilibrium set, given the initial distribution of non-adopters m_0 (and for finite T the terminal valuation D_T), is a lattice. We can compute the value of the extreme equilibria (i.e., the smallest and the largest) by iterating on $N^{k+1} = \mathcal{F}(N^k; D_T, m_0)$ for $k = 0, 1, \ldots,$ starting from $N^0(t) = 1$ or from $N^0(t) = 0$, for all t. The theorem ensures that the limit of this iterative process converges to a fixed point. Moreover, if the two sequences converge to the same limit, then the equilibrium is unique.

Two remarks are in order about part (i). First, as mentioned above, we implement the computations of the extreme equilibrium using the discretized version of the model for a finite T. Second, while this theorem shows that an equilibrium exists in the continuous time case, the theorem does not show that the fixed point N or $\bar{x}$ are continuous functions of time. As a consequence, the theorem does not establish the existence of a classical solution of the p.d.e.'s discussed in Section 2.2. Nevertheless for all the numerical examples of the extreme equilibrium that we computed using the discretized model we have found no behavior that resembles jumps in the path of N or $\bar{x}$. A type of equilibrium with a single jump is described in the next section as a "delayed adoption equilibrium". There we establish conditions under which the technology adoption can be shifted to an arbitrary future period.[5]

[5]We thank one referee to prod us to extend Theorem 1 to both discrete and continouos time as well as T

Part (ii) of the theorem focuses on the "high-adoption" equilibrium and establishes a useful comparative statics result: considering a larger θ, or a "smaller" m_0 (more agents endowed with the app at time zero), leads to more adoption.

3.2 Delayed Adoption Equilibria

In this section, we consider a family of multiple equilibria that are identical except for being shifted over time, with arbitrary delays in the onset of adoption. This family of equilibria arises under two main assumptions. The first assumption requires a prominent role of complementarities, namely parameters that satisfy $U\theta_0 < \rho c$, so that the adoption decisions depend on adoption decisions by "others". The second assumption is a restriction on the initial condition, namely that no agent starts with the technology at time zero. We have the following proposition.

Proposition 1. Assume that $U\theta_0 < \rho c$, that $m_0(x) = 1/U$, and that $T = \infty$. Assume that there is an equilibrium with $(N, \bar{x})$ such that $N(0) > 0$ and $\bar{x}(0) < U$. Let $t_0 > 0$ but otherwise arbitrary. Then there is an equilibrium $(N', \bar{x}')$ with $N'(0) = 0$ and $\bar{x}'(t) = U$ for $t \in [0, t_0)$, and with $N'(t) = N(t - t_0)$ and $\bar{x}'(t) = \bar{x}(t - t_0)$ for all $t \in [t_0, \infty)$.

A few comments are in order. First, the interpretation of $U\theta_0 < \rho c$ is clear. It says that using technology when nobody else uses it (even if x were kept at its highest value forever) does not compensate the adoption cost. Second, the delay in adoption featured in the $(N', \bar{x}')$ equilibrium is arbitrary and hence depends on an extreme amount of coordination of agents agreeing on when to adopt. Third, an equilibrium with $t_0 > 0$ cannot be the highest one, since it is dominated by the equilibrium with zero delay. Fourth, the initial equilibrium $(N, \bar{x})$ can be the highest one.

We note that the previous proposition assumes that the equilibrium $(N, \bar{x})$ starts with $N(0) > 0$ and $\bar{x}(0) < U$. These assumptions imply that, in the corresponding equilibrium with a delay $t_0 > 0$, there is a downward jump in $\bar{x}'$ and an upward jump in N' at time t_0. The next proposition shows that the assumption that $N(0) > 0$ and $\bar{x}(0) < U$ is without loss of generality.

Proposition 2. Assume that $U\theta_0 < \rho c$ and that $m_0(x) = 1/U$. Then there is no equilibrium in which $\bar{x}(0) = U$, $\bar{x}(t) < U$ for $t > 0$, and $\bar{x}(t)$ is continuous at $t = 0$.

The proposition establishes that if there exist an equilibrium with adoption, i.e. one where $\bar{x}(t) < U$, and the economy starts with zero adoption, then the threshold $\bar{x}(t)$ must

finite / infinite. We thank the editor and one referee for suggesting us to investigate the possibility of jumps when there are strong strategic complementarities.

display a jump, which is necessary to create a significantly large mass of adopters to initiate the technology diffusion. Given a calibrated value of ρ, we estimate the parameters θ_0 and c, among others, in Section 8: for our estimated values, we find that $\theta_0 U > \rho c$, that is, the opposite inequality from the one that assumed in the propositions above.

3.3 No-Adoption Equilibrium

The setup may feature an equilibrium with zero adoption, i.e., $\bar{x}(t) = U$ for all t. For simplicity we focus on the case where $T = \infty$. This case is particularly easy because agents' decisions are in a corner. We characterize the basin of attraction for such equilibrium, i.e., we find a threshold for the number of adopters $\underline{N}$, such that a no-adoption equilibrium exists if and only if at $t = 0$ the mass of agents with the technology is smaller than $\underline{N}$.

PROPOSITION 3. A no-adoption equilibrium with $\bar{x}(t) = U$ and $N(t) = N(0)e^{-\nu t}$ for all $t \geq 0$ exists if and only if $1 - \int_0^U m_0(x)dx \leq \underline{N}$, where

$$\frac{\rho c}{U} = \theta_0 \left[1 + g(\eta U)\right] + \underline{N} \frac{\rho \theta_n}{\rho + \nu} \left[1 + g(\eta' U)\right] \tag{13}$$

$$\eta \equiv \sqrt{\frac{2\rho}{\sigma^2}}, \; \eta' \equiv \sqrt{\frac{2(\rho + \nu)}{\sigma^2}} \text{ and } g(y) \equiv \frac{\operatorname{csch}(y) - \coth(y)}{y} \in (-\tfrac{1}{2}, 0) \; . \tag{14}$$

Note that $\underline{N} > 0$ if and only if $\frac{\rho c}{U} > \theta_0 \left[1 + g(\eta U)\right]$. Moreover, if $\underline{N} > 0$ we have:
(i) $\underline{N}$ is an increasing function of σ, satisfying

$$\frac{\rho + \nu}{\rho \theta_n} \left(\frac{\rho c}{U} - \theta_0 \right) \; \leq \; \underline{N} \; \leq \; \frac{\rho + \nu}{\rho \theta_n} \left(2\frac{\rho c}{U} - \theta_0 \right), \tag{15}$$

where the lower (upper) boundary is reached as $\sigma \to 0$ ($\sigma \to \infty$).
(ii) $\underline{N}$ is a decreasing function of θ_n.

An immediate corollary of this proposition is that $m_0(x) = 1/U$ is an invariant distribution provided that $\underline{N} \geq 0$, i.e., if the economy starts with no adoption, then it may remain in that equilibrium forever (no adoption is a stationary equilibrium). That $\underline{N} > 0$ requires θ_0 to be small is easily understood: if θ_0 is large agents with a high x will find it profitable to adopt regardless of what the others choose. Likewise, that $\underline{N} > 0$ is increasing in σ implies that if agents are hit by larges shocks the no-adoption equilibrium is more likely to occur. This result follows because, for a given U, a large σ makes the reversion to the mean faster, lowering the benefit of adoption. Finally, if θ_n is large then it is more profitable to coordinate on high N and the basin of attraction of the no-adoption equilibrium is smaller.

3.4 Stationary Equilibria

In this section we let $T = \infty$ and analyze the stationary equilibria of the continuous time model. We look for an initial condition m_0, such that the distribution is invariant, so that both $\bar{x}(t) = \bar{x}_{ss}$ and $N(t) = N_{ss}$ are constant through time. We will show that convergence to the stationary equilibrium must be gradual, i.e., that it is not possible to "jump" to the stationary equilibrium given a generic initial condition in the model where $\sigma > 0$.[6]

A stationary equilibrium is given by two constant values of N_{ss} and $\bar{x}_{ss}$ that solve the time-invariant version of the partial differential equations presented in Section 2.2. From a mathematical point of view the equilibrium is a fixed point. Given N_{ss}, $\tilde{D}$ and $\bar{x}_{ss}$ solve:

$$\begin{aligned}
\rho \tilde{D}(x) &= x(\theta_0 + \theta_n N_{ss}) + \frac{\sigma^2}{2}\tilde{D}_{xx}(x) \text{ if } x \in [0, \bar{x}_{ss}] && \text{Value of Adoption} \\
\tilde{D}_x(0) &= 0 && \text{Reflecting} \\
\tilde{D}(\bar{x}_{ss}) &= c && \text{Value Matching} \\
\tilde{D}_x(\bar{x}_{ss}) &= 0 && \text{Smooth Pasting .}
\end{aligned}$$

Conversely, given $\bar{x}_{ss}$, the density $\tilde{m}$ solves

$$\begin{aligned}
0 &= -\nu \tilde{m}(x) + \nu \frac{1}{U} + \frac{\sigma^2}{2}\tilde{m}_{xx}(x) && \text{KFE if } x \leq \bar{x}_{ss} \\
\tilde{m}(\bar{x}_{ss}) &= 0 \text{ and } \tilde{m}_x(0) = 0 && \text{Exit and Reflecting .}
\end{aligned}$$

Notice that the (stationary) equilibrium $\tilde{m}(x)$ and $\bar{x}_{ss}$ solve the fixed point

$$N_{ss} = 1 - \int_0^{\bar{x}_{ss}} \tilde{m}(s)dx.$$

We begin by solving for $\tilde{D}(x)$ and $\bar{x}_{ss}$ given a value for N_{ss} (see Appendix A.1 for details). Using the solution for $\tilde{D}$ we can solve for $\mathcal{X}_{ss} : [0,1] \to [0,U]$, a function that gives the *optimal* stationary threshold as a function of a given N_{ss}. The monotonicity properties of the function $\tilde{D}$ on the parameters N_{ss}, θ_n, c and θ_0 give the following characterization of the threshold $\mathcal{X}_{ss}$.

LEMMA 4. The function $\mathcal{X}_{ss}$ is decreasing in N_{ss}, strictly so at the points where $0 < \bar{x}_{ss} < U$. Fixing a value of N_{ss}, the function $\mathcal{X}_{ss}$ is strictly increasing in c, strictly so at the points where $0 < \bar{x}_{ss} < U$. Fixing a value of N_{ss}, the function $\mathcal{X}_{ss}$ is strictly decreasing in θ_0

[6] An immediate jump to the stationary equilibrium might instead occur in a model with $\sigma = 0$ (See the Online appendix J of Alvarez et al. (2023b)).

and θ_n at the points where $0 < \bar{x}_{ss} < U$. Moreover, we have the following expansion: $\mathcal{X}_{ss}(N_{ss}) = \frac{\rho c}{\theta_0 + \theta_n N_{ss}} + \frac{\sigma}{\sqrt{2\rho}} + o(\sigma)$.

Since the function $\mathcal{X}_{ss}(N_{ss})$ is decreasing in N_{ss}, it has an inverse, $\mathcal{X}_{ss}^{-1}$, given by:

$$\mathcal{X}_{ss}^{-1}(\bar{x}_{ss}) = \frac{1}{\theta_n}\left[\frac{\rho c}{\left(\bar{x}_{ss} + \bar{A}_1 e^{\eta \bar{x}_{ss}} + \bar{A}_2 e^{-\eta \bar{x}_{ss}}\right) - \frac{\left(1+\eta\left(\bar{A}_1 e^{\eta \bar{x}_{ss}} - \bar{A}_2 e^{-\eta \bar{x}_{ss}}\right)\right)\left(e^{\eta \bar{x}_{ss}} + e^{-\eta \bar{x}_{ss}}\right)}{\eta\left(e^{\eta \bar{x}} - e^{-\eta \bar{x}_{ss}}\right)}} - \theta_0\right] \text{ where}$$

$$\bar{A}_1 \equiv \frac{1}{\eta}\frac{\left(1 - e^{-\eta U}\right)}{\left(e^{-\eta U} - e^{\eta U}\right)},\ \bar{A}_2 \equiv \frac{1}{\eta}\frac{\left(1 - e^{\eta U}\right)}{\left(e^{-\eta U} - e^{\eta U}\right)} \text{ and } \eta \equiv \sqrt{2\rho/\sigma^2}. \quad (16)$$

Note that, from the expansion given in Lemma 4, fixing $\bar{x}_{ss}$, then $\mathcal{X}_{ss}^{-1}(\bar{x}_{ss})$ is increasing in σ in a neighborhood of $\sigma = 0$. Provided that $\theta_n > 0$ we have

$$\mathcal{X}_{ss}^{-1}(\bar{x}_{ss}) \approx \frac{1}{\theta_n}\left(\frac{c\rho}{\bar{x}_{ss} - \sigma/\sqrt{2\rho}} - \theta_0\right).$$

Next we can solve the Kolmogorov forward equation for $\tilde{m}(x)$, given a barrier $\bar{x}_{ss}$ subject to an exit point and to the boundary conditions coming from the reflecting barriers. We denote the corresponding value of the fraction that have adopted as $\mathcal{N}_{ss}(\bar{x}_{ss})$. Solving this equation we obtain

$$\mathcal{N}_{ss}(\bar{x}_{ss}) = 1 - \frac{\bar{x}_{ss}}{U} + \frac{\tanh\left(\gamma \bar{x}_{ss}\right)}{U\gamma} \text{ where } \gamma \equiv \sqrt{2\nu/\sigma^2}. \quad (17)$$

Inspection of equation (17) yields the following characterization of $\mathcal{N}_{ss}$.

Lemma 5. Fix $\gamma > 0$, then $\mathcal{N}_{ss}(\bar{x})$ is strictly decreasing in $\bar{x}_{ss}$. Fixing $\bar{x} > 0$, then $\mathcal{N}_{ss}$ is strictly increasing in γ, and hence strictly decreasing in σ. Moreover, we have the expansion: $\mathcal{N}_{ss}(\bar{x}) = 1 - \frac{\bar{x}_{ss}}{U} + \frac{\sigma}{U\sqrt{2\nu}} + o(\sigma)$.

As is intuitive, the value of $\mathcal{N}_{ss}(\bar{x}_{ss})$ is *decreasing* in the level of the barrier $\bar{x}$. The system given by equation (16) and equation (17) determines $\bar{x}_{ss}$ and N_{ss}. In particular, a stationary equilibrum is described by the pair $\{\bar{x}_{ss}, N_{ss}\}$, which solves

$$N_{ss} \equiv \mathcal{N}_{ss}(\bar{x}_{ss}) = \mathcal{X}_{ss}^{-1}(\bar{x}_{ss}).$$

Next, we summarize the behavior of the stationary equilibrium for small values of σ. We label the stationary equilibrium with superscripts $\{H, L\}$ to hint at the associated High or Low level of adoption, so that $\bar{x}^H < \bar{x}^L$. Indeed setting $\sigma = 0$ in the two expansions given in the previous two lemmas one obtains a quadratic equation for $\bar{x}_{ss}/U$ whose solution, whenever in $(0,1)$, gives the two interior steady states.

Proposition 4. Assume that $\nu > 0$ and that the parameters θ_0, θ_n, c and ρ are such that there are two interior stationary equilibria in the deterministic case of $\sigma = 0$, and label them as $\bar{x}^H_{ss} < \bar{x}^L_{ss}$. Then, (i) there exists a $\bar{\sigma} > 0$ such that for all $\sigma \in (0, \bar{\sigma})$ there are two interior stationary equilibria with $\bar{x}^H_{ss} < \bar{x}^L_{ss}$. (ii) The threshold for each stationary equilibrium is continuous with respect to σ at $\sigma = 0$. (iii) The sign of the comparative static differs across stationary equilibria, with

$$\frac{\partial \bar{x}^H_{ss}}{\partial c} > 0 > \frac{\partial \bar{x}^L_{ss}}{\partial c} \quad \text{and} \quad \frac{\partial \bar{x}^L_{ss}}{\partial \theta_0} > 0 > \frac{\partial \bar{x}^H_{ss}}{\partial \theta_0}.$$

The proposition shows that the high adoption stationary equilibrium behaves in an intuitive way, with more adoption (a lower $\bar{x}^H_{ss}$) associated with a smaller adoption cost (c), or with a larger intrinsic value of the technology (θ_0). The comparative statics for the low adoption stationary state are just the opposite: adoption is higher as the adoption cost increases. The latter (unrealistic) feature, and the unstable nature of the low adoption equilibrium (see the next section), will lead us to focus on the high adoption equilibrium in our quantitative analysis.

Figure 1: Stochastic Stationary Equilibria: Density of non-adopters: $\tilde{m}(x)$

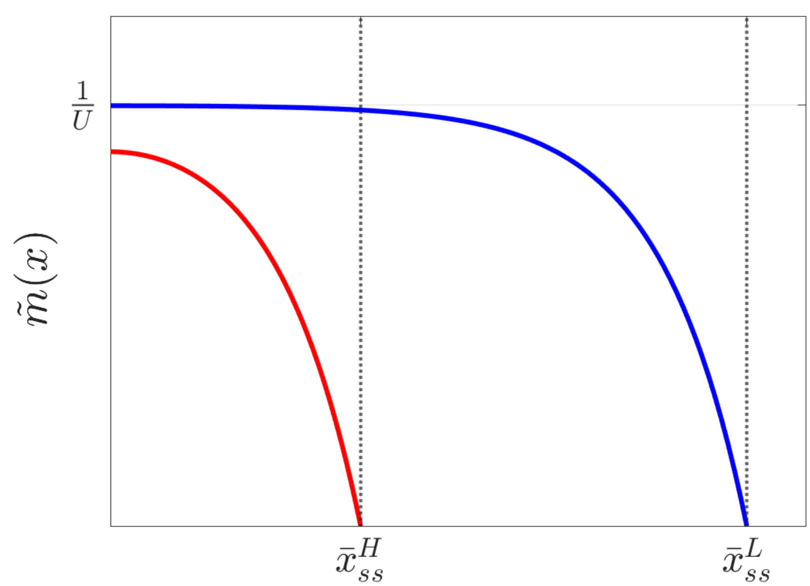

High and Low Adoption Stationary Equilibria

Figure 1 shows the densities of the invariant distribution of the high- and low-adoption equilibria. A notable feature of the stationary distribution of non-adopters is that, provided

$\sigma > 0$, the distibution features agents with "low x, namely with $x(t) < \bar{x}_{ss}$, who have the technology. These are agents who adopted the technology in the past (for some $t' < t$ when $x(t') > \bar{x}(t')$, and whose x decreased over time. As a result, $\tilde{m}(x) < 1/U$ when $\sigma > 0$, and the density of non-adopters below $\bar{x}_{ss}$ is not uniform. Given that the density takes time to adjust, the stochastic model features dynamics in the adoption of a new technology: for instance if the economy starts with $m_0 = 1/U$, it takes time to move from the initial to the invariant distribution, as agents adopt when $x(t) > \bar{x}(t)$ and it takes time for the $x's$ to crawl back below the stationary threshold. In other words, this form of the invariant distribution and the fact that agents follow a threshold rule implies that there is no equilibrium where at some finite time t the economy jumps to the steady state. Instead, as mentioned above, such a jump must occur in a model where x is heterogeneous across agents and $\sigma = 0$ for a large set of initial conditions (see Proposition 20 in Appendix J in Alvarez et al. (2023b)).

4 Stability of Stationary Equilibria

In this section we analyze the local stability of the stationary equilibria. We explore the question by perturbing the stationary distribution of adopters, using techniques from the Mean Field Game literature developed in Alvarez, Lippi and Souganidis (2023a). For this purpose, we use the equilibrium Definition 2. This dynamical system is infinite-dimensional because the state, at every time t, is given by the entire density $m(x, t)$.

The objective is to consider the stationary equilibrium $\tilde{m}$ and ask if, starting from a condition m_0 close to $\tilde{m}$, the economy converges to $\tilde{m}$. As the system is infinite-dimensional, many "deviations" are possible. Any initial condition can be described by $m_0(x) = \tilde{m}(x) + \epsilon\omega(x)$, for some ω satisfying $\int_0^U \omega(x)dx = 0$. The sense in which the analysis is local is that we differentiate the system with respect to ϵ and evaluate it at $\epsilon = 0$. The alert reader will notice that the local dynamics of a system in $\mathbb{R}^q$ are encoded in a $q \times q$ matrix. The analogous infinite dimensional object is a linear operator that will be presented below.

We begin the analysis with the approximation of $\bar{x}(t) = \mathcal{X}(N)(t)$. That is, we study how perturbing the aggregate path of adoption N leads to adjusting the decision rule for threshold path $\bar{x}$. To do this, we take the directional derivative (Gateaux) with respect to an arbitrary perturbation n of a constant path N. In particular, we consider paths defined by $N(t) = N_{ss} + \epsilon\, n(t)$ around the stationary value N_{ss}. We denote this Gateaux derivative by $\bar{y}$, so that $\bar{x}(t) \approx \bar{x}_{ss} + \epsilon\bar{y}(t)$.

Lemma 6. Fix a stationary equilibrium with interior $\bar{x}_{ss}$, and its corresponding N_{ss}. Let D_T be equal to the stationary value function $\tilde{D}$ corresponding to that stationary equilibrium. Let $n : [0, T] \to \mathbb{R}$ be an arbitrary perturbation. Then

$$\bar{y}(t) \equiv \lim_{\epsilon \downarrow 0} \frac{\mathcal{X}(N_{ss} + \epsilon n; \tilde{D})(t) - \mathcal{X}(N_{ss}; \tilde{D})(t)}{\epsilon}$$
$$= \frac{\theta_n}{\tilde{D}_{xx}(\bar{x}_{ss})} \int_t^T G(\tau - t) n(\tau) d\tau, \tag{18}$$

where

$$G(s) \equiv \sum_{j=0}^{\infty} c_j e^{-\psi_j s} \geq 0 \,,\; \psi_j \equiv \rho + \frac{\sigma^2}{2} \left(\frac{\pi(\frac{1}{2} + j)}{\bar{x}_{ss}} \right)^2 \text{ and } c_j \equiv 2 \left(1 - \frac{\cos(\pi j)}{\pi(j + \frac{1}{2})} \right) ,$$

where $\tilde{D}_{xx}(\bar{x}_{ss}) < 0$ is the second derivative of the stationary value function:

$$\tilde{D}_{xx}(\bar{x}_{ss}) = \frac{\rho c - \bar{x}_{ss} \left[\theta_0 + \theta_n N_{ss} \right]}{\sigma^2/2} \,,\; N_{ss} = 1 - \frac{\bar{x}_{ss}}{U} + \frac{\tanh(\gamma \bar{x}_{ss})}{\gamma U} \text{ and } \gamma = \sqrt{\frac{2\nu}{\sigma^2}}.$$

Thus, we can write $\bar{x}(t) = \bar{x}_{ss} + \epsilon \bar{y}(t) + o(\epsilon)$. Note that G is positive and D_{xx} is negative, so the effect of the future path on the current value is negative, which is consistent with the property that $\mathcal{X}$ is decreasing. Also note that it is proportional to θ_n, so if $\theta_n = 0$, then the threshold will be constant. Thus, the approximation of $\bar{x}(t)$ depends on the perturbation of the path of N from t to T, given by $n(s)$ for $s = [t, T]$. The proof of the proposition is obtained by jointly differentiating with respect to ϵ the system defined by D and $\bar{x}$ in equation (9) and equation (10). This yields a new p.d.e., and new boundary conditions. The expression for $\bar{y}$ is obtained once we solve this new p.d.e., see the proof in Appendix C.1.

Now we turn to the perturbation for the fraction of the adopters, as a function of the threshold path and of a perturbation of the initial condition. We approximate $N(t) = \mathcal{N}(\bar{x}, m_0)(t)$ by taking the directional derivative (Gateaux) with respect to an arbitrary perturbation $\bar{y}$ of a constant path $\bar{x}$ and a perturbation ω on the stationary density $\tilde{m}$. In particular, we consider paths defined by $\bar{x}(t) = \bar{x}_{ss} + \epsilon \, \bar{y}(t)$ around the stationary threshold x_{ss}, and $m_0(x) = \tilde{m}(x) + \epsilon \, \omega(x)$. We will denote this Gateaux derivative by n.

Lemma 7. Fix the interior threshold $\bar{x}_{ss}$ of a stationary equilibrium and its corresponding N_{ss}, and let $\tilde{m}$ be the corresponding invariant distribution of non-adopters. Let $\omega : [0, \bar{x}_{ss}] \to \mathbb{R}$ be an arbitrary perturbation to the distribution, and let $\bar{y} : [0, T] \to \mathbb{R}$ be an arbitrary perturbation of the threshold. Then

$$n(t) \equiv \lim_{\epsilon \downarrow 0} \frac{\mathcal{N}(\bar{x}_{ss} + \epsilon \bar{y}; \tilde{m} + \epsilon w)(t) - \mathcal{N}(\bar{x}_{ss}; \tilde{m})(t)}{\epsilon}$$
$$= n_0(\omega)(t) + \frac{\tilde{m}_x(\bar{x}_{ss}) \sigma^2}{\bar{x}_{ss}} \int_0^t J(t - \tau) \bar{y}(\tau) d\tau \tag{19}$$

where
$$J(s) = \sum_{j=0}^{\infty} e^{-\mu_j s} \text{ with } \mu_j = \nu + \frac{1}{2}\sigma^2 \left(\frac{\pi(\frac{1}{2}+j)}{\bar{x}_{ss}} \right)^2 \tag{20}$$
$$n_0(\omega)(t) \equiv - \sum_{j=0}^{\infty} \frac{\bar{x}_{ss}}{\pi(\frac{1}{2}+j)} \frac{\langle \varphi_j, \omega \rangle}{\langle \varphi_j, \varphi_j \rangle} e^{-\mu_j t}, \tag{21}$$

$$\varphi_j(x) \equiv \sin\left(\left(\frac{1}{2} + j \right) \pi \left(1 - \frac{x}{\bar{x}_{ss}} \right) \right) \text{ for } x \in [0, \bar{x}_{ss}] \tag{22}$$
$$\frac{\langle \varphi_j, \omega \rangle}{\langle \varphi_j, \varphi_j \rangle} = \frac{2}{\bar{x}_{ss}} \int_0^{\bar{x}_{ss}} \varphi_j(x)\omega(x)dx \text{ and } \tilde{m}_x(\bar{x}_{ss}) = -\frac{\gamma}{U} \tanh(\gamma \bar{x}_{ss}).$$

Thus, we can write $N(t) = N_{ss} + \epsilon n(t) + o(\epsilon)$. This formula encodes the effect of two perturbations: ω and $\bar{y}$. The former is the perturbation on the initial condition m_0, whose effect is in the term $n_0(\omega)(t)$. We note that $n_0(\omega)(t)$ is the effect at time t on the path $N(t)$ triggered by a perturbation of the initial condition keeping the threshold rule $\bar{x}$ fixed. The function $n_0(\omega)$ can be further reinterpreted by considering the limiting case of a perturbation ω given by a distribution concentrated at $x = \hat{x} \leq \bar{x}_{ss}$, i.e., a Dirac's delta function as $\omega(x) = \delta_{\hat{x}}(x)$. In this case,

$$n_0(\delta_{\hat{x}})(t) = - \sum_{j=0}^{\infty} 2 \frac{\sin\left(\left(\frac{1}{2}+j\right) \pi \left(1 - \frac{\hat{x}}{\bar{x}_{ss}}\right) \right)}{(\frac{1}{2}+j)\pi} e^{-\mu_j t}.$$

The effect of the perturbation, $\bar{y}$, on the path of the threshold, $\bar{x}(s)$, is captured by the second term in equation (19). This term gives the effect at time t on the path $N(t)$ of a perturbation of the threshold rule $\bar{x}$, keeping the initial condition $\tilde{m}$ fixed. Also, consistent with our general result for $\mathcal{N}$, the effect of the threshold is negative, as $J > 0$ and $\tilde{m}_x(\bar{x}_{ss}) < 0$.

For future reference it is useful to understand the behavior of $n_0(t)$ as function of time. In particular, the rate at which the perturbation ω to the initial distribution converges back to the stationary distribution, while keeping $\bar{x}(t) = \bar{x}_{ss}$. This rate is given by the value of $\mu_0 = \nu + \frac{\sigma^2}{8} \left(\frac{\pi}{\bar{x}_{ss}} \right)^2$, i.e., the dominant eigenvalue.[7]

The next step is to use the last two lemmas to derive one equation for the linearized equilibrium as a function of the perturbed initial distribution $m_0(x) = \tilde{m}(x) + \epsilon\omega(x)$. We combine equation (18) and equation (19) to arrive to a single linear equation that $n(t)$ must solve as a function of ω.

Theorem 2. Fix an interior threshold $\bar{x}_{ss}$ for a stationary state, with its corresponding

[7] The proof is in Appendix C.2 and resembles the one for the previous proposition.

N_{ss}, and let $\tilde{m}$ be the corresponding invariant distribution of non-adopters. Let $m_0(x) = \tilde{m}(x) + \epsilon\omega(x)$. Let D_T be equal to the value function $\tilde{D}$ corresponding to that stationary equilibrium. The linearized equilibrium solves

$$n(t) = n_0(\omega)(t) + \Theta \int_0^T K(t,s)n(s)ds, \tag{23}$$

where $n_0(\omega)(t)$ is given in Lemma 7 and $\Theta \equiv \frac{\tilde{m}_x(\bar{x}_{ss})\sigma^2\theta_n}{\bar{x}_{ss}\tilde{D}_{xx}(\bar{x}_{ss})} > 0$. The kernel K is given by

$$K(t,s) = \sum_{i=0}^{\infty}\sum_{j=0}^{\infty} c_j e^{-\mu_i t - \psi_j s}\left[\frac{e^{(\mu_i+\psi_j)\min\{t,s\}} - 1}{\mu_i + \psi_j}\right] > 0. \tag{24}$$

Moreover, $\text{Lip}_K \equiv \sup_t \int |K(t,s)|ds \leq \left(\frac{\bar{x}_{ss}^2}{\sigma^2}\right)^2$. Furthermore, if $\Theta\,\text{Lip}_K < 1$ there exists a unique bounded solution to equation (23) which is the limit of

$$n = \left[I + \Theta\mathcal{K} + \Theta^2\mathcal{K}^2 + \ldots\right] n_0(\omega) \;\; \text{where} \;\; \mathcal{K}(g)(t) \equiv \int_0^T K(t,s)g(s)ds, \tag{25}$$

and where $\mathcal{K}^{j+1}(g)(t) \equiv \int_0^T K(t,s)\,\mathcal{K}^j(g)(s)\,ds$ for any bounded $g : [0,T] \to \mathbb{R}$.

A few remarks are in order. First, note that K depends on θ_n as μ_j, ψ_j are a function of $\bar{x}_{ss}$, which is itself a function of θ_n. The coefficient Θ depends on θ_n directly and indirectly via $\bar{x}_{ss}$. Hence equation (23) depends on which stationary equilibrium we focus on. Second, if we discretize time so that $t \in \{\Delta_t(j-1) : j = 1, \ldots, J\}$ for $\Delta_t = \frac{T}{J-1}$, as done in Section 2.1, then the operator $\mathcal{K}$ is a $J \times J$ matrix with elements $K(t_i, t_j)$, and n_0, n are $J \times 1$ vectors, so that equation (23) becomes the linear equation $n = n_0 + \Theta\mathcal{K}\,n$. Third, the fact that $\Theta\mathcal{K} > 0$ implies that the terms $\Theta\mathcal{K} + \Theta^2\mathcal{K}^2 + \ldots$ in equation (25) give the amplification over and above n_0, due to the time-varying path of the barrier $\bar{x}$.

Figure 2 illustrates the stability of the high and low adoption equilibria, respectively, in Panels (a) and (b). Each panel considers two shocks that displace a small mass of agents away from the invariant distribution of non-adopters and endows them with the app. The blue line depicts the case where the app is given to agents with low benefit, namely with $x \approx 0$, while the red line considers a perturbation where the app is given to agents with a high benefit, namely with $x \approx \bar{x}_{ss}$.[8] Two remarks are due. First, the high adoption equilibrium is locally stable, as displayed in Panel (a): for all shocks considered, the system returns to its invariant distribution. We also note that the half life of the shock is much shorter when

[8] Parameters used for illustration: $\theta_0 = 26.32$; $\theta_n = 5.72 \cdot \theta_0$; $U = 1$; $\nu = 0.0278$; $r = 0.05$; $\sigma = 0.032$; $c = 2 \cdot 10.54 \cdot \theta_0$.

Figure 2: Perturbation of Stationary Equlbria

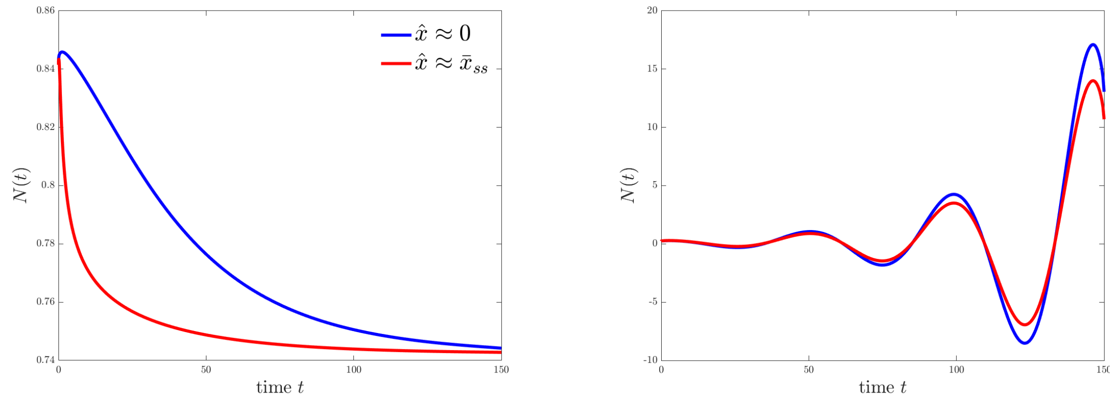

(a) High Adoption Stat. Eqbm.

(b) Low Adoption Stat. Eqbm.

the perturbation assigns the app to agents with a high benefit ($x \approx \bar{x}_{ss}$), as these agents were going to get the app soon anyways. Second, Panel (b) reveals that the low adoption equilibrium is unstable: the dynamics of the system following a perturbation are explosive, i.e., the sequence in equation (25) does not converge so that the system does not return to the invariant distribution after the shock. To appreciate the explosive nature of the path nearby the low activity stationary equilibrium, notice the difference in the scales of the two panels.

5 The Planning Problem

This section sets up the planning problem, characterizes of its solution, and shows how it can be decentralized as an equilibrium with a subsidy.[9] The planner solves a non-trivial dynamic problem since the state of the economy is an entire distribution.

At time zero the planner solves:

$$\max_{\{\bar{x}(t)\}} \left\{ \int_0^\infty e^{-rt} \int_0^U \underbrace{(1/U - m(s,t))}_{\text{Density of adopters}} \underbrace{s\,(\theta_0 + \theta_n N(t))}_{\text{Flow benefit}} \, ds \, dt \right.$$
$$\left. - \underbrace{\int_0^\infty e^{-rt} c\,(N_t(t) + \nu N(t))\, dt}_{\text{Flow of adoption cost: gross new adoptions}} \quad \right\}$$

[9]Appendix D.1 characterizes the stationary solution of this problem. Appendix D.5 uses a linearized version of the problem to analyze dynamics around its invariant distribution, an exercise that is akin to the one of Section 4.

subject to

$$N(t) = 1 - \int_0^{\bar{x}(t)} m(s,t) ds \quad \text{for all } t$$

$$m_t(x,t) = -\nu \left(m(x,t) - 1/U \right) + \frac{\sigma^2}{2} m_{xx}(x,t) \qquad \text{for } x \in (0, \bar{x}(t)) \text{ and all } t \geq 0 \quad \text{KFE}$$

$$m(x,t) = 0 \quad \text{for } x \in [\bar{x}(t), U] \text{ and all } t \geq 0 \quad \text{Adoption}$$

$$m_x(0,t) = 0 \quad \text{for all } t \geq 0 \quad \text{Reflecting}$$

$$m(x,0) = m_0(x) \text{ for all } x \ . \quad \text{Initial condition}$$

The objective function of the planner integrates the lifetime utility of agents using as a weight the discount factor e^{-rt} for the cohort born at t. The first term contains the utility flow of the agents who use the technology. The second term subtracts the cost of adoption, where $c(N_t(t) + \nu N(t))$ is the gross flow cost of adoption at time t. This flow cost is driven by the inflow of new adopters $N_t(t)$ and by the replacement of dead agents (who had adopted in the past) by newborns.[10] The first constraint defines $N(t)$, the second constraint is the KFE for the density of non-adopters, m. As before, the density is zero to the right of $\bar{x}(t)$, there is an exit point at $\bar{x}(t)$, and there is a boundary condition from the reflection at zero.

At each time t the planner decides a threshold $\bar{x}(t)$ that determines adoption, taking as given the initial condition $m_0(x)$ and the law of motion of the density m (affected by the choice of $\bar{x}$). To characterize the solution, we write the Lagrangian for this problem. We denote the Lagrange multiplier of the KFE equation by $e^{-rt}\lambda(x,t)$ and replace $N(t)$ and $N_t(t)$ by the corresponding definition. To derive the p.d.e's for non-adopters, we first adapt the planning problem to a discrete-time discrete-state problem using a finite-difference approximation. In this set up, we allow for a more general policy, i.e., not necessarily a threshold rule. We obtain the first order conditions for a problem in finite dimensions and take limits to find the corresponding p.d.e's, summarized in the following proposition.[11]

Lemma 8. A planner's problem is given by $\{\bar{x}(t), \lambda(x,t), m(x,t)\}$ such that adoption occurs for $x \geq \bar{x}(t)$, and the Lagrange multiplier λ, and the density of non-adopters m solve

[10] At every moment there is an inflow ν of newborns without the app. A fraction $1 - \frac{\bar{x}(t)}{U}$ of the newborns immediately pays the cost c and adopts, see Appendix D.2 for details.

[11] We provide details of this derivation in Appendix D.3.

the p.d.e. for non-adopters:

$$\begin{aligned}
\rho\lambda(x,t) &= x\Big(\theta_0+\theta_n[1-\int_0^{\bar{x}(t)} m(s,t)ds]\Big)+\theta_n\Big(\tfrac{U}{2}-\int_0^{\bar{x}(t)} m(s,t)s\,ds\Big) \qquad (26)\\
&\quad+\tfrac{\sigma^2}{2}\lambda_{xx}(x,t)+\lambda_t(x,t) \text{ for } x\le\bar{x}(t) \text{ and } t\ge 0\\
\lambda(x,t) &= c \text{ for } x\ge\bar{x}(t) \text{ and } t\ge 0\\
\lambda_x(\bar{x}(t),t) &= 0 \text{ for } t\ge 0 \qquad (27)\\
\lambda_x(0,t) &= 0 \text{ for } t\ge 0
\end{aligned}$$

and

$$\begin{aligned}
m_t(x,t) &= \nu\big(1/U-m(x,t)\big)+\tfrac{\sigma^2}{2}m_{xx}(x,t) \text{ for } x<\bar{x}(t) \text{ and } t\ge 0\\
m(x,t) &= 0 \text{ for } x\ge\bar{x}(t) \text{ and } t\ge 0\\
m_x(0,t) &= 0 \text{ for } t\ge 0\\
m(x,0) &= m_0(x) \text{ for all } x\ .
\end{aligned}$$

This lemma has two important consequences. First, it allows us to compute the solution of the planning problem following similar steps as the ones used to compute the equilibrium in Section 3.1. Second, it indicates how to decentralize the optimal allocation as an equilibrium. Define $Z(t)\equiv\frac{U}{2}-\int_0^{\bar{x}(t)} m(s,t)s\,ds\ge 0$ and note that this non-negative magnitude is the difference between the average x in the population, $U/2$, and the average x among those who have not adopted the technology (the integral term). Comparing the p.d.e. for the Lagrange multiplier λ in equation (26) with the p.d.e. for D that characterizes the equilibrium in equation (9), we see that these equations only differ in the flow term $\theta_n Z(t)$. Thus, if agents who adopt the technology are given a flow subsidy $\theta_n Z(t)$ every period after they have adopted (independent of the app's usage), then the planner allocation is an equilibrium. Clearly, this is equivalent to a once and for all payment to agents adopting at t equal to $\theta_n\int_t^\infty e^{-\rho(s-t)}Z(s)ds$. Note that $\theta_n Z(t)$ contains the inframarginal valuation of the technology for those that use it, so the subsidy's work by correcting the externality associated with the individual adoption. We summarize this discussion in the following theorem.

Theorem 3. Fix an initial condition m_0 and the solution to the planner's problem $\{\bar{x},\lambda,m\}$. The planner's allocation coincides with an equilibrium with the same initial conditions and a time-varying flow subsidy paid to adopters given by $\theta_n Z(t)$, where

$$Z(t)\equiv\tfrac{U}{2}-\int_0^{\bar{x}(t)} m(s,t)\,s\,ds \quad \text{for all } t\ge 0 \qquad (28)$$

The subsidy $\theta_n Z$ is independent of x.

For future reference, we define $Z \equiv \mathcal{Z}(\bar{x}; m_0)$ as the solution of the path for Z defined in equation (28). In particular, given $\bar{x}$ and m_0, we solve for m using the KFE and then compute Z.

Consider the path $\bar{x}$ that solves the p.d.e. $\rho\lambda(x,t) = x(\theta_0 + \theta_n N(t)) + \theta_n Z(t) + \frac{\sigma^2}{2}\lambda_{xx}(x,t) + \lambda_t(x,t)$ with the three boundaries given in equation (27) given the paths of N and Z and terminal condition $\lambda(x,T) = \lambda_T(x)$. Let $\bar{x} = \mathcal{X}^P(N, Z; \lambda_T)$ denote the functional, defined as the $\mathcal{X}$ in Section 2.1, where the superscript P denotes the planning problem. Note that, using the definitions for $\mathcal{X}^P, \mathcal{Z}$ and $\mathcal{N}$ the planner's problem must satisfy the fixed point $\bar{x}^* = \mathcal{H}(\bar{x}^*, \lambda_T, m_0)$ where $\mathcal{H}(\bar{x}; \lambda_T, m_0) \equiv \mathcal{X}^P(\mathcal{N}(\bar{x}; m_0), \mathcal{Z}(\bar{x}; m_0); \lambda_T)$. We can use the analysis used in Section 3, based on monotonicity, to characterize the solution to this fixed point problem, and to compute it.

Figure 3 illustrates how the application of the optimal subsidy leads to a high adoption equilibrium. In Panel (a) of the figure, we plot the time path of the share of adopters, $N(t)$, for the planning problem, using the stationary equilibrium distribution of non-adopters as the initial distribution (i.e., $m_0(x) = \tilde{m}(x)$). Let denote by N_{ss} the value of the equilibrium steady state. In the planning problem, the path of $N(t)$ jumps immediately from N_{ss} (at the time the subsidy is introduced) and gradually converges to the stationary distribution for the planning problem.[12] Panel (b) shows the time path of the optimal subsidy to implement the optimal, $Z(t)$, which starts at the value $Z(0) = \frac{U}{2} - \int_0^{\bar{x}_H} \tilde{m}(s) s ds$ and increases over time. In this example, the high-adoption equilibrium has partial adoption, i.e., $N_{ss} < 1$, but the efficient allocation, as can be seen in panel (a), converges to almost full adoption of the technology.

Figure 3: Planning Problem: $m_0(x) = \tilde{m}(x)$

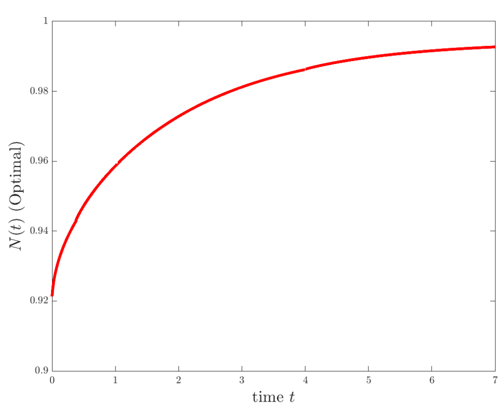

(a) $N(t)$ (Optimal)

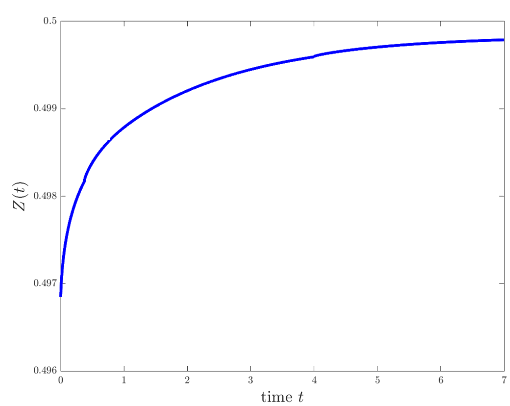

(b) $Z(t)$

[12] In this example, $N_{ss} = 0.42$.

6 Adding Information Diffusion

We incorporate learning about the existence of the technology through a simple extension of the celebrated Bass's (1969) model of information diffusion. We do so both for completeness and to accommodate the evidence on knowledge on the technology presented in Section 8. We first describe how the learning process works, then present the extension to the baseline model, and conclude with two simpler special cases: $\theta_n = 0$ or $\sigma = 0$.[13]

Dynamics of Learning. We assume newborn agents are initially uninformed and become informed by randomly matching with informed agents. The set $I(t)$ of informed agents is divided into $N(t)$ agents who have adopted and $M(t)$ agents who are informed but have not adopted, so that $I(t) = N(t) + M(t)$. Both types of informed agents transmit information, so the dynamics of $I(t)$ are independent of agents' adoption decisions. The law of motion of m needs to be modified to include the inflow of informed agents as in a random diffusion model:

$$m_t(x,t) = \frac{\sigma^2}{2} m_{xx}(x,t) + \frac{\beta_0}{U} I(t)(1 - I(t)) - \nu m(x,t) \text{ all } t \geq 0 \text{ and } x \in [0, \bar{x}]$$
$$m(x,t) = 0 \text{ all } t \geq 0 \text{ and } x \in [\bar{x}, U]$$
$$m_x(0,t)\,\sigma = 0 \text{ all } t \geq 0$$

where $I(t)$ denotes the fraction of the population informed about the technology, and the parameter β_0 is the parameter governing the number of meetings per unit of time between those informed, $I(t)$, and those uninformed, $1 - I(t)$. The term $\frac{\beta_0}{U} I(t)(1 - I(t))$ represents the flow of agents per unit of time who learn about the app. The time path for $I(t)$ has a closed form solution: it is initially convex, then becomes concave, and converges to $I_{ss} = 1 - \nu/\beta_0$. The entire path of I depends only on the parameters (ν, β_0) and on the initial value of $I(0)$. If $I(0)$ is sufficiently small, then $I(t)$ remains low for an extended period of time. See Appendix E for additional details.

Baseline model ($\sigma > 0, \theta_n > 0$). Above, we described the law of motion of m for a given path of $\bar{x}$. We now describe the determination of $\bar{x}$, which applies only to informed agents. The variational inequality governing the adoption decision (i.e., net value of adoption $a(x,t) - c$ and the net optimal value $v(x,t)$) are the same as in the model with strategic complementarities, since adoption can occur only after agents become aware of the technology. Thus, the only change lies in the law of motion of m described above. Theorem 1 and

[13] We thank the editor and one referee for suggesting us to investigate the possibility of jumps and the role of $\sigma = 0$.

Theorem 3 extend in a straightforward manner to this case. Also, as $t \to \infty$, so that $I(t) \to I_{ss}$ this model has stationary states which are a scaled version of the ones for the model without learning.

Next we analyze the interaction between learning and strategic complementarity. The following proposition considers a case of strong strategic complementarity with an initial condition in which no agent has the technology. In this case, for sufficiently small $I(0)$, there is a period where there is no adoption followed by a downward jump in $\bar{x}$ and an upward jump in N (coming from zero):

Proposition 5. *Assume that $U\theta_0 < \rho c$. Then for any $\tau > 0$, there is a $I(0)$ small enough, such that $\bar{x}(t) = U$ for $t \in [0, \tau)$.*

The intuition behind Proposition 5 is straightforward: until there is a critical mass of informed agents, $I(\tau)$, adoption cannot occur. Recall that for any $\tau > 0$ and $\epsilon > 0$, there exists an $I(0)$ small enough such that $I(t) < \epsilon$ for all $t \in [0, \tau)$. It then follows that adoption is bounded above by $N(t) \leq I(t)$, and, given the assumptions of the proposition, even agents with the highest x find it optimal to delay adoption (since $U\theta_0 < \rho c$). Indeed the proof is almost a line-by-line adaptation of the proof of Proposition 2. Figure E3 in the appendix uses parameters and initial conditions as in Proposition 5 to illustrate that there is no adoption as long as $I(t) < I(\tau)$. This figure also shows that the multiple equilibria, involving the further delay in adoption described in Proposition 2, extends to the model with learning.[14]

We conclude this section by considering two simpler versions of the learning model.

Learning without strategic complementarity. This setup, developed in detail in Appendix E.1, considers a "pure" learning model without complementarities ($\theta_n = 0$) in which adoption benefits are random ($\sigma > 0$). Informed agents can pay the cost c and adopt. We show that the optimal decision for informed agents is given by a time-invariant threshold $\bar{x}$, which is independent of the network size. This invariance implies that there is no selection in the adoption of the technology: early and late adopters are similar agents in terms of their x, differing only in the timing of when they learn about the technology. This prediction contrasts with the evidence on selection discussed in Section 7.2.1. The model has a unique constrained-efficient equilibrium with a logistic S-shaped path for N when the initial share of informed agents is small. Along the equilibrium path, the dynamics of $N(t)$ are fully determined by the dynamics of $I(t)$.

[14] With learning, the continuation of the highest equilibrium differs from that of the equilibrium with delay. This difference arises because, when $\sigma > 0$, agents who adopted before the delay continue to contribute to the size of the network even if their x falls below the adoption threshold.

Learning with complementarities and fixed types. A learning model with complementarities ($\theta_n > 0$) is analyzed in details in Appendix E.2, where the idiosyncratic benefit of adoption x is deterministic ($\sigma = 0$). In this model the adoption benefit depends on the size of the network $N(t)$, while the agent type x is immutable. We focus on the case where either none or very few agents start with the technology –defined properly in Appendix E.2. In this case, due to the network effects ($\theta_n > 0$), the optimal decision for informed agents is a monotone, time-varying threshold $\bar{x}(t)$, with an associated monotone (increasing or decreasing) path for $N(t)$. The model may feature multiple constrained-efficient equilibria, and the dynamics of the equilibrium path $N(t)$ are again determined by the dynamics of $I(t)$.

The simplicity of the model with $\sigma = 0$ has both advantages and disadvantages. One advantage is that it allows for a complete characterization of the critical delay threshold $I(\tau)$ given in Proposition 5. Another advantage is that the equilibrium path starting from low adoption can be computed explicitly. On the other hand, the reason the equilibrium can be computed explicitly is that is is essentially static. In particular, the level of $N(t)$ at the highest equilibrium is just a function of $I(t)$. This implies that if we consider an "MIT" shock that removes the technology from a group of agents who had previously adopted, these agents immediately re-adopt, and the equilibrium returns to its pre-shock position, as described in Proposition 16. Instead, in a model with $\sigma > 0$, re-adoption occurs gradually.

7 Application: SINPE, A Digital Payments Platform

In May 2015, the Central Bank of Costa Rica (BCCR) launched SINPE Móvil (hereafter, SINPE), a digital platform that enables users to make money transfers using their mobile phones.[15] To utilize SINPE, users must have a bank account at a financial institution and link it to their mobile number. According to the BCCR, the primary objective of SINPE was to become a mass-market payment mechanism that could reduce the demand for cash as a method of payment. As such, SINPE was originally designed for relatively small transfers, which are not subject to any fee as long as they do not exceed a daily sum. The maximum daily amount transferred without a fee varies by bank; for most users, it is approximately \$310, although some banks have lower limits of \$233 and \$155.[16] The average transaction size in SINPE is about \$50, and has slowly decreased over time, as shown in Figure G2.

While, in theory, firms are allowed to adopt SINPE and conduct transactions within the app, in practice, transactions involving firms represent less than 5% of all payments. This

[15]SINPE is an acronym for the initials of "National Electronic Payment System" (*Sistema Nacional de Pagos Electrónicos*) in Spanish.

[16]Respectively, these limits in dollars correspond with approximately 200,000; 150,000; and 100,000 Costa Rican colones. These amounts correspond with 2021 limits and exchange rates.

motivates us to study adoption through the lens of our model while focusing on peer-to-peer transactions where small agents trade with each other, rather than one with a few non-atomistic players (large firms). Appendix G.2 presents details on the transactions by user type and between networks and discusses the slow adoption for person-to-business(P2B) and business-to-business (B2B) payments.

7.1 Data

This section describes the battery of administrative datasets used in the paper. First, we leverage data on SINPE transactions. Our data on SINPE usage is comprehensive: For each user in the country, we have official records on the *exact date* when she adopted the technology, along with records on each transaction made. In particular, for each transaction, the data records the *amount transacted* along with the individual identifier of *the sender and the receiver* of the money. Records also include the sender's and the receiver's bank. Importantly, this information is available, not only for individuals, but also for firms.

We also leverage information on networks of coworkers for each formally employed individual, along with their income. Matched employer-employee data is obtained from the Registry of Economic Variables of the Central Bank of Costa Rica, which tracks the universe of formal employment and labor earnings. The data include *monthly* details on each employee, including her earnings and employment history spanning SINPE's lifetime (2015-2021).[17] With this information, we can identify which people are working at the same firm in a given month to construct networks of coworkers which can be matched to SINPE records. Networks of coworkers vary at a monthly frequency as people change employers.

While our baseline analysis focuses on coworkers networks, we complement its statistics with those of other network types, namely, networks of neighbors and relatives. We construct networks of neighbors for all adult citizens in the country leveraging data from the National Registry and the Supreme Court of Elections. The data consist of official records on the residence of each citizen.[18] Data on nationwide family networks is available from the National Registry and makes it possible to reconstruct each person's family tree.[19] The data includes individual identifiers that can be linked to SINPE. The same data source provides details on individual demographics. Finally, we leverage data on corporate income tax returns from the

[17]It is worth noting that informal workers are a relatively small share of all workers in Costa Rica (27.4%), which is significantly below the Latin American average of 53.1% (ILO, 2002).

[18]Records include each person's district of residence, with 488 total districts, and also include the voting center which is closest to the citizen's residence, with 2,059 centers in total. We leverage the latter to get a more precise notion of a person's neighborhood. See Méndez and Van Patten (2025) for further details.

[19]We find that the average number of first-degree, second-degree, and third-degree relatives is 6 (median 5), 8 (median 7), and 14 (median 11), respectively.

Ministry of Finance for the universe of formal firms. The data contains typical balance sheet variables since SINPE's inception, and includes details on each firm's sector and location. Summary statistics on each type of network are reported in Table G1.

7.2 From Model to Data

As described in the previous section, we obtained (i) transaction-level data including information on the senders and receivers who took part in each transaction since the app's inception, and (ii) individual-level data on networks from official sources. Further, crucially, we can link identifiers in (i) and (ii). We leverage this substantial data effort to construct measures of networks (N) for *each* individual and to obtain individual-level measures of adoption at the extensive and intensive margins. Figure G3 shows the diffusion path of the technology for the median network.[20]

Our baseline analysis focuses on networks of coworkers—the network for which we can more credibly identify network changes that are plausibly orthogonal to changes in app usage. This will enable us to document evidence of selection (x) and cleanly identify θ_n, which governs the strength of the strategic complementarities and will be crucial for the policy analysis and the estimation of the optimal subsidy. We will also emphasize changes in the intensive margin of adoption, which can be mapped to our model, as particularly informative for teasing out the role of strategic complementarities relative to other potential drivers, such as learning.

7.2.1 Evidence of Selection at Entry

Through the lens of our model, early adopters—who started using the technology even when the network was small—should be more intense users (with higher x). Consistent with this notion, we document that early adopters have distinct characteristics as compared with users who adopted later. For this exercise, and throughout the entire paper, we classify an individual as an adopter starting from the time when she first used the app. First, as shown in Figure 4, we find that early adopters have a higher average wage as compared with individuals who adopted later (Panel (a)), and are on average more high-skill (Panel (b)).[21] Early adopters are also younger, on average, than later adopters, as shown in Figure G5.

[20] We classify networks (i.e., neighborhoods, families, firms) according to their level of adoption. In particular, we calculate the share of individuals within a network who had adopted SINPE by December 2021, the last period available in our data set. We then compute percentiles of this share across networks to generate a distribution.

[21] We classify an occupation as high-skill if it requires education or training beyond a high-school diploma. The dashed vertical line in each figure denotes the beginning of the pandemic, which just as in Figure G1

Figure 4: Average Wage and Skill at the Time of Adoption

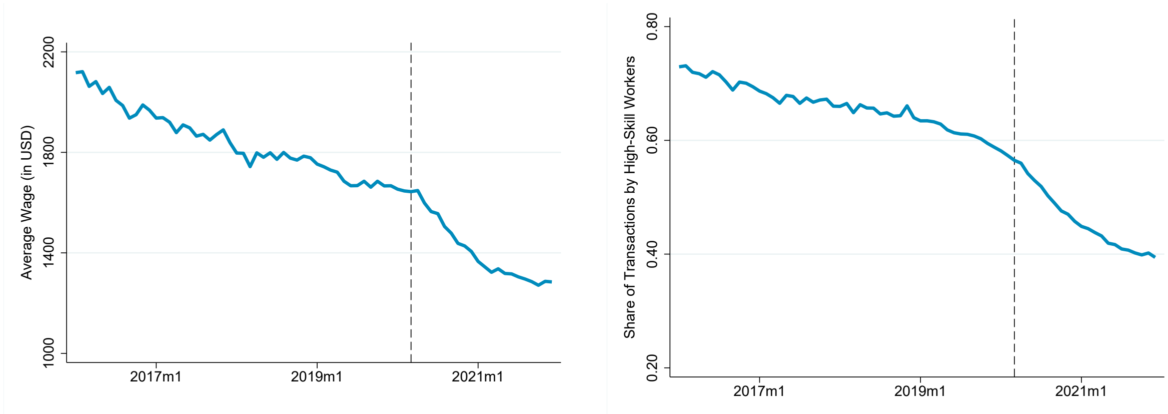

(a) Wage (b) High Skill

Notes: Panel (a) shows the cross-sectional distribution of SINPE users' monthly wages in USD. Panel (b) shows the cross-sectional distribution of SINPE users' skills. High skill users are those that are in an occupation that requires more than a high school degree. Both panels show averages weighted by the number of transactions of each user. Both figures include a vertical dashed line to mark the start of the COVID-19 pandemic (March 2020).

Second, we can more closely bring the model to the data by interpreting the flow benefit of agents who adopt the technology as being proportional to how intensively they use SINPE. Suppose SINPE users choose the intensity with which they use the app. Specifically, suppose ξ_t is the probability of a transaction per unit of time, maximizing the following expression:

$$\xi_t^*(x_t, N_t) = \arg\max_{\xi_t} \frac{1+p}{p}\left[\beta(x_t, N_t)\xi_t - \frac{\xi_t^{1+p}}{1+p}\right],$$

where $p > 0$ so that the problem is convex and $\beta(x_t, N_t) > 0$. The first order condition describes the optimal intensity with which the technology is used: $\xi_t^*(x_t, N_t) = \beta(x_t, N_t)^{1/p}$, and we can choose the function $\beta(x_t, N_t)$ such that the indirect utility function gives the specified flow benefit, i.e:

$$[\theta_0 + \theta_n N_t]x_t = \max_{\xi_t} \frac{1+p}{p}\left[\beta(x_t, N_t)\xi_t - \frac{\xi_t^{1+p}}{1+p}\right] \text{ for all } x_t \in [0, U] \text{ and } N_t \in [0, 1]. \quad (29)$$

The solution is given by $\beta(x_t, N_t) = [(\theta_0 + \theta_n N_t)\, x_t]^{\frac{p}{p+1}}$; combining this expression with the first-order condition and taking logs with obtain:

$$\ln \xi_t^* = \frac{1}{1+p} \ln\left[(\theta_0 + \theta_n N_t)\right] + \frac{1}{1+p} \ln x_t. \quad (30)$$

did not have a major impact on overall trends.

Given the discreteness of the number of transactions in the data, ξ_t^* is interpreted as the mean of a Poisson distribution; transactions each period are drawn from a Poisson probability distribution with mean ξ_t^* (i.e., $\mathtt{T}_t \sim \text{Poisson}(\xi_t)$). Thus, if we were to remove the network×time variation from the logarithm of the number of transactions, then they would proxy for $\ln x_t$, as through the lens of the model only the idiosyncratic variation would remain. The model also predicts that users with a higher x would adopt the technology earlier. Thus, we can obtain a relation between intensity of usage ($\mathtt{T}_{it}^n$) and the share of user i's network who had adopted the technology at the time *when she first used the app* ($N_{i,entry}^n$):

$$\ln \mathtt{T}_{it}^n = \gamma + \zeta N_{i,entry}^n + \lambda_t^n + \nu_{it}^n,$$

where $n \in$\{neighbors, coworkers, relatives\} and $\mathtt{T}_{it}^n$ is defined as number of transactions of user i each month t. Our model predicts that $\zeta < 0$, as users who adopted the app ("entered") when the network was smaller should have a higher idiosyncratic taste for the app and use it more intensively—note that the inclusion of the network-time fixed effect, λ_t^n, prevents this relationship from being mechanical.

We estimate $\hat{\zeta}$ to be -2.7 when defining a network as a neighborhood. This relationship is shown in Column (1) of Table 1, and while suggestive, points to the presence of selection at entry. The relation is also robust to defining networks using coworkers and relatives, as shown in Columns (2) and (3) in Table 1. The relation also holds if, instead of the total number of transactions, we consider the value of transactions as our dependent variable, as reported in Table G3.

Table 1: Number of Transactions and Size of Network at Entry

Dependent variable: Number of Transactions (IHS)

	(1)	(2)	(3)
Size of Coworkers' Network at Entry	-1.300***		
	(0.043)		
Size of Neighbors' Network at Entry		-2.730***	
		(0.025)	
Size of Family Network at Entry			-1.181***
			(0.006)
Observations	16,138,736	34,409,818	14,700,288
Network×Time/Cohort FE	Yes	Yes	Yes
Adjusted R-squared	0.304	0.234	0.199

Notes: The dependent variable in this estimation is the number of transactions each month for each user transformed using the inverse hyperbolic sine function. Coefficients describe the effect of increasing the share of an individual's network who had adopted the app at the time when she used it for the first time. All regressions control for network size (in levels) and use data from May 2015, when the technology launched, to December 2021. Standard errors, clustered by network, are in parenthesis.

7.2.2 Estimating the Strength of the Strategic Complementarities

The core idea behind strategic complementarities is that usage benefits increase with the size of a user's network. Recall the expression in equation (30). Under this interpretation of the model, the intensity with which the application is used, which is observable in the data (e.g., number or value of transactions), is proportional in logs to the flow benefit of adopting the application as described in the model. After taking the first order Taylor expansion of $\ln(\theta_0 + \theta_n N_t)$ around N^* and plugging it into equation (30), we obtain:

$$\ln \mathtt{T}_t \approx \frac{1}{1+p}\Big(\ln(\theta_0 + \theta_n N^*) + \frac{\theta_n (N_t - N^*)}{\theta_0 + \theta_n N^*} + \ln x_t \Big). \tag{31}$$

Moreover, taking first differences, it follows that:

$$\Delta \ln \mathtt{T}_t = \beta \Delta N_t + \nu_t, \tag{32}$$

where $\beta \equiv \frac{1}{1+p}\frac{\vartheta}{1+\vartheta N^*}$, $\vartheta \equiv \frac{\theta_n}{\theta_0}$, and $\nu_t \equiv \frac{1}{1+p}\Delta \ln x_t$. Further, if $p \approx 0$, then $\vartheta = \frac{\beta}{1-N^*\beta}$. Thus, throughout all the tables in this section, we can evaluate N^* at its mean value to recover ϑ from each β; these are our coefficients of interest since strategic complementarities in the adoption of the technology exist if $\beta > 0 \iff \vartheta > 0 \iff \theta_0 > 0$ and $\theta_n > 0$. Note that equation (32) is in *differences*, therefore, *any individual or network characteristics which are time invariant will cancel out.*

With these expressions, one can first naively run an OLS specification. We do so in Appendix G.4 and find a significant correlation between the intensity of app usage and the share of individuals in the user's network who have adopted it. This correlation remains robust across various network definitions, usage intensity measures, and specifications. Then, we show that the impact of network size on usage intensity persists even after employing a leave-one-out instrument to address potential endogeneity concerns and measurement errors. Additionally, this relationship is unaffected when accounting for selection through a balanced panel of adopters. However, to quantify the model, one ultimately needs to take a stand on the causal impact of changes in the number of adopters; we do so by focusing on mass layoffs.

Usage After a Mass Layoff (Intensive Margin of Adoption). This strategy focuses on the network of coworkers and implements both (i) a mover design, where we follow workers displaced during mass layoffs to examine the effect of network changes on the intensive and extensive margins of adoption and (ii) an analysis of stayers, in which we instead focus on workers who remained at a firm after a mass layoff.[22] The main hypothesis of the movers exercise is that workers, who were displaced during a mass layoff and who ended up at firms where a larger share of colleagues had SINPE (larger N), have more incentives to use the app via the effect of strategic complementarities. Similarly, the idea behind the analysis of stayers is that workers who remain at a firm that, for instance, laid off most of its SINPE-using employees (smaller N), have now less incentives to use the app.

We first analyze the impact of a mass layoff on movers' usage. To do so, we focus on workers who were fired during a mass layoff and consider only displaced workers *who had already adopted and had used SINPE at least once by the time they were fired.* We then examine how the intensity with which they use the app changes depending on the change in the share of coworkers who had SINPE at their old and new firm. As explained before, it is possible to derive the relationship in equation (58) from our theoretical model, which speaks to the technology's usage intensity. Thus, we consider:

$$\begin{aligned}\Delta \ln \mathtt{T}_i =& \alpha + \zeta \Delta N_i^{coworkers} + \gamma \Delta \ln wage_i + \psi \Delta \ln size_i + \varphi \, date\,hired_i + \\ & \omega \Delta Covid_i + \delta \lambda_{ic} + \nu \ln \sum_{t=0}^{move} \mathtt{T}_{ti} + \nu \sum_{t=0}^{move} (\ln \mathtt{T}_{t,\text{ new firm}} - \ln \mathtt{T}_{t,\text{ old firm}}) + \epsilon_i, \end{aligned} \tag{33}$$

where $\Delta \ln \mathtt{T}_i$ refers to the change in monthly intensity with which individual i used SINPE within 6 months *after* arriving at her new firm compared with 6 months *before* being fired; $\Delta N_i^{coworkers}$ is the change between the share of coworkers who had adopted at the old and the new employer; $\Delta \ln wage_i$ corresponds with the change in the average wage (in logs) across 6 months before the layoff and after the rehiring; $\Delta \ln size_i$ is the change in the number of workers at each firm; $date\,hired_i$ is a time fixed effect corresponding with the month in which individual i was hired by the new firm; $\Delta Covid_i$ controls for the change in the cumulative COVID-19 cases (transformed using the inverse hyperbolic sine function) in the individual's neighborhood across the 6 months before the layoff and after the rehiring; λ_{ic} controls for cohort (i.e., the date when individual i adopted SINPE); $\ln \sum_{t=0}^{move} \mathtt{T}_i$ is the sum of all historical transactions made by agent i since she adopted the app, and $\sum_{t=0}^{move} (\ln \mathtt{T}_{t,\text{ new firm}} - \ln \mathtt{T}_{t,\text{ old firm}})$ is the difference in the (log) historical transactions made by workers at the new firm and the old firm up until the move occurred, which aims to control for factors—other than strategic complementarities—which might facilitate adoption at the new vs. the old firm.[23]

[22] To define a mass layoff, we follow Davis and Von Wachter (2011) and identify establishments with at least 50 workers that contracted their monthly employment by at least 30% *and* had a stable workforce before this episode and did not recover in the following 12 months. Given we also analyze stayers, we implement a few additional refinements. Details are provided in Appendix G.5.1.

[23] Results are robust to also including a dyadic interaction controlling for industry before and after the

This is our preferred specification for several reasons. First, the results are likely not driven by learning about the app since (i) workers had already adopted the app when they were fired—and we define "adoption" as making at least one transaction—so they were at least aware of the app's existence and had used it before; (ii) we control for tenure in the app (i.e., the cohort when the user adopted) and for the historical number of transactions in the app, which as shown before correlate with observables like age, skill, and wage. These variables aid in controlling for characteristics that are particularly relevant for intensity of usage and are also useful to addressing learning to better use the app after adopting. Second, of course, the choice of the new firm after a mass layoff is not exogenous, but this does not pose a measurement problem as long as sorting is not (both): (i) stronger after a mass layoff—note that there is no reason why this might be the case, especially as results hold even when we focus on job-to-job transitions, where workers had little time to find a new job after being fired exogenously—and (ii) not controlled for by the cohort of the mover, which proxies for her idiosyncratic characteristics, and difference in the historical transactions at the new vs. the old firm. The latter control, in particular, helps us rule out stories where, for instance, workers select into firms where people use the app more intensively for reasons other than strategic complementarities (like demographics or the internet speed at the firm).

Table 2: Intensity of Usage and Changes in Coworkers' Network After a Mass Layoff

Dependent Variable: Δ Number of transactions (IHS)

	(a) Movers			(b) Stayers		
	(1)	(2)	(3)	(4)	(5)	(6)
$\Delta N_i^{coworkers}$	2.646***	1.406***	1.283***	3.284***	0.952**	0.971**
	(0.203)	(0.268)	(0.294)	(0.237)	(0.443)	(0.435)
$\Delta \ln wage_i$		0.383***	0.385***		0.203**	0.132
		(0.070)	(0.077)		(0.087)	(0.103)
$\Delta Covid_i$		0.168**	0.167***		-0.010	-0.012
		(0.027)	(0.032)		(0.025)	(0.024)
Observations	917	917	917	2,236	2,236	2,236
Time FE	No	Yes	Yes	No	Yes	Yes
Cohort FE/Historical T	No	No	Yes	No	No	Yes
Adjusted R-squared	0.153	0.244	0.262	0.093	0.122	0.184

Notes: The unit of observation is the individual. We run regressions using data on mass layoffs which occurred between May 2015, when the technology was introduced, until December 2021. While time and cohort fixed-effects' inclusion varies across columns, all other independent variables in equation (33) are present across columns. Standard errors are in parentheses.

Panel (a) of Table 2 displays our results using the number of transactions per user as our

move.

dependent variable. Changes in the intensity of usage depend positively and significantly on the change in the share of adopters at the old and new firm. Panel (a1) of Figure 5 displays the marginal effect of these network changes following the specification described by Column (2) of Table 2. As this panel shows, not only is the relationship between usage and network changes positive, but also whenever a worker moves to a firm with a lower adoption rate, her usage decreases (i.e., the change on the vertical axis is negative), a relationship that would be hard to reconcile with a pure learning story.[24]

Column (3) controls for cohort, i.e., date of adoption, which aims to mitigate any effect of more experienced users behaving differently than beginners. Column (3) also controls for the total historical transactions made, which in a similar spirit as cohort, intends to mitigate any effect resulting from learning how to use the app from others. Interestingly, as compared with Column (3), adding these controls does not change the coefficient of interest almost at all. This result aligns with the following intuition: at the intensive margin—once users have already adopted and used the app—a learning story is less plausible, as reflected by ζ not changing after controlling for cohort and historical usage.

The analysis can be taken to an even more detailed level if, instead of considering all transactions in the left-hand-side variable, we focus only on those which had a coworker as a counterpart. This subsample allows us to better identify changes in usage intensity which are a direct consequence of the arguably exogenous changes in the network of coworkers. Reassuringly, results are remarkably similar to those using all transactions, as shown in Panel (a2) of Figure 5.

A similar analysis can be conducted based on *stayers*. Namely, we focus on workers who remain at a firm after it experienced a mass layoff. Their change in N will therefore depend on how the composition of SINPE adopters changed after the mass layoff. We then consider a regression similar to equation (33), except for the last regressor which would be zero in this case.[25] Results based on stayers are reported in Panel (b) of Table 2 and Panel of Figure 5. Remarkably, although the movers design is based on a very different sample than the analysis based on stayers, the estimated coefficients in our preferred specifications, in columns (3) and (6) of Table 2, are statistically equal.

Adoption After a Mass Layoff. Lastly, we analyze changes in the extensive margin of adoption. For movers, we consider the change in the probability of adoption for displaced

[24] The marginal effect considering the value of transactions as dependent variable, as opposed to the number of transactions, is reported in Figure G7. We also report the distribution of network changes in Figure G8 and the absence of pretrends in Figure G9.

[25] An additional control equal to the change in the average wage at the workers' firm delivers statistically equal results, both for the intensive and extensive margin analyses.

Figure 5: Marginal Effect of Network Changes on Usage Intensity

(a) Movers

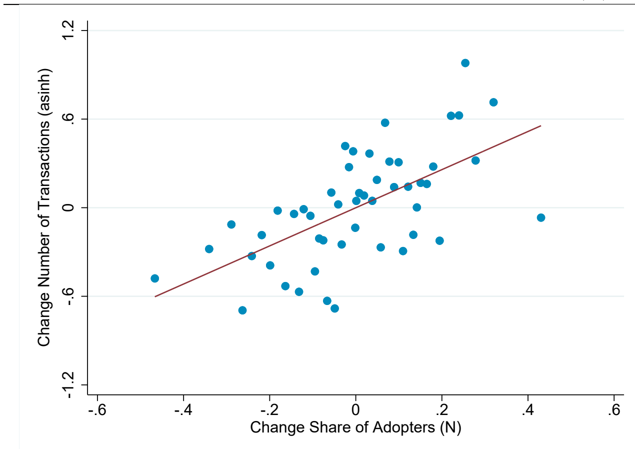

(a1) All transactions

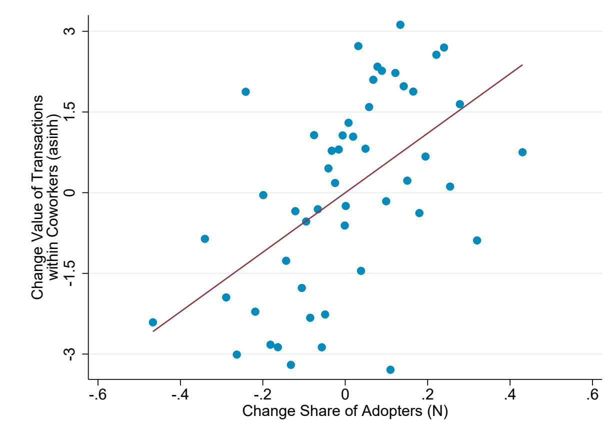

(a2) Transactions with coworkers only

(b) Stayers

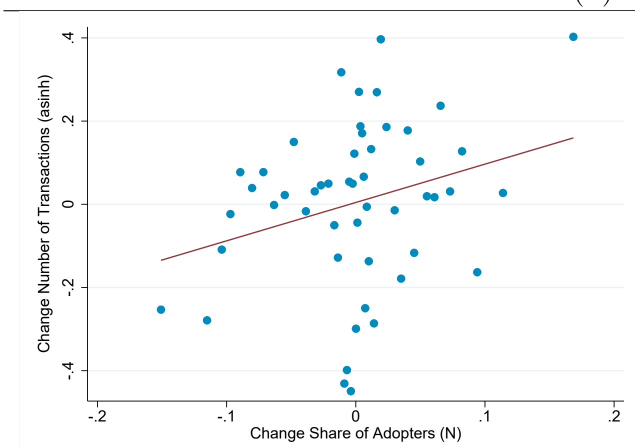

(b1) All transactions

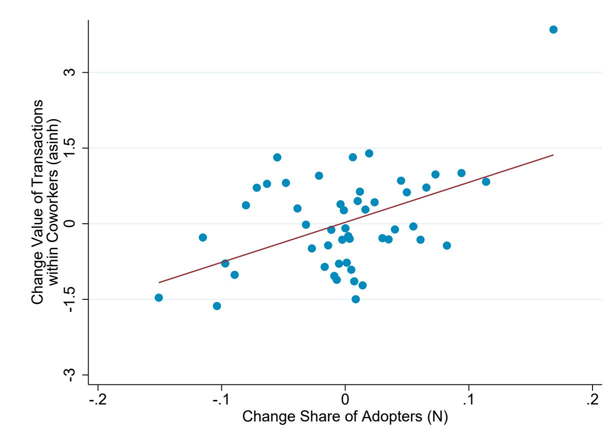

(b2) Transactions with coworkers only

Notes: Panel (a1) plots the marginal effect of $\Delta N_i^{coworkers}$ in the specification described by Column (3) of Table 2, while Panel (b1) plots the marginal effect of $\Delta N_i^{coworkers}$ in the specification described by Column (6) of Table 2. Bars denote 95% confidence intervals. The dependent variable in this estimation is the number of transactions (transformed using the inverse hyperbolic sine function) on each period for each user. Panels (a2) and (b2) are similar, but differ as the dependent variable in these estimations is the number of transactions *which have a coworker as a counterpart* (transformed using the inverse hyperbolic sine function) on each period for each user. Results are robust to winsorizing the top and bottom 5th percentiles of the distribution of network changes.

workers *who had not adopted the app by the time they were rehired*, and how it depends on the change in the share of coworkers who had SINPE at their old and new firm. We consider:

$$
\begin{aligned}
Adopt_i =& \alpha + \zeta \Delta N_i^{coworkers} + \gamma \Delta \ln wage_i + \psi \Delta \ln size_i + \varphi \, date\, hired_i + \\
& \omega \Delta Covid_i + \nu \sum_{t=0}^{move} \left(\ln \mathtt{T}_{t,\text{ new firm}} - \ln \mathtt{T}_{t,\text{ old firm}} \right) + \epsilon_i,
\end{aligned} \tag{34}
$$

where $Adopt_i$ equals one if individual i adopted SINPE within 6 months after arriving at her new firm, and zero otherwise. Other variables are defined in the same way as in equation (33). For stayers, we instead consider the probability of adoption for workers who were not fired by a firm which underwent a mass layoff and how it depends on the change in the composition of workers who had SINPE, before and after the mass layoff took place. We then use a regression similar to equation (34), except for the last regressor which would be zero.

Panels (a1) and (b1) of Figure 6 estimate equation (34) using a logit model. The marginal effects of changes in network adoption are reported in brackets. The analysis of movers in panel (a1) consistently finds that workers who, after a mass layoff, were hired by firms where the rate of SINPE adoption was higher than their previous employer's are more likely to adopt SINPE than their counterparts who moved to firms where the change in their coworkers' rate of adoption was smaller. Reassuringly, panel (b1) also finds that workers who experienced an increase in the share of adopters among their peers were more likely to adopt SINPE themselves. The marginal effect of $\Delta N_i^{coworkers}$, under the specification described by Column (3) in each table, is shown in panels (a2) and (b2). These marginal effects are monotonous and, as expected, are present only when the change in the share of adopters is positive, regardless of the subsample considered.

Figure 6: Adoption Probability and Changes in Coworkers' Network After a Mass Layoff

(a) Movers

	(1)	(2)	(3)
$\Delta N_i^{coworkers}$	8.321***	5.197***	5.018***
	(0.161)	(0.207)	(0.208)
	[0.450]	[0.376]	[0.350]
$\Delta \ln wage_i$		0.076	0.040
		(0.055)	(0.056)
$\Delta Covid_i$			0.097***
			(0.026)
Observations	10,176	8,035	8,035
Time/Cohort FE	No	Yes	Yes
Pseudo R2	0.507	0.529	0.530

(a1) Changes in Adoption Probability

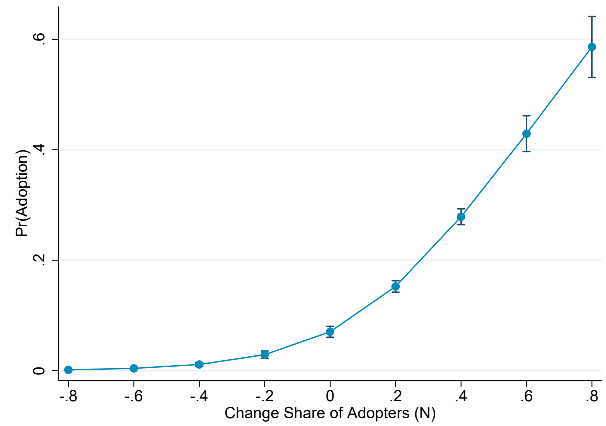

(a2) Marginal Effect of Network Changes

(b) Stayers

	(1)	(2)	(3)
$\Delta N_i^{coworkers}$	11.249***	6.489***	6.418***
	(0.172)	(0.356)	(0.357)
	[0.716]	[0.442]	[0.437]
$\Delta \ln wage_i$		-0.104	-0.098
		(0.089)	(0.089)
$\Delta Covid_i$			0.037*
			(0.020)
Observations	24,329	23,005	23,005
Time/Cohort FE	No	Yes	Yes
Pseudo R2	0.330	0.356	0.357

(b1) Changes in Adoption Probability

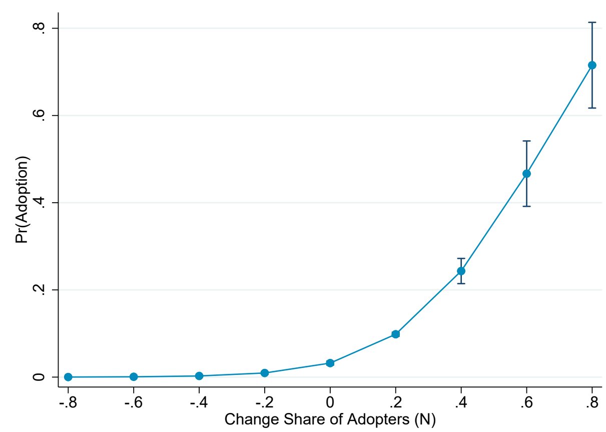

(b2) Marginal Effect of Network Changes

Notes: Panels (a1) and (b1): The unit of observation is the individual. We run regressions using data on mass layoffs that occurred between May 2015, when the technology was introduced, and December 2021. Standard errors are in parentheses. Marginal effects for the main variable of interest are reported in brackets. Panels (a2) and (b2): The figures plot the marginal effect of $\Delta N_i^{coworkers}$ in the specification described by column (3) of panels (a1) and (b1), respectively, in this figure. Vertical bars denote 95% confidence intervals.

8 Quantitative Performance and Optimal Subsidy

In this section, we calibrate the version of the model that incorporates learning, described in Section 6, and evaluate its performance relative to SINPE data. The model with only strategic complementarities assumes that all individuals are informed about the technology at all times. However, according to the 2017 Survey of Payment Methods conducted by the Central Bank of Costa Rica, only about 4% of adults reported knowing about SINPE Móvil more than two years after its launch. The model that incorporates learning helps align the model with this fact and, as a result, with the smooth and relatively flat path of $N(t)$ during the first few years after launch shown in Figure 7. In what follows, we describe our calibration procedure in detail.

Calibration. We interpret the flow benefit of agents who adopt the technology as being proportional to how many transactions they conduct, and assuming a convex adjustment cost (i.e., $p > 0$). U can be normalized without loss of generality (we use the normalization $U = 1$), so the problem features seven independent parameters: $\nu, r, \theta_n, \theta_0, \sigma$, p, and c. The model with learning has an additional parameter, β_0, and an initial condition for the informed population, $I(0)$.

The parameters ν, r, β_0 and are calibrated externally. We set ν to 0.0278 to match the rate at which agents stop using SINPE; namely, the average fraction of agents in 2019-2021 who had adopted SINPE but did not conduct a single transaction in the app within a year. We use the last three years of the data, when the adoption rate is higher, to focus on periods closer to a stationary equilibrium. We set the discount factor r to be consistent with a 5 percent annual interest rate. This value is a lower bound for r, which can admit higher values if we assume agents expect new technologies to arrive in the future and replace SINPE. The values of ν and r imply $\rho = r + \nu = 0.0778$. Lastly, we calibrate β_0 and $I(0)$ jointly to match two observed moments of the information diffusion process.[26] We use the expression for the diffusion path (see Proposition 11) and set $I(t_1) = 0.0409$ at $t_1 = 31$ months (based on the 2017 Survey of Payment Methods) and $I(t_2) = 0.9617$ at $t_2 = 117$ months (based on a follow-up survey conducted by the Central Bank in 2025). This yields $\beta_0 \approx 1.0717$ and $I(0) \approx 0.00287$, implying that about 0.29% of workers were informed about SINPE Móvil at the time of its launch.

The parameters $\theta_n, \theta_0, \sigma, p$ are calibrated using simulated methods of moments (SMM). Intuitively, we aim to choose parameters that make the model consistent with the distribution of transactions in the data and the mass layoff exercise. To achieve this, in the data, we focus on workers at firms active from 2019 to 2021 with more than 5 employees. We take advantage of having closed form solutions for the steady state. Thus, we concentrate on firms close to a stationary equilibrium, specifically those whose N (fraction of employees with the app) changed by less than 0.1 percentage points in 2021. We then compute moments from the empirical distribution of transactions over the years 2020-2021 and simulate the model, replicating the same characteristics as our empirical sample. In addition, we simulate a

[26]Diffusion models typically impose a small seed value for the unobserved initial condition. Observing two moments of the information diffusion path allows us to identify it jointly with the transmission rate. Standard SIR applications assume $I(0) = \varepsilon N$ for small ε (e.g., Acemoglu et al., 2021; Alvarez et al., 2021).

sample of firms that replicates the characteristics of those subject to a mass layoff. We do this to run the same estimation, presented in Section 7.2.2, in the simulated data to obtain information on the parameters governing the strength of the strategic complementarities. We then choose the parameters that minimize the distance between the moments in the data and the model. We provide more details of our strategy below.

Simulation. We begin by simulating the model for a monthly panel of agents. Our simulation takes as given the values of ν, r, and β_0, since they are calibrated externally, and $N_{ss} = 0.90$, which is obtained from our sample of firms close to a stationary equilibrium. Initial conditions $x(0)$ are drawn from the stationary distribution of adopters. To find this distribution, we first find $\bar{x}_{ss}$ using the following equation:

$$N_{ss} = (1 - \tfrac{\nu}{\beta_0}) \left[1 - \bar{x}_{ss} \left(1 - \frac{\tanh(\gamma \bar{x}_{ss})}{\gamma \bar{x}_{ss}}\right)\right].$$

Then, given $\bar{x}_{ss}$, we find the distribution of adopters using the stationary distribution of non-adopters:

$$\tilde{m}(x) = (1 - \tfrac{\nu}{\beta_0}) \left(1 - \frac{\cosh(\gamma x)}{\cosh(\gamma \bar{x}_{ss})}\right) \qquad \text{where } \gamma = \sqrt{2\nu}/\sigma$$

using that $N_{ss} = I_{ss} - M_{ss}$ and $I_{ss} = (1 - \frac{\nu}{\beta_0})$. We simulate a panel of 5,000 individuals.[27] In the simulation, agents die at rate ν and they become inactive in the application just as in the data. The process of x follows a Brownian motion, independent across agents, with variance per unit of time σ, no drift, and reflecting barriers at $x = 0$ and $x = 1$. Since x is unobserved and what is observed are transactions, as before, we interpret the flow benefit of agents who adopt the technology as being proportional to how intensively they use SINPE. Thus, we compute: $\xi_t = [\theta_0(1 + \vartheta N_{ss})x_t]^{\frac{1}{1+p}}$, where $\vartheta \equiv \frac{\theta_n}{\theta_0}$, to find the number of transactions $\mathtt{T}_t$ by drawing them from a Poisson probability distribution $\mathtt{T}_t \sim \text{Poisson}(\xi_t)$.

Mass Layoff. We also simulate a panel of workers at firms with the same characteristics as those experiencing mass layoffs in the data. Specifically, as presented in Table G12, we simulate a sample of 292 firms with 94 employees each. We focus on workers who remain at a firm after it has experienced a mass layoff (i.e., stayers).[28] We then examine how the intensity with which they use the app changes depending on the change in the share of coworkers who had SINPE after a mass layoff. The change in N depends on how the composition of adopters

[27] Our estimates are not sensitive to simulating a larger sample of users.

[28] Table 2 shows that the estimated impact of a mass layoff on usage is statistically equal for movers and stayers.

changes after the mass layoff, which involves randomly choosing and removing a fraction of workers from each firm undergoing a mass layoff. We choose the magnitude of these mass layoffs to match the average size of these events in the data (i.e., 57%). We then run the same regression that is implemented in Table 2. First, we calculate the number of transactions before and after the mass layoff event. Then, we regress the change in monthly transactions within six months of the mass layoff event on the change in the share of coworkers who had adopted the app before and after the event. The estimated coefficient is a moment that we target in our calibration

Calibrated Moments. We target the following five moments: the mean number of transactions, the median number of transactions, the absolute value of changes in transactions, the coefficient of the mass layoffs regression, and the autocorrelation of the number of transactions. As done throughout the paper, all the targeted data moments are calculated after controlling for COVID-19 cases. Parameters $\theta_n, \theta_0, \sigma$, and p are chosen to minimize the sum of the norms of the percent deviations of simulated moments from target moments.[29] Table 3 reports the empirical and simulated moments.[30]

Table 3: Moments: Distribution of Transactions

Parameter	Value	Std. Dev.	Moment	Data	Model
σ	0.032	0.002	Mean Number of Transactions	6.88	6.84
θ_0	26.32	4.726	Median Number of Transactions	6.08	6.64
p	0.0059	0.0009	Absolute Value Changes in Transactions	3.48	2.76
$\vartheta \equiv \frac{\theta_n}{\theta_0}$	5.722	0.343	Coefficient Mass Layoffs Regression	0.97	0.96
			Autocorrelation of Transactions	0.97	0.95

Intuitively, the mean and median number of transactions provide information about θ_0 and p, as shown by equation (30). The dispersion in the changes of transactions and the autocorrelation of transaction provide relevant information to pin down σ; a lower variance decreases the absolute value of the changes in transactions but increases the autocorrelation coefficient. Lastly, equation (32) shows that the coefficient of the mass layoffs regression informs the estimation of θ_n.[31] Targeting this moment allows us to leverage the rich vari-

[29] This is, $\min \sum_i^5 \omega(i) \frac{|\text{Model}(i) - \text{Data}(i)|}{\text{Data}(i)}$, where Model$(i)$ is a simulated i-th moment and Data(i) is a target value of i-th moment. We assign half of the weight to each of the mean and median of transactions since they provide similar information for the calibration. We assign twice as much weight to the coefficient of the mass layoff regression, as this moment provides information about the strength of the strategic complementarities.

[30] We simulate the model 200 times and use the average values of the moments from the simulated data. In the model and the data, we calculate the autocorrelation of the average transactions over two years to minimize the impact of measurement error in the autocorrelation coefficient. The standard deviation of the parameters is obtained from the the SMM variance-covariance matrix, which is obtained from calculating the derivative of the criterion function with respect to each parameter.

[31] The learning model in Appendix E cannot capture the patterns observed after mass layoffs, as it features

ation across networks to inform the model estimation. Importantly, by running the same regression in both the data and the model, we do not rely on approximating the relationship between the change in transactions and the change in the share of users who have adopted the technology around the stationary equilibrium to obtain information about ϑ, as done in equation (32). Overall, Table 3 shows that the model is quantitatively consistent with the empirical distribution of transactions.[32]

Cost of Adopting. Lastly, the adoption cost, c, can be obtained from the solution of the stationary problem for adoption, given a value of $\bar{x}_{ss}$ and the parameters $\theta_n, \theta_0, \sigma$. In particular, we use the following equation:[33]

$$c = \bar{\theta}_{ss}\left[\bar{x}_{ss} + \bar{A}_1 e^{\eta\bar{x}_{ss}} + \bar{A}_2 e^{-\eta\bar{x}_{ss}} - \left(\frac{1}{\eta}\frac{e^{\eta\bar{x}_{ss}} + e^{-\eta\bar{x}_{ss}}}{e^{\eta\bar{x}_{ss}} - e^{-\eta\bar{x}_{ss}}}\right)\left(1 + \eta\bar{A}_1 e^{\eta\bar{x}_{ss}} - \eta\bar{A}_2 e^{-\eta\bar{x}_{ss}}\right)\right]$$

where $\eta = \sqrt{2\rho}/\sigma$, $\bar{A}_1 = \frac{1}{\eta}\frac{\left(1-e^{-\eta}\right)}{(e^{-\eta}-e^{\eta})}$, $\bar{A}_2 = \frac{1}{\eta}\frac{(1-e^{\eta})}{(e^{-\eta}-e^{\eta})}$ and $\bar{\theta}_{ss} \equiv \frac{\theta_0+\theta_n N_{ss}}{\rho}$. The calibrated parameters imply an adoption cost of $c = 10.54 \cdot \theta_0$, which rationalizes the observed steady-state adoption level N_{ss}. Importantly, our procedure does not assume a unique interior steady state—only that the economy is observed at a steady state. The calibration relies entirely on aggregation and the statistical properties of the unobservable state variable x and transactions, without using agents' optimality. As a result, parameters are estimated via the method of moments, and the adoption cost c is then computed using the firm's first-order condition to ensure consistency with the steady state.

Results. Using the estimated parameters, we simulate the dynamic model to obtain the adoption path predicted by the model. Panel (a) of Figure 7 compares the path of adoption in the model and in the data. The solid red line indicates the diffusion of the technology in the median firm and the dashed lines represent the 10^{th} and 90^{th} percentiles after controlling for COVID-19 cases.[34] The figure shows that both the speed and the level of adoption generated by the model are consistent with those in the data. Panel (b) shows the path of $I(t)$, $N(t)$ and $\bar{x}(t)$. The path of $I(t)$ shows that most people are informed about the technology within the first 7 years; in the stationary distribution, approximately 97.4% (i.e., $I_{ss} = 1 - \frac{\nu}{\beta_0}$) of the population knows about the application and 89% of the workers the median firm adopt the

random diffusion of the technology. Consequently, after adoption, an agent's flow benefit does not depend on the network size N (i.e., $\theta_n = 0$).

[32] A sensitivity analysis of the relevant parameters can be found in Appendix H.

[33] All details on the derivation of this equation can be found in Appendix A.1.

[34] We adjust the adoption path in the data to control for the pandemic. To do so, we estimate the impact of COVID-19 cases on the number of new users and subtract the predicted number of pandemic-driven new users from the cumulative number of users.

application as shown by the path of $N(t)$. Importantly, the declining path of $\bar{x}(t)$ indicates that, consistent with our empirical evidence, the model features selection: agents that benefit the most from the technology adopt first. This contrasts with the model that features only learning, which shows no selection in the adoption of the technology.[35]

Figure 7: Path of Adopters (Short-Run and Long-Run)

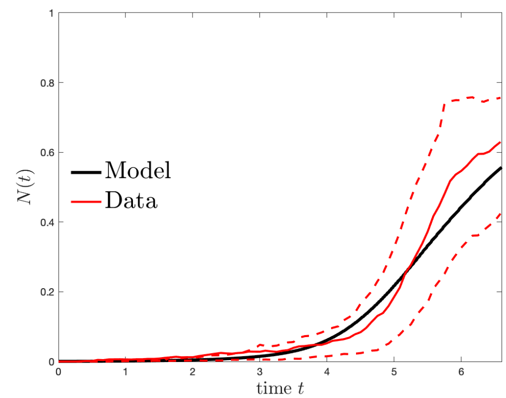

(a) Model vs Data

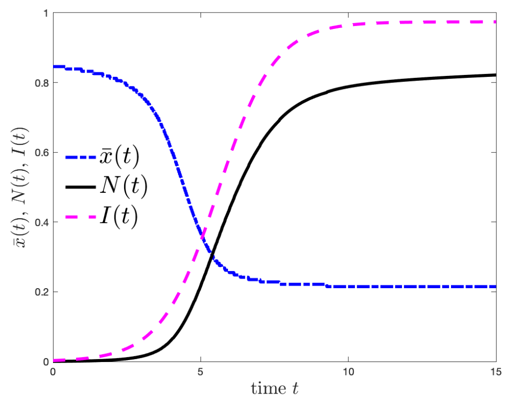

(b) Long-Run Path

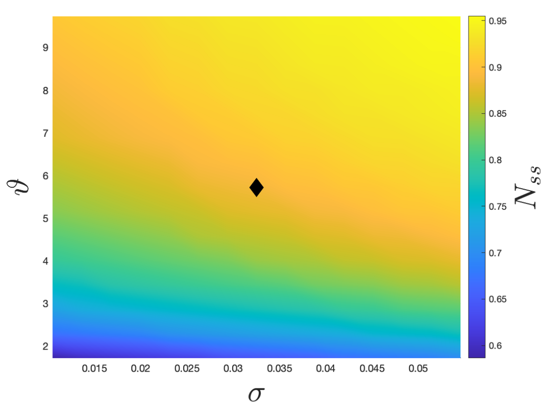

(c) Comparative Statics: N_{ss}

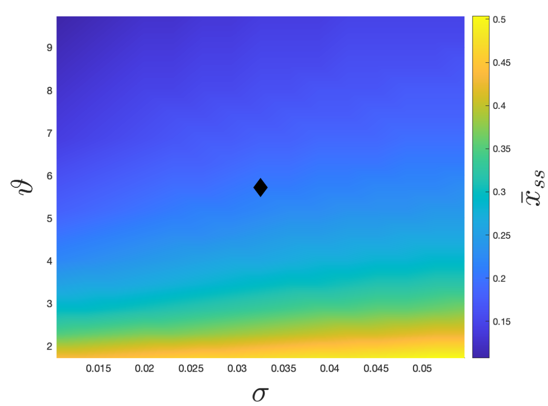

(d) Comparative Statics: $\bar{x}_{ss}$

Notes: Panel (a) compares the path of adopters in the model and in the data. The solid red line shows the patterns of diffusion of the technology in the median firm, where the percentile is calculated in the last period of the sample using the share of individuals that had adopted the technology. The dashed red lines show the 10^{th} and 90^{th} percentiles. Panel (b) shows the share of informed agents, $I(t)$, the share of adopters, $N(t)$, and the levels of $\bar{x}(t)$ predicted by the model under our baseline calibration. Panel (c) and (d) show how N_{ss} and $\bar{x}_{ss}$ change with ϑ and σ, keeping the rest of the parameters constant. The black diamonds indicate the levels of ϑ and σ in our baseline calibration.

Panels (c) and (d) of Figure 7 display the values of N_{ss} and $\bar{x}_{ss}$ in the stationary equilibrium as we vary ϑ and σ, while holding others constant. These panels illustrate the

[35] Appendix H.3 presents a version of the model without strategic complementarities and only learning (i.e., $\vartheta = 0$). In this case, the path of $\bar{x}(t)$ is completely flat. Figure H7 shows the paths of $N(t)$ and $\bar{x}(t)$ for different speeds of information diffusion; namely, different values of β_0. It shows that selection occurs in the model even when the speed of information diffusion is very high.

comparative statics of the stationary equilibrium derived in Section 3.4. Panel (c) shows how the stationary level of adoption changes with ϑ and σ (a black diamond denotes N_{ss}'s level in the baseline calibration). As ϑ increases, so does the strength of the strategic complementarities, and not surprisingly, N_{ss} increases as ϑ rises. The effect of σ is more subtle and results from two opposing forces. On the one hand, higher σ decreases N_{ss} since agents have a higher option value of waiting to adopt. On the other hand, higher σ increases N_{ss}, since it implies a smaller density of non-adopters below $\bar{x}_{ss}$. In our calibration the latter effect dominates and N_{ss} increases with σ. Panel (d) displays a similar exercise for $\bar{x}_{ss}$. It shows that strategic complementarities ϑ play an important role in decreasing the adoption threshold. Moreover, given ϑ, a higher σ increases $\bar{x}_{ss}$

Variation Across Networks. The model is consistent with both high and low adoption networks of firms, each implying a different path of adopters in equilibrium. Specifically, we calibrate the model by targeting moments from individuals at firms whose level of adoption is either above the median (high adoption) or below the median (low adoption).[36] We target the same data moments computed for different samples of workers, specifically those working at firms whose average level of adoption is either above the median, $N_{ss}^{high} = 0.95$, or below the median, $N_{ss}^{low} = 0.73$, and we assume the same coefficient for the mass layoffs regressions in both calibrations. We estimate a higher level of strategic complementarities (i.e., higher ϑ) in networks with high adoption and a higher convexity in the cost of conducting transactions in low adoption networks (i.e., higher p). Panel (a) of Figure 8 shows the path of adopters in the two calibrated networks (high and low adoption) relative to the data, indicating that these calibrated versions of the model are consistent with the 10^{th} and 90^{th} percentiles of adoption in the data. Panel (b) show the path of $\bar{x}(t)$, which indicates the strength of the strategic complementarities in each of the calibrated networks. In the high adoption network, 96% of the population adopts the application. In the low adoption network, only 73% of the population adopts in the stationary equilibrium.

Optimal Subsidy. Panel (a) of Figure 9 shows the optimal adoption path in the model with complementarities (blue line) relative to the high-adoption equilibrium (black line). During the first three years after the launch of the technology, the optimal level of adoption is similar to that of the equilibrium without subsidy. Afterward, the optimal path of adopters from the planning problem is higher. In fact, by the beginning of 2020, it is equal to the total number of informed agents in the economy—over 12 percentage points higher than the levels of adoption observed in the data—and by the end of 2021, it is over 15 percentage points

[36] The details of the calibration can be found in Appendix H.2.

Figure 8: Variation Across Networks: Path of Adopters

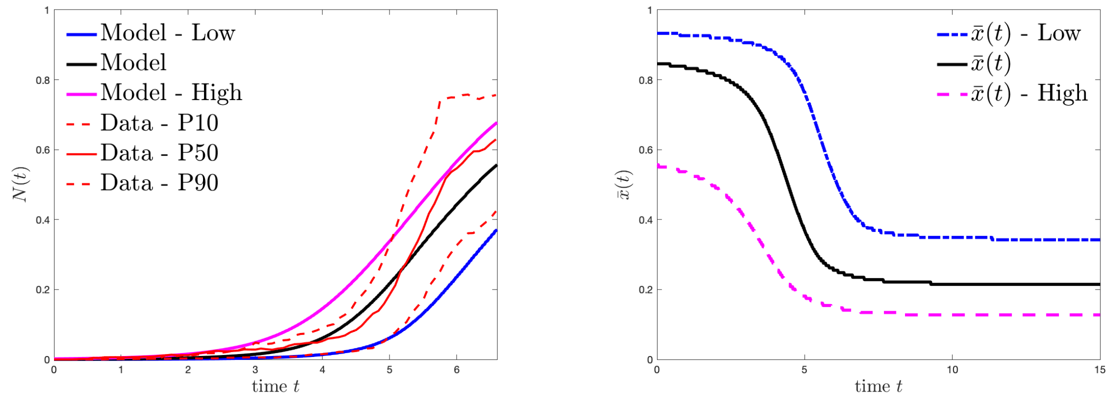

(a) Model vs Data

(b) Path of $\bar{x}(t)$

Notes: Panel (a) compares the path of adopters in the model and in the data. The solid red line shows the diffusion patterns of the technology in the median firm, and the solid black line shows the diffusion patterns in the benchmark calibration of the model. The dashed red lines indicate the 10^{th} and 90^{th} percentiles of adoption in the data. The solid magenta line shows the path of adopters in the model calibrated for high adoption, and the solid blue line shows the path of adopters in the low adoption calibration. Panel (d) shows the levels of $\bar{x}(t)$ under each of the calibrations, respectively.

higher. Panel (b) shows the path of the optimal subsidy.[37] As the share of adopters increases, so does the externality. Thus, the optimal subsidy, which is the same across agents, increases over time. To see why, notice the optimal flow subsidy in equation (28) can be written as

$$\theta_n Z(t) = \theta_n N(t) \times \mathbb{E}(x|\text{adopted}),$$

where the expectation over x is taken over the set of agents that have adopted the technology (see Theorem 3). The first term $\theta_n N$ captures the size of the adoption externality, i.e., the additional benefits for agents that adopt the technology. Thus, the subsidy increases as more agents adopt. Conversely, $\mathbb{E}(x|adopted)$ decreases as more agents adopt, since the marginal adopter has lower idiosyncratic benefits from adopting the technology. Intuitively, the planner internalizes that subsidizing agents with low x also benefits the rest of the agents, even if the subsidy to incentivize these agents to adopt is large. The first component of the optimal subsidy dominates and eventually pushes the economy to universal adoption. The optimal subsidy contrasts with that of a pure learning model, which is constrained efficient and where the optimal subsidy to adopt the technology is zero. Importantly, *the planner is also constrained by the share of people who are informed*; otherwise, while the subsidy would still be increasing and the same across agents, there would be a "jump" in the subsidy level

[37] Figure 9 shows the subsidy $\theta_n Z(t)$ as a ratio of the net flow benefits (i.e., $(\theta_0 + \theta_n N(t))\mathbb{E}(x|adopted)$). In the invariant distribution, the subsidy-to-benefit ratio is approximately 0.84.

Figure 9: Planning Problem: Solution and Optimal Subsidy

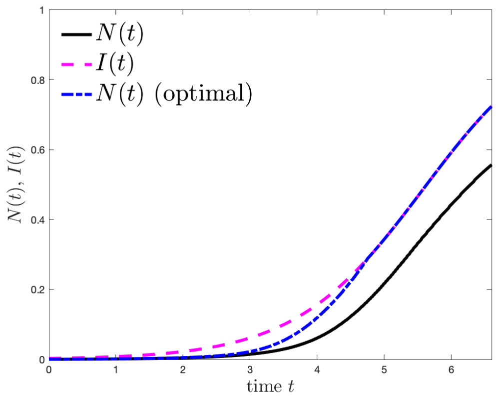

(a) Optimal Path of Adopters

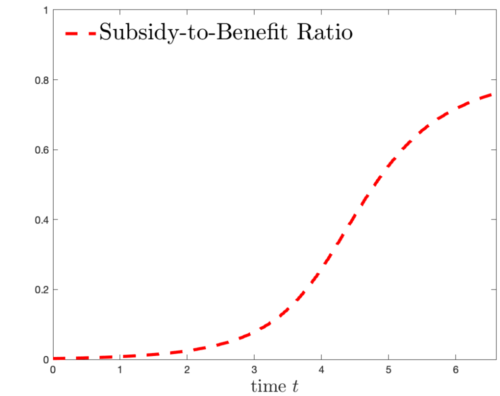

(b) Optimal Subsidy

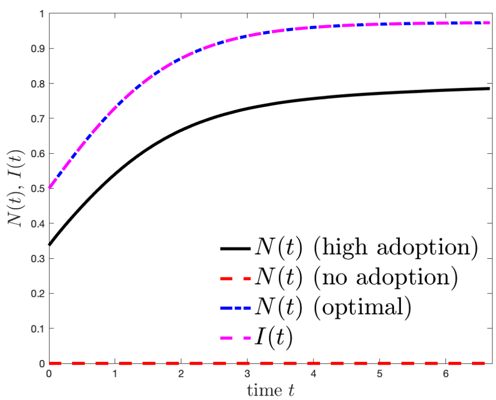

(c) Multiplicity: High Cost

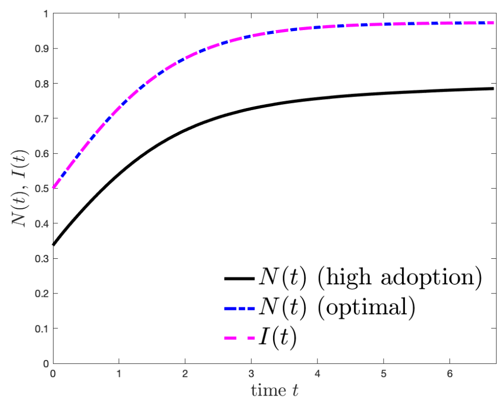

(d) No Multiplicity: Low Cost

Notes: Panel (a) shows the share of informed agents, $I(t)$, the share of adopters, $N(t)$, and the optimal levels of adoption, $N(t)$ (optimal), according to the solution of the planning problem. Panel (b) shows the path of the ratio between the optimal subsidy $\theta_n Z(t)$ and the flow benefit of the average adopter, $Z(t)(\theta_0 + \theta_n N(t))$. Panel (c) shows the share of informed agents, $I(t)$, the share of adopters, $N(t)$, and the optimal levels of adoption, $N(t)$ (optimal), according to the solution of the planning problem for a high adoption cost and 70% of the population informed 7 months after the launch of the technology. Panel (b) shows the same variables for a lower adoption cost and 70% of the population informed 7 months after the initial launch. The initial distribution in both panels is $m_0(t) = 1/U$.

as soon as the application is launched, as depicted in panel (b) of Figure 3.[38]

In Appendix H.2 we estimate the model using variation across different networks. Our findings indicate that the model aligns with both high and low adoption networks of firms, each implying different paths of adopters in equilibrium and different optimal adoption paths in the planning problem. Consistent with our benchmark calibration, all versions of the model show that the optimal subsidy pushes the economy toward universal adoption. Figure H6 shows that only for lower levels of ϑ does the planner prescribe lower adoption levels.

[38] Figure H8 shows the optimal adoption paths and the respective subsidy-to-benefit ratios for different speeds of information diffusion.

Multiplicity. Our model can be used to study economies with higher adoption costs featuring multiple equilibria. We consider an economy with higher adoption cost c and higher fraction of the population informed about the technology at launch. This example is motivated by a recent experience in El Salvador, where 70% of the population knew about a payment app introduced by the government (i.e., Chivo Wallet) 7 months after its initial launch.[39] Panel (c) shows the possible paths of adopters $N(t)$ for this economy. It shows that, when the adoption cost is larger (in this case 15% higher than in Costa Rica), the low-adoption equilibrium where nobody adopts the technology is not ruled out; for the same initial conditions, there is an equilibrium with high adoption and one with no adoption. Panel (d) shows the same paths for a lower adoption cost. Our model allows for the study and quantification of policies that eliminate the no-adoption equilibrium even if the optimal subsidy is not implemented. In this case, a large enough permanent subsidy can lower the adoption cost, solve the coordination failure, and send the economy to the high adoption equilibrium, i.e., from Panel (c) to (d).[40]

9 Conclusion

Understanding the adoption process of a technology and the transition from low to high adoption is challenging, especially in the presence of strategic complementarities. This paper develops a new dynamic model of technology adoption that allows us to model this transition. The model provides a framework to generate gradual adoption through a novel mechanism—waiting for others to adopt—and allows us to derive predictions that can be tested empirically. We solve for the social planner's problem. The planner in our setup controls the entire distribution of adopters across time. The presence of strategic complementarities enriches the problem and allows us to link our results to the "big push" literature, as they imply that small subsidies can lead to large changes in adoption given the multiplicity of equilibria. In our framework, the optimal subsidy increases over time but it is flat across people, thus, easily implementable. The methodology we develop can be useful for a wide set of multidimensional dynamic problems, and can be applied to the study of any technology that features strategic complementarities, learning, or both.

Our application analyzes new electronic methods of payment, which are particularly relevant today and are undergoing a digital transformation. This revolution has been echoed

[39] The app allows users to digitally trade both bitcoin and dollars.

[40] The Salvadorean government did in fact implement a similar subsidy. As an incentive to adopt, citizens who downloaded Chivo Wallet received a $30 bitcoin bonus from the government. Our model suggests that the subsidy was not large enough to rule-out the no-adoption equilibrium given the low levels of adoption of Chivo Wallet reported by Alvarez, Argente and Van Patten (2022).

by a growing interest from monetary authorities to promote and develop digital payment platforms, both in developed and developing countries. Using individual- and transaction-level data on SINPE, a national electronic payment system adopted by a large fraction of the adult population in Costa Rica, along with extensive data on the networks of each user, we document that strategic complementarities play an important role in the adoption of this technology. SINPE also provides a rich environment to calibrate the model, which allows us to estimate the optimal time-varying adoption subsidy and the degree of selection into adoption across time. These results have implications for the launch and implementation of payment technologies with similar features such as CBDCs.

References

Acemoglu, D., Chernozhukov, V., Werning, I., Whinston, M.D., 2021. Optimal Targeted Lockdowns in a Multigroup SIR Model. American Economic Review: Insights 3, 487–502.

Alfaro-Urena, A., Castro-Vincenzi, J., Fanelli, S., Morales, E., 2023. Firm export dynamics in interdependent markets. Technical Report. National Bureau of Economic Research.

Alvarez, F., Argente, D., Lippi, F., 2021. A simple planning problem for covid-19 lock-down, testing, and tracing. American Economic Review: Insights 3, 367–382.

Alvarez, F., Lippi, F., Souganidis, P., 2023a. Price setting with strategic complementarities as a mean field game. Econometrica 91, 2005–2039.

Alvarez, F.E., Argente, D., Lippi, F., Mendez, E., Van Patten, D., 2023b. Strategic Complementarities in a Dynamic Model of Technology Adoption: P2P Digital Payments. Working Paper 31280. National Bureau of Economic Research.

Alvarez, F.E., Argente, D., Van Patten, D., 2022. Are Cryptocurrencies Currencies? Bitcoin as Legal Tender in El Salvador. Technical Report 29968. National Bureau of Economic Research.

Argente, D., Gonzalez Alvarez, P., Méndez, E., Van Patten, D., 2025. Drivers of Digital Payment Adoption: Lessons from Brazil, Costa Rica, and Mexico. Working Paper.

Arkolakis, C., Eckert, F., Shi, R., 2023. Combinatorial discrete choice: A quantitative model of multinational location decisions. Technical Report. National Bureau of Economic Research.

Auclert, A., Rigato, R.D., Rognlie, M., Straub, L., 2022. New Pricing Models, Same Old Phillips Curves? Technical Report. National Bureau of Economic Research.

Bass, F.M., 1969. A New Product Growth for Model Consumer Durables. Management Science 15, 215–227.

Benhabib, J., Perla, J., Tonetti, C., 2021. Reconciling Models of Diffusion and Innovation: A Theory of the Productivity Distribution and Technology Frontier. Econometrica 89, 2261–2301. doi:https://doi.org/10.3982/ECTA15020.

Bigio, S., Mendez, E., Van Patten, D., 2025. A theory of Payments-Chain Crises. Technical Report. National Bureau of Economic Research.

Bilal, A., 2023. Solving heterogeneous agent models with the master equation. Technical Report. National Bureau of Economic Research.

Björkegren, D., 2018. The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda. The Review of Economic Studies 86, 1033–1060.

Buera, F.J., Hopenhayn, H., Shin, Y., Trachter, N., 2021. Big Push in Distorted Economies. Working Paper 28561. National Bureau of Economic Research.

Calvo, G.A., 1983. Staggered prices in a utility-maximizing framework. Journal of Monetary Economics 12, 383–398.

Comin, D.A., Cirera, X., Cruz, M., 2025. Technology Sophistication Across Establishments. Technical Report. National Bureau of Economic Research.

Crouzet, N., Gupta, A., Mezzanotti, F., 2023. Shocks and Technology Adoption: Evidence from Electronic Payment Systems.

Davis, S.J., Haltiwanger, J., 1992. Gross Job Creation, Gross Job Destruction, and Employment Reallocation. The Quarterly Journal of Economics 107, 819–863.

Davis, S.J., Von Wachter, T.M., 2011. Recessions and the Cost of Job Loss. Technical Report. National Bureau of Economic Research.

Duarte, A., Frost, J., Gambacorta, L., Koo Wilkens, P., Shin, H.S., 2022. Central banks, the monetary system and public payment infrastructures: lessons from Brazil's Pix. Available at SSRN 4064528 .

Griliches, Z., 1957. Hybrid Corn: An Exploration in the Economics of Technological Change. Econometrica 25, 501–522. URL: http://www.jstor.org/stable/1905380.

ILO, 2002. Women and men in the informal economy: A statistical picture.

Jia, P., 2008. What happens when wal-mart comes to town: An empirical analysis of the discount retailing industry. Econometrica 76, 1263–1316.

Mansfield, E., 1961. Technical Change and the Rate of Imitation. Econometrica 29, 741–766. URL: http://www.jstor.org/stable/1911817.

Matsuyama, K., 1991. Increasing Returns, Industrialization, and Indeterminacy of Equilibrium. The Quarterly Journal of Economics 106, 617–650.

Méndez, E., Van Patten, D., 2025. Voting on a trade agreement: Firm networks and attitudes towards openness. Review of Economic Studies , rdaf005.

Milgrom, P., Shannon, C., 1994. Monotone Comparative Statics. Econometrica 62, 157–180.

Topkis, D.M., 1978. Minimizing a Submodular Function on a Lattice. Operations research 26, 305–321.