Network Externalities and Technology Adoption: Lessons from Electronic Payments

Gautam Gowrisankaran
Department of Economics, University of Minnesota
and Federal Reserve Bank of Minneapolis
gautam@econ.umn.edu

Joanna Stavins
Research Department
Federal Reserve Bank of Boston
joanna.stavins@bos.frb.org

This Version: August 12, 1999

Abstract:

We seek to determine the presence and causes of network externalities for the automated clearinghouse (ACH) electronic payments system, using a monthly panel data set on individual bank adoption of ACH. We construct a model of ACH usage that shows how to separately identify network externalities from technological advancement and peer-group effects. We find significant evidence of network effects and find evidence that these network effects are not internalized. Moreover, a large part of these network effects are due to informational problems. Sunk costs of adoption appear to be low.

We thank Dan Ackerberg, V.V. Chari, April Franco, Austan Goolsbee, Ed Green, Tom Holmes, Pete Klenow, Matt Mitchell, Andrea Moro, members of the Regulated Industries Workteam at the Federal Reserve Bank of Cleveland, and seminar participants at UCLA, Duke and the Federal Reserve Banks of Boston and Cleveland for helpful comments. We thank Marie Willard for able research assistance. Gowrisankaran acknowledges financial support from the Financial Services Research Group at the Federal Reserve Bank of Cleveland and from the Office of the Vice-President for Research at the University of Minnesota. The views expressed herein are solely those of the authors and not those of the Federal Reserve System.

Section 1: Introduction

The goal of this paper is to analyze the extent and sources of network externalities for electronic payments markets, using data on bank adoption. A good is characterized by a network externality when an increase in the number of users of the good increases the value to other users, even after controlling for price and other characteristics of the good. Electronic payments markets have some characteristics of network industries—parties directly involved in a payment transaction have to agree on the method of the payment, and their financial institutions have to coordinate technologies and standards.[1] If a bank decides to adopt a particular electronic payment technology, this benefits other banks that use the technology, because banks that already participate can then directly exchange payments with one more institution. Moreover, because electronic payments products are technologically intensive, they may be characterized by informational networks, where the value of the good increases with more users because user familiarity lowers costs. If present, network externalities may give rise to a market failure, where the good is underprovided.

We examine network externalities in the electronic payments industry by using data from the Federal Reserve on one form of electronic payments, the automated clearinghouse (ACH). ACH is an electronic payments system typically used for small recurring payments between consumers and businesses, such as automatic payroll deposit or mortgage deduction. Currently about 75 percent of ACH transactions are processed by the Federal Reserve System. The Federal Reserve processes ACH payments for financial institutions, which in turn sell their ACH services to businesses and individuals. For an ACH transaction to take place, both the originating and receiving banks must have adopted the ACH technology.[2] Banks can adopt the technology as originators and/or as recipients. We perform our estimation using a 34-month panel (March 1995 to December 1997) of the number of ACH transactions for all the individual financial institutions that purchased ACH services from the Federal Reserve.

Our results have some potentially important policy implications for payment systems

[1] Other authors, such as Roberds (1998) and Weinberg (1997) have discussed network effects in payments markets.
[2] Even the financial institutions that correspondent banks to process their ACH transactions have to adjust their infrastructure and invest in staff training.

policy. In an age when computers and technology have become prevalent, only a tiny fraction of payments are completed using electronic payments systems. There are at least two possible explanations for this: either current electronic products are simply less preferred at the current prices to cash and checks for most types of transactions, or there are network externalities. It is only if network externalities are present that electronic payment products are being underused at their current prices. Thus, the two explanations have very different policy implications: the first one calls for laissez-faire policies including market pricing for electronic payments products, while the second one suggests that there may be a need for policy interventions such as aggressive marketing efforts, pricing below marginal cost, or government subsidies for new electronic payments technologies.

The results of our study are also relevant to policy decisions for other high-technology industries such as VCRs, e-mail, or banking on the Internet. These industries all have some network aspects that imply that subsidies to encourage their usage might be welfare enhancing, and are similar to ACH in that they are high-technology industries. By using micro-level panel data that simply does not exist for these industries, our study may shed light on the extent of network externalities for other industries. Thus, our work can also function as an empirical case study that may illustrate the importance of network externalities for other industries.

Indeed, in spite of the substantial theoretical work on network externalities,[3] there have been comparatively few empirical analyses of network effects.[4] A central reason for the lack of empirical analyses is the lack of data. In most industries, the only source of data is time series information, such as monthly sales and price information. However, for technologically intensive goods, price and costs are generally decreasing over time due to technological advances while quantity is increasing over time. One cannot identify whether the increasing quantities are due to the network benefit from having more users or simply due to the lower prices. The empirical studies that have sought to examine network externalities using time-series data have been beset

[3] For instance, Farrell and Saloner (1985) and Katz and Shapiro (1986) have examined the equilibrium adoption of new technologies with network externalities.

[4] We consider network externalities for the case of homogeneous networks, such as the FAX or ACH standard. A whole different literature analyzes network externalities for industries where different goods may give different levels of network compatibility, such as spreadsheets, mainframe computers, and ATM machines. See Gandal (1994), Greenstein (1993) or Saloner and Shephard (1995) for examples of these, respectively.

by this identification issue.[5] Some recent studies of network externalities have instead made use of regional geographical cross-sectional data.[6] Cross-sectional data has its own set of problems: it is difficult to disentangle whether regional correlations in the pattern of usage are due to network externalities or simply to regional variations in preferences, sometimes called peer group effects.[7]

In contrast to these existing studies, we make use of panel data with many observations at any time period, and geographic data to measure the distance between banks. We construct our tests by making use of a simple theoretical model. In our model, customers have a higher utility from ACH usage the more other users there are, and cannot pay other users to use ACH, which leads to a network externality. Banks choose ACH adoption decisions in a simultaneous equilibrium, acting as agents for their customers. The model has testable implications for both network effects and externalities.

We test for network effects by regressing adoption decisions of individual banks on adoption decisions of nearby banks, after controlling for peer group effects and technological improvement. Since we have panel data, we control for peer group effects with fixed effects for each bank, and for technological improvement with time fixed effects. As we use fixed effects, we will find network externalities only when *increases* in usage *above the general time trend* are correlated for nearby banks. Thus, we can separately identify network externalities from increases in usage due to technological improvement (which cannot be done using time series data) and from regional variations in preferences (which cannot be done using cross-sectional data). We find significant evidence of network effects.

We also test for externalities by regressing the market-level usage of ACH on a measure of market concentration. The intuition is that in monopoly markets, the monopolist bank should be able to internalize any network effect, and eliminate any externality. Then, because a network externality is always positive, ACH usage should be increasing in the level of market

[5] See Cabral and Leite (1992), Economides and Himmelberg (1995) and Park (1997). One recent study (Gandal, Kende and Rob (1997)) does attempt to separate the two effects with time series data, by noting that the network benefit for CD players is due to the number of CD software titles available, and not to the quantity of CD players.
[6] For instance, Goolsbee and Klenow (1998) examine network externalities in computers and Rysman (1998) considers network externalities for Yellow Pages telephone books.
[7] This identification problem is common across many fields. For instance, peer group effects are commonly observed in educational outcomes, but it is not clear whether these are caused by variations in preferences or by network effects.

concentration, after controlling for market characteristics. We control for market characteristics via total assets and fixed effects. We find significant evidence of externalities.

In addition to testing for the presence of network externalities, we would like to better characterize the possible sources of the network externalities. There are at least two possible explanations of network externalities in electronic payments. First, with a fixed cost of adoption, potential users may obtain an increase in utility from ACH usage from simply having more people to trade with. We call this a usage externality. Second, if knowledge of how to best use ACH is shared among banks, then ACH may be more valuable with more users, for informational rather than trade reasons. We call this an informational externality.

We can separate informational from usage externalities because we can separately observe two types of ACH transactions: origination and receipt transactions. Because originator banks can only exchange payments with recipient banks, an increase in the number of originators causes only an increase in the informational value of ACH, while an increase in the number of recipients causes both an informational and a usage value increase. We regress the decision to adopt ACH as an originator on both the number of banks that use ACH as an originator and the number that use ACH as a recipient. We find that banks are more likely to adopt ACH as an originator when there are more ACH recipient banks, while the effect of the number of originator banks varies in sign depending on the specification. Thus, we infer that usage externalities are present but that informational externalities may be important as well.

We are also interested in finding out whether sunk costs are an important determinant of ACH usage. In the presence of sunk costs, the decision of banks to adopt ACH will depend on current adoption decisions and expectations of future adoption decisions of other banks in the network, and on the banks' own past decisions to adopt. However, it will not depend on other banks' past decisions to adopt. We test for sunk costs by regressing current decision to use ACH on lagged and future own and competitor usage decision. The lagged and future adoption decisions of competitors enter negatively, which suggests that any sunk cost from adoption is small.

The remainder of this paper is divided as follows. In Section 2, we describe the data set that we use to test for network externalities. Section 3 discusses our base model and its testable implications for network externalities. Section 4 provides results. Section 5 provides extensions to the model where we test for informational externalities and for sunk costs. Section 6 concludes.

Section 2: Data

Our principal data set is the Federal Reserve's monthly billing data that provides information on individual financial institutions that processed their ACH payments through Federal Reserve Banks.[8] We observe data for the period of March 1995 through December 1997. We have two data sets: one lists the billing information for transaction originations, and the other lists the billing information for transaction receipts. ACH transactions can be one of two types: credit or debit. A credit transaction is initiated by the payer; for instance, direct deposit of payroll is originated by the employer's bank, which transfers the money to the employee's bank account. A debit transaction is originated by the payee; for example, utility bill payments are originated by the utility's bank, which initiates the payment from the customer's bank account. For each financial institution in the data set, we have the ACH volume processed through the Federal Reserve each month and the total amount that the Federal Reserve charged for processing that volume. We also have the American Banking Association (ABA) number that allows us to link this data with other publicly available banking data.

The Federal Reserve is currently the dominant provider of ACH services. The Federal Reserve handled approximately 75 percent of the roughly 3.3 billion on-others commercial ACH transactions processed in 1996.[9] The remaining share of the on-others market is handled by three private sector ACH providers: Visa, New York Automated Clearing House, and American Clearing House (formerly Arizona Clearing House). There are some network linkages between the different ACH providers. For instance, the Federal Reserve processes ACH items originated by members of the private networks and vice versa. However, for lack of data, we deal only with ACH transactions that are billed through the Federal Reserve, and treat Federal Reserve ACH as the relevant network for the good.

---

[8] We thank the Federal Reserve's Retail Payments Product Office for making this data set available to us.

[9] NACHA and Federal Reserve estimates. Government transactions constituted another 600 million.

In addition to the ACH billing data, we used a number of publicly available databases. First, we linked the Federal Reserve data with the quarterly Call Reports database. The Call Reports database provides information on bank assets, deposits, name, and the zip code of the headquarters, for all banks that are registered with the Federal Deposit Insurance Corporation (FDIC). We interpolated the Call Reports database from the quarterly to the monthly level. Several banks opened and closed during our sample period. We kept these banks in the sample for the quarters in which they were open. For months in between quarters, we kept a bank only if it was open at the start and end of the quarter.

One data problem that we encountered is that a large fraction of the American Bankers' Association (ABA) numbers—an identifier in the ACH billing data collected by the Federal Reserve—were not in the Call Reports database. Most of the ABA numbers that did not match are credit unions or thrifts.

The Call Report data on assets and deposits are reported by the FDIC certificate number. Banks with a given FDIC certificate number may use one or more ABA identifiers when billing the Federal Reserve for ACH services. Thus, we aggregated the Federal Reserve ACH volume up to the FDIC number level. We then excluded all banks with assets of less than $10 million for all months in the sample and all remaining credit unions. We were left with approximately 11,000 banks over the 34-month sample period.

We also merged our data set with the annual Summary of Deposits database, from June of 1995-97. This database provides the zip code and total deposits for each bank branch, at the FDIC certificate level. We used this database to construct the Herfindahl-Hirschman Index of market concentration (HHI) for each MSA and each non-MSA county, necessary for some of our tests. We chose the Summary of Deposits database in order to be able to construct a branch-level measure of HHI, which gives amore accurate picture than a headquarters-level measure. HHI is defined as the sum of squares of the market shares. We based our measures of market shares for a given bank on its total deposits at all branches in that MSA. After constructing the HHI, we also interpolated this to the 34 months in our billing data.

In order to find the MSA or county (for banks not in MSAs) for each bank in the Call Reports and Summary of Deposits database, we used the Census Geocorr database, that translates from zip codes to MSAs and counties. We used the 1995 county and MSA mappings, and chose the highest weighted MSA or non-MSA county for each zip code. The Census Geocorr database is incomplete, and many of the zip codes from our bank databases are not reported there. In this case, we searched for the zip code in the Geocorr database with a centroid nearest to the missing one. If this nearest zip code was within 10 kilometers of the missing one, we adopted its MSA/county information. Otherwise, we treated that observation as missing and dropped it. Overall, this algorithm was very successful at identifying MSAs/counties, with only a handful of missing observations.

Lastly, we needed to find the distance between zip codes, in order to test for network effects as well as to use the Geocorr databases, as above. We used the Census Tiger database to find the latitude and longitude of zip code centroids, and used the standard great circle formula to find the distance between centroids.

One factor that can affect usage of ACH is its price. Prices that the Federal Reserve charges banks for ACH processing are set at a fixed rate and adjusted periodically. Figure 1 displays a time series of these prices. Note that the intraregional per-item prices (that is, prices for ACH items exchanged between banks located within the same Federal Reserve District) did not change throughout our sample period. At the same time, the interregional prices declined from $0.014 in 1995 to $0.009 in 1997. Because prices are set by fiat and do not respond to changes in local demand, they may be viewed as exogenous. We do not have any information on the prices that banks charge to their customers. In addition to per-transaction costs, banks must pay an ACH participation fee of $25 per month. Also, banks that offer ACH generally maintain a Fedline connection for ACH as well as other electronic payment services.

Financial institutions can adopt ACH as either originators or recipients or both. We assume that a bank has adopted ACH as an originator in a given month if it originated at least one ACH transaction. Our analysis is based on the banks in our sample that adopted ACH in a given month. Table 1 lists the number of banks of different asset sizes that have adopted ACH as originators or recipients or both, while Table 2 lists the number of exitors and entrants in each month. An exitor is defined as a bank that has adopted ACH in the current month, but not in the next month. One can see from Table 1 that the number of banks that use ACH is growing during our sample period, and that more banks use ACH as recipients than as originators. Nonetheless, there is a substantial number of banks that do not use ACH.

In addition to the large number of entrants, there are many exitors. This suggests two things: first, that sunk costs of adoption are low, which we examine more fully in Section 5. Second, that there are banks that technically offer ACH, but have no transactions in a given month. We tried thresholds for adoption other than zero, such as ten or twenty monthly transactions. These different thresholds did not reduce the number of exitors. Moreover, this noise in the data is unlikely to affect the conclusions except to diminish the power of our tests.

Section 3: The Model and Testable Implications

3.1 The Model

We postulate a simple, static base model to test for the presence of network externalities. The goal is to define a model that can explain how network externalities (i.e. interdependent preferences for usage by the customers of banks) translate into different adoption patterns by banks. While our base model is restrictive, our testable implications will be robust to much more general models, as we discuss at the end of this section.

We define the network for a bank to be the set of all other banks with headquarters within 30 kilometers of the current bank.[10] Let J be the number of banks in the network at time t.

Each bank has a fixed number of customers $I_{j,t}$ at time t, and acts as an agent for its customers. Thus, we do not consider market power, or allow for customers to choose between banks.[11] In every period t, banks must decide whether or not to adopt the network good, ACH origination. If a bank adopts the good, it will price at cost. Since there are no marginal profits, a bank values the surplus of its customers, net of the fixed cost of adoption, when deciding whether or not to adopt.

We model a two-stage discrete choice game at each period t. First, banks simultaneously choose whether or not to adopt ACH. Let the set of 0-1 adoption decisions for banks be called

[10] We discuss the implications of this assumption below.

[11] There is survey evidence from the Survey of Small Business Finances and the Survey of Consumer Finances that customers choose their local financial institutions for all their financial needs, including ACH services. (See Kwast *et al.* (1997) for details.)

$\left(A_{1,t},\ldots,A_{J,t}\right)$. Following the set of adoption decisions, customers simultaneously decide whether to use checks or ACH. Let $Usage_{i,j,t}$ denote the 0-1 choice of ACH usage, by customer i of bank j at time t. A customer can only use ACH at time t if $A_{j,t}=1$, i.e. if its bank has adopted the technology.

Let us first consider the customer usage decision, for customer i of bank j at time t. This is the second stage of the game, and is a subgame conditional on $\left(A_{1,t},\ldots,A_{J,t}\right)$.

We can normalize the utility of checks to be $u_{i,j,t,CHECK}=0$. To illustrate the utility of ACH, we define a few more terms: let $W_{i,j,t}$ be the set of individual customer characteristics. Note that some of these characteristics may be observable to the econometrician and some of them may be unobservable. Let $Usage_{\_i,j,t}$ be the 0-1 usage decisions for all customers of banks that are in same network as bank j. Additionally, let $\beta_{i,j,t}$ be a vector of parameters, specific to each customer. Lastly, let P_t be the market price of ACH at time t. Recall that this price is set by the Federal Reserve System, and hence is constant across banks for a given time period.

We can then write the net utility of using ACH for this customer as:

(1) $$u_{i,j,t,ACH}=f\left(W_{i,j,t},Usage_{\_i,j,t},P_t,\beta_{i,j,t}\right),$$

for some function f. We define that a network benefit exists when $\partial u_{i,j,t,ACH}/\partial Usage_{\_i,j,t}$ is positive. In this case, the utility of customer i from using ACH depends on the set of other customers who are using ACH, even after controlling for the product’s price and exogenous attributes. Note that we assume that any network benefit is positive so that utility is increasing in $Usage_{\_i,j,t}$. Also note that $\beta_{i,j,t}$ varies by customer, since each customer will value usage by different customers to different degrees, and hence places different weights on different usages.

Conditional on a vector of bank adoption decisions $\left(A_{1,t},\ldots,A_{J,t}\right)$, consumers choose whether or not to use ACH. The choice of usage is determined simultaneously by all consumers as a Nash equilibrium. Define a vector $U=\left(Usage_{1,1,t},\ldots,Usage_{I_1,1,t},\ldots,Usage_{1,J,t},\ldots,Usage_{I_J,J,t}\right)$ of usage strategies for consumers at every bank in the network. As consumers can only use ACH if their bank has adopted the technology, U will be a Nash equilibrium if and only if:

$$\text{(2)} \qquad \text{Usage}_{i,j,t} = \begin{cases} 1, \text{ if } A_{j,t} = 1, \text{ and } f\left(W_{i,j,t}, \text{Usage}_{\_i,j,t}, P_t, \beta_{i,j,t}\right) > 0 \\ 0, \text{ otherwise} \end{cases}, \forall i, j, t .$$

Thus, from (2), a given consumer will choose to use ACH if, conditional on the set of other consumers who use ACH, her utility of using ACH, from (1), is positive.

Correspondingly, we can define the consumer surplus to bank j's customers from this vector of strategies U as $CS_{j,t}\left(U, W_{\_,j,t}, P_t, \beta_{\_,j,t}\right)$, where "$_{\_,j,t}$" refers to the set of characteristics for all consumers of that bank. The consumer surplus is defined as the area under the Marshallian demand curve, which is to say that it is the integral of the demand function from the current price to infinity. As we are modeling a discrete choice demand curve, the quantity demanded is either 0 or 1, and can be represented by an indicator function. Thus, $CS_{j,t}\left(U, W_{\_,j,t}, P_t, \beta_{\_,j,t}\right)$ can be written as:

$$\text{(3)} \qquad CS_{j,t}\left(U, W_{\_j,t}, P_t, \beta_{\_j,t}\right) = \sum_{i=1}^{I_{j,t}} \int_{P_t}^{\infty} \left(\text{Usage}_{i,j,t}\right) f\left(W_{i,j,t}, \text{Usage}_{\_i,j,t}, p, \beta_{i,j,t} > 0\right) dp .$$

Note that the indicator function $\text{Usage}_{i,j,t}$ is used in (3), because, if bank j has not adopted, usage will be zero, even if consumers would gain from usage.

While existence of equilibrium can easily be shown by construction, one important feature about network externality models is that there may be multiple Nash equilibria, characterized by tipping behavior, as in Farrell and Saloner (1985). For instance, if the magnitude of the network externality is sufficiently large, there may be a Nash equilibrium where everyone is using the good, and another where no one is using the good. While the equilibrium is not unique, we can show that there is a unique Pareto dominating Nash equilibrium, and that usage in this equilibrium is higher than for any other Nash equilibrium:

Proposition 1: Assume that $\partial u_{i,j,t,ACH} / \partial Usage_{\_i,j,t}$ is always strictly positive. Then, given adoption decisions $A = (A_{1,t}, \ldots, A_{J,t})$, there exists a unique Nash equilibrium $U(A)$ of the usage game such that U Pareto dominates all other Nash equilibria and U is Pareto optimal over the set of possible usage decisions. Moreover, if $A' = (A'_{1,t}, \ldots, A'_{J,t})$ satisfies $A' \geq A$ [12] then $U(A') \geq U(A)$ and $CS_{j,t}(U(A'), W_{\_,j,t}, P_t, \beta_{\_,j,t}) \geq CS_{j,t}(U(A), W_{\_,j,t}, P_t, \beta_{\_,j,t}), \forall j$.

Proof: This proposition follows because the network externality is assumed to be positive and so, if an additional customer uses ACH, everyone will be weakly better off. To show this, we first construct a Pareto optimal Nash equilibrium of this game, which we call U. Then, we show that U Pareto dominates all other Nash equilibria.

To construct U, start with a strategy profile U^0 where every customer from a bank that has adopted ACH is using ACH. Then, construct U^1 by setting usage to 0 for all customers who do not want to use ACH even if everyone is using ACH, as in U^0. Then, construct U^2 by removing from U^1 everyone who does not want to use ACH, given that the set of other people using ACH is specified by U^1. Note that $U^2 \leq U^1 \leq U^0$. Repeat this process until we find some N such that $U^N = U^{N+1}$. Such an N must exist because there are only a finite number of customers. Let $U^N = U$. Then, U is a Nash equilibrium, because no one who stopped using ACH at some earlier stage U^n would want to start, given that the network externality is positive, and everyone who uses ACH wants to use it. Also, U is a Pareto optimal set of usage strategies, because the only way to make customers who use ACH better off is to induce some subset of the customers who do not use it to use it. But, if this subset were to use ACH, then at least one customer in the subset would have a negative utility of using ACH, by construction.

Now, suppose that there is another Nash equilibrium V that is not Pareto dominated by U. Then, as there are some customers who are better off under V than under U, there is a set of customers who use ACH under V but not under U. Let W be a set of usage decisions constructed by taking the union of people who use ACH under either U or V. Everyone who uses ACH under V but not under U has a positive utility from W and a zero utility from U.

[12] Recall that adoption is a vector of 0's and 1's. Thus, the inequality means that the set of customers that use ACH

Everyone else has utility at least as high, from W as from U. Thus, W Pareto dominates U, contradicting our earlier proof that U is a Pareto optimal strategy.

Lastly, note that if $A' \geq A$, then in our construction of the Pareto dominating equilibrium, we start with a higher value of U^0 as more customers can use the good; equivalently $U^0(A') \geq U^0(A)$. At each stage n, then, it follows that $U^n(A') \geq U^n(A)$ since the network benefit from those new customers to the existing customers is positive. Thus, $U(A') \geq U(A)$. Since the network benefit is positive, customers from existing banks under A will only gain from this increased usage. Customers from new banks will earn non-negative surplus, versus the zero surplus that they earned previously. Thus, $CS_{j,t}(U(A'), W_{-,j,t}, P_t, \beta_{-,j,t}) \geq CS_{j,t}(U(A), W_{-,j,t}, P_t, \beta_{-,j,t}), \forall j$.

While we showed that there is a unique Pareto dominating Nash equilibrium, there is still an externality because customers cannot compensate other customers for their ACH usage. Thus, while this equilibrium is Pareto optimal over the set of usage strategies, it is possible to obtain a Pareto improvement in welfare levels if customers can pay other customers to use ACH.

We now aggregate to the bank level in order to examine the choice of adoption decisions. Let the fixed costs for bank j of adopting ACH be $FC_{j,t} - \varepsilon_{j,t}$, where $FC_{j,t}$ is observed and $\varepsilon_{j,t}$ is unobserved to the econometrician. We assume that $\varepsilon_{j,t}$ is observed to all firms in the network. Let $U(A)$ denote a set of second-stage Nash equilibria, one for each vector of adoption strategies. The surplus to the bank is the surplus of its consumers net of its fixed costs. Thus, we can write the surplus from adopting ACH to bank j as:

(4) $$\Pi_{j,t}((A_{-j,t}, 1), W_{-,j,t}, P_t, \beta_{-,j,t}) = CS_{j,t}(U(A_{-j,t}, 1), W_{-,j,t}, P_t, \beta_{-,j,t}) - FC_{j,t} + \varepsilon_{j,t},$$

where $(A_{-j,t}, 1) \equiv (A_{1,t}, \ldots, A_{j-1,t}, 1, A_{j+1,t}, \ldots, A_{J,t})$. If bank j adopts ACH in period t, it will realize the surplus in (4), otherwise it will receive a surplus of 0. Thus, a strategy profile $(A_{1,t}, \ldots, A_{J,t})$ induces a subgame perfect equilibrium on the game if and only if:

in the second instance is a subset of the set of customers that use ACH in the first instance.

(5) $$A_{j,t} = \{\Pi_{j,t}((A_{-j,t},1), W_{-,j,t}, P_t, \beta_{-,j,t})\} = \{CS_{j,t}(U(A_{-j,t},1), W_{-,j,t}, P_t, \beta_{-,j,t}) - FC_{j,t} + \varepsilon_{j,t}\}$$

where { } is the indicator function.

From (5), one can see that the bank adoption problem is itself an indirect network externality type problem. Moreover, if customers choose the Pareto dominating Nash equilibrium, then from Proposition 1, an additional ACH adoption will only have a positive effect in terms of consumer surplus, and hence in terms of the bank's "utility" function from (4). Thus, if consumers choose their usage decisions based on the Pareto dominating Nash equilibrium, (4) will then be of exactly the same functional form as (1). We can then apply Proposition 1 recursively to show that the adoption problem from (4) will also have a unique Pareto dominating subgame perfect equilibrium.

As a result, the indirect adoption network externality game will have as a property that equilibrium profits from adopting are increasing in other firms' ε's, which we can test with data. To formalize:

Proposition 2: Suppose customers choose the unique Pareto dominating Nash equilibrium of usage and banks choose the unique Pareto dominating subgame perfect equilibrium of adoption. Then, an increase in $\varepsilon_{j,t}$ weakly increases the equilibrium adoption $A_{-j,t}$ for other firms.

Proof: First, as above, with the Pareto dominating equilibrium of the usage game, the reduced-form "utility" to bank k expressed in (4) is increasing in other banks' adoption decisions. Thus, we can construct the unique Pareto dominating equilibrium for this problem, in the same way as in Proposition 1.

From (4), the reduced-form utility from adoption for bank j is increasing in $\varepsilon_{j,t}$. Now, let us go through the similar construction of the equilibrium given in Proposition 1 for the adoption game, for two values of $\varepsilon_{j,t}$, $\varepsilon^1_{j,t}$ and $\varepsilon^2_{j,t}$, where $\varepsilon^1_{j,t} < \varepsilon^2_{j,t}$. If firm j adopted ACH for the first value, then it would adopt it for the second value. To see this, if j did not adopt it for $\varepsilon^1_{j,t}$, then there is some stage n in our equilibrium construction where it dropped ACH for $\varepsilon^1_{j,t}$. Now, at that stage n, if $\varepsilon^2_{j,t}$ is sufficiently high, it may continue to adopt ACH. If this is the case, then, at some later stage, this may cause some other firm k that would have not adopted ACH to adopt it, instead. This may, in turn, cause other firms to adopt it. At no stage could it be the case that a firm that would have adopted ACH for the first value would choose not to adopt ACH for the second value. Thus, the increase in $\varepsilon_{j,t}$ weakly increases the Pareto dominating equilibrium adoption for other banks.

Given Proposition 2, after controlling for exogenous factors $\beta_{\_,j,t}$, P_t and $W_{\_,j,t}$, the unique Pareto dominating equilibrium will generate a positive partial correlation in adoption decisions for a given network. Using data on multiple networks over multiple times, we can use this insight to estimate a network effect.

3.2 Testing for Network Effects

We test for network effects by examining whether we find clusters in adoptions, as above. We make one major assumption regarding the error terms: that the unobserved components of utility ($\varepsilon_{j,t}$ and any unobserved components of $W_{\_,j,t}$) are not correlated across banks in a market. As we discuss later, this assumption is necessary to identify network effects, in the absence of any exogenous variation in usage.

Because we control for fixed effects for banks and time periods, this restriction still allows for peer group effects and for time specific shocks. Thus, if one bank or region uses more ACH than another bank or region because its residents or customers are more technologically sophisticated, this will be captured by a high fixed effect and be perfectly consistent with our assumption. Also, if ACH usage is increasing over time, because of decreased price or increased acceptance of electronic payment methods, this will result in increasing monthly fixed effects over time. However, if there are shocks that are correlated across nearby banks that are both local and time-specific, then that will not be consistent with our model. Lastly, we note that we use ACH origination data rather than receipt data because this restriction is not as credible for receipts: for instance, if a large employer starts a direct deposit program, then this may cause many banks to simultaneously choose recipient adoption, exactly because of a concurrent change in the unobserved components, caused by the actions of this one large employer.

Given our assumption, we can test for network externalities by modeling the profit function (4). To illustrate, define $\#(j,t)$ as the number of other banks in the network of bank j that use ACH at time t, i.e. $\#(j,t) = A_{1,t} + \ldots + A_{j-1,t} + A_{j+1,t} + \ldots + A_{J,t}$. Recall that the network is defined to be the number of banks within 30 kilometers of bank j. Define also $\#_Q(j,t)$ as the volume of usage by other banks in the network of bank j at time t, which is a proxy for the number of customers using ACH. Then, our base testing equation is:

$$\text{(6)} \qquad A_{j,t} = \{\Pi_{j,t} > 0\} = \{\lambda_1 \#(j,t) + \lambda_2 \#_Q(j,t) + \alpha_j + \delta_t + \gamma_1 x_{j,t} + \varepsilon_{j,t} > 0\},$$

where α is a vector of bank fixed effects, δ is a vector of monthly fixed effects (one of which is omitted), and $\varepsilon_{j,t}$ is a logit error term. The bank fixed effect α_j captures fixed costs, characteristics of the bank's customers from $W_{\_,j,t}$ and network preferences from $\beta_{\_,j,t}$. The time fixed effect δ_t captures price at time t and other factors that might evolve with time, such as technological awareness. The variable $x_{j,t}$ captures factors from $W_{\_,j,t}$ or $\beta_{\_,j,t}$ that may change over time and bank, such as assets and deposits of the bank and its competitors. The two measures of adoption, $\#(j,t)$ and $\#_Q(j,t)$ measure the potential network benefit $CS_{j,t}(U(A_{\_j,t},1), W_{\_,j,t}, P_t, \beta_{\_,j,t})$.

Equation (6) is our basic test for network effects. We will interpret the data as providing significant evidence of network effects if and only if λ_1 or λ_2 is significantly positive.

The justification for using (6) as a test of network externalities is as follows. Suppose that there are no network externalities. Then, $A_{j,t}$ is a function of its exogenous characteristics, and the fixed cost error term. By the assumption that the unobserved components are uncorrelated across banks, $\#(j,t)$ and $\#_Q(j,t)$ will be exogenous to bank j's error terms. Thus, with no network externalities, if we perform a regression of ACH adoption on exogenous characteristics and adoption decisions of other banks, the estimated coefficients on λ_1 or λ_2 will be consistent, and should converge in probability to zero.

Now suppose that there are network externalities. In this case, $(A_{1,t},\dots,A_{J,t})$ will be determined via a simultaneous Nash equilibrium as in (5). Suppose also that firms and customers choose the Pareto dominating Nash equilibrium. Then, by Proposition 3.2, we know that a high $\varepsilon_{j,t}$ will yield a high $A_{j,t}$ and a high $A_{-j,t}$, after controlling for the exogenous factors. Then, as the ε's are iid, this will induce a positive partial correlation between $A_{j,t}$ and $\#(j,t)$ and $\#_Q(j,t)$, in which case λ_1 or λ_2 will converge in probability to significantly positive values. This means that if we performed a simple likelihood ratio test, and found that we could reject that $\#(j,t)$ is zero with 95 percent probability, then, because of the endogeneity we know that there is at least a 95 percent probability that the coefficient is not zero. Thus, we can use the simple test statistics, with the small caveat that they may underestimate the probability of finding network effects, even if they are present.

If banks and/or their customers choose different equilibria from the Pareto dominating ones, then it is not always the case that this positive partial correlation will still be observed. For instance, we could find pathological subgame perfect equilibria, where fewer customers use ACH the more banks adopt. In general, though, as long as the equilibrium satisfies that, as adoption increases (in the case of usage) and as ε increases (in the case of adoption) the equilibrium does not have less users,[13] a positive partial correlation will still be generated. Indeed, much of our identification will come from the fact that some networks will have coordinated on the Pareto dominating equilibrium and other networks will still largely be using checks.

A couple of explanations about our identification of network externalities may be useful at this point. First, note that we do not observe any exogenous variation in the adoption decisions of competitors: in contrast, usage levels for the network good are determined simultaneously, based on the error realizations. Thus, the regression of adoption on adoption in (6) will be beset by the classic simultaneity problem, as above. Therefore, λ_1 cannot be interpreted as a consistent estimate of the *structural* parameter on $\#(j,t)$ in the utility function. Nonetheless, as shown above, a partial correlation can be interpreted as evidence of network externalities. While recovery of the structural parameters would require an equilibrium model, one can see that the coefficients on usage will be asymptotically biased upwards, because part of any observed correlation of high usage levels across banks will be due to high realization of an error term, and not to causation.

Second, recall that we assumed that $\varepsilon_{j,t}$ is uncorrelated for banks in a network. If the ε's are in fact correlated within a network, then even in the absence of network externalities, adoption decisions would be endogenous. With a positive correlation, we would estimate a positive value for α. Thus, we make the assumption that the ε's are uncorrelated because this is the only way to identify network externalities with simultaneously determined usage data. Moreover, any correlation in the ε's may be interpreted as a type of network externality. For instance, some Federal Reserve Banks had regional advertising campaigns to promote ACH, which may have led to higher usage levels for banks in those regions at the times of the advertising campaigns.[14] To the extent that these campaigns increased the usage level for banks in specific regions at certain times, this is a type of network externality from the advertising. However, these campaigns will cause informational externalities as opposed to pure network externalities. We discuss in Section 5 how to use our data to separate informational externalities from pure network externalities.

3.3 Testing for Externalities

If there is more than one bank in a network, then as discussed in Section 3.1, one bank's adoption of the network good will increase the utility for customers of other banks in the network. Implicit in our Nash equilibrium definition is the assumption that banks cannot compensate other banks for the benefit that they cause by adopting ACH. This is why the network effect leads to an externality, where the network good may be underproduced.

---

[13] Note that we know that such an equilibrium exists, because Proposition 1 shows that the Pareto dominating equilibrium satisfies this condition.

An alternative approach to testing for network effects is to test for externalities. From (5), we can see that each bank chooses the adoption decision that maximizes the surplus of its customers, conditional on other banks' policies. The socially optimal adoption policy would be to choose the adoption decision that maximizes the sum of each banks welfare conditional on other banks' actions, i.e., to choose a vector of strategies $\left(A_{1,t},\ldots,A_{J,t}\right)$ such that:

(7) $$A_{j,t} = \left\{\Pi_{1,t}\left(\left(A_{\_j,t},1\right), W_{\_,1,t}, P_t, \beta_{\_,1,t}\right) + \ldots + \Pi_{J,t}\left(\left(A_{\_j,t},1\right), W_{\_,J,t}, P_t, \beta_{\_,J,t}\right)\right\}.$$

For networks in which there is only one bank, (5) and (7) will be exactly the same expression. Hence, for monopoly networks, the bank will be able to fully internalize any adoption network effect. Note that the bank will not necessarily be able to implement the optimal usage level conditional on adoption since customers are still not able to compensate other customers for the network externality that their usage generates. However, conditional on usage, the probability of adoption, over realizations of the fixed cost error ε, will be optimal.

We can show that a two-firm oligopoly with the same customers and the same fixed costs as the monopoly will have a weakly lower adoption probability than the monopoly and hence a lower adoption probability than is socially optimal. This is true regardless of the choice of equilibrium. To illustrate, consider a given set of customers, and two industry structures, monopoly and two-firm oligopoly. Under monopoly, the optimal adoption decision will be implemented, as above.

Now for a two-firm oligopoly formed by splitting the set of customers into two banks, bank 1 and bank 2, the set of realizations of ε under which either bank 1 or 2 adopts is a subset of the set under which the monopoly adopts. To see this, for ε's for which the monopoly adopts, the equilibrium surplus to bank 1 from adopting is weakly less than the surplus to the monopolist, because bank 1 will have to bear the same fixed costs, but will only internalize the surplus of its customers (which will be the same as under the monopoly) but not the surplus of bank 2's customers. Thus, bank 1 may or may not adopt, even if bank 2 adopts. In this case, every

[14] The information was provided by the Federal Reserve Bank of St. Louis. We were unable to find data on exact advertising expenditures.

subgame perfect equilibrium will have no adoption by bank 1. For ε's for which the monopoly did not adopt, the equilibrium gains to either firm from adopting are less than (if the other firm does not adopt) or equal to (if the other firm does adopt) the gains to the monopoly from adopting, and the costs are the same. Hence there is no subgame perfect equilibrium where either firm adopts.

Thus, with a two-firm oligopoly, the expected probability of being at a bank that uses ACH is less than for a monopoly. The same logic shows that as we move the customers into finer and finer partitions of banks, the expected probability of being at a bank that uses ACH becomes smaller and smaller.

We use this logic to construct a network-level test for externalities. The test uses as the network the set of banks within each MSA or non-MSA county. We use MSA or non-MSA county instead of the 30-kilometer boundary from earlier to avoid the problem of overlapping networks. Let $HHI_{k,t}$ denote the MSA/county-level market concentration.

For the test, we construct a dependent variable $y_{k,t}$ for MSA k and month t that is the fraction of customers at banks that have adopted ACH; i.e. $y_{k,t} = \sum_{j=1}^{J} I_{j,t} A_{j,t} \Big/ \sum_{j=1}^{J} I_{j,t}$. We proxy for number of customers $I_{j,t}$ using each bank's deposits. Then, we regress:

(8) $$y_{k,t} = \lambda_3 HHI_{k,t} + \gamma_2 x_{k,t} + \delta_t + \alpha_k + \varepsilon_{k,t} .$$

We choose market size, measured by dummy variables created by partitioning banks into 100 classes, depending on their size in terms of assets or deposits, and/or MSA-level fixed effects as our controls $x_{k,t}$ and α_k. As in (6), δ is a set of monthly fixed effects and α is a set of MSA fixed effects. If we find that λ_3 is a significantly positive predictor of y, this is evidence of externalities, which we attribute to network externalities.

We can use a similar logic to construct a test at the bank level. Consider the following experiment. Suppose that we fix bank 1, and split the customers of a second bank into banks 2 and 3. Now, banks 2 and 3 will be less likely to adopt ACH than before, as above. Because of this, with network externalities, bank 1 will also be less likely to adopt ACH than before. Thus, with network externalities, bank 1's adoption decision will be influenced by the market concentration of its network, after controlling for its own size and attributes.

To construct this test of externalities, we again use the MSA/county definition of the network. We perform a regression:

(9) $$A_{j,t} = \{\lambda_4 HHI_{j,t} + \gamma_3 x_{j,t} + \delta_t + \alpha_j + \varepsilon_{j,t}\},$$

where we use a logistic error specification for $\varepsilon_{j,t}$. The HHI variable that is used is the HHI for the MSA/county in which the bank is located. We again use asset dummies and/or bank level fixed effects as our controls $x_{k,t}$ and α_k. If we find a significantly positive value of λ_4, this is evidence that banks are more likely to adopt in an industry that is more concentrated, which we attribute to network externalities.

Since ACH is a very small percentage of bank transactions, it is plausible to view HHI as exogenous to our structural error terms ε. In this case, we can view the coefficient on market concentration as a reduced-form parameter that indicates the true effect that a change in market concentration would have on ACH adoption. This allows us to directly use our coefficient estimates as evidence on the presence of network externalities.

We note also that our HHI measure of market concentration is based on the Summary of Deposits. As this data is collected annually, and interpolated to the monthly level, it forms a rough measure of the market concentration in any given month. Because this measure is rough, we also estimate (8) without the MSA fixed effects and (9) without the bank fixed effects, α, in order to find an estimate that is not based solely on within-market changes in concentration. We still use monthly fixed effects to control for price variation in all the specifications.

Finally, note that we might be tempted to construct different bank-level tests for externalities. However, these would not be appropriate. To illustrate, we might regress adoption on $\#(j, t)$ and the HHI market concentration. However, even with network externalities, it is unclear whether the effect of concentration will be negative or positive, because adoption is an endogenous variable. Alternatively, if we regress adoption at the bank level on market concentration without controlling for the bank's own size then the implications of network externalities are also unclear, since the bank's own size affects its decision to adopt.

3.4 Robustness of the Results

We have presented a simple static model of network externalities, to illustrate the basis for our test. Since we are testing for network externalities and not structurally estimating their size, there are several assumptions about our model that could be relaxed and still yield the same testable implications. We detail each of these below.

Size of network

We assumed a network size of 30 kilometers for the tests of network effects, and that the network is the MSA for the tests of externalities. Certainly, banks exchange payments with other banks with headquarters that are outside this limited area, and thus a bank's true network may include banks that are not in our network. If the size of the network is larger, then we will not capture the full extent of the network effect with our parameter on usage. Since usage of ACH is increasing over our sample period, the non-captured network effect will be confounded with the time fixed effects. Nonetheless, if there is a network externality, a portion of it is likely to be local, and this will be captured by our estimation process. If the size of the network is smaller than 30 kilometers, then again, our measures of the network will be imprecise, but any network effect will be captured. Thus, our tests are robust to different sizes of the network. To the extent that our network size understates a bank's true network size, our externality regression estimates will understate the deviation of oligopoly ACH usage from its optimal level; these estimates will nonetheless be valid as a test of network externalities.

Market power

We assumed that banks are perfect agents for their customers, and that customers are captive to one institution. Suppose now that customers are captive to one institution, but that banks choose to maximize profits rather than maximize surplus and cannot price discriminate.[15] In this case, as firms have market power, the price that banks charge their customers for ACH usage will be different from the price that they pay to the Fed, and will be decreasing in the

[15] If banks can perfectly price discriminate, then they would again want to maximize total surplus and their adoption decisions would be identical to those of our base model.

elasticity of demand of their customer base. Thus, a bank's adoption decision will reflect both its customers' surplus from ACH and the bank's ability to extract this surplus from its customers. However, if customers are still captive to one institution, then other banks' adoption decisions will not cause a bank to change its adoption decision. While adoption decisions may be correlated because market power will be similar for nearby banks, this effect of market power on adoption will be picked up by the bank fixed effects. Thus, this type of market power will not induce any correlation among the adoption decisions, in the absence of network externalities. Our tests for network effects will be robust to this type of market power.

Now suppose that there is some common pool of customers, and that banks compete for these customers. We would then postulate a three stage game: in the first stage, firms simultaneously make adoption decisions; in the second stage, firms simultaneously set prices (assuming some product differentiation) or quantities; in the final stage, consumers choose their bank and whether or not to use ACH. In this case, a high value of ε_j will cause firm j to be more likely to adopt. This will cause the residual demand for firm k to be less than it otherwise would be, which will cause firm k to be less likely to adopt. This is the opposite sign from a network externality. Thus, if we do not see a positive correlation in adoption decisions, this may be due to confounding between market power and network externalities. However, a positive correlation cannot be explained by market power, because the sign of the effect is the opposite.

Market power will also have the opposite effect from network externalities in the externality regressions. In particular, equilibrium quantity for an oligopoly declines in the number of firms, in most models of production. Thus, any market power effect will result in a lower demand for ACH the more concentrated the market, which is the opposite conclusion from our base model. Just as for the network effect regression, then, a positive correlation between market power and ACH usage must be interpreted as a network externality. Also, any market power will again result in an underestimate of the optimal ACH usage level.

Dynamics

One assumption that we made is that banks and consumers make adoption and usage decisions every period. If there are sunk costs of ACH adoption, then the decision to adopt or use ACH may depend on past usage decisions. Additionally, if there are sunk costs and agents are not myopic, adoption decisions will depend on future expectations of the network benefit from usage.

Even if there are important dynamic interactions as above, in the absence of network externalities, there will be no partial correlation between the adoption decisions. Moreover, with network externalities, a portion of the adoption decision will not depend on sunk costs and will depend solely on current network benefits. The implication is that if and only if there are network externalities, will there be a correlation between current adoption decisions. Thus, our tests are valid tests of some of the cross-sectional implications of the dynamic adoption model. In Section 5, we explore more fully some of the dynamics implications.

One other thing relevant to dynamics is that we do not model any serial correlation in the error structure. However, if there is serial correlation but we are not accounting for it, then our results may be inconsistent.

Section 4: Results of the Model

We present two sets of basic results. Our first set tests for network effects using (6), and our second set of results tests for externalities using (8) and (9).

Recall that our tests for network effects from (6) include a fixed effect α_j for each bank. As we have 11,000 banks in our data set, it is not computationally feasible to estimate a dummy variable for each bank. Fortunately, Chamberlain (1980) provides a conditional likelihood method of consistently estimating a logit model with fixed effects. Chamberlain shows that one can estimate the conditional probability of seeing a particular choice sequence, where we condition on the total number of times that the good was chosen. By conditioning in this manner, the fixed effects α can be divided out from the conditional likelihood, in the same way that fixed effects can be differenced out in a linear model. We make use of Chamberlain's model to estimate our regression of network effects. Note that Chamberlain's method gives consistent but not efficient estimates, since it is only a conditional maximum likelihood estimator.

The results from (6) are presented in Table 3. We present a few variants of the results, to check for robustness. All regressions include a full set of 34 monthly dummies δ. Models 1 - 3 include fixed effects for each bank, calculated using Chamberlain's conditional maximum likelihood estimator. As Chamberlain's results give consistent but not efficient estimates, we calculated standard maximum likelihood results in Models 4 - 6. For these results, we included size dummies, but not bank-level fixed effects. Models 1 and 4 give results of a regression of adoption on $\#(j,t)$. In the second column (Models 2 and 5), adoption is regressed on $\#_Q(j,t)$; the third column gives a regression of adoption on $\#(j,t)$ and $\#_Q(j,t)$.

We find significant evidence of network externalities: All the models show a strongly significant value of λ_1, the effect of adoption by other banks on own adoption. Note that this is evidence not just of potential network externalities, but that unexploited network externalities exist during our sample period, twelve years after ACH was already established. Regressions with just volume of other banks (Models 2 and 5) shows mixed evidence of λ_2, the effect of volume by other banks on own adoption. With both adoption and volume (Models 3 and 6), λ_2 is estimated to be negative. This suggests two things: first, the functional form of the network benefit between adoption and volume is non-linear and complicated. Second, since banks value other banks' adoptions rather than volume, the network externality at the consumer level may be less than at the bank level.

The size of the coefficients on usage varies significantly between the conditional maximum likelihood and standard maximum likelihood estimators. For instance, $\#(j,t)$ is estimated to be 0.174 in Model 1, but 0.005 in Model 4. However, as we do not interpret these coefficients as structural estimates of the effect of usage, this is not necessarily problematic. Moreover, in a linear model, a fixed-effects estimator is likely to exacerbate any endogeneity problem, and the same is likely to be true for fixed-effects logit versus the regular logit.

Table 4 presents our tests for externalities. We include six specifications, of which three are network level regressions based on (8) and three are bank level regressions based on (9). For both tests, we include standard regressions with controls for assets, as well as a fixed effects regression. The fixed effects regression for the network level regression (Model 3) is computed using the standard linear fixed effects model, while the fixed effects regression for bank level regressions (Model 6) is computed using Chamberlain's conditional logit model. All specifications include a full set of monthly dummies. For the network level tests (Models 1 - 3), we include dummy variables that control for MSA/county level characteristics, including total assets, deposits and population. For the bank level tests (Models 4 - 6), the dummy variables, of assets and deposits, are at a bank level.

We find significant evidence of externalities. For the network level tests (Models 1 - 3), both non fixed-effects regressions show a significantly positive effect of HHI on usage. The fixed effects regression shows an effect that is positive, although not significant. Thus, a higher value of HHI (i.e. less competition) is associated with a higher percentage of banks using ACH. For the bank level tests (Models 4 - 6), we again find significant evidence of externalities. The fixed effects model and one of the non-fixed effects models both show that an increase in HHI significantly increases a bank's probability of using ACH. The non-fixed effects model with controls for assets and deposits (Model 5), shows a positive, though not significant, effect.

Section 5: Extensions of the Model

5.1 Separating Informational and Usage Externalities

Our base model in Section 3 analyzed only a generic network externality, where it was assumed that customers' utility increases with other customers' use of ACH. We did not characterize the underlying causes of this externality. In this subsection, we seek to better characterize the causes of this externality.

We consider two possible sources for the network externality in (1). First, we might think of a customer as a small business that has to make payments to a random exogenously specified set of other customers each month. Suppose that the customer has to pay a fixed cost CFC for making an ACH payment. Then, its ex-ante reduced-form utility would be increasing in the number of other users of ACH, as in (1). We can write this as:

(10) $$u_{i,j,t,ACH} = \sum_{h \in \_i} \Pr(\text{i transacts with h}) \cdot \text{Usage}_{h,j,t} \cdot \left(AVC^{Check} - AVC^{ACH}\right) - CFC_{i,t} \ .$$

We assume that ACH has a lower average variable cost than checks, which empirical studies have found.[16] Costs include the price paid to the bank for using ACH as well as transaction costs and record keeping. In this case, a customer will have an increase in utility from other customers' use of ACH because the benefit from using ACH, which is the lower average variable cost, is spread over more units of fixed costs. We call this first type of network externality a *usage externality*.

Second, another reason for a network externality is that the costs of usage might be lower with more customers. This would be true if customers or banks learn from each other about the best methods of using ACH, or if usage is more accepted with more users.[17] In this case, we can write the utility from ACH as:

$$(11) \qquad u_{i,j,t,ACH} = (\#\text{ of ACH transactions}) \cdot \left(AVC^{Check} - AVC^{ACH}\left(Usage_{-i,j,t}\right)\right) - CFC_{i,t}\left(Usage_{-i,j,t}\right).$$

Note that in (11), we assume that the cost of using ACH is lower even with a constant volume of output, as the number of other customers using ACH increases. We call this type of externality an *informational externality*. Informational externalities can also exist at the bank level. In this case, bank fixed costs FC from (4) would be increasing in the number of other banks that have adopted ACH.

Potentially, both informational and usage externalities can exist for an industry. In this case, utility from ACH usage would incorporate elements from both (10) and (11). Formally, we can define:

$$(12) \qquad \begin{aligned} u_{i,j,t,ACH} = &\sum_{h \in -i} \Pr(\text{i transacts with h})h \cdot Usage_{h,j,t} \cdot \left(AVC^{Check} - AVC^{ACH}\left(Usage_{-i,j,t}\right)\right) \\ &- CFC_{i,t}\left(Usage_{-i,j,t}\right) \end{aligned}$$

[16] See, for instance, Wells (1997).

[17] For instance, we have heard reports that small businesses sometimes do not want to use ACH because in a tax audit, an ACH receipt is poorly understood by IRS inspectors relative to a cancelled check.

Informational externalities are thought to exist for many technologically intensive goods, such as computers,[18] cellular phones or even prescription medications.[19] While these two externalities appear quite similar, it is important to distinguish between them. The two types of externalities have different policy implications for ACH. For instance, the only way to internalize a usage externality is to subsidize usage of ACH. However, user training or advertising may lower costs and hence internalize an informational externality. Moreover, we may have some strong prior characterization about the relative levels of these two externalities for different industries. For instance, VCRs are largely driven by usage externalities, since customers care about the number of software titles available, not because more customers give more information about the good. In contrast, computer spreadsheets may be driven by informational externalities, since many of them offer strong cross-program compatibility and superior exogenous characteristics, but have quite low market shares.[20]

We are in a unique position to separate informational from usage externalities, because we observe separate ACH origination and receipt adoption. As customers can only originate payments to customers of recipient banks, the usage externality in (12) is determined exclusively by the number of recipient customers. In contrast, the informational externality is determined by both origination and recipient customers, and probably more by origination customers, because the type of transaction is more similar.

As in Section 3, this insight will translate up to the bank level, to generate an analog of (5). Define $\#^{R}(j,t)$ to be the number of other ACH recipient banks in the network for bank j at time t. Then, we test for informational versus usage externalities by constructing a similar test of adoption as in (6):

$$(13) \qquad A_{j,t} = \{\Pi_{j,t} > 0\} = \{\lambda_5 \#(j,t) + \lambda_6 \#^{R}(j,t) + \alpha_j + \delta_t + \gamma x_{j,t} + \varepsilon_{j,t} > 0\}.$$

Similarly to (6), we cannot interpret the λ coefficients in (13) as structural coefficients. Nonetheless, if there is no usage externality and information is obtained solely from other

[18] For instance, Goolsbee and Klenow (1998) analyze this type of network externality for computer usage.
[19] See Berndt, Pindyck, and Azoulay (1999).
[20] See Gandal (1994).

origination adopters, then $\#^R(j,t)$ will be exogenous and hence λ_6 will be zero. If there is no informational externality, then the coefficient on λ_5 may still be non-negative, because it will be endogenously determined, in the system of equations with origination and receipt adoption. However, such an effect will be a minor second-order effect, and thus we would expect its value to be close to zero. Hence, significantly positive values of λ_6 will be interpreted as evidence of usage externalities, and significantly positive values of λ_5 as evidence of informational externalities.

The results from a regression of (13) are given in Table 5. We choose two specifications: a Chamberlain conditional maximum likelihood with fixed effects for each bank (Model 1), and a maximum likelihood with dummy variables for asset size classes (Model 2).

We find significant evidence of usage externalities: our estimated values of λ_6 are 0.081 or 0.009 depending on the model chosen, and significantly positive. In contrast, we find mixed evidence of informational externalities. Depending on the specification, λ_5 is either significantly positive or negative. This suggests that both usage and informational externalities are important parts of the story for ACH.

5.2 Sunk Costs of Adoption

Thus far, we have examined a static decision process. In the presence of sunk costs of adoption, lagged usage will affect the current decision of usage. Moreover, if customers and banks are not myopic, then in the presence of sunk costs, expected future usage of ACH will affect current usage decisions.

We construct a simple extension of our model in (1), where decisions are sunk for some number of months, m. Suppose that every period, a fraction $1/m$ of customers can make their decisions as in (1). Each customer makes its decision knowing that it must stick to the same technology for m periods. Each period, banks choose their adoption decision as in (5). However, there may be sunk costs of adoption, in the sense that costs will be lower if the bank has adopted ACH previously.

Given this model, a static implication is that usage will be correlated cross-sectionally, as we discuss in Section 3. However, there are other dynamic implications that we can test. One implication is that current usage and hence current adoption will be correlated with lagged adoption. Another implication is that customers care about future usage by other customers, since usage decisions are made to maximize expected surplus over the m month period.

In order to examine the importance of dynamics, we perform a regression of adoption on past and future adoption decisions. Specifically, we estimate:

(14) $$A_{j,t} = \{\lambda_7 \#(j,t) + \lambda_8 \#(j-6,t) + \lambda_9 \#(j+6,t) + \lambda_{10} A_{j,t-6} + \lambda_{11} A_{j,t+6} + \delta_t + \gamma x_{j,t} + \varepsilon_{j,t} > 0\}.$$

Equation (14) differs from our base model (6) in two ways: first, it includes 6-month lagged and future decisions, and second, we do not include the bank-level fixed effects, α. We choose 6 month lagged and future values and not earlier lags in order to avoid problems of serial correlation, which will confound any structural analysis of sunk costs. The reason for not using bank-level fixed effects is that one cannot consistently estimate a Chamberlain-style logit model with lagged dependent variables, using the standard conditional maximum likelihood estimator. We do control for MSA-level fixed effects and assets and deposits in (14).

We present two specifications of (14) in Table 6. We first control for all four lagged and future usages (Model 1) and also control for only $A_{j,t-6}$ and $\#(j+6,t)$ (Model 2). We choose the second specification, as these are the only two that should enter the structural adoption decision, in theory.

We find that own future and lagged adoption (λ_{10} and λ_{11}) are significantly positive. This suggests that there are large sunk costs of adoption. However, as we are not controlling for bank-level fixed effects in this specification, this may be partly or fully due to bank characteristics.

We also find that competitor lagged and future usages (λ_8 and λ_9) are significantly negatively correlated with the adoption decision. Current competitor usage (λ_7) is still significantly positive and much larger in magnitude than these two. If the error terms are not serially correlated after six months, then we can view the lagged and future decisions as exogenous. Thus, the results suggest that banks are likely to use ACH in months where other banks are using ACH. However, the network externalities do not appear to induce persistence in usage, or any sort of long-run effect. Indeed the negative sign suggests that perhaps the banks are coordinating on a different standard instead of ACH.

6. Conclusions

In this paper, we displayed a simple theoretical model of technology usage with network externalities, and used this model, together with detailed panel data, to test for network externalities in the Federal Reserve automated clearinghouse (ACH) payment system. Our tests reveal strong and consistent evidence of network effects. This finding of network effects occurs even though we are able to control for improvements in technology, price changes and variations in preferences from individual banks and customers, using our detailed data. The conclusion is robust to a variety of specifications. Moreover, we also find strong evidence that these network effects are not internalized and that the externalities are large in magnitude.

We extend our model to analyze whether informational or usage externalities are more important. The evidence suggests that both informational externalities and usage externalities are important. We also consider dynamic decision-making, but find no evidence that there is any sunk component of the network benefit.

As our results strongly indicate the presence of network externalities, we draw two policy implications. First, it is indeed possible that ACH is underused relative to its socially optimal usage level. As some part of the network externality is due to informational externalities, Federal Reserve policy should attempt to encourage usage through informative advertising and other policies that increase information about ACH. Second, other high-technology industries may also be characterized by informational externalities. Further policy recommendations will have to depend on estimating the underlying structural parameters of the network benefit in the utility function. As this paper lays out a theoretical framework for examining network externalities, it also suggests a future research strategy whereby the size of the network effects can be directly estimated in order to draw further policy conclusions.

References

Berndt, Ernst R., Robert S. Pindyck, and Pierre Azoulay (1999). "Network Effects and Diffusion in Pharmaceutical Markets: Antiulcer Drugs," NBER Working Paper #W7024 (March).

Cabral, Luis and Antonio Leite (1992). "Network Consumption Externalities: The Case of Portuguese Telex Service," in *The Economics of Information Networks*. Amsterdam: Elsevier.

Chamberlain, Gary (1980). "Analysis of Covariance with Qualitative Data," *Review of Economic Studies* 47: 225-38.

Economides, Nicholas and Charles Himmelberg (1995). "Critical Mass and Network Size with Application to the U.S. Fax Market," Discussion Paper EC-95-11, NYU Stern School of Business.

Farrell, J. and G. Saloner (1985). "Standardization, Compatibility, and Innovation," *RAND Journal of Economics* 16: 70-83.

Gandal, Neil (1994). "Hedonic Price Indexes for Spreadsheets and an Empirical Test for Network Externalities." *RAND Journal of Economics* 25: 160-70.

Gandal, Neil, Michael Kende and Rafael Rob (1997). "The Dynamics of Technological Adoption in Hardware/Software Systems: The case of Compact Disc Players," Sackler Institute of Economics Studies Working Paper 21-97.

Goolsbee, Austan and Peter Klenow (1998). "Evidence on Network and Learning Externalities in the Diffusion of Home Computers." Mimeo, University of Chicago, GSB.

Greenstein, Shane (1993). "Did Installed Base Give an Incumbent Any (Measurable) Advantages in Federal Computer Procurement." *RAND Journal of Economics* 24; 19-39.

Katz, M.L. and C. Shapiro (1986). "Technology Adoption in the Presence of Network Externalities," *Journal of Political Economy* 94: 822-41.

Kwast, Myron L., Martha Starr-McCluer, and John D. Wolken (1997). "Market Definition and the Analysis of Antitrust in Banking," *The Antitrust Bulletin* 42: 973-95.

Park, Sangin (1997). "Quantitative Analysis of Network Effects in Competing Technologies: the VCR Case," mimeo, SUNY at Stony Brook.

Roberds, William (1998). “The Impact of Fraud on New Methods of Retail Payment,” *Federal Reserve Bank of Atlanta Economic Revie*w: 42-52.

Rysman, Marc (1998). “Competition Between Networks: A Study of the Market for Yellow Pages” Mimeo, University of Wisconsin.

Saloner, Garth and Andrea Shephard (1995). “Adoption of Technologies with Network Effects: An Empirical Examination of the Adoption of Automated Teller Machines.” *RAND Journal of Economics* 26: 479-501.

Weinberg, John A. (1997). “The Organization of Private Payments Networks.” *Federal Reserve Bank of Richmond Economic Quarterly* 83: 25-43.

Wells, Kirsten E. (1996). “Are Checks Overused?” Federal Reserve Bank of Minneapolis *Quarterly Review* (Fall): 2-12.

Table 1: Number of banks using ACH each month, by ACH function and bank size

Month	Originators						Receivers					
	Using ACH			Not using ACH			Using ACH			Not using ACH		
	small	med	large	small	med	large	small	med	large	small	med	large
Mar-95	3309	1979	509	4305	779	228	6761	2321	527	853	437	210
Jan-96	3393	2089	537	4060	798	232	6664	2454	565	789	433	204
Jul-96	3587	2104	542	3909	749	218	6823	2471	569	673	382	191
Jan-97	3955	2317	579	3464	614	180	6858	2662	663	561	269	96
Feb-97	3976	2316	586	3451	603	177	6872	2649	667	555	270	96
Mar-97	4074	2302	584	3364	600	185	6891	2639	671	547	263	98
Apr-97	4142	2328	593	3282	580	184	6875	2641	681	549	267	96
May-97	4142	2336	598	3293	564	176	6875	2638	675	560	262	99
Jun-97	4210	2342	593	3275	525	164	6918	2616	659	567	251	98
Jul-97	4288	2346	592	3194	524	165	6912	2618	658	570	252	99
Aug-97	4313	2330	596	3178	531	161	6924	2605	661	567	256	96
Sep-97	4352	2305	595	3158	544	155	6947	2595	657	563	254	93
Oct-97	4385	2369	593	3085	515	162	6907	2634	658	563	250	97
Nov-97	4298	2354	596	3164	530	167	6892	2634	663	570	250	100
Dec-97	4368	2372	588	3096	516	169	6875	2647	657	589	241	100

Note: Authors' calculations based on the sample. Small banks are banks with assets greater than $10 million and less than $100 million, medium with assets greater or equal to $100 million and below $500 million, and large with assets greater or equal to $500 million.

Table 2: Number of banks that enter and exit ACH origination each month

Month	Entrants	Exitors	Month	Entrants	Exitors
Mar-95		588	Aug-96	595	418
Apr-95	420	404	Sep-96	457	360
May-95	599	521	Oct-96	551	383
Jun-95	507	548	Nov-96	432	392
Jul-95	446	379	Dec-96	397	332
Aug-95	548	475	Jan-97	454	353
Sep-95	482	515	Feb-97	380	331
Oct-95	476	467	Mar-97	413	317
Nov-95	530	510	Apr-97	420	337
Dec-95	529	430	May-97	350	351
Jan-96	520	426	Jun-97	420	296
Feb-96	516	451	Jul-97	377	323
Mar-96	508	418	Aug-97	336	312
Apr-96	496	496	Sep-97	325	265
May-96	463	496	Oct-97	360	408
Jun-96	460	434	Nov-97	309	264
Jul-96	492	383	Dec-97	344	

Note: Authors’ calculations based on the sample.

Table 3: Test of network effects

Unit of observation: Bank / month
Dependent variable: 1 if a bank offers ACH origination, 0 otherwise.
Monthly dummies for all 34 months are included in regressions.
Bank fixed effects are included in models 1-3.
Asset size category dummies[21] are included in models 4-6.

	Estimation with bank-level fixed effects (Conditional maximum likelihood estimation)		
Model:	Model 1: Effect of usage	Model 2: Effect of volume	Model 3: Effect of both
Regressor: $\#(j,t)$ (usage)	0.174*** (0.005)		0.172 (0.005)
Regressor: $\#_Q(j,t)$ (volume)		1.23×10^{-7}*** (1.20×10^{-8})	-1.19×10^{-7}*** (1.54×10^{-8})
Cond. log likelihood:	-62,811	-63,700	-62,782
	Estimation with asset size category dummies (Maximum likelihood estimation)		
	Model 4: Effect of usage	Model 5: Effect of volume	Model 6: Effect of both
Regressor: $\#(j,t)$ (usage)	0.005*** (0.0002)		0.016*** (0.0003)
Regressor: $\#_Q(j,t)$ (volume)		-1.95×10^{-8}*** (1.17×10^{-9})	-8.70×10^{-8} (1.75×10^{-9})***
Regressor: Assets (Thousands of $)	-1.86×10^{-7}*** (1.22×10^{-8})	-1.74×10^{-7}*** (1.22×10^{-8})	-1.65×10^{-7}*** (1.23×10^{-8})
Regressor: $Assets^2$	2.35×10^{-15}*** (2.32×10^{-16})	2.31×10^{-15}*** (2.30×10^{-16})	2.29×10^{-15}*** (2.27×10^{-16})
Regressor: Deposits	3.28×10^{-7}*** (1.76×10^{-8})	3.14×10^{-7}*** (1.75×10^{-8})	2.99×10^{-7}*** (1.76×10^{-8})
Regressor: $Deposits^2$	-5.22×10^{-15}*** (4.41×10^{-16})	-5.15×10^{-15}*** (4.37×10^{-16})	-5.10×10^{-15}*** (4.31×10^{-16})
Log likelihood:	-205,899	-206,082	-204,674

** significant at the 10 percent level
*** significant at the 1 percent level

[21] Assets were partitioned into 100 equal-length intervals. Dummy variables were set to 1 if the bank's assets fell within a given interval. The intervals are (<$20M, $20-$40M, ..., $1.8B-$2B, >=$2B). This same system is used in subsequent tables.

Table 4: Test of externalities

Unit of observation: MSA or non-MSA county / month (Models 1 – 3). Bank / month (Models 4 – 6).
Dependent variable: percent of deposits at banks that use ACH (Models 1 – 3). 1 if a bank offers ACH origination, 0 otherwise.(Models 4 – 6).
Monthly fixed effects for all 34 months are included in regression.

Estimation of network level % of assets using ACH on HHI			
	Model 1: Dummies for asset size categories	Model 2: Dummies for asset, population, and deposit categories	Model 3: Fixed effects estimation
Regressor: MSA/county HHI	0.150*** (0.008)	0.176*** (0.008)	0.069 (0.043)
Regressors: Total assets Total assets squared Total deposits Total deposits squared	Not included	Not included	Included in regression
R^2:	0.13	0.15	0.011
Estimation of bank level usage on HHI			
	Model 4: Dummies for asset size categories	Model 5: Dummies for asset and deposit categories	Model 6: Fixed effects (conditional logit) estimation
Regressor: MSA/county HHI	0.070*** (0.026)	0.030 (0.026)	0.821** (0.480)
Regressors: Assets Assets squared Deposits Deposits squared	Not included	Not included	Included in regression
Log likelihood:	-203,538	-206,440	-63,692 (cond. likelihood)

Note: Population dummies are at intervals of 40,000 people; asset and deposit dummies are at intervals of $20M. Asset and deposits dummies are for the MSA/county, for models 1 and 2, and for the bank, for models 4 and 5.

Table 5: Tests for Informational vs. Usage Externalities

Unit of observation: Bank / month
Dependent variable: 1 if a bank offers ACH origination, 0 otherwise.
Monthly fixed effects for all 34 months are included in regressions.

Model:	Model 1: Bank fixed-effects (Conditional ML estimation)	Model 2: Asset category fixed-effects (ML estimation)
Regressor: $\#(j,t)$ (orig. usage)	0.149*** (0.006)	-0.007*** (0.001)
Regressor: $\#^{R}(j,t)$ (recip. usage)	0.081*** (0.007)	0.009*** (0.001)
Regressor: Assets (Thousands of $)		-1.85×10^{-7}*** (1.22×10^{-8})
Regressor: Assets^2		2.34×10^{-15}*** (2.33×10^{-16})
Regressor: Deposits		3.26×10^{-7}*** (1.76×10^{-8})
Regressor: Deposits^2		-5.18×10^{-15}*** (4.41×10^{-16})
Log likelihood:	-62,736 (cond. likelihood)	-205,848

Table 6: Tests for sunk costs of decision-making

Unit of observation: Bank / month starting at month 7 (Sep-96) and ending at month 28
Dependent variable: 1 if a bank offers ACH origination, 0 otherwise.
Monthly fixed effects for all 22 months and for asset-size categories are included in regressions.

	Model 1: Control for all lagged usage	Model 2: Control for own lagged and other future usage
Regressor: $\#(j,t)$ (current usage)	0.068*** (0.003)	0.030*** (0.002)
Regressor: $A_{j,t-6}$ (lagged own usage)	2.681*** (0.016)	3.570*** (0.014)
Regressor: $\#(j,t-6)$ (lagged other firms' usage)	-0.014*** (0.002)	
Regressor: $A_{j,t+6}$ (future own usage)	2.842*** (0.016)	
Regressor: $\#(j,t+6)$ (future other firms' usage)	-0.052*** (0.003)	-0.029*** (0.002)
Regressor: Assets (Thousands of $)	$-9.69\text{x}10^{-9}$ ($1.73\text{x}10^{-8}$)	$-4.72\text{x}10^{-8}$*** ($1.58\text{x}10^{-8}$)
Regressor: Assets^2	$1.79\text{x}10^{-16}$ ($1.85\text{x}10^{-16}$)	$5.50\text{x}10^{-16}$*** ($1.94\text{x}10^{-16}$)
Regressor: Deposits	$5.79\text{x}10^{-8}$** ($2.68\text{x}10^{-8}$)	$1.15\text{x}10^{-7}$*** ($2.32\text{x}10^{-8}$)
Regressor: Deposits^2	$-8.22\text{x}10^{-16}$** ($4.12\text{x}10^{-16}$)	$-1.62\text{x}10^{-15}$*** ($3.86\text{x}10^{-16}$)
Log likelihood:	-60,565	-78,467

Figure 1: Per-item origination fees for Federal Reserve ACH processing, 1985-1997