

Network Externalities and Technology Adoption: Lessons from Electronic Payments

Gautam Gowrisankaran

Department of Economics
University of Minnesota,
Federal Reserve Bank of San Francisco,
and NBER
gautam@econ.umn.edu

Joanna Stavins

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

This Version: May 3, 2002

Abstract:

We seek to analyze the extent and sources of network externalities for the automated clearinghouse (ACH) electronic payments system using a quarterly panel data set on individual bank adoption and usage of ACH. We provide three methods to identify network externalities using this panel data. The first method identifies network externalities from the clustering of ACH adoption. The second method identifies them by examining whether banks in areas with higher market concentration or larger competitors are more likely to adopt ACH. The third method identifies them by examining whether the ACH adoption by small branches of large banks affects the adoption by local competitors. Using fixed effects and panel data these methods separately identify network externalities from technological advancement, peer-group effects, economies of scale and market power. We find evidence that the network externalities are moderately large.

We thank Dan Akerberg, V.V. Chari, April Franco, Austan Goolsbee, Ed Green, Justine Hastings, Tom Holmes, Pete Klenow, Matt Mitchell, Andrea Moro, Rob Porter, two anonymous referees and seminar participants at UCLA, Duke, Johns Hopkins, NYU/Stern, and the Federal Reserve Banks of Boston, Cleveland, New York and San Francisco for helpful comments. We thank Anita Todd for editorial assistance and Marie Willard for 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, or the Federal Reserve Banks of Boston or San Francisco.

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 and usage. 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 deposits or mortgage deductions. During our sample period, the Federal Reserve System processed approximately 75 percent of ACH transactions. 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. We perform our estimation using an 11-quarter panel (1995:Q2 to 1997:Q4) of the number of ACH transactions for each of the individual financial institutions that purchased ACH services from the Federal Reserve.

Our results have some potentially important implications for payment systems policy. In an age when computers and technology have become prevalent, only a tiny fraction of payments

¹ Other authors, such as Roberds (1998) and Weinberg (1997) have discussed network effects in payments markets.

are completed using electronic payments systems. There are at least two possible explanations for this: current electronic products may simply be preferred less, at the current prices, to cash and checks for most types of transactions; or network externalities may exist. Only if network externalities exist are electronic payment products being underused at their current prices. Thus, the two explanations have very different policy implications: the first calls for laissez-faire policies, including market pricing for electronic payments products, while the second suggests that there may be a need for policy interventions such as aggressive marketing efforts or pricing below marginal cost.

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 for ACH, which simply do not exist for these other industries, our study may shed light on the extent of network externalities for other industries. Thus, our work also functions 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,² there have been comparatively few empirical analyses of network effects.³ 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

² For instance, Farrell and Saloner (1985) and Katz and Shapiro (1986) have examined the equilibrium adoption of new technologies with network externalities.

³ We consider network externalities for the case of homogeneous networks, such as fax machines or e-mail. A separate 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 Gandall (1994), Greenstein (1993) or Saloner and Shephard (1995) for examples of these, respectively.

by this identification issue.⁴ Some recent studies of network externalities have instead made use of regional geographical cross-sectional data.⁵ 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.⁶

We make use of panel data with many observations at any time period, geographic data to measure the distance between banks and a simple theoretical model of technology adoption for ACH. The model specifies that banks in a network simultaneously choose whether or not to adopt ACH, based on the preferences of their customers and the adoption decisions and ACH volumes of other banks. With network externalities, a customer's value from using ACH is increasing in the number of customers at other banks that are using ACH. This leads to an interdependence in preferences across banks and hence to a simultaneity in the equilibrium adoption decisions of banks.

Because of the simultaneity of adoption decisions, it is not straightforward to identify network externalities. However, this paper develops three novel methods of identification. The first method is to examine the partial correlations of adoption decisions for banks within a network. The second method is to use excluded exogenous variables based on bank size to control for the endogenous adoption decisions. The third method is to create a quasi-experimental source of variation, using the adoption decisions of small, remote branches of banks. By using panel data and detailed geographic variation, these methods are robust to some of the pitfalls from other methods of identifying network externalities. We now briefly summarize the three methods and results.

Our first method is to regress the decision to adopt ACH on the adoption decisions of nearby banks, controlling for bank and time fixed effects. As we use fixed effects, we identify

⁴ See Cabral and Leite (1992), Economides and Himmelberg (1995) and Park (1997). One recent study (Gandal, Kende and Rob (2000)) 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.

⁵ For instance, Goolsbee and Klenow (1998) examine network externalities in computers and Rysman (1998) considers network externalities for Yellow Pages telephone books.

⁶ 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. Manski (1993) has denoted this inability to estimate the treatment effect as the reflection problem.

network externalities when *increases* in usage *above the general time trend* are correlated for nearby banks. The identifying assumption that justifies this test is that the time-bank specific profit shocks are uncorrelated across banks in a network after controlling for fixed effects. Even though this covariance restriction is nonlinear, we show that we can use simple fixed effects methods to test for the presence of network externalities, though not to estimate their magnitudes. This test is similar to tests done by cross-sectional studies that look for clustering of usage. However, since we use panel data, we control for peer group effects with fixed effects for each bank (which cannot be done using cross-sectional data) and for technological improvement with time fixed effects (which cannot be done using time series data). Performing this test, we find significant evidence of network externalities. More structural methods can also be used to find the magnitude of the network externalities with this assumption.

Our second method is to use a bank's size, measured by assets or deposits, as an exogenous variable, that can be excluded from other banks' adoption decisions and, hence, can control for the bank's adoption decision. The rationale is that, all else being equal, a larger bank is more likely to adopt ACH because it can spread the fixed cost of adoption over more customers, but that its size will not *directly* affect the adoption decisions of other banks in its network. A related exogenous variable to size is market concentration. Concentration, which is a function of sizes, is a particularly useful way of encapsulating this source of exogenous variation, since our model predicts that the equilibrium effect of concentration on adoption is positive. The reason for this is that a monopoly network, or to a lesser extent an oligopoly network, is able to internalize the network externality and, hence, is more likely to adopt ACH. We exploit this variation in two ways. First, we regress the decision to adopt ACH on the market-level concentration after controlling for the bank's own size. We find a positive and significant effect, consistent with the presence of network externalities and can reject market power and economies of scale as the reasons for the result. Second, we regress the decision to adopt ACH on the adoption decisions of other banks in the network, using the deposits of these banks and a Herfindahl-Hirschman Index (HHI) as instruments for their adoption decisions. We find a positive, significant, and large structural effect of other banks' adoption decisions on adoption.

Our final method is to create a sample where the adoption decisions of nearby banks can plausibly be considered exogenous and, using this sample, regress the decision to adopt ACH on the adoption decisions of nearby banks. To create such a sample, we exploit the fact that banks with multiple branches either adopt ACH for all of their branches or for none. This creates a quasi-experimental source of variation for local banks in networks where all the other banks are small branches of large banks. The rationale is that if a branch is a sufficiently small part of the entire company, a large bank will not change its ACH adoption strategy in response to the local bank's decision of whether or not to adopt ACH. In such cases, the large banks' adoption decisions can be viewed as exogenous to the error terms of the small bank, even in the presence of network externalities. Defining "small branches" as branches with less than 5 percent of the total deposits, we estimate the structural adoption decision on a subsample of local banks in networks with only small branches. We find network externalities that are moderately large though smaller than those found using our instrumental variables method above. For instance, an exogenous increase of 10 percentage points in the deposit-weighted percent of banks adopting would cause a representative bank to increase its probability of adoption by 4.4 percentage points. Using these estimates of the structural parameters, we also compute the Nash equilibria adoption decisions for stylized networks of different sizes. We find that the Pareto-best Nash equilibria for networks with two or more banks have an adoption probability 20 percentage points lower than the first-best outcome.

The remainder of this paper is divided as follows. In Section 2, we describe the data. Section 3 discusses our model and its implications for identifying network externalities. Section 4 provides results. Section 5 concludes.

Section 2: Data

Our principal data set is the Federal Reserve's billing data that provides information on individual financial institutions that process their ACH payments through Federal Reserve Banks.⁷ We observe quarterly data for 1995:Q2 through 1997:Q4. The data set lists the billing

⁷ We thank the Federal Reserve's Retail Payments Product Office for making this data set available to us.

information for all ACH transaction originations. ACH transactions can be one of three types: credit, debit or return. A credit transaction is initiated by the payer; for example, 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. Return transactions, which are much less common than credit or debit transactions, occur in cases of billing irregularities or disputes. For each financial institution in the data set, we have the ACH volume processed through the Federal Reserve each quarter as well as the American Banking Association (ABA) number that allows us to link these data with other publicly available banking data.

A financial institution that plans to offer Federal Reserve ACH services to its customers has to bear initial fixed costs. To transmit ACH information to the Fed, financial institutions other than the largest banks generally obtain a FedLine connection. Most banks use the FedLine connection for other financial services as well. The startup FedLine costs include buying a dedicated PC, modem, printer, and a special security card for encryption. The card costs approximately \$1,000. The above costs are shared among all FedLine-related services. In addition, banks buy special software for ACH processing from a vendor, at a cost of approximately \$2,000. After the installation, banks pay a monthly FedLine fee of between \$75 and \$2,000, depending on the type of connection. In addition, banks have to bear the costs of training their personnel. The Federal Reserve provides training at a cost of approximately \$150 per day. Except for employee training costs, these capital costs are unlikely to be sunk upon exit. Some of the largest banks do not use FedLine to transmit ACH information to the Fed, but instead use a direct computer interface connection between their mainframe computer and a Fed host mainframe.

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 (between two different banks) commercial ACH transactions processed in 1996.⁸ The remaining share of the

⁸ National ACH Association (NACHA) and Federal Reserve data. Government transactions constituted another 600 million.

on-others market was handled by three private sector ACH providers: Visa, New York Automated Clearing House (now called Electronic Payments Network), and American Clearing House (formerly Arizona Clearing House). There are some network linkages among 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.

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 Consolidated Reports of Condition and Income (Call Reports) database. The Call Reports database provides information on bank assets, deposits, name, ABA numbers, and the zip code of the headquarters for all banks that are registered with the Federal Deposit Insurance Corporation (FDIC). Several banks opened and closed during our sample period. We kept these banks in the sample for the quarters in which they were open.

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 numbers when billing the Federal Reserve for ACH services. Thus, we aggregated the Federal Reserve ACH volume up to the FDIC number level. One data problem that we encountered is that a substantial fraction of the ABA numbers from the ACH billing data were not in the Call Reports database. Most of the ABA numbers that did not match are credit unions or thrifts. We further excluded all banks with deposits of less than \$10 million for all quarters in the sample and all remaining credit unions. We were left with approximately 11,000 banks over the 11-quarter sample period.

We also merged our data set with the annual Summary of Deposits database for 1995-1998. This database, which is issued for the second quarter, provides the zip code and total deposits for each bank branch at the FDIC certificate level. The Summary of Deposits database is useful because, unlike the Call Reports, it provides branch-level information. However, since it is at the annual, not the quarterly, level it is more difficult to identify fixed effects when using it. We kept all banks from the Summary of Deposits data that are also in our Call Report sample.

In order to find the metropolitan statistical area (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 \$0.010 in 1997. In addition to these per-transaction fees, banks also paid a file fee of \$1.75 or \$6.75 depending on the number of transactions in the file. 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.

We assume that a bank has adopted ACH in a given quarter if it originated at least one ACH transaction. Our analysis is based on the banks in our sample that adopted ACH in a given quarter. Table 1 lists the number of banks of different deposit sizes that have adopted ACH. One can see from Table 1 that the number of banks that use ACH grew during our sample period. Nonetheless, the number of banks that did not use ACH is substantial. Larger banks are more

likely to adopt ACH than smaller banks, consistent with them spreading the fixed cost of adoption over more customers.

Table 2 lists the number of exiters and entrants in each quarter. An exiter is defined as a bank that has adopted ACH in the current quarter, but not in the next quarter. In addition to the large number of entrants, there are many exiters. This suggests two things. First, the sunk costs of adoption are low. Second, some of the exit may be due to mismeasurement, for instance, banks that technically offer ACH but have no transactions in a given quarter. Because we were concerned about the potential mismeasurement, we asked managers from the Federal Reserve Bank of Minneapolis ACH department to review a small subsample of banks that appeared to both enter and exit. From our discussions with the managers, we believe that about one fourth of the subsample are banks that stop originating ACH transactions and another fourth are exiting or merging banks. The remaining half the banks that enter and exit appear to be banks that originate a small number of ACH transactions, often return items, and so do not have volume in some quarters. Because some of these banks may not have performed any ACH origination transactions in a given period but may have offered ACH services if an interested client was available, it is debatable whether these banks should be considered to be entering and exiting. To alleviate concerns about this potential misspecification of our definition of entry and exit, we provide robustness checks of the results with an alternate definition of entry and exit in Section 4.1.

Tables 3 and 4 provide some evidence on whether the data support network externalities, using our first and second methods of identifying network externalities respectively, noted above. Both tables make use of the fraction of banks adopting ACH within an MSA or a non-MSA county. We weigh the fraction by the deposits for all branches within the MSA, using the Summary of Deposits data. Table 4 also uses a HHI concentration measure of banks within an MSA/county, which we create from the branch-level Summary of Deposits data.

Table 3 provides information on whether *changes* in adoption decisions are correlated for banks within an MSA/county, in keeping with a fixed effects identification of the clustering of ACH adoption. The table shows the percent of banks adopting in the middle and last periods conditional on the adoption fraction in the first period. One can see sizeable evidence of

correlations of changes in adoption decisions. For instance, in markets for which 0-19% or 20-39% of the banks adopt ACH in the first period, if more banks adopt ACH in the middle or last period, then it is most likely that the vast majority (80-100%) of them adopt. In the absence of within-network correlations, we would expect a much more gradual change in adoption decisions within a network.

Table 4 provides information on the relation between HHI and adoption decisions for MSAs/counties. We show the variation within narrow population ranges to control for economies of scale from larger-sized markets. The table shows that, within a population range, more concentrated markets have a larger fraction of banks adopting ACH. For instance, in the 20,000-40,000 population range, banks with an HHI of 0.1-0.2 have a 0.766 adoption fraction, while banks with an HHI of 0.4-1.0 have a 0.862 adoption fraction. Additionally, the table shows that larger MSAs/counties have less concentrated banking structures and more ACH adoption, consistent with larger banks adopting more frequently, and suggesting the importance of controlling for size.

Section 3: The Model and Testable Implications

3.1 The Model

We develop a simple static model of technology adoption that illustrates how network externalities, that is, interdependent preferences for ACH usage by banks or their customers, translate into testable implications. While our base model is restrictive, its testable implications will be robust to several alternatives, as we discuss below in Section 3.5.

Consider a network with J banks, each of which must simultaneously decide whether or not to adopt ACH. Let (A_1, \dots, A_J) denote the vector of 0-1 bank adoption decisions. Each bank has a fixed set of customers and acts as an agent for its customers. Thus, we do not consider

market power or allow for customers to choose among banks.⁹ Following the adoption decisions, the customers simultaneously decide whether or not to use ACH. A customer that is at a bank that has not adopted ACH cannot use ACH.

We model the customer decision as a discrete choice between checks and ACH; let $Usage_{i,j}$ equal 1 if consumer i at bank j chooses ACH and 0 if the consumer chooses a check.¹⁰ Normalize the utility of checks to 0. Then, let the utility of using ACH be:

$$(1) \quad u_{i,j}(Usage_{-i,j}) = \gamma^c X_{i,j} + f(\beta^c, Usage_{-i,j}),$$

where $X_{i,j}$ are observable variables that affect the utility from using ACH, $Usage_{-i,j}$ denotes the usage decisions of other customers, both at bank j and at other banks, and β^c and γ^c are consumer-level parameters. Consumers will use ACH if the utility is positive and if their bank has adopted ACH:

$$(2) \quad Usage_{i,j}(Usage_{-i,j}) = \{A_j = 1 \text{ and } u_{i,j} > 0\},$$

where $\{\cdot\}$ is an indicator function. With network externalities, we expect that β^c will be a positive coefficient vector, so that increased usage by other customers increases the utility from using ACH.

As we have a two-stage game, we consider subgame perfect equilibria. Let us first condition on the bank adoption decisions and examine the second-stage usage decision.

Existence of Nash equilibrium for this subgame can easily be shown by construction. One important feature is that there may be multiple equilibria, characterized by tipping behavior, as in Farrell and Saloner (1985). For instance, with large network externalities, there may be a Nash equilibrium where all customers use ACH, and another equilibrium where no one uses ACH. Although equilibrium is not necessarily unique, it turns out that there is one equilibrium that Pareto dominates all other equilibria. This equilibrium has a greater level of ACH adoption than

⁹ We justify this assumption with 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.)

¹⁰ This assumption is an approximation to the real world, where there are other payment mechanisms, such as cash, debit cards and credit cards. Moreover, in the real world, consumers may use services, such as CheckFree, where the actual payment mechanism is unknown to the consumer.

any other equilibrium, is weakly increasing in bank ACH adoption, and is easy to compute. Formally,

Proposition 1: Assume that $\partial u_{i,j}/\partial \text{Usage}_{-i,j}$ is strictly positive. Then, given adoption decisions $A = (A_1, \dots, A_J)$, there exists a unique Nash equilibrium $\text{Usage}^P(A)$ of the usage game such that $\text{Usage}^P(A)$ Pareto dominates all other Nash equilibria. Moreover, if $A' = (A'_1, \dots, A'_J)$ satisfies $A' \geq A$ ¹¹ then $\text{Usage}^P(A') \geq \text{Usage}^P(A)$.

Proof: See Appendix.

Turning now to the first stage, we can write the profits from adopting ACH as a function of the other banks' adoption decisions, conditioning on the Pareto-dominating second-stage equilibrium:

$$(3) \quad \pi_j(A_{-j}) = g(X_j, \gamma, A_{-j}, \beta, \varepsilon_j),$$

where X_j are observable bank-specific variables that affect the profits from adoption, principally assets, deposits, price, and time and bank fixed effects, A_{-j} are the adoption decisions of other banks in the network, ε_j are unobservable bank-specific variables, β are the parameters on A_{-j} and γ are parameters on X_j and on interactions. We normalize the profit from not adopting ACH to 0 and choose a linear functional form of (3) for tractability. The adoption decision can then be expressed as:

$$(4) \quad A_j = \{\pi_j(A_{-j}) > 0\} = \{\gamma X_j + \beta h(A_{-j}) + \varepsilon_j > 0\},$$

where $h(A_{-j})$ encapsulates the network externality. We specify two functional forms for $h(A_{-j})$. First, we define $\#(j)$ to be the fraction of other banks in the network that have adopted ACH. Second, we define $\#^Q(j)$ to be ACH volume (in terms of number of transactions) among other banks in the network divided by the assets of these banks. The latter measure is intended to proxy for the percent of transactions performed with ACH. Given that $h(A_{-j})$ is increasing in

¹¹ Recall that adoption is a vector of 0s and 1s. Thus, the inequality means that the set of customers who use ACH in the second instance is a subset of the set of customers who use ACH in the first instance.

adoption, network externalities imply that β is positive. We use combinations of these two measures for the different specifications. Because the measures are not defined for networks with one bank, we excluded these networks for all specifications that use (4).

If consumers choose the Pareto-best equilibrium, then ACH adoption by an additional bank k will weakly increase both the profits to bank j from adopting ACH and the equilibrium ACH usage among customers of bank j . This follows from Proposition 1: adoption by bank k will lead to a weakly higher usage level overall, which will make all customers weakly better off. As we assume that banks capture the surplus of its customers, this will weakly increase profits. Thus, even if the network externality is at the consumer level, this will lead to an indirect network externality at the bank level. Moreover, by applying Proposition 1 to the bank problem, it follows that if we restrict consumers to always play the Pareto dominating equilibrium of their subgame, then the bank game will have a unique Pareto dominating equilibrium.

It is entirely possible that the network externality is at the consumer level and is due to increased acceptance of ACH or to learning across consumers. However, it is also possible that the network externality is at the bank level. In this case, it would be due to the direct effect of having more banks with which to exchange ACH payments and to it being cheaper for banks to exchange ACH payments if they can learn about procedures from other banks in the network. If there is a direct network externality component of this form, this will again lead to β being positive. Hence, our estimation will pick up both consumer- and bank-level network externalities and we cannot distinguish between them. However, we might hypothesize that $\beta(j)$ indicates more of a bank level externality and that $\beta^Q(j)$ indicates more of a consumer-level externality.

Our main equation of empirical interest is (4). Thus, we focus on estimating the network externality using the bank as the unit of observation. We have three different methods of identification which we discuss below.

3.2 Identifying Network Externalities from Clustering

Our first method of identifying network externalities is to test for clustering of adoption and usage decisions within a network using (4) and our panel data of bank adoption and volume.

Because we use panel data, X_j includes bank and time fixed effects, as well as assets, deposits and their squares. We assume that a bank is in the network with bank j if its headquarters is within 30 kilometers of the headquarters of bank j .

We make one major assumption for this identification: we assume that the ε 's are not correlated across banks in a network. 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 consistent with this assumption. Also, if ACH usage is increasing over time because of decreased price or increased acceptance of electronic payment methods, the quarterly fixed effects will increase 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.¹² We expect that the ε_j 's will capture unobserved time-varying factors such as an employer that newly decides to permit or encourage ACH for payroll. Importantly, we use the ACH origination volume throughout. If, instead, we had used the recipient volume, then such an employer adoption would generate correlations in recipient volume, even in the absence of network externalities, since different employees would receive the payments at different banks.

Even though this identifying assumption is nonlinear, we develop a test for network externalities based on it that uses standard fixed effects methods. Specifically, we test for network externalities by directly estimating (4) using either ACH adoption or ACH volume (in terms of number of transactions) per assets as the dependent variable.

With the above assumption, we can interpret the data as providing significant evidence of network effects if and only if the estimated β 's (the coefficients on other banks' ACH adoption and volume) are significantly positive. To see this, note that with no network externalities, the true value of β is 0. In this case, A_j is a function of bank j 's observed and unobserved characteristics but not of A_{-j} . Given the assumption that the ε 's are uncorrelated within the

¹² An example of a local time-specific effect would be a regional advertising campaign that stimulated demand in a particular area.

network, $h(A_{-j})$ will be exogenous to ε_j . Thus, with no network externalities, if we perform a regression of ACH adoption on exogenous characteristics and adoption decisions of other banks, the estimated β will be consistent, and should converge in probability to zero.

Now suppose that there are network externalities. In this case, A will be determined via a simultaneous Nash equilibrium. Since a high ε_k will cause bank k to be more likely to adopt, which will in turn cause bank j to be more likely to adopt, $h(A_{-j})$ will be endogenous. Thus, we have a classic simultaneity problem and the estimated β will not converge to the structural parameter but will instead reflect the equilibrium correlation between ACH adoption decisions. While the estimate of β cannot be interpreted as the structural parameter, the equilibrium correlation will be positive, and generally larger in magnitude than the structural parameter. Thus, if we find a significantly positive estimate of β , we can interpret this as evidence of network externalities.

It is important to note that this estimation of (4) will identify network externalities but not structurally estimate their magnitudes. Equilibrium methods can be used to recover the magnitudes of the structural β parameters with the same identification assumption; see Akerberg and Gowrisankaran (2002) for details.

3.3 Identifying Network Externalities from Excluded Bank Size Variables

Our second method of identification is to use X_{-j} , the exogenous characteristics of other banks in the network, as “instruments” to control for their endogenous adoption decisions A_{-j} , as suggested by Manski (1993). Specifically, in (4), X_{-j} does not enter into bank j 's adoption decisions, but does enter into the adoption decisions for other banks in the network. This suggests that X_{-j} , and functions of it, can be used as instruments for A_{-j} . The natural X_{-j} to use are functions of bank size, measured by assets and deposits. Since ACH is a very small percentage of bank transactions, it is plausible to view the sizes of other banks, in terms of assets

and deposits, as exogenous to the error term ε from (4), which is the assumption needed for identification.

There is an economic intuition behind using size as excluded exogenous variables. All else being equal, a bank's size should affect its decision to adopt ACH because larger banks can spread the fixed costs of adopting ACH over more customers. However, while bank j 's profits from adoption might be affected by bank k 's adoption decision and ACH usage, there is no reason to think that its profits from adoption would be directly affected by the size of bank k . Hence, bank k 's size will be correlated with bank k 's adoption decision, but will not directly enter into bank j 's adoption decision, making it a good instrument for bank k 's adoption decision. Importantly, since we control for a bank's own size, this identification will not confound network externalities with economies of scale.

A related exogenous variable to bank size is market concentration. Concentration, a function of bank sizes, is a useful way of encapsulating this source of variation because the model predicts that, with network externalities, the equilibrium effect of concentration on adoption is positive. To see this, consider a network with a number of banks. Each bank will adopt ACH based on whether its profits from adopting are positive, but it will not internalize the externality that its adoption has on profits for the other banks in the network. Since the externality is positive, this will yield underadoption relative to the monopoly adoption probability, since the monopoly can fully internalize the externality. Formally, we can show that if we merge the banks into a monopoly holding company with fixed costs from adopting at each of its member banks equal to the fixed costs before the merger, then the set of banks that adopt ACH is weakly greater.

We use size and concentration to identify network externalities with two separate regressions. First, we perform a reduced form regression of bank adoption on market concentration, where we use HHI as the measure of concentration. Specifically, we estimate:

$$(5) \quad A_j = \{\gamma X_j + \alpha \text{HHI}_j + \varepsilon_j > 0\}.$$

The HHI variable is constructed from deposits in the MSA/county in which the bank is located, as discussed in Section 2. We include time fixed effects and asset dummies (at intervals of \$20 million) as well as assets, deposits and their squares in X_j . In some versions of (5) we include

bank fixed effects. However, this is problematic since the variation in asset levels (and, hence, HHI) is small over the sample. It is important to remember that, unlike (4), (5) is a reduced form regression, where HHI_j is a function of X_j and X_{-j} . Nonetheless, we interpret a positive coefficient on α as evidence that concentrated markets are more likely to adopt ACH, which we attribute to network externalities.

Second, we estimate (4), using size and concentration as instruments for $h(A_{-j})$. Even though adoption is a discrete variable, we need to assume a linear functional form to apply instrumental variables. Hence, we use a linear probability model variant of (4):

$$(6) \quad A_j = \gamma X_j + \beta h(A_{-j}) + \varepsilon_j.$$

For our instruments, we choose mean assets for all banks within 30 kilometers of bank j as well as different MSA/county HHI variables. We calculate the HHI using two different data sources, the annual branch-level Summary of Deposits and the quarterly bank-level Call Reports. For both data sources, we define one HHI that includes bank j (as in (5)) and one that excludes bank j , resulting in a total of five instruments. We use time fixed effects and asset dummies in X_j and also try specifications with bank fixed effects. We perform (6) on the sample of banks with at least one other bank in the network.

3.4 Identifying Network Externalities from Quasi-Experimental Variation in Adoption

Our third method of identification is to find a quasi-experimental source of variation for A_{-j} in (4). While we cannot observe exogenous variation in adoption decisions in general, it is possible to observe exogenous variation for certain markets because of the fact that banks with multiple branches either adopt ACH for all of their branches or for none. To see why this creates a source of exogenous variation, consider the case of a small town with one small bank j and some small branches of large banks. No large bank is going to adopt ACH over all its branches in response to an adoption from the small bank, and hence its decisions are unaffected by A_j . Since A_k for large banks $k \neq j$ are unaffected by ε_j , they are exogenous variables. Essentially,

the adoption strategies of large banks form a “natural experiment” in these networks, which allows us to trace out the structural reaction function for the local bank.

To exploit this source of variation, we create a sample of local banks for which all the bank branches within 30 kilometers are small branches of large banks, where we define “small” as less than 5 percent of the total deposits of that bank. We exclude from this sample banks without at least one other bank branch in their network, since no meaningful inference can be drawn from this method for banks that are alone in their network.¹³

Using this sample, we estimate a version of (4). We do not observe the volumes of large banks’ ACH transactions specific to the network of bank j , hence we include $\#(j)$, the fraction of other banks that adopt, for the network benefit $h(A_{-j})$. Moreover, for this estimation, we compute $\#(j)$ by weighting it by the local deposits of the bank branches because of possible size disparities. While we cannot allow for bank fixed effects in X_j because of the small amount of data, we do include time fixed effects, deposits and deposits squared. In some specifications, we include deposits dummies at intervals of \$2 million.

Note that, in addition to assuming that adoption decisions for large multi-market banks are exogenous, we are again making the identifying assumption that the branch-level deposit sizes and branch entry decisions of these banks are exogenous. We feel that this is justified because ACH is a very small percentage of bank transactions. Also, even if the adoption decisions and deposits are exogenous, there may be more than one reason for adoption to increase in response to an exogenous change in another bank’s adoption. We will attribute it entirely to network externalities.

3.5 Robustness of the Results

We have presented a simple static model of network externalities, to illustrate the basis for our tests and estimation. There are several assumptions about our model that could be relaxed

and still yield the same testable implications, although relaxing these assumptions would yield different estimated magnitudes for the network externalities. We detail each of these below.

Size of network

We assumed a network size of 30 kilometers and also constructed HHI at the MSA/county level, as is common practice. Certainly, banks exchange payments with other banks with headquarters that are outside this limited area—indeed, the majority of bank payments are non-local. Thus, a bank’s true network may include banks that are outside of the 30-kilometer radius. If the size of the network is larger, then we will not capture the full extent of the network effect with our clustering regression. Since ACH usage 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 the specified size, then again, our measures of the network will be imprecise, but any network effect will be captured. Thus, our clustering identification is robust to different sizes of the network.

Market power

We assumed that banks capture all the surplus from customers and that customers are captive to one institution. This is consistent with either a perfectly price-discriminating bank or a bank that prices at marginal cost. We now consider a more general model where customers can switch between banks and where banks are profit-maximizers who set a fixed ACH price and, hence, are unable to capture all of the surplus from their customers. For our identification based on quasi-experimental variation, market power should not affect the coefficient estimates, since local market power in a small market is unlikely to affect overall adoption by a large bank. However, it is worth considering in more detail the implications of market power on our first two methods of identification.

¹³ We have tried variants where we include such networks and assume that $A_{-j} = 1$ for such cases, consistent with a monopoly network being able to internalize externalities as above. The results are not sensitive to the inclusion of such networks.

Consider first the implications of market power on our clustering method of identification. If customers can switch banks, then there may be reasons besides network externalities for adoption decisions to be correlated. To understand the equilibrium pattern of correlations, we need to assess whether adoption decisions are strategic complements, which would lead to positive clustering, or strategic substitutes, which would lead to negative clustering. In general, models of technology adoption without network externalities imply that adoption decisions are strategic substitutes; see Spencer and Brander (1983) for instance. The reason for this is that, if one bank adopts the technology, this lowers the value to other banks from adopting the technology, since their potential customer base will be smaller. Thus, standard models would have the opposite prediction from network externalities. An exception to this occurs when banks use ACH adoption to generate revenue for some other service as a loss-leader. However, there is no evidence of this type of behavior for ACH and, indeed, ACH volume is probably too small to affect other business that much. Moreover, while such models would predict a positive correlation on adoption decisions, they would predict a negative correlation on volumes, because it remains true that the customer base shrinks following entry.

Now consider the implications of market power on the excluded concentration method of identification. Suppose that firms in concentrated markets can charge a higher price for ACH. Then, even in the absence of network externalities, firms in concentrated markets might be more willing to adopt ACH because they can capture more of the consumer surplus from their adoption, all else being equal.¹⁴ This market power effect could be confounded with network externalities. However, the two models differ in one way. A firm with market power will have a lower equilibrium quantity conditional on adoption than a competitive firm because it is able to enter with less customers by charging a higher price per customer. This is the opposite implication from network externalities: with market power, equilibrium quantity conditional on adoption will be decreasing in concentration, while with network externalities, it will be increasing. We test for this by examining whether the ACH volume per assets conditional on adoption is increasing or decreasing in concentration. This result was obtained assuming that banks can choose only one price for ACH. Even if banks in concentrated markets could perfectly

price discriminate, they would still only obtain the same quantity conditional on adoption as competitive markets, which set the price at marginal cost. The overall point is, again, that it is difficult to find models besides network externalities that show an unmitigated positive effect from network externalities.

Dynamics

One assumption that we make 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 of different banks. 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 among current adoption decisions. Thus, our identification will capture the cross-sectional implications of the dynamic adoption model. In contrast, our discrete-choice estimation methods are not robust to serial correlation in the error structure, since we do not model serial correlation.

Section 4: Results of the Model

We present results corresponding to each of our three methods of identification, and an analysis of equilibrium adoption based on the estimated parameters.

4.1 Results Using Clustering Method of Identification

¹⁴ Studies such as Bresnahan and Reiss (1991) have attributed the higher probability of entry per capita by monopolists to this market power reason.

In Table 5 we present the estimates of the clustering tests using (4). We present results from six models. The first four use adoption as the dependent variable, while the last two use volume per assets. We specify the first four models with a logit and the last two models as a linear model. All of the models include bank fixed effects, time fixed effects, assets, assets squared, deposits and deposits squared.

As our data set contains approximately 11,000 banks, it is not computationally feasible to estimate a dummy variable for each bank. Instead, we estimate Models 1–4 by using Chamberlain’s (1980) conditional likelihood method. This method maximizes a *conditional* likelihood function, the probability of a particular choice sequence for a bank conditioning on the total number of times that the bank chose to adopt ACH. Chamberlain shows that the fixed effects do not enter into this conditional likelihood function and that this method gives consistent but not efficient estimates. For Models 5–6, we use a linear fixed effects specification, for which the fixed effects also do not enter into the criterion function, but which is numerically equivalent to estimating the 11,000 bank fixed effects using OLS.

We find that the fraction of other banks that adopt, $\#(j)$, has a significant and very positive correlation with a bank’s own adoption decision (Models 1 and 3).¹⁵ Because of potential data concerns with our definition of ACH adoption, we also experimented with an alternate definition of adoption, where we assume that a bank adopts ACH on the first date at which it has 20 transactions and exits if its transactions subsequently drop to 0. These results, presented in Model 4, show a coefficient on $\#(j)$ that is very similar to that from Models 1 and 3. 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.

It is not possible to interpret these coefficients as structural estimates. Moreover, with Chamberlain’s conditional logit model, it is not possible to estimate the bank fixed effects and, hence, one cannot find the magnitude of the partial correlation. Nonetheless, we can infer something about the magnitudes from the standard logit probability derivative formula,

¹⁵ We also separately estimated this model with linear probability and duration model specifications. Both models gave qualitatively similar results.

$d\text{pr}(A_j)/d\#(j) = \beta_{\#(j)} \text{pr}(A_j)(1 - \text{pr}(A_j))$. If we assume that $\text{pr}(A_j) = 0.5$, then from Model 1, a 10 percentage point increase in the fraction of other banks adopting ACH is associated with a $0.5 \times 0.5 \times 1.874 \times 0.1 = 0.047$ or 4.7 percentage point increase in the probability that bank j is using ACH,¹⁶ with a standard error of 0.2 percentage points. If $\text{pr}(A_j) \neq 0.5$, this implies a smaller magnitude than 4.7%. Since the mean adoption probability is 67.2%, the actual magnitude is likely to be somewhat smaller than 4.7%.

In contrast, the evidence on the effect of the number of transactions is much weaker (Models 2, 5, 6). While the three models show a positive impact of $\#(j)$ or $\#^Q(j)$ on ACH adoption or volume per assets, this effect is not significant in any of the specifications. Moreover, the magnitude of the effect can be assessed by noting that the mean number of ACH transactions per volume in our sample is normalized to be 0.031. Thus, using Model 3, a doubling in the volume per assets per volume by a mean bank is associated with an increase of $0.142 \times 0.031 \times 0.5 \times 0.5 = 0.0013$ or 0.13 percentage points in the adoption probability of a bank with $\text{pr}(A_j) = 0.5$. This number is economically quite small, in addition to being insignificant.

The fact that adoption is quantitatively much more important than the volume of transactions, as evidenced both by the magnitude and insignificance of coefficients, suggests that banks value other banks' adoption decisions rather than volume. Thus, the network externality at the consumer level seems to be less important than at the bank level. While informational externalities may be important among bank customers, banks are affected by direct usage externalities because they exchange transactions among themselves. As the network externality at the bank level is probably due more to direct usage externalities than informational externalities, this suggests that direct usage externalities are present.

4.2 Results Using Excluded Size and Concentration Method of Identification

Table 6 presents our results based on the concentration method of identification. All specifications include time fixed effects, assets, assets squared, deposits and deposits squared.

¹⁶ Note that these figures are approximate, since the formula applies only to the case of an infinitesimal change.

We start with reduced form regressions of HHI on concentration, which are in Models 1 and 2. Model 1 estimates a logit specification with categorical controls for bank size, while Model 2 includes bank fixed effects, using Chamberlain's conditional logit specification again.

Model 1 reveals significant evidence of network externalities. A higher value of HHI (i.e. less competition) is associated with a higher probability of using ACH. The magnitude of the effect is moderately large. Considering again a bank with a 50 percent probability of adoption, the model predicts that an increase in market concentration from three equal-sized banks to two equal-sized banks will lead to an increase in the adoption probability by roughly $0.5 \times 0.5 \times 0.1667 \times 0.392 = 0.0163$ or 1.63 percentage points, with a standard error of 0.2 percentage points. Model 2 shows a positive effect that is similar in magnitude to Model 1, though not significant. The likely reason for the lack of significance is that the amount of variation in the data is small.

As mentioned in Section 3.5, a potential alternative explanation for this observed pattern in the data is market power. In order to investigate whether market power can explain the pattern, Model 3 presents the results from a regression of ACH volume per assets on HHI, using as the sample banks with a positive volume. We implement this with a linear regression with bank fixed effects. We find that conditional on adoption, volume of ACH is significantly increasing in HHI. This result is not consistent with market power instead of network externalities explaining the increased ACH adoption in concentrated markets and, thus, it reinforces our findings that network externalities exist.

Lastly, in Models 4 and 5 we estimate instrumental variable specifications of (4) using bank size and concentration as instruments for the adoption fraction $\#(j)$. Model 4 presents the results with categorical controls for bank size while Model 5 specifies bank fixed effects. The specifications both show a significantly positive impact of $\#(j)$ on adoption, with a magnitude of 0.858 without bank fixed effects and 0.930 with fixed effects. In Model 4 we also present robust standard errors to control for the fact that the error terms in a linear probability model are heteroskedastic; robust standard errors do not substantially affect the conclusions. In addition to being significant, the estimated magnitudes of the coefficients are large: they imply that a 10

percentage point increase in the adoption fraction of other banks will increase the adoption probability by roughly 9 percentage points.

4.3 Results Using Quasi-Experimental Variation in Adoption Method of Identification

Table 7 presents our results for our estimation based on the subsample of banks where all other banks in the network are small branches of large banks. We have included two probit specifications. Both specifications include a full set of time fixed effects, bank deposits and bank deposits squared; Model 2 includes categorical variable controls for deposits. Since we are using the Summary of Deposits database, we do not have assets for this estimation. We also included a linear probability model specification to compare the results to the instrumental variables results from Table 6. For both estimation results, we have given regular standard errors as well as robust standard errors that account for clustering based on the fact that adoption decisions for the same bank in different periods may be correlated.

Recall that we construct the sample by keeping local banks with only small branches of large banks in their network, in order to allow for only exogenous variation in the adoption decision. Thus, our sample is small, with 237 observations in Model 1, of which 158 are kept in Model 2. Nonetheless, the models show significant evidence of positive and moderately large network externalities. For instance, if bank j has an adoption probability of 50 percent, then Model 1 implies that an exogenous increase of 10 percentage points in the deposit-weighted ACH adoptions of other banks would raise bank j 's adoption probability by $\phi(\Phi^{-1}(0.5)) \times 1.094 \times 0.1 = 0.044$ or 4.4 percentage points with a standard error of 1.4 percentage points using the robust method, where ϕ and Φ are the standard normal density and cumulative distribution function respectively. Lastly, Model 3 shows that a linear probability model yields a magnitude of 0.511 which is somewhat smaller than the magnitude of 0.9 from the instrumental variables regressions in Table 6 Models 4 and 5. The smaller size is likely due to the fact that our sample consists of banks in isolated, rural areas, which are likely to have a lower than average value from technologically intensive electronic payment mechanisms.

4.4 Implications of Network Externalities on Equilibrium Adoption

In order to further interpret the economic magnitude of the network externalities, we simulate equilibria for representative networks using (4) and the estimated parameters from Table 7, Model 1. Note that even though we used a specific quasi-experimental source of variation to estimate the parameters, we can use the estimated parameters to compute equilibria for any network configuration since the estimated parameters are structural. We compute equilibria for networks with N local banks, where each bank had the Table 7, Model 1 sample mean deposit size and existed at the time period 1996:Q2.

The equilibrium in a network depends on the realization of the unobservables $\varepsilon_1, \dots, \varepsilon_j$. For this simulation, we assume that these unobservables are i.i.d. across firms in a network. We are interested in finding the expected value of network adoption decisions, over the distribution of these unobservables. Hence, we simulate over the unobservables, and solve for the equilibrium adoption decisions for each vector of unobservables. The iterative technique described in the proof of Proposition 1 provides a natural method to solve for the Pareto-worst and -best equilibria given any vector of unobservables. To get an accurate picture of the adoption probabilities, we compute the mean adoption probability over 10,000 simulation iterations, with different draws on the unobservables for each iteration. We also compute the first-best adoption decisions for these draws, which is the adoption decisions that would be chosen if the firms could function as a perfectly colluding cartel. This can be computed similarly to the Nash equilibrium computation, with the alternate assumption that firms value the profits of all firms in the network when making their adoption decisions.

Table 8 reports the mean adoption probability for $N = 1, \dots, 6$, under the Pareto-worst, Pareto-best and first-best outcomes. For $N = 1$ we set $\#(j) = 1$. Note that the three outcomes are the same for this case since the monopoly will internalize any externality. We find that the Pareto-best Nash equilibrium outcome with $N = 1$ has banks adopting ACH 62 percent of the time, and that this falls monotonically in N to 44 percent with $N = 2$ and to 39 percent by $N = 6$. For $N > 1$, the Pareto-worst Nash equilibrium has a lower adoption probability that

increases from 30 to 33 percent as N increases from 2 to 6. In contrast, the first-best outcome involves a mean adoption frequency that is much higher, around 67 percent for $N > 1$.

While it is not possible to identify which equilibrium is observed in the data from this estimation, a central result is that the first-best adoption levels are much higher than any Nash equilibrium levels for $N > 1$, suggesting an economically important network externality. As a caveat, it is worth noting that these results pertain specifically to small banks. Most of the large banks have adopted ACH, and the results show no significant effect from having more volume at a bank that has already adopted.

We have not computed the deadweight losses from the different Nash equilibria relative to the Pareto optimum. Since we have recovered the structural parameters of the model, one could easily compute these deadweight losses. However, they would not be expressed in any meaningful units. To compute the welfare measures in dollars, one would ideally want to have some estimated coefficient, such as price, that is directly expressed in dollar terms. In our estimation, price variation will be captured by the time fixed effects, and, hence, price is not usable in this regard. Nonetheless, if one is interested in dollar welfare measures, one can conceivably match the estimated fixed cost in the model to the fixed costs reported in Section 2. This would then allow us to scale the parameters to dollar terms and recover the dollar losses from the network externalities.

5. Conclusions

In this paper, we analyzed a simple theoretical model of technology adoption with network externalities and used this model, together with detailed panel data, to test for and estimate the magnitude of network externalities in the Federal Reserve automated clearinghouse (ACH) payment system.

We develop three methods to identify network externalities. The first method examines whether there is clustering in ACH adoption and is identified by the assumption that the shocks to banks' profits are uncorrelated for banks in a network. The second method uses sizes and concentration of other banks as instruments for their adoption decisions and is identified by the

assumption that the sizes of other banks in a network are excluded exogenous variables in the ACH adoption decision. The third method uses a subsample of the data with small branches of large banks, where the variation in adoption decisions can plausibly be thought of as exogenous. Because we have panel data, we use time fixed effects and bank fixed effects or size categories to control for differences in technology, price changes, scale effects and variations in preferences among individual banks and customers.

All three identification methods reveal significant evidence of network externalities. The conclusion is robust to a variety of alternative hypotheses. Moreover, there is some evidence that the network externalities exist at the bank level rather than at the individual customer level, suggesting that direct usage externalities may be more important than informational externalities. The magnitude of the network externalities is estimated to be moderately large. We find that the Pareto-best Nash equilibrium has a substantially lower adoption probability than the first-best outcome.

As our results strongly indicate the presence of network externalities, we draw two policy implications. First, it does appear that ACH is underused relative to its socially optimal usage level. Federal Reserve policy should attempt to encourage ACH adoption and usage. Second, other high-technology industries may also be characterized by network externalities. Further policy recommendations will have to depend on estimating costs and computing the welfare effects of different counterfactual policy environments. As this paper lays out a theoretical framework for examining network externalities, it suggests a future research strategy to accomplish this goal.

Appendix

Proof of Proposition 1:

We prove existence of a Pareto dominating Nash equilibrium by constructing a Nash equilibrium $Usage^P(A)$ (P for short), and then proving that it Pareto dominates all other Nash equilibria. We then show that P is increasing in A.

To construct P , start with a strategy profile $Usage_0^P(A)$ (P^0 for short) where every customer from a bank that has adopted ACH is using ACH. Then, construct P^1 by setting usage to 0 for all customers who do not want to use ACH even if the usage decisions are given by P^0 . Similarly, construct P^2 by removing from P^1 everyone who does not want to use ACH, given that the usage decisions are given by P^1 . Note that $P^2 \leq P^1 \leq P^0$. Repeat this process until we find some $N(A)$ such that $P^N = P^{N+1}$. Such an $N(A)$ must exist because there are only a finite number of customers. Let $P = P^{N+1}$. By construction, no customer would want to unilaterally deviate from this strategy profile and, hence, P is a Nash equilibrium.

Now, suppose that there is another Nash equilibrium $Usage^Q(A)$ (Q for short) that is not Pareto dominated by P . Then, some customer must be better off under Q than under P which implies that some consumer (not necessarily the same one) must be using ACH under Q but not under P , so that it is not the case that $P \geq Q$. Consider the last stage i such that $P^i \geq Q$. Such an i must exist since $P^0 \geq Q$. Consider a consumer who stopped using ACH between P^i and P^{i+1} but is using ACH under Q . Given our construction of P^{i+1} , it would be individually optimal for this consumer to stop using ACH under Q , which contradicts the assumption that Q is a Nash equilibrium.

Lastly, note that if $A' \geq A$, then $Usage_0^P(A') \geq Usage_0^P(A)$. It then follows that $Usage_1^P(A') \geq Usage_1^P(A)$ and, by recursion, that $Usage_i^P(A') \geq Usage_i^P(A)$ for all stages i . Since $Usage^P(A') = Usage_{\max\{N(A), N(A')\}}^P(A') \geq Usage_{\max\{N(A), N(A')\}}^P(A) = Usage^P(A)$, the last part of the proposition holds.

References

Akerberg, Daniel and Gautam Gowrisankaran (2002). "Quantifying Equilibrium Network Externalities in the ACH Banking Industry." Mimeo, UCLA.

- Bresnahan, Timothy and Peter Reiss (1991). "Entry and Competition in Concentrated Markets." *Journal of Political Economy* 99: 977-1009.
- 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, Joseph and Garth 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 (2000). "The Dynamics of Technological Adoption in Hardware/Software Systems: The Case of Compact Disc Players." *RAND Journal of Economics* 31: 43-61.
- 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, Michael and Carl 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.
- Manski, Charles (1993). "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies* 60: 531-42.
- 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 Review*: 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.
- Spencer, Barbara J. and James A. Brander (1983). "International R&D Rivalry and Industrial Strategy." *Review of Economic Studies* 50, 707-722:
- Weinberg, John A. (1997). "The Organization of Private Payments Networks." Federal Reserve Bank of Richmond *Economic Quarterly* 83: 25-43.

Table 1: Number and percent of banks using ACH origination each quarter, by bank size

Quarter	Number Using ACH			Percent Using ACH		
	small	medium	large	small	medium	large
95:Q2	4208	1811	440	53.1%	75.1%	73.8%
95:Q3	4235	1831	448	53.7%	75.3%	73.9%
95:Q4	4273	1893	462	54.5%	76.9%	75.1%
96:Q1	4383	1911	468	55.8%	77.9%	75.2%
96:Q2	4559	1901	449	57.7%	78.0%	76.1%
96:Q3	4600	1948	480	58.3%	79.8%	81.8%
96:Q4	4732	2032	482	60.5%	81.1%	80.3%
97:Q1	4857	2042	473	62.0%	81.8%	79.4%
97:Q2	5061	2057	470	64.1%	83.8%	82.0%
97:Q3	5138	2057	470	65.1%	83.3%	82.6%
97:Q4	5169	2112	465	65.9%	83.9%	81.9%

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

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

Quarter	Entrants	Percent of total	Exiters	Percent of total
95:Q2	--	--	407	6.3%
95:Q3	462	7.1%	414	6.4%
95:Q4	528	8.0%	397	6.0%
96:Q1	531	7.9%	388	5.7%
96:Q2	535	7.7%	493	7.1%
96:Q3	612	8.7%	319	4.5%
96:Q4	537	7.4%	328	4.5%
97:Q1	454	6.2%	256	3.5%
97:Q2	472	6.2%	266	3.5%
97:Q3	343	4.5%	254	3.3%
97:Q4	335	4.3%	--	--

Authors' calculations based on the sample.

Table 3: Fraction of banks adopting ACH in 1996:Q3 and 1997:Q4, by 1995:Q2 adoption

		Fraction of banks adopting in 1995:Q2				
		0-19%	20%-39%	40-59%	60-79%	80-100%
Fraction of banks adopting in 1996:Q3 given 1995:Q2 level	0-19%	0.618	0.097	0.046	0.012	0.003
	20%-39%	0.097	0.312	0.060	0.030	0.003
	40-59%	0.077	0.175	0.337	0.067	0.019
	60-79%	0.069	0.156	0.189	0.445	0.109
	80-100%	0.139	0.260	0.368	0.445	0.866
Fraction of banks adopting in 1997:Q4 given 1995:Q2 level	0-19%	0.376	0.033	0.010	0.005	0.002
	20%-39%	0.105	0.208	0.014	0.019	0.005
	40-59%	0.097	0.136	0.214	0.077	0.028
	60-79%	0.093	0.143	0.221	0.271	0.129
	80-100%	0.329	0.481	0.540	0.629	0.837

Authors' calculations based on the sample. An observation is the deposit-weighted mean fraction of banks adopting ACH in one MSA or non-MSA county in one quarter. The table indicates the fraction of observations in a given cell in 1996:Q3 and 1997:Q4 conditional on the 1995:Q2 cell.

Table 4: Fraction of banks adopting ACH, by population and concentration

	Population of MSA/county					
HHI	20,000 - 40,000		100,000 - 200,000		400,000 - 1 million	
	Mean	No. of obs.	Mean	No. of obs.	Mean	No. of obs.
0 - 0.1	0.591	3	0.661	11	0.679	4
0.1 - 0.2	0.766	108	0.797	67	0.912	34
0.2 - 0.3	0.810	247	0.878	38	0.828	8
0.3 - 0.4	0.850	131	0.935	6	--	0
0.4 - 1.0	0.862	81	0.972	1	--	0

Authors' calculations based on the sample. An observation is the deposit-weighted mean fraction of banks adopting ACH in one MSA or non-MSA county in 1996:Q3. The table indicates the number of observations in each population/concentration cell and the fraction of observations in each cell out of the total number of observations for that population range.

Table 5: Identification of network externalities using clustering

		Regressor: Fraction adopting ACH [#(j)]	Regressor: ACH volume per assets [# ^Q (j)]	Number of observations	Conditional log likelihood / R ² (within)
Chamberlain's conditional logit	Model 1: Adoption on #(j)	1.874*** (0.096)	--	Tot: 104,302 Used: 35,739	-12,882.9
	Model 2: Adoption on #(j) and # ^Q (j)	1.873*** (0.096)	0.142 (0.109)	Tot: 104,302 Used: 35,739	-12,882.0
	Model 3: Adoption on # ^Q (j)	--	0.172 (0.108)	Tot: 104,302 Used: 35,739	-13,090.1
	Model 4: Adoption on #(j) with alternate definition of entry/exit	2.073*** (0.135)	--	Tot: 104,302 Used: 33,964	-9,218.2
Linear fixed effects	Model 5: Volume per assets on #(j)	0.0014 (0.0057)	--	104,302	0.0012
	Model 6: Volume per assets on #(j) and # ^Q (j)	0.0013 (0.0057)	0.0075 (0.0052)	104,302	0.0012

*** significant at the 1 percent level

Unit of observation: Bank / quarter for banks with at least one other bank within 30 kilometers. Quarterly dummies for all 11 quarters, bank fixed effects, assets, deposits and their squares are included in all regressions.

Table 6: Identification of network externalities using concentration and size

	Regressor: MSA/county level concentration [HHI_j]	Regressor: Fraction adopting ACH [$\#(j)$]	Number of observations	Log likelihood / R^2 (within)
Model 1: Adoption on HHI_j and bank size dummies (Logit estimation)	0.392*** (0.046)	--	109,321	Log L= -64,392.6
Model 2: Adoption on HHI_j and bank fixed effects (Chamberlain's conditional logit estimation)	0.295 (0.642)	--	Tot: 109,321 Used: 37,908	Cond Log L= -13,876.70
Model 3: Volume per assets on HHI_j and bank fixed effects, for banks with positive volume (Linear FE estimation)	0.171 (0.053)	--	73,255	$R^2=0.002$
Model 4: Adoption on $\#(j)$ and bank size dummies (Linear IV estimation)	--	0.858*** (0.057) Robust: (0.063)	104,302	--
Model 5: Adoption on $\#(j)$ and bank fixed effects (Linear IV FE estimation)	--	0.930*** (0.072)	104,302	--

*** significant at the 1 percent level

Unit of observation: Bank / quarter. Models 4 and 5 exclude banks with no other bank within 30 kilometers. Quarterly dummies for all 11 quarters, assets, deposits and their squares are included in all regressions. Bank size dummies for Models 1 and 4 measure assets at intervals of \$20 million up to \$2 billion.

Table 7: Identification of network externalities using quasi-experimental variation in adoption

	Regressor: Fraction adopting ACH [#(j)]	Number of obs.	Log likelihood / R ²
Model 1: Adoption on #(j) (Probit estimation)	1.094 (0.411)*** Robust: (0.354)***	237	Log L= -138.4
Model 2: Adoption on #(j) and bank size dummies (Probit estimation)	1.623 (0.679)** Robust: (0.788)**	158	Log L= -84.7
Model 3: Adoption on #(j) (Linear estimation)	0.511 (0.180)*** Robust: (0.291)*	237	R ² = 0.467

*** significant at the 1 percent level

** significant at the 5 percent level

* significant at the 10 percent level

Unit of observation: Bank / quarter for local banks with no other local banks and at least one small branch of a large bank in the network. Quarterly dummies for all 11 quarters, assets, deposits and their squares are included in all regressions. Bank size dummies for Models 2 and 3 measure deposits at intervals of \$2 million up to \$100 million.

Table 8: Simulation of network equilibria


Number of firms, N	Pareto-worst Nash equilibrium	Pareto-best Nash equilibrium	First-best (perfect cartel) outcome
1	0.62	0.62	0.62
2	0.30	0.46	0.66
3	0.32	0.43	0.67
4	0.32	0.41	0.67
5	0.32	0.40	0.67
6	0.33	0.39	0.67

Each cell gives the expected percent of banks adopting under each environment.

We assume that there are N symmetric firms.

Deposits=\$33 million, Time = 1996:Q2, and parameters are those in Table 7, Model 1.

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



Note: In May 1997, volume-based pricing was introduced, with price set to 0.9 cents per item for files with less than 2500 items and 0.7 cents per item for files with 2500 or more items.