

Network externalities and technology adoption: lessons from electronic payments

Gautam Gowrisankaran*
Joanna Stavins **

Abstract:

We analyze the extent of network externalities for the automated clearinghouse (ACH) electronic payments system using a panel data set on bank adoption and usage of ACH. We develop three methods. The first examines the clustering of ACH adoption. The second examines the impact of market concentration and the size of competitors on ACH adoption. The third examines the impact of ACH adoption by small branches of large banks on local competitors. 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.

* Department of Economics, Yale University, P.O. Box 208264, New Haven, CT 06511 USA and National Bureau of Economic Research, gautam_gowrisankaran@nber.org.

** Research Department, Federal Reserve Bank of Boston, 600 Atlantic Avenue, Boston, MA 02106 USA, joanna.stavins@bos.frb.org.

Acknowledgements:

We thank Dan Akerberg, V.V. Chari, April Franco, Austan Goolsbee, Ed Green, Justine Hastings, Tom Holmes, Pete Klenow, Jamie McAndrews, Matt Mitchell, Andrea Moro, Rob Porter, two anonymous referees, and various seminar participants for helpful comments. We thank Anita Todd for editorial assistance and Marie Willard for research assistance. Gowrisankaran acknowledges financial support from the Federal Reserve Banks of Cleveland, New York and San Francisco 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, Cleveland, New York, or San Francisco.

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 already use the technology because they 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 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% of all 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 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. Our study also illustrates the types of data and methods that would identify network externalities for other high-technology industries.

In spite of substantial theoretical work on network externalities,² there are comparatively few empirical analyses of network effects, primarily due to the lack of data.³ In most industries, there is only 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 based on time series data have been beset by this identification issue; see Cabral and Leite (1992), Economides and Himmelberg (1995) and Park (2003).⁴ Some recent studies of network externalities have instead made use of regional geographical cross-sectional data.⁵ Cross-sectional data have their 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 examines the partial correlations of adoption decisions for banks within a network. The second uses excluded exogenous variables based on bank size to control for the endogenous adoption decisions. The third creates 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.

The remainder of this paper is divided as follows. Section 2 describes the data. Section 3 describes the model. Sections 4 through 6 provide tests and results for our three identification strategies. Section 7 concludes.

2. Data

Our principal data set is the Federal Reserve's billing data that provide 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 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 use the total 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 generally obtain a FedLine connection, although some of the largest banks instead use a direct computer interface between their mainframe computer and a Fed host mainframe. Most banks use their 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 ACH processing software for 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.

The Federal Reserve is currently the dominant provider of ACH services. The Federal Reserve handled approximately 75% of the roughly 3.3 billion on-others (between two different banks) commercial ACH transactions processed in 1996.⁸ The remaining share of the 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. 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 use 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 keep 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 aggregate the Federal Reserve ACH volume up to the FDIC number level. One data problem is that a substantial fraction of the ABA numbers from the ACH billing data is not in the Call Reports database. Most of the ABA numbers that do not match are credit unions or thrifts which we exclude. We further exclude all banks with deposits of less than \$10 million for all quarters in the sample and all remaining credit unions. We are left with approximately 11,000 banks over the 11-quarter sample period.

We also merge 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 for some of our methods 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 keep all banks from the Summary of Deposits data that are also in our Call Report sample.

As is generally done for banks, we define a market as the metropolitan statistical area (MSA) or county (for banks not in MSAs). We find bank MSA/counties using the Census Geocorr database, which

translates from zip codes to MSAs and counties. We use the 1995 county and MSA mappings and choose 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 search for the zip code in the Geocorr database with a centroid nearest to the missing one. If this nearest zip code is within 10 kilometers of the missing one, we adopt its MSA/county information. Otherwise, we treat that observation as missing and drop it. Overall, this algorithm is successful at identifying MSAs/counties, with only a handful of missing observations.

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

One factor that can affect usage of ACH is price. The Federal Reserve charges fixed per-transaction prices to banks for ACH processing. These prices are adjusted periodically. 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 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, which is consistent with their ability to spread 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 offer ACH but have no transactions in a given quarter. Because we were concerned about 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 is banks that stop originating ACH transactions and another fourth is exiting or merging banks. The remaining half that enter and exit appears 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 in Section 4.

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 use the fraction of banks adopting ACH within an MSA or a non-MSA county. We weight the fraction by the deposits for all branches within the MSA, using the Summary of Deposits data. Table 4 also uses an 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. There is 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 .1-.2 have a .766 adoption fraction, while banks with an HHI of .4-1.0 have a .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.

3. 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.

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 captures a fixed proportion of the consumer surplus. Thus, we do not explicitly consider oligopoly pricing, wherein the proportion would likely vary based on the number of firms, or allow for

customers to choose among banks.⁹ Following adoption decisions, customers simultaneously decide whether or not to use ACH. A customer 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:

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

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:

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

where $\{\cdot\}$ is an indicator function. With network externalities, we expect that $u_{i,j}$ will be weakly increasing in every component of $Usage_{-i,j}$, and strictly increasing for some components, in particular in the usage of other customers at the same bank.

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 others. This equilibrium has a greater level of ACH adoption than any other equilibrium, is weakly increasing in bank ACH adoption, and is easy to compute. Formally,

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

Proof: See the 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:

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

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 of A_{-j} , and γ are parameters of X_j and interactions. We normalize the profit from not adopting ACH to 0 and choose a linear version of (3) for tractability. The adoption decision can then be expressed as:

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

where $h(A_{-j})$ encapsulates the network externality. Our principal purpose is to obtain inference on β in (4).

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 $\#^o(j)$ 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 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 exclude 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 a fixed proportion of the surplus of their 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. Banks may coordinate on the Pareto dominating Nash equilibrium, for instance, by informally agreeing on it at a banking trade conference.¹²

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 having more banks with which to exchange ACH payments and due to it being cheaper for banks to exchange ACH payments if they can learn about procedures from other banks in the network. Hence, our estimation will pick up both consumer- and bank-level network externalities and we cannot distinguish between them.

We assume a network size of 30 kilometers and also construct HHI at the MSA/county level. Certainly, banks exchange payments with other banks with headquarters outside this limited area — indeed, the majority of bank payments are non-local. Thus, any estimated β coefficient from (4) will capture only a portion of the network effect. Since ACH usage is increasing over our sample period, the non-captured network effect will be confounded with the time fixed effects.

4. Identifying network externalities from clustering

Our first method of identifying network externalities tests 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 identification, 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 the fixed effect. 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, shocks that are correlated across nearby banks and that vary across time will not be consistent

with our model.¹³ We expect that the ε_j 's will capture unobserved time-varying factors such as employers that permit or encourage ACH for payroll. Importantly, we use the ACH origination volume throughout. If, instead, we had used the recipient volume, then 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.

We test for network externalities with this identification assumption using standard fixed effects methods. Specifically, we estimate (4) using either ACH adoption or ACH volume per assets as the dependent variable. We can interpret the data as providing evidence of network effects if and only if the estimated β 's (the coefficients on other banks' ACH adoption and volume) are significantly positive.

The reason for this is that with no network externalities, the value of β is 0. In this case, A_j is a function of bank j 's characteristics but not of A_{-j} . Given the assumption that the ε 's are uncorrelated within the network, $h(A_{-j})$ will not be correlated with ε_j . This implies that the estimated β will be consistent and should converge in probability to zero. With network externalities, 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, the estimated β will reflect the equilibrium correlation between ACH adoption decisions. While the estimate of β cannot be interpreted as the structural parameter, we can interpret a significantly positive estimate of β as evidence of network externalities.¹⁴

Table 5 presents the estimates of the clustering tests for 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 linear models.

As our data set contains approximately 11,000 banks, it is not computationally feasible to estimate the bank fixed effects by including 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 chooses 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, which is numerically equivalent to estimating the 11,000 bank dummy variables 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 2).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 2. Note that this evidence shows not just the potential for network externalities, but also the existence of unexploited network externalities, twelve years after ACH was 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 magnitudes from the logit probability derivative formula, $dpr(A_j)/d\#(j) = \beta_{\#(j)}pr(A_j)(1 - pr(A_j))$. If we assume that $pr(A_j) = .5$, then from Model 1, a 10 percentage point increase in the fraction of other banks adopting ACH is associated with a 4.7 percentage point increase in the probability that bank j is using ACH,¹⁶ with a standard error of .2 percentage point. Since the mean adoption probability is 67.2%, the actual magnitude is likely to be smaller than 4.7 percentage points.

In contrast, the evidence on the effect of the number of transactions is much weaker (Models 2, 3, and 6). While the three models show a positive impact of $\#(j)$ or $\#^o(j)$ 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 .031. Thus, using Model 3, a doubling in the volume per assets by a mean bank is associated with an increase of .13 percentage point in the adoption probability of a bank with $pr(A_j) = .5$. This number is economically quite small.

The fact that adoption is quantitatively much more important than the volume of transactions, as evidenced by both the coefficient magnitude and insignificance, indicates a much greater partial correlation across banks' adoption decisions than between bank adoption and volume. However, since adoption and volume are endogenous and interrelated, it is difficult to assess the relative magnitudes of the structural coefficients on adoption and volume from these regressions. This again suggests the need for a more structural analysis to recover the sources of the network externalities.

The evidence of network externalities in Table 5 is robust to allowing for some types of market power and dynamics. First, our base model assumes that banks capture a fixed proportion of the surplus from customers and that customers are captive to one institution. Consider a more general model where customers can switch between banks and where banks choose ACH price in addition to adoption. If customers can switch banks, then there may be correlations in adoption decisions for reasons other than network externalities. 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 the customer base shrinks following entry.

Second, we assume 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. Adoption decisions will also depend on expectations of the future network benefit from usage. Even if there are important dynamic interactions, in the absence of network externalities or correlations in the unobservables over time and across banks, there will be no partial correlation between the adoption decisions of different banks. Moreover, with network externalities, a portion of the adoption decision will depend on current network benefits and not on sunk costs. The implication is that the correlation among current adoption decisions is due to network externalities. In contrast, our estimation methods are not robust to serial correlation in the error structure.

5. 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. Importantly, since we control for a bank's own size, this identification will not confound network externalities with economies of scale provided that size is correctly measured.

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 adoption are positive, but it will not internalize the positive externality that its adoption has on profits for the other banks in the network. Thus, if we merge the banks into a monopoly holding company with fixed costs at each of its member banks equal to the fixed costs before the merger, then the set of banks that adopts weakly increases.

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:

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

The HHI variable is constructed from deposits in the MSA/county in which the bank is located, as discussed in Section 2. We include asset dummies (at intervals of \$20 million) in X_j as well as time fixed effects, assets, assets squared, deposits, and deposits squared. We interpret a positive coefficient on α as evidence that concentrated markets are more likely to adopt ACH, which we attribute to network externalities.

Table 6, Models 1 and 2 present these results. 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.

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. 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 1.63 percentage points, with a standard error of .2 percentage point, for a bank with a 50% probability of adoption. Model 2 shows a positive effect that is similar in magnitude to Model 1, though not significant. The insignificance is likely due to the fact that the within-bank variation in asset levels (and, hence, HHI) is small.

Note that our base model assumes that banks capture a fixed proportion of the surplus from customers. With market power, firms in concentrated markets are likely to capture a higher proportion of the surplus by charging a higher price. 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.¹⁷ This market power effect could be confounded with network externalities. However, there is one difference: 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 fewer customers by charging a higher price per customer. In contrast, with network externalities, equilibrium quantity conditional on adoption will be increasing in concentration. Model 3 presents the results from a linear regression of ACH volume per assets on HHI, using as the sample banks with a positive volume and including 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 and, thus, reinforces our findings that network externalities exist.

Second, we estimate (4) using size and concentration as instruments for $h(A_{-j})$. Even though adoption is a discrete variable, we assume a linear functional form to apply instrumental variables.¹⁸ Hence, we use a linear probability model variant of (4):

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

Because of our results in Section 4, the only network benefit that we include is $\#(j)$. 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 include the same controls in X_j as in (5). We perform (6) on the sample of banks with at least one other bank in the network. 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 .858 without bank fixed effects and .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; the conclusions are not substantially affected. The estimated coefficients are large: they imply that a 10 percentage point increase in the adoption fraction of other banks increases the adoption probability by roughly 9 percentage points.

6. 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). There is quasi-experimental variation for certain markets because banks with multiple branches either adopt ACH for all of their branches or for none. We exploit this variation by creating a subsample of small towns with one small, local bank j and some small branches of large banks. No large bank is going to adopt ACH over all of its branches in response to 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. Thus, the adoption decisions of large banks form a “natural experiment” in these networks which allows us to trace out the structural reaction function for the local bank.

We estimate a version of (4) on 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% 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.¹⁹ As we do not observe the volumes of large banks’ ACH transactions specific to the network of bank j , we define the network benefit $h(A_{-j})$ to be $\#(j)$, the fraction of other banks that adopt. We weight the fraction by deposits to account for size differences across banks.

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. However, 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.

Table 7 presents the results from 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. We do not allow for bank fixed effects in X_j because of the small amount of data. We also include a linear probability model specification to compare with the instrumental variables results in Table 6. We report 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.

By construction, our sample is small, with 237 observations in Model 1, of which 158 are kept in Model 2. Nonetheless, there is evidence of significant, positive, and moderately large network externalities. For instance, if bank j has an adoption probability of 50%, then Model 1 implies that an exogenous increase of 10 percentage points in the deposit-weighted adoption fraction of other banks would raise bank j 's adoption probability by 4.4 percentage points with a standard error of 1.4 percentage points using the robust method. Last, Model 3 shows that a linear probability model yields a magnitude of .511 which is somewhat smaller than the magnitude of .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 electronic payment mechanisms.

In order to further interpret the economic magnitude of the network externalities, we simulate equilibria for representative networks using (4) and the estimated structural parameters from Model 1. The iterative technique described in the proof of Proposition 1 provides a natural method to solve for the

Pareto-worst and -best equilibria given parameters, observables, and unobservables. We compute equilibria for networks with N local banks, where each bank had the Model 1 sample mean deposit size, existed at the midpoint of the sample, and had an iid unobservable ε_j . For accuracy, we base our results on 10,000 simulation draws. We also compute the first-best adoption decisions, which would be chosen by a perfectly colluding cartel. These decisions can be computed similarly but 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. The Pareto-best Nash equilibrium outcome with $N=1$ has banks adopting ACH 62% of the time, and this falls monotonically in N . For $N > 1$, the Pareto-worst Nash equilibrium has a lower adoption probability that increases as N increases from 2 to 6. In contrast, the first-best outcome has a much higher mean adoption frequency.

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 to small banks. Most 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, one could compute these deadweight losses. However, they would not be expressed in any meaningful units. To compute the welfare measures in dollars, one would want some estimated coefficient, such as price, that is expressed in dollar terms. In our estimation, price variation will be captured by the time fixed effects. Nonetheless, to obtain dollar measures, one could match the estimated fixed cost to the fixed costs reported in Section 2, scale the parameters to dollar terms, and recover the dollar losses from the network externalities.

7. Conclusions

In this paper, we used a simple theoretical model of technology adoption with network externalities, 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. 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 indicate that ACH appears to be underused relative to its socially optimal level due to network externalities, Federal Reserve policy should encourage ACH adoption and usage. In addition, this paper lays out a theoretical framework for examining network externalities, suggesting a research strategy for estimating network effects and their consequences in other industries.

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 , which further implies 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. *Q.E.D.*

References

- Akerberg, D. and Gowrisankaran, G. "Quantifying Equilibrium Network Externalities in the ACH Banking Industry." Mimeo, University of Minnesota, 2003.
- Bresnahan, T. and Reiss, P. "Entry and Competition in Concentrated Markets." *Journal of Political Economy*, Vol. 99 (1991), pp. 977-1009.
- Cabral, L. and Leite, A. "Network Consumption Externalities: The Case of Portuguese Telex Service." C. Antonelli, ed., *The Economics of Information Networks*. Amsterdam: Elsevier, 1992.
- Chamberlain, G. "Analysis of Covariance with Qualitative Data." *Review of Economic Studies*, Vol. 47 (1980), pp. 225-38.
- Economides, N. and Himmelberg, C. "Critical Mass and Network Size with Application to the U.S. Fax Market." Discussion Paper EC-95-11, NYU Stern School of Business, 1995.
- Farrell, J. and Saloner, G. "Standardization, Compatibility, and Innovation." *RAND Journal of Economics*, Vol. 16 (1985), pp. 70-83.
- Gandal, N. "Hedonic Price Indexes for Spreadsheets and an Empirical Test for Network Externalities." *RAND Journal of Economics*, Vol. 25 (1994), pp. 160-70.
- _____, Kende, M., and Rob, R. "The Dynamics of Technological Adoption in Hardware/Software Systems: The Case of Compact Disc Players." *RAND Journal of Economics*, Vol. 31 (2000), pp. 43-61.
- Goolsbee, A. and Klenow, P. "Evidence on Network and Learning Externalities in the Diffusion of Home Computers." *Journal of Law and Economics*, Vol. 45 (2002), pp. 317-344.
- Greenstein, S. "Did Installed Base Give an Incumbent Any (Measurable) Advantages in Federal Computer Procurement." *RAND Journal of Economics*, Vol. 24 (1993), pp. 19-39.
- Katz, M. and Shapiro, C. "Technology Adoption in the Presence of Network Externalities." *Journal of Political Economy*, Vol. 94 (1986), pp. 822-41.
- Kwast, M.L., Starr-McCluer, M., and Wolken, J.D. "Market Definition and the Analysis of Antitrust in Banking." *The Antitrust Bulletin*, Vol. 42 (1997), pp. 973-95.
- Manski, C. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies*, Vol. 60 (1993), pp. 531-42.
- Park, S. "Quantitative Analysis of Network Externalities in Systems Competition: the VCR Case." Mimeo, SUNY at Stony Brook, 2003.
- Roberds, W. "The Impact of Fraud on New Methods of Retail Payment." Federal Reserve Bank of Atlanta *Economic Review*, Vol. 83 (1998), pp. 42-52.

- Rysman, M. "Competition between Networks: A Study of the Market for Yellow Pages." Mimeo, Boston University, 2002.
- Saloner, G. and Shephard, A. "Adoption of Technologies with Network Effects: An Empirical Examination of the Adoption of Automated Teller Machines." *RAND Journal of Economics*, Vol. 26 (1995), pp. 479-501.
- Spencer, B.J. and Brander, J.A. "International R&D Rivalry and Industrial Strategy." *Review of Economic Studies*, Vol. 50 (1983), pp. 707-722:
- Weinberg, J.A. "The Organization of Private Payments Networks." Federal Reserve Bank of Richmond *Economic Quarterly*, Vol. 83 (1997), pp. 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	4,208	1,811	440	53.1%	75.1%	73.8%
95:Q3	4,235	1,831	448	53.7%	75.3%	73.9%
95:Q4	4,273	1,893	462	54.5%	76.9%	75.1%
96:Q1	4,383	1,911	468	55.8%	77.9%	75.2%
96:Q2	4,559	1,901	449	57.7%	78.0%	76.1%
96:Q3	4,600	1,948	480	58.3%	79.8%	81.8%
96:Q4	4,732	2,032	482	60.5%	81.1%	80.3%
97:Q1	4,857	2,042	473	62.0%	81.8%	79.4%
97:Q2	5,061	2,057	470	64.1%	83.8%	82.0%
97:Q3	5,138	2,057	470	65.1%	83.3%	82.6%
97:Q4	5,169	2,112	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%	.618	.097	.046	.012	.011
	20%-39%	.097	.312	.060	.030	.011
	40-59%	.077	.175	.337	.067	.023
	60-79%	.069	.156	.189	.445	.086
	80-100%	.139	.260	.368	.445	.870
Fraction of banks adopting in 1997:Q4 given 1995:Q2 level	0-19%	.376	.033	.010	.005	.008
	20%-39%	.105	.208	.014	.019	.007
	40-59%	.097	.136	.214	.077	.026
	60-79%	.093	.143	.221	.271	.101
	80-100%	.329	.481	.540	.629	.858
Fraction of 1995:Q2 banks at each level		.102	.061	.112	.169	.556

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

HHI	Population of MSA/county					
	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 - .1	.591	3	.661	11	.679	4
.1 - .2	.766	108	.797	67	.912	34
.2 - .3	.810	247	.878	38	.828	8
.3 - .4	.850	131	.935	6	--	0
.4 - 1.0	.862	81	.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*** (.096)	--	Tot: 104,302 Used: 35,739 -12,882.9
	Model 2: Adoption on $\#(j)$ and $\#^Q(j)$	1.873*** (.096)	.142 (.109)	Tot: 104,302 Used: 35,739 -12,882.0
	Model 3: Adoption on $\#^Q(j)$	--	.172 (.108)	Tot: 104,302 Used: 35,739 -13,090.1
	Model 4: Adoption on $\#(j)$ with alternate definition of entry/exit	2.073*** (.135)	--	Tot: 104,302 Used: 33,964 -9,218.2
Linear fixed effects	Model 5: Volume per assets on $\#(j)$.0014 (.0057)	--	104,302 .0012
	Model 6: Volume per assets on $\#(j)$ and $\#^Q(j)$.0013 (.0057)	.0075 (.0052)	104,302 .0012

*** significant at the 1% 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)	.392*** (.046)	--	109,321	$Log L =$ -64,392.6
Model 2: Adoption on HHI_j and bank fixed effects (Chamberlain's conditional logit estimation)	.295 (.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)	.171 (.053)	--	73,255	$R^2 = .002$
Model 4: Adoption on $\#(j)$ and bank size dummies (Linear IV estimation)	--	.858*** (.057) Robust: (.063)	104,302	--
Model 5: Adoption on $\#(j)$ and bank fixed effects (Linear IV FE estimation)	--	.930*** (.072)	104,302	--

*** significant at the 1% 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 (.411)*** Robust: (.354)***	237	$\text{Log } L =$ -138.4
Model 2: Adoption on $\#(j)$ and bank size dummies (Probit estimation)	1.623 (.679)** Robust: (.788)**	158	$\text{Log } L =$ -84.7
Model 3: Adoption on $\#(j)$ (Linear estimation)	.511 (.180)*** Robust: (.291)*	237	$R^2 =$.467

*** significant at the 1% level

** significant at the 5% level

* significant at the 10% 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	.62	.62	.62
2	.30	.46	.66
3	.32	.43	.67
4	.32	.41	.67
5	.32	.40	.67
6	.33	.39	.67

Each cell gives the expected fraction of banks that adopt in an industry with N firms in 1996:Q3, each with deposits of \$33 million and iid draws of ε_j . The parameters are those in Table 7, Model 1.

¹ Roberds (1998) and Weinberg (1997) discuss network effects in payments markets.

² Farrell and Saloner (1985) and Katz and Shapiro (1986) examine the equilibrium adoption of new technologies with network externalities.

³ We consider network externalities for 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 Gandal (1994), Greenstein (1993), or Saloner and Shephard (1995) respectively.

⁴ 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 (2002) examine network externalities in computers and Rysman (2002) 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.

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

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

⁹ We justify this latter 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, Starr-McCluer, and Wolken (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.

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.

¹² Another possibility is that the banks coordinate on the first-best outcome. We prefer the Nash assumption because absent side payments, which are likely infeasible and illegal, banks would not be able to fully internalize the externality. The first-best outcome will generate the same partial correlation implications as the Nash equilibrium assumption, implying that our correlation tests will be robust to this choice of assumption. However, our tests based on market concentration would not be valid if banks can coordinate on the first-best outcome.

¹³ An example of a local time-specific effect would be a regional advertising campaign that stimulated demand in a particular area.

¹⁴ Equilibrium methods can be used to recover the magnitudes of the structural β parameters with the same identification assumption; see Akerberg and Gowrisankaran (2003).

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

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

¹⁷ Bresnahan and Reiss (1991) attribute the higher probability of entry per capita by monopolists to this market power reason.

¹⁸ It is possible to estimate a nonlinear functional using the method of simulated moments or simulated maximum likelihood; see Akerberg and Gowrisankaran (2003).

¹⁹ 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.