NET Institute*

www.NETinst.org

Working Paper #03-06

October 2003

Quantifying Equilibrium Network Externalities in the ACH Banking Industry

by

Daniel A. Ackerberg
Department of Economics, University of Arizona and NBER

Gautam Gowrisankaran
John M. Olin School of Business, Washington University in St. Louis and NBER

* The Networks, Electronic Commerce, and Telecommunications ("NET") Institute, http://www.NETinst.org, is a non-profit institution devoted to research on network industries, electronic commerce, telecommunications, the Internet, "virtual networks" comprised of computers that share the same technical standard or operating system, and on network issues in general.

Quantifying Equilibrium Network Externalities in the ACH Banking Industry[1]

Daniel A. Ackerberg
Department of Economics
University of Arizona
and NBER
daniel_ackerberg@nber.org

and

Gautam Gowrisankaran
John M. Olin School of Business
Washington University in St. Louis
and NBER
gautam_gowrisankaran@nber.org

This Version: October 1, 2003

Abstract

We seek to estimate the causes and magnitudes of network externalities for the automated clearinghouse (ACH) electronic payments system, using a panel data set on individual bank usage of ACH. We construct an equilibrium model of consumer and bank adoption of ACH in the presence of a network. The model identifies network externalities from correlations of changes in usage levels for banks within a network, from changes in usage following changes in market concentration or sizes of competitors and from adoption decisions of banks outside the network with small branches in the network, and can separately identify consumer and bank network effects. We structurally estimate the parameters of the model by matching equilibrium behavior to the data, using simulated maximum likelihood and a data set of localized networks, and use a bootstrap to recover confidence intervals. The parameters are estimated with high precision and fit various moments of the data reasonably well. We find that most of the impediment to ACH adoption is due to large consumer fixed costs of adoption. The deadweight loss from the network externality is moderate: the optimal number of ACH transactions is about 16% higher than the equilibrium level.

[1] We acknowledge funding from the NET Institute, and thank Steve Berry, Jinyong Hahn, Andrea Moro, Klaas van't Veld and seminar participants at numerous institutions for helpful comments.
.

1. Introduction

The goal of this paper is to estimate the size and importance of network externalities for the automated clearinghouse (ACH) banking industry using an equilibrium model of ACH usage. ACH is an electronic payment mechanism developed by the Federal Reserve and used by banks. ACH is a network: banks on both sides of a transaction must adopt ACH technology for an ACH transaction to occur. Network externalities are thought to exist in many high-technology industries. Examples include fax machines, where network effects may exist because two separate parties must communicate for a transaction to occur and computers, where network effects may exist because information on how to use new technology is costly. ACH shares the network features of fax machines, computers and other technological goods, and hence network externalities may exist for ACH.

If present, network externalities typically cause underutilization of the network good. When the network externality is positive, Nash equilibria can be Pareto ranked, and it is possible that the industry is stuck in a Pareto inferior equilibrium, characterized by even less usage than the Pareto best equilibrium. The underutilization is particularly relevant for the case of ACH. In an age when computers and technology have become prevalent, most payments continue to be performed with checks and cash. By estimating the magnitude of network externalities, we can further understand the causes of such externalities, uncover how much the usage of ACH differs from the socially optimal level, and find out whether markets are stuck in Pareto inferior equilibria. Moreover, by estimating an equilibrium model, we can evaluate the welfare and usage consequences of policies such as government subsidies of the network good.

This work extends previous research on estimating network externalities for ACH (Gowrisankaran and Stavins, 2002). That study postulated a simple game which resulted in bank adoption of ACH being a function of the adoption decisions of other banks in the network, as well as of a bank's characteristics, such as its size. The interdependence of preferences for ACH adoption leads to a simultaneity in the equilibrium adoption decisions of banks, making identification of the network externalities potentially difficult. Thus, the study proposed three methods to identify network externalities: examining whether adoption is clustered (after controlling for bank fixed effects), using excluded exogenous variables based on bank size to control for endogenous adoption decisions, and exploiting the quasi-experimental variation from the adoption decisions of small, remote branches of banks. Each of the three strategies revealed significant and positive network externalities, even after controlling for factors such as economies of scale and market power.

This paper builds on the previous research, by specifying and structurally estimating an equilibrium model of technology adoption for ACH in the presence of network externalities. The estimation uses similar data to the earlier work, and hence is identified from the same sources. However, our use of structural estimation has several advantages. First, we estimate a functional form for the network model that is directly consistent with the underlying theory of consumer utility maximization. Most importantly, this allows us to identify whether the network effects are arising at the consumer or bank level. Additionally, this allows us to efficiently combine data on bank adoption of ACH and volume conditional on adoption, and to handle networks with one bank in a logical way.[2] Second, we can recover the magnitudes of the network externalities, in a way that uses the power from the combination of all three methods of identification.[3] Third, the structural model leads very naturally to welfare and policy analysis. Note that the empirical distinction between consumer and bank level network externalities is very important here. With a subsidy to promote adoption, for example, one would want to know whom to subsidize, banks or consumers. Lastly, the structural estimation methods that we develop here are novel and

---

[2] In contrast, the earlier work could either model adoption or bank volume as dependent variables. There were measurement issues in using the adoption variable, but it is difficult to model the quantity choice outside of a structural model. Moreover, the previous work had to exclude networks with one bank, because the network variables, which are based on the fraction of other banks adopting ACH, were not defined for this case.

[3] The earlier work was able to recover magnitudes of the network externalities for some of the individual specifications. However, these specifications were somewhat limited and problematic. For instance, the quasi-experimental source of variation identified the magnitudes of network externalities, but only for a very small data set (0.2% of the total observations) of rural banks. The instrumental variables specification identified the network externalities but at the cost of imposing a linear functional form for the discrete adoption variable. The work on treatment effects (e.g. Heckman and Robb (1987) and Angrist and Imbens (1994)) suggests that linear probability models (and their associated heteroskedasticity) cause significant problems identifying causal effects using IV. The correlation source of variation could not be used at all to identify the magnitude of the network externalities without structural methods. All of the identification of the magnitudes used a reduced-form profit function for banks that was not consistent with the underlying consumer preferences.

contribute to the literature on structural estimation of simultaneous games and network games in particular.[4]

Our model of technology adoption is as follows. We consider a localized, repeated static market with a given set of banks and consumers that are tied to the bank. Each consumer must make a fixed number of transactions to other consumers evenly distributed throughout the network; transactions can be made using either checks or ACH. While all banks and consumers accept checks, some may not have adopted ACH. Some banks are local to the market while others are branches of big banks based outside the network. In each time period, local banks decide whether to adopt ACH capabilities, based on whether the marginal profits from ACH transactions conditional on adoption are greater than the fixed costs of adoption; the decisions of non-local banks are made exogenously and known to the local banks. Following bank adoption, each consumer at each bank that has adopted ACH chooses whether or not to adopt ACH. If the consumer adopts ACH, she must pay a fixed cost of adoption, but then can, and by assumption will, use ACH for her transactions to those consumers that have also adopted ACH. We model the consumer fixed costs of adoption with random effects to control for correlated preferences. The fact that ACH transactions can only be made to other individuals who have adopted implies that, in equilibrium, consumers are more likely to adopt if more consumers have adopted ACH. Similarly, banks are more likely to adopt ACH if more consumers and banks are expected to adopt. There may be multiple equilibria, and the model would not be valid without an assumption on the observed selection of equilibrium.[5] Because the network game is supermodular, there exist Pareto-best and -worst equilibria.[6] We assume that the world is characterized by some frequency of best and worst equilibria.

Our model is straightforward to solve for a given vector of parameters and draws on econometric unobservables, but it is not possible to solve the likelihood function analytically.

[4] For instance, Brock and Durlauf (2001) discusses identification for social interaction games, which are conceptually identical to network externality games. Topa (2001) structurally estimates a social interaction model using a GMM procedure. We develop a simulated maximum likelihood (SML) estimation procedure for our model, that can be used to estimate these types of games.
[5] See Heckman (1978).
[6] See Milgrom and Shannon (1994).

Thus, we estimate the model using simulated maximum likelihood. To estimate the model, we solve for the subgame perfect equilibrium of the model for each market conditional on random effects and evaluate the likelihood of the simulated equilibrium predictions. We then numerically search for the parameter vector whose simulated equilibrium predictions maximize the likelihood. We recover confidence intervals for the parameters using the bootstrap. As our endogenous variable, the number of ACH transactions at each bank, is a truncated continuous variable, we need to smooth the simulated likelihood, which we do by postulating a measurement error in the reported quantities.[7]

In our model, the network effects are captured by four parameters: the consumer and bank fixed costs of adoption, and the consumer and bank per-transaction benefits from adoption. For both banks and consumers, the fixed costs and benefits are only identified up to a common proportion. Our data identifies these two ratios via three separate mechanisms, similar to Gowrisankaran and Stavins (2003). The first source of identification is covariance restrictions. We assume that after controlling for bank and market characteristics with random effects, unobservables affecting adoption are independently distributed across banks in a given market. Thus, the estimation will find network externalities from this source if, after controlling for the random effects, the pattern of adoption within a network displays correlations consistent with network externalities. The second source is exclusion restrictions, based on the fact that the sizes of other banks do not enter into a bank's adoption decision. The estimation will find network externalities from this source if, for example, concentrated markets experience more ACH adoption. The third source of identification is the variation in adoption decisions by large, non-local banks. We assume that the adoption decisions of these banks are exogenous, and not made in response to equilibrium conditions in the market, but allow the customers at that bank to make their usage decisions in equilibrium.

By using data on both bank adoption and the proportion of transactions completed with ACH, our model can separately identify the bank network effect from the consumer network

[7] Keane and Wolpin (2000) use a similar technique.

effect. To see this, note that if there were no consumer fixed cost, then the proportion of transactions completed with ACH would simply be the square of the transaction-weighted fraction of banks adopting. As the consumer fixed cost grows, this proportion will fall, conditional on bank adoption, indicating a more important role to the consumer externality relative to the bank externality.

The remainder of this paper is divided as follows. Section 2 describes the model. Section 3 describes the data. In Section 4, we detail our estimation procedure, including the computation of the equilibrium and identification of the parameters. Section 5 contains results and Section 6 concludes.

2. The Model

We propose a simple static model of network externalities at a geographically local level. Consider a localized network of J banks in market m at time t, each with a given number of customers. The timing of our game is as follows. In the first stage, banks simultaneously decide whether or not to adopt the ACH technology. Let $A_{mt} = \left(A_{1mt}, \ldots, A_{Jmt}\right)$ be a set of indicator functions representing these adoption decisions. In the second stage, consumers decide whether to adopt ACH for their individual transactions. For a particular transaction (between two consumers) to be made through ACH, both consumers' banks must have adopted ACH, and both consumers themselves must have adopted the technology. Assume that all econometric unobservables are common knowledge to all firms and are unobservable only to the econometrician. Lastly, we assume that a consumer at bank j only knows the adoption decision of bank j, when making her adoption decision. However, in equilibrium, the consumer will have conjectures about the decisions of other banks. We proceed by first analyzing consumer decisions conditional on A_{mt}. Then we move to the first stage and analyze equilibrium bank decisions.

Since ACH transactions are a small percentage of a bank's total business, we assume that the bank's consumer base and deposits are exogenous to our model of ACH usage. Denote the deposits under bank j's control as x_{jmt}. Assume that the number of total (both ACH and check) transactions that bank j's consumers engage in at time t is proportional to these deposits x_{jmt}, i.e.,

$$(1) \qquad T_{jmt} = \lambda x_{jmt}.$$

We assume that the demand for transactions is perfectly inelastic, and hence that prices of transactions do not enter into (1). We feel that this is a reasonable assumption because the demand for transactions is in fact likely to be fairly inelastic and because ACH is a small proportion of transactions.

While we assume that the total number of transactions that consumers make is a constant fraction of deposits, we do model the proportion of these transactions that are made through ACH, which we denote as T^{ACH}_{jmt}. We assume that each bank j has a set of consumers each of whom needs to make N transactions in period t.[8] By definition, if bank j *has not* adopted ACH technology, these N transactions must be made through paper checks. If bank j *has* adopted ACH, the consumer does have the option of using ACH.

Consider consumer i's adoption decision conditional on her bank having adopted ACH. We assume that the consumer obtains net utility:

$$(2) \qquad V_{ACH} - V_{CHK} = \tilde{\beta}_1 + \tilde{\beta}_2 (p^{ACH}_t - p^{CHK}_t)$$

[8] There are a number of dimensions in which this is a stylized model of consumer behavior – in particular the fact that consumers all make an identical number of transactions. This is necessary as we have no consumer level data on behavior.

from making an ACH (versus check) transaction, where p_t^{ACH} and p_t^{CHK} represent the prices of ACH and check transactions respectively. Note that prices do not vary cross-sectionally, as they are set nationally by the Federal Reserve.

We assume that the consumer's transaction partners are allocated randomly among consumers of banks in the network, that the number of consumers is large enough to treat consumers as atoms, and that the net utility from an ACH transaction is positive. An ACH transaction can only occur if both the originating and receiving consumers have adopted ACH. Since the net utility from using ACH is assumed positive, any pair of consumers who have both adopted ACH will use ACH to process their transaction. Thus, if u_{mt} denotes the equilibrium fraction of consumers who adopt ACH, then the equilibrium probability that a transaction is made with ACH must satisfy:

$$(3) \qquad u_{mt}^2 = \frac{\sum_j T_{jmt}^{ACH}}{\sum_j T_{jmt}}.$$

Thus, u_{mt} is the square root of the total proportion of ACH (vs. check) transactions in the entire market.

Using the above definitions, we can write consumer i's net expected utility from adopting ACH (vs. not adopting) as

$$(4) \qquad EU_{ijmt} = N \cdot u_{mt} \cdot (V_{ACH} - V_{CHK}) + F_{ijmt}$$

where F_{ijmt} denotes the negative of the fixed costs of adopting. From (4), expected utility is the number of transactions that the consumer will make (N) times the probability that each transaction will be with another consumer who has adopted ACH (u_{mt}), times the utility gain from those ACH transactions ($V_{ACH} - V_{CHK}$) minus the fixed costs of adopting.

In our empirical work, we want to allow for very general unobserved correlation in ACH transactions across markets, firms, and time, to separately identify the network benefits of ACH from differences in consumer fixed costs. To allow for this, we specify F_{ijmt} as:

$$(5) \qquad F_{ijmt} = \beta_0 + \beta_3 t + \alpha_{jmt} + \varepsilon_{ijmt},$$

where t is a time trend, β_0 and β_3 are parameters to estimate, α_{jmt} is a normally distributed bank level econometric unobservable, and ε_{ijmt} is an iid consumer level logit error. We then allow α_{jmt} to be both correlated across time for consumers of a given firm and to be correlated among all consumers in a given network - specifically, we let

$$(6) \qquad \alpha_{jmt} = \alpha^A_{jmt} + \alpha^B_{jm} + \alpha^C_{m} + \alpha^D_{mt},$$

where $\alpha^A_{jmt} \sim \text{iid } N(0,1)$,[9] $\alpha^B_{jm} \sim \text{iid } N\left(0, \sigma^2_{\alpha^B}\right)$, $\alpha^C_{m} \sim \text{iid } N\left(0, \sigma^2_{\alpha^C}\right)$, $\alpha^D_{mt} \sim \text{iid } N\left(0, \sigma^2_{\alpha^D}\right)$ and where α^A, α^B, α^C and α^D are all independent of each other.[10]

Substituting from (5) into (4), we obtain:

---

[9] As adoption is a discrete decision, the variance of 1 is a normalization.

[10] As we detail the rest of the model, one might note that there are a number of places in the model where one might include a flexible unobservable structure like α_{jmt} in (6). This includes consumers' marginal benefits, in consumers' fixed costs, in banks' marginal profits, in banks' fixed costs). Because we essentially have one dependent variable in our analysis (number of ACH transactions), we felt that from an identification perspective it was only prudent to

$$\text{(7)} \quad \begin{aligned} EU_{ijmt} &= \beta_0 + N \cdot u_{mt} \cdot (\tilde{\beta}_1 + \tilde{\beta}_2 (p_t^{ACH} - p_t^{CHK})) + \beta_3 t + \alpha_{jmt} + \varepsilon_{ijmt} \\ &= \beta_0 + \beta_1 u_{mt} + \beta_2 p_t^{ACH} u_{mt} + \beta_3 t + \alpha_{jmt} + \varepsilon_{ijmt}, \end{aligned}$$

where β_1 and β_2 are newly defined parameters, defined by $\beta_1 = N \cdot \left(\tilde{\beta}_1 - \tilde{\beta}_2 p_t^{CHK}\right)$ and $\beta_2 = N \cdot \tilde{\beta}_2$.[11] By integrating out over the logit error ε_{ijmt},[12] we get the probability that a consumer at bank j in market m in time t adopts ACH as:

$$\text{(8)} \quad P_{jmt} = \frac{\exp\left(\beta_0 + \beta_1 u_{mt} + \beta_2 p_t^{ACH} u_{mt} + \beta_3 t + \alpha_{jmt}\right)}{1 + \exp\left(\beta_0 + \beta_1 u_{mt} + \beta_2 p_t^{ACH} u_{mt} + \beta_3 t + \alpha_{jmt}\right)}.$$

Using again the assumption that there are a large number of consumers at each bank, P_{jmt} is the exact proportion of consumers who adopt ACH. Then, the equilibrium number of ACH transactions at bank j must satisfy:

$$\text{(9)} \quad T_{jmt}^{ACH}\left(A_{mt}\right) = A_{jmt} T_{jmt} P_{jmt}\left(A_{mt}\right) u_{mt}\left(A_{mt}\right).$$

Note that if bank j does not adopt ACH, $A_{jmt} = 0$ and $T_{jmt}^{ACH} = 0$. If bank j does adopt, the number of ACH transactions is equal to the total number of transactions (T_{jmt}) times the proportion of the banks' customers who adopt ACH (P_{jmt}) times the proportion of those customers' transactions that are with other customers in the market who have adopted ACH (u_{mt}).

We next turn to optimal bank adoption decisions conditional on the above model of transaction choice. Recall that in the first stage, banks simultaneously decide whether to adopt ACH technology. Denote the marginal cost to the bank of an ACH and a check transaction as mc_t^{ACH} and mc_t^{CHK}, respectively. Assume that there is a per-period fixed cost FC of adopting

include one set of flexible unobservables. The reason we put them in consumer fixed costs is because this was the specification that appears to fit the data best.

[11] We fold p_t^{CHK} into u_{mt} in (7) because we do not have data on the price of checks.

ACH technology. Importantly, this is a per-period cost, not a one time sunk cost of adoption. As such, there are no dynamic optimization issues and firms simply maximize per-period profits.[13]

Banks compare profits from adopting ACH to profits from not adopting ACH. The increment in profits from adopting is the number of ACH transactions times the difference in margins, minus the fixed cost of adoption. This increment is:

$$\text{(10)} \quad \begin{aligned} \Pi_{jmt}\left(T_{jmt}^{ACH}\right) &= T_{jmt}^{ACH}\left[\left(p_{jmt}^{ACH} - mc_{jmt}^{ACH}\right) - \left(p_{jmt}^{CHK} - mc_{jmt}^{CHK}\right)\right] - FC \\ &= T_{jmt}^{ACH} \times \text{markup} - \left(\overline{FC} + \alpha_{jmt}^{E}\right) \end{aligned}$$

where fixed costs are divided into a common component ($\overline{FC}$) and an idiosyncratic component (α_{jmt}^{E}). As with the consumer fixed cost ε_{ijmt}, we normalize α_{jmt}^{E} to have a standard logistic distribution. We estimate both $\overline{FC}$ and markup.

Bank j will adopt ACH at time t if and only if $\Pi_{jmt}\left(T_{jmt}^{ACH}\right) > 0$. We can see that adoption will depend on other banks' decisions through T_{jmt}^{ACH}, which is a function of the equilibrium network adoption u_{mt}. An equilibrium $\left(A_{1mt} \ldots, A_{JmT}, T_{1mt}^{ACH}, \ldots, T_{Jmt}^{ACH}\right)$ requires that all banks' adoption decisions are optimal conditional on all other banks adoption decisions, i.e.

$$\text{(11)} \quad A_{jmt} = \left\{\Pi_{jmt}\left(T_{jmt}^{ACH}\left(A_{1mt}, \ldots A_{j-1m,t}, 1, A_{j+1,mt} \ldots, A_{Jmt}\right)\right) > 0\right\}, \ \forall j,$$

where $T_{jmt}^{ACH}(\cdot)$ satisfies (3), (8) and (9).

Some customers in our model will have accounts at branches of banks whose headquarters are outside the network. We assume that these banks make their adoption decisions without considering the conditions in the network; i.e. their adoption decisions are exogenous to

[12] Note that since we do not have consumer level data and include a flexible α_{jmt}, the assumption that the logit errors are iid is essentially WLOG.

[13] There is some evidence of this nature of fixed costs in our data as we see a number of banks switching from adoption to non-adoption between periods. See Gowrisankaran and Stavins (2001) for details.

the unobservables in the market. However, conditional on adoption, customers at those banks choose their adoption decisions using the same criteria as banks whose headquarters are in the network. Thus, if a bank with headquarters outside the network chooses to adopt ACH, the probability of its customer adopting ACH will be given by (8). As in Gowrisankaran and Stavins (2003), the non-local banks will provide an important source of identification.

There are often multiple equilibria of this network adoption game. To see this, note that on one hand, if every customer is using the network good, then any one customer is likely to want to use it. On the other hand, if no customer is using it, then any one customer is likely to not want to use it. This same logic is also true at the bank level. Because the value from another bank or customer adopting ACH is higher if the bank is itself adopting ACH, the adoption game is supermodular. Several properties follow from supermodularity.[14] These properties can easily be proved directly,[15] and do not depend on continuity but only on this monotonicity property. First, there exists at least one pure strategy subgame perfect equilibrium. Second, there exist one subgame perfect equilibrium that Pareto dominates all others and one (not necessarily distinct) subgame perfect equilibrium that is Pareto inferior to all others. Third, the proof of the second property is constructive, and it provides a very quick way to compute the Pareto-best and -worst subgame perfect equilibria. This last property is particularly important for estimation purposes.

To ensure an internally consistent specification, we need to specify the selection of equilibrium.[16] We want to estimate a specification that is consistent with the presence of multiple equilibria, and that can allow us to estimate whether markets tend to be in good or bad equilibria. Since we observe several separate networks, we want to allow for the possibility that some networks are in a good equilibrium while others are in a bad equilibrium. Hence, we assume that there is some frequency that any given network is in the Pareto-best equilibrium, with a corresponding frequency of being in the Pareto-worst equilibrium. We estimate the frequency as a parameter. Formally, let

[14] See Milgrom and Shannon (1994).
[15] See Gowrisankaran and Stavins (2003).
[16] Heckman (1978) shows that this type of simultaneous equations model is not well-specified without some such assumption.
[17] Our method of estimating models with multiple equilibria is similar to the method used by Moro (2000) who treats the equilibrium choice as a parameter.

$\omega_m \sim \text{iid } U(0,1)$. We assume that the market will be in the Pareto-best equilibrium if and only if $\omega_m < \exp(\Omega)/(1+\exp(\Omega))$, where Ω, the probability of being in the Pareto-best equilibrium, is a parameter that we estimate.[18] Note that we do not allow for the equilibrium to vary within a network across time.

3. Data

Our principal data set is the Federal Reserve's billing data that provides information on individual financial institutions that processed their ACH payments through Federal Reserve Banks.[19] We observe quarterly data on the number of transaction originations by bank for the period of 1995 Q2 through 1997 Q4. ACH transactions can be one of two types: credit or debit. A credit transaction is initiated by the payer; for instance, direct deposit of payroll is originated by the employer's bank, which transfers the money to the employee's bank account. A debit transaction is originated by the payee; for example, utility bill payments are originated by the utility's bank, which initiates the payment from the customer's bank account. For each financial institution in the data set, we have the ACH volume processed through the Federal Reserve each month and the total amount that the Federal Reserve charged for processing that volume. We also have the American Banking Association (ABA) number that allows us to link this data with other publicly available banking data.

The Federal Reserve is currently the dominant provider of ACH services. The Federal Reserve handled approximately 75 percent of the roughly 3.3 billion on-others commercial ACH transactions processed in 1996 and approximately 70 percent in 1998.[20] The remaining share of the on-others market was handled by three private sector ACH providers: Visa, New York

[18] Our method of estimating models with multiple equilibria is a generalization of the method used by Moro (2002) who estimates the equilibrium as a parameter. The difference is that we estimate the frequency of being in either equilibrium as a parameter, since we observe several regional markets, while Moro (2002) only has one market per year.

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

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

Automated Clearing House (now called Electronic Payments Network), and American Clearing House (formerly Arizona Clearing House). There are some network linkages between the different ACH providers. For instance, the Federal Reserve processes ACH items originated by members of the private networks and vice versa. However, for lack of data, we deal only with ACH transactions that are billed through the Federal Reserve, and treat Federal Reserve ACH as the relevant network for the good.

In addition to the ACH billing data, we use a number of publicly available databases to augment our data. First, we linked the Federal Reserve data with the quarterly Call Reports database. The Call Reports database provides information on bank assets, deposits, name, and the zip code of the headquarters for all banks that are registered with the Federal Deposit Insurance Corporation (FDIC). Several banks opened and closed during our sample period. We kept these banks in the sample for the quarters in which they were open.

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

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

In our model, we define a network to be a set of banks that are geographically close. Thus, we needed to find the distance between zip codes. We used Census information to find the latitude and longitude of zip code centroids, and used the standard great circle formula to find the distance between centroids.

Our estimation procedure is based on the assumption that a bank's network is geographically local. Our basic definition of a network is the set of banks whose headquarters are within 30 kilometers of the headquarters of a given bank. Because we are solving for an equilibrium of the adoption game, we need to also include all the banks that are within 30 kilometers of the banks that are within 30 kilometers, and all the banks that are near these banks, etc. We performed this process in order to separate our data set of 11,000 banks into mutually exclusive networks. Each network is self-contained, in the sense that every bank headquarters that is within 30 kilometers of any bank headquarters in the network is also in the network, and no bank headquarters in the network is within 30 kilometers of any bank headquarters outside the network.

One significant data problem is that many banks have become national in scope. As the relevant network for these banks is likely to be national, our model would not be particularly meaningful for these banks. Thus, we kept in our sample only banks that are in small markets. Specifically, we kept all networks with 10 or less bank headquarters total during every time period of our sample. From this set, we excluded networks where any one bank had more than 20 percent of its deposits outside the network, or where in aggregate, 10 percent of deposits for local banks were outside the network. We were left with a sample of 456 mutually exclusive networks comprising 878 local banks, observed over 11 time periods.

Figure 1 displays a map of New England with the networks from this region marked with asterisks, in order to give some idea about typical networks. One can see that these networks are comprised of small, isolated towns, such as Lewiston, ME and Nantucket, MA.

As described in Section 2, we use information from banks with branches in the network but headquarters outside the network, but model them separately from banks with headquarters in the network. We include in our sample 661 bank branches from banks outside the network.

Table 1 gives some specifics on the networks at every time period, broken down by the number of banks with headquarters in the networks. Approximately half of the network time-periods in our sample – 2730 in all – are composed of only one local bank. Another quarter of the network time-periods have two banks. However, there are large numbers of network time-periods with up to 10 local banks. Banks in our sample tend to be small banks, with assets of around \$100 million. The percentage of firms using ACH appears to be quite consistent across network size, although banks in smaller networks have fewer ACH transactions.

Table 2 examines the non-local banks in these markets. Of note is the large number of outside banks. For instance, in markets with one local bank, the average number of outside banks is 2.74. Although the sizes of non-local bank branches and local banks are similar in terms of deposits, non-local banks adopt ACH much more frequently. This is due to the fact that the non-local banks are, on average, much larger than the local banks, and than their local branches.

Table 3 gives some specifics on the changes in ACH usage over our sample period. We can see that the fraction of banks using ACH increased during our sample period. Moreover, there appears to be a large fraction of networks where every bank uses ACH – more than one would expect without correlations in usage.

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 2 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 \$0.01 in 1997. In May 1997, the Federal Reserve implemented a two-tier price system of \$0.009 for banks with less than 2500 transactions per file and \$0.007 for banks with more than 2500 transactions per file. We ignore the \$0.007 price because we do not have data on the number of transactions per file (only monthly totals) and because most of the banks in our sample are sufficiently small as to only pay the higher rate. Because prices are set by fiat and do not respond to changes in local demand, they may be viewed as exogenous. We do not have any information on the prices that banks charge to their customers. In addition to per-transaction costs, banks must file fees of \$1.75 per small file and \$6.75 per file per large file and pay an ACH participation fee of \$25 per month. Also, banks that offer ACH generally maintain a Fedline connection for ACH as well as other electronic payment services.

4. Estimation

Our model is based on a vector of unknown parameters $\left(\lambda, \beta, \Omega, \overline{FC}, \text{markup}, \sigma\right)$ and econometric unobservables (α, ω).[21] For ease of notation, let us group the unknown parameters together as θ. Our estimation algorithm seeks to recover θ from the data. In this section, we describe our estimation algorithm, including the computation of equilibria, and explain how the parameters of the model are identified.

4.1 Estimation Algorithm

Let us start by defining the data for one network in a given time period. For each bank, our data contain observed predetermined variables, namely its local deposits x_{jmt}, price p_t, time t, and its local/non-local status. For branches of non-local banks, our data also contain their observed ACH adoption decisions A_{jmt}, that we assume to be pre-determined. Our data also contain the observed endogenous variable T_{jmt}^{ACH}, for local banks only.

Now consider a given parameter vector $\theta' = \left(\lambda', \beta', \Omega', \overline{FC}', \text{markup}', \sigma'\right)$. For this parameter vector, densities for fixed-cost unobservables α are defined and it is thus possible to simulate them for a given vector over time. Given a vector of simulation draws on α and exogenous data, we can easily compute the Pareto-worst and -best subgame perfect equilibrium of the industry.[22] Conceptually, we can then match the weighted sum of the two predicted equilibria to the data, where the weights depend on the equilibrium selection parameter Ω'.

We estimate the model by using simulated maximum likelihood. To understand the estimation procedure, consider the likelihood for market m for a market with J banks of which

---

[21] Note that the consumer level unobservables ε are aggregated up in the model.
[22] In Section 4.2, we provide details on the computation of the Nash equilibrium.

the first $\hat{J}$ are local and the remaining are branches of non-local banks. Defining $X_m = \left\{x_{1mt}, \ldots, x_{Jmt}, A_{\hat{J}+1mt}, \ldots, A_{Jmt}, t, p_t\right\}_{t=1}^{T}$ to be the exogenous data for this market, the likelihood is:

$$(12) \qquad L_m(\theta') = P\left(\left\{T_{1mt}^{ACH}, \ldots, T_{\hat{J}mt}^{ACH}\right\}_{t=1}^{T} \middle| X_m, \theta'\right),$$

where "P" indicates the probability density function. The likelihood specifies the joint probability of the actions of banks and consumers in market m. Note that this is the joint probability for the banks and consumers in market m *over all time periods*, which is necessary because the unobservables are potentially correlated across time.

As noted above, it is quite easy to simulate data from our model. However, it would be virtually impossible to evaluate (12) analytically. This suggests the use of simulation estimation. One approach to simulation estimation would be to use GMM, simulating data from the model and finding parameters that make moments of this simulated data as close as possible to moments of the observed data. However, the complicated correlation structure of the model (correlation within banks across time, within markets across time, and across banks in a given market) makes it hard to write down a concise set of moments to match. For instance, we would need a large number of covariance moments. Therefore, we use simulated maximum likelihood (SML).

Many recent papers use SML to estimate structural models.[23] SML is attractive to use with fully-specified structural models, because one does not have to worry that the specification of the estimator is influencing the estimated parameters. However, there is a significant technical problem in using SML to estimate our model: our dependent variables have a continuous component to them, namely the number of ACH transactions. Therefore, straightforward

[23] See Keane and Wolpin (1997) or Rust (1987) for instance.

simulation will necessarily result in likelihood zero events, as our simulated transactions will never match the observed transactions.

To solve this problem, we add measurement error to the outcome variables of the model as in Keane and Wolpin (2000). This gives the observed data a non-zero likelihood, which makes it feasible to estimate.[24] There are several issues that determined our functional form choice for the measurement error. First, it is conceptually difficult to define an appropriate measurement error process because 0 transactions is a common outcome (accounting for 33.5% of local bank observations), but the number of transactions can be large. We want a log functional form of the measurement error (i.e. proportional measurement error) to account for the observations with many transactions but a linear functional form to account for measurement error in the observations with few transactions. Moreover, there are a number of banks with a positive but very small number of transactions (21.7% of local bank observations have one or more but less than ten transactions during a quarter). Many of these observations are likely due to banks that have not adopted ACH, but are processing a return item, or initiating a transaction as a one-time favor to a specific customer. Lastly, it is difficult to separately identify the measurement error process from the structural parameters.

To account for these different factors, we add two probabilities to our data generating process – first, a probability that the bank reports positive transactions when the bank has not adopted ACH, and second the reverse, i.e. a probability that the bank reports zero transactions when the bank has in fact adopted. After a careful look at the number of transactions data for evidence on the degree of the spurious adoption story above, we set the first probability (i.e. the probability of reporting positive transactions when a bank hasn't adopted) to 20%. Again in a specification motivated by examining the data, we assume that in this 20% case, then the reported number of transactions follows an exponential decay process, with a decay factor of 0.7.[25] In a more arbitrary way, we set the second probability (the probability of reporting zero

[24] Note that this is analogous to kernel smoothing a simulated likelihood.
[25] We arrived at the 20% figure by examining the number of banks reporting strangely low numbers of transactions relative to the banks that had clearly adopted. We then chose the 0.7 figure to match the distribution of these spurious observations.

transactions when there are in fact a positive number of transactions) to 5%. We intend to check our model for robustness with respect to this parameter. When there are positive transactions, we assume a log normal measurement error process for the case when positive transactions are actually reported, which happens 95% of the time. We normalize the log normal to a minimum of 1000, to get proportional measurement error for high numbers of transactions and more than proportional measurement error for low numbers of transactions. Mathematically, then, if $T_{jmt}^{ACH} > 0$, then

$$\text{(13)} \quad \begin{aligned} &\log\left(\text{observed } T_{jmt}^{ACH} + 1000\right) = \log\left(\text{actual } T_{jmt}^{ACH} + 1000\right) + e_{jmt}, \text{ with prob. } 0.95, \\ &\text{observed } T_{jmt}^{ACH} = 0, \text{ with prob. } 0.05, \end{aligned}$$

while if $T_{jmt}^{ACH} = 0$, then

$$\text{(14)} \quad \begin{aligned} &\text{observed } T_{jmt}^{ACH} = 0, \text{ with prob. } 0.8, \\ &\text{observed } T_{jmt}^{ACH} = n > 0, \text{ with prob. } 0.2 \times \left(1 - 0.7\right) \times 0.7^{n-1}. \end{aligned}$$

We assume that $e_{jmt} \sim N\left(0, \sigma_e^2\right)$. The only measurement error parameter that we estimate is σ_e^2. The simulated likelihood function that we maximize now has observed T_{jmt}^{ACH}, and not T_{jmt}^{ACH}, as the dependent variable. Formally, the simulated log likelihood for the market m with NS simulation draws can be written as:

$$\text{(15)} \quad \begin{aligned} &\log \ L_m\left(\theta'\right) \\ &= \log\left(P\left(\left\{\text{observed } T_{1mt}^{ACH}, \ldots, \text{observed } T_{\hat{J}mt}^{ACH}\right\}_{t=1}^{T} X_m, \theta'\right)\right) \\ &\approx \log\left(\frac{1}{NS}\sum_{s=1}^{NS}\left[\begin{aligned} &\tfrac{\exp(\Omega')}{1+\exp(\Omega')}\prod_{j,t}^{\hat{J},T} P\left(\text{observed } T_{jmt}^{ACH} \middle| \text{simulated } T_{jmt,s}^{ACH}\left(X_m, \theta', \omega_m = 1, \alpha_{m,s}\right)\right) \\ &+\left(1 - \tfrac{\exp(\Omega')}{1+\exp(\Omega')}\right)\prod_{j,t}^{\hat{J},T} P\left(\text{observed } T_{jmt}^{ACH} \middle| \text{simulated } T_{jmt,s}^{ACH}\left(X_m, \theta', \omega_m = 0, \alpha_{m,s}\right)\right) \end{aligned}\right]\right), \end{aligned}$$

where the formulas for the conditional densities of observed T_{jmt}^{ACH} in (15) can be evaluated using (13) and (14), and simulated $T_{jmt,s}^{ACH}$ is computed by solving for the appropriate subgame perfect equilibrium of the model. Note that since there are only two equilibria, we do not simulate ω_m, but instead integrate over the two equilibria, weighting the Pareto-best equilibrium with probability $\exp(\Omega')/(1+\exp(\Omega'))$.

Recall that a number of our local markets contain bank branches of large banks from outside the market. As adoption decisions of these banks are likely at a regional or national level, we treat them as exogenous to our model. However, we *do* model the adoption decisions of local consumers of those branches. As we do not observe the number of ACH transactions at these branches, these consumer decisions do not directly enter the likelihood function (15). Nonetheless, these decisions do indirectly enter the likelihood function through their effects on the adoption decisions of local banks and the consumers of these local banks.

4.2 Computation of Equilibrium

In order to compute the likelihood function (15), we need to evaluate simulated $T_{jmt,s}^{ACH}\left(X_m, \theta', \omega_m = 0, \alpha_{m,s}\right)$, which involves solving for the Pareto-best or -worst subgame perfect equilibrium of the model conditional on a vector of pre-determined variables and econometric unobservables.

In general, estimation of Nash equilibria can be very computationally intensive. This computational intensity is a large part of the reason why structural models are notoriously difficult to estimate. In our case, it is computationally simple to solve for both subgame perfect equilibria. The underlying reason for this is that the network externality is assumed to always be positive, which makes the game supermodular. Because of this, the optimal reaction functions will always be a monotone mapping of the previous stage reaction functions. This is also the basis of the proof that there is a Pareto-best subgame perfect equilibrium given in Gowrisankaran and Stavins (2003), Proposition 1.

Thus, we solve for the Pareto-best subgame perfect equilibrium by using the following iterative process on adoption of banks and consumers. We start the first iteration by assuming that all banks and consumers use ACH, i.e. $\left(A^1_{1mt}=1,\ldots,A^1_{Jmt}=1\right)$ and $\left(P^1_{1mt}=1,\ldots,P^1_{Jmt}=1\right)$. In the second iteration we consider each bank in turn. For bank j, we find the consumer adoption decisions given the adoption decisions in the first iteration, except with the assumption that bank j has adopted ACH.[26] We then determine whether bank j would find it profitable to adopt given this level of usage, and enter this as the new strategy. We repeat this process for each bank. This results in a vector $\left(A^2_{1mt},\ldots,A^2_{Jmt},P^2_{1mt},\ldots,P^2_{Jmt}\right)$ where each level is weakly less than in the first iteration. We repeat this process until convergence; convergence is guaranteed by this monotonicity property. As in Gowrisankaran and Stavins (2003), we can show that the limiting values $\left(A^N_{1mt},\ldots,A^N_{Jmt},P^N_{1mt},\ldots,P^N_{Jmt}\right)$ form a Pareto-best subgame perfect equilibrium. Correspondingly, if we start the first iteration by assuming that no one is using ACH, i.e. $\left(A^1_{1mt}=0,\ldots,A^1_{Jmt}=0\right)$ and $\left(P^1_{1mt}=0,\ldots,P^1_{Jmt}=0\right)$ and then iterate to convergence, the algorithm will converge to the Pareto-worst Nash equilibrium.

We can also use variants of this algorithm to solve for the outcomes when local banks internalize the network externality and when consumers internalize the externality, both of which we report. For the case of banks internalizing the externality, we solve for the bank adoption decisions differently, assuming that banks value the difference in profits from *all* banks resulting from their adoption decision. For the case of the consumers internalizing the externality, we need to solve for the optimal cutoff fixed cost for each consumer, which differs from the non-cooperative case, even conditional on other agents' actions.

Because of the monotonicity of the reaction functions, our algorithm converges to the appropriate Nash equilibrium very quickly. For instance, to evaluate one parameter iteration with 10 simulation draws, we require computing a Nash equilibrium for the roughly 500 markets over

11 time periods with 10 different simulation draws and 2 equilibria. It takes about 3 seconds to solve for these 100,000 equilibria on a modern workstation.

To compute confidence intervals for the parameter estimates, we use bootstrap methods, which are robust to most misspecification, given that observations are iid. For our bootstrap method, we resample the data with replacement from the original data set, treating a network over time as the unit of observation. We then recompute the maximum likelihood estimates using the new data set, and repeat this 50 times to obtain accurate confidence intervals for the parameters.

4.3 Identification

We now explain what identifies the important parameters of our model. We focus on the fixed costs and marginal benefits of ACH adoption at both the bank and consumer level. These are the parameters that govern the extent of the externalities associated with ACH. As noted in the motivation, we have three sources of identification. Our formal model of equilibrium allows us to combine all the sources of identification into one estimation procedure that uses all these assumptions.

If we knew the marginal benefits of ACH adoption (for banks the relative markup for ACH, for consumers the relative utility from an ACH transaction), the levels of adoption decisions would identify the fixed costs of adoption. In other words, the observed proportion of bank adoption would identify $\overline{\mathrm{FC}}$, and the observed proportion of consumer decisions would identify β_0.[27] Thus we focus on identification of the marginal benefits of adoption.

A first source of identification comes from the (assumed exogenous) adoption decisions of large, non-local banks. Consider a local bank j and its consumers. As the adoption of non-local banks exogenously increases, the equilibrium adoption rate of bank j's consumers will

[26] Recall that consumers of bank j observe the decisions of bank j before making their adoption decisions.

[27] Note that there is a selection issue here, since we do not observe the proportion of consumers adopting for banks that do not adopt.

increase. The extent of this increase will identify the relative utility that consumers gain from an ACH transaction. The exogenous increase in adoption of non-local banks will also increase the probability that bank j adopts ACH. Note that this comes from two sources – first the fact that bank j's consumers are more apt to adopt, and second that the number of transactions each of those consumers would make conditional on adoption would increase. Put another way, even if the adoption decisions of bank j's consumers were not affected by the increase in non-local adoption, bank j would still be more likely to adopt (due to the increase in transactions for existing adopters). The extent of the increase in bank j's adoption (or adoption probability across banks in different markets) will identify the relative markup from an ACH transaction.

A second source of identification comes from the assumption that market structure (i.e. the deposit sizes of banks) is exogenous. Consider banks in two sets of markets – the first set (A) consists of monopolies, the second (B) duopolies. With network externalities, note that banks in the A markets should be more likely to adopt, as the bank level externality is completely internalized. As the relative markup banks obtain from ACH transactions increases, we should see bigger differences in adoption probabilities between the two sets of markets. Thus, differences in adoption probabilities across different types of markets should identify the relative markup. A last source of identification of network externalities we examine comes from correlation in adoption decisions, both at the bank and consumer level. We can test for the robustness of both of these sources of identification. The second source can be eliminated by allowing the relative markup of ACH to checks to differ based on market power. The third source is only used if we set $\alpha_{mt}^{D} = 0$; thus, our base model does not allow for this source of identification.

All of these sources of identification will yield different effects depending on whether the network externality is at the consumer or bank level. To see this, note that if there are no consumer fixed costs of adoption, we can precisely predict the ACH volume conditional on bank decisions: it is the square of the fraction of banks that adopt. As consumer fixed costs increase, there will be less transactions conditional on a set of bank adopters, particularly when the set is small. Thus, the extent to which numbers of ACH transactions increase moving from markets with different numbers of bank adopters will identify the consumer level parameters, while the extent to which bank adoption changes will identify the bank level parameters.

Lastly, the equilibrium selection parameter Ω will be identified by differences in usage given different industry structures. For instance, as the number of firms increases, the increasing externality should make it more likely that there is a Pareto-worst equilibrium that is distinct from the Pareto-best equilibrium. Thus, we can identify the equilibrium selection parameters by examining whether there is increased unexplained variance in behavior for networks with more than one bank that does not exist for networks with one bank. Note that if we saw a high variance in the usage levels in *all* markets, this could be evidence of high variances of the random effects α, not necessarily multiple equilibria.

5. Results and Implications

Using the simulated maximum likelihood developed in Section 4, we have estimated structural parameters for our base model and various specifications. We first present the results and then present policy experiments.

5.1 Base results

Table 4 gives base parameter values. For the base specification, we allow for market-specific time-varying random effects. Thus, we are only using our first two sources of identification.

Most of the parameters listed in Table 4 appear to be reasonable. For instance, the coefficient on time trend is positive, suggesting that there is increased acceptance of technological goods and that a portion of the network externality is from outside the 30 kilometer area of our model. The ACH price coefficient is negative. On the consumer side, both consumer fixed costs and marginal benefits are positive and the ratio appears reasonable. On the other hand, for banks, the estimated mean fixed costs of adoption seems very small in comparison to the normalized net markup.

The correlation parameters are interesting in that the firm specific (constant over time) and market specific (constant over time) random effects appear to be considerably more important than the time varying effects α^{A}_{jmt} and α^{D}_{mt}. This appears to suggest that there is not much varying over time in these markets, at least with regard to unobservables. The last important parameter is the equilibrium selection parameter Ω. The estimated value of the parameter, -0.461, suggests that approximately 39% of markets are in the Pareto best equilibrium.

Virtually all of the parameters in Table 4 are precisely estimated. Only one of the thirteen parameters, $\sigma_{\alpha^{D}}$, is not significantly different from zero at the 1% level. The high precision of the estimates is characteristic of structural estimation models.

Table 5 examines the fit of the model. Since we assume that the data is characterized by measurement error, we compare the predictions of the model with measurement error to the data. The model matches the percentage of banks adopting in the data very precisely. The model predicts somewhat more ACH transactions than we observe in the data, and predicts a somewhat lower standard deviation of the number of ACH transactions across banks. However, the relative increase in the number of transactions between the start and end of our sample period mirrors the data.

We also report various correlations as generated by the model and the data. The model captures the correlations between bank deposits and the number of ACH transactions quite well. The model does fairly well at capturing the cross-time correlation in bank adoption decisions and in the number of ACH transactions. However, the model somewhat overpredicts the correlation between adoption decisions and number of ACH transactions for banks in a network at a given time.

Although the parameter estimates are interesting in of themselves, it is much more valuable to examine the impact of the parameters on the estimated equilibrium. This is done in Table 6. We look at 3 statistics of the estimated equilibrium – the percentage of banks adopting ACH, the percentage of consumers adopting, and percentage of overall transactions done through ACH. At the estimated parameters, 79.4% of banks are adopting, 24.6% of consumers are adopting, and 8.3% of all transactions are ACH. Note that according to our model, in a given market the percentage of ACH transactions should equal the percentage of consumers adopting squared. The reason that this is not true in the data as a whole is due to the convexity of this function – there are some markets where lots of consumers adopt and some where very few do.

The second row of Table 6 examines what our model predicts if there were no mean bank fixed costs of adoption. The difference between this and the first row is indicative of the level of the network externalities at the bank level. Although many more banks adopt ACH, the differences in transactions processed with ACH are small. This is due to our small estimated bank fixed cost of adoption, which implies that the holdup from consumers not using ACH is not due to their banks. On the other hand, when we eliminate the consumer mean fixed cost of adoption, there are big changes in the equilibrium proportion of consumers that adopt ACH. Consumer adoption increases to 54.5%, a number that is still far less than 100% due to the random component of fixed costs. In response to this expected adoption by consumers, banks also increase adoption, to 98.2%. In this equilibrium, 33.1% of all transactions are done using ACH. These estimates suggest that consumer fixed costs are the primary impediments to ACH adoption.

The next two rows of Table 6 examine the existence of multiple equilibria at our estimated parameter values by forcing either the Pareto-worst or the Pareto-best equilibria. The results across the two equilibria are very similar, though not identical. This suggests that at our estimated parameters, multiple equilibria are not a significant issue.

Lastly, we investigate what would happen if some of these externalities could be internalized. There is no natural way to compare consumer utility to firm profits. As a result, we cannot solve for the first best outcome, i.e. if a social planner controlled all agents in the economy. However, we can investigate what happens if all the local banks coordinated decisions to maximize joint profits, or if all consumers coordinate to maximize joint utility. Results are in the last two columns of Table 6. As might be expected, joint profit maximization of all the local banks does not change matters much. On the other hand, joint utility maximization of consumers does increase adoption, causing the number of ACH transactions to rise by about 16 percent, from 8.3% of all transactions to 9.6% of all transactions. While it is difficult to assess the tradeoff between consumer and bank utility, these two results taken together suggest that complete joint surplus maximization would also yield about 16 percent more transactions than the equilibrium outcome. Thus, while the world is not at the first-best usage level, even with complete surplus maximization most transactions are completed with checks.

5.2 Policy Experiments

The above discussion suggests that network externalities are really biting at the consumer level rather than the firm level. It appears to be consumer fixed costs which are limiting the adoption of ACH. In contrast, bank fixed costs are small and not significantly preventing ACH use. This suggests that government policy, particularly at the consumer level, might increase welfare. We examine this possibility in Table 7.

The first column of Table 7 again examines properties of the estimated equilibrium. In addition to statistics on consumer and bank adoption, we report welfare measures – the sum of firm profits and the sum of consumer utilities. We have no way of converting these measures into dollars, so it is important to realize that these measures are not comparable to each other. Consumer utility is measured in "utils", and profits are measured in "profit units."

The second two columns essentially repeat two of the experiments of the prior section. We remove, sequentially, consumer and bank mean fixed costs through a government subsidy. Rows 6 and 7 of the table report the cost to the government (in profit units and utils respectively) of these policies. Rows 8 and 9 report the total profit units (bank profits – government cost in profit units) and total utils (consumer utils – government cost in utils) resulting from these policies respectively. Again, the consumer subsidy is far more effective at increasing ACH usage. Bank fixed costs are simply not large enough to prevent adoption. Also note the extremely large benefits to banks from this consumer subsidy, as they are able to make considerably more variable profits.

Considering a subsidy of mean fixed costs is rather arbitrary, as there are distributions of these fixed costs at both the bank and consumer level. Columns 4 and 5 consider very large subsidies to banks and consumers, subsidies large enough to get virtually everyone to adopt (note that since there are some non-local banks who do not adopt, we cannot get all consumers to adopt). Even in this case, there are extremely limited changes with the bank subsidies.

We cannot make any conclusive evaluations of the above policies. This is because in all cases, either total utils or total profits go down as a result of the policy. Since we have no way of relating the increases in profits to the decrease in utils (or vice-versa), we cannot conclude a policy is welfare improving. Note why, for example, with the mean consumer fixed cost subsidy, total utils go down. This is because this subsidy is too large, generating inefficient adoption decisions by consumers. With a smaller, more efficient, consumer subsidy, we might hope to keep even or increase total utility, as well as total profits. Column 7 exhibits results from the smallest consumer subsidy (approximately) that does this, 25% of their mean fixed cost. With this subsidy, total utils are unchanged, but firm (and total) profits increase by more than 50%. This subsidy therefore unambiguously increases welfare in the market. Note that this is likely not an optimal policy - to determine that we would need to devise a way to compare profits to utility.

6. Conclusions

In this paper, we have estimated a structural equilibrium model of network externalities in the ACH banking industry in order to estimate the causes and magnitudes of network externalities for this industry. Our parameter estimates are precisely estimated and fit the data reasonably well.

We find that bank fixed costs from ACH adoption are low and do not explain why ACH is not more widely used. In contrast, consumer fixed costs of ACH adoption are substantial, and are a major explanation for the lack of ACH usage. Thus, changes that lower the consumer fixed cost of ACH adoption will encourage adoption and usage of ACH. As electronic payment technologies become more widely accepted and used at the consumer level, we will expect to use vastly more ACH transactions.

Although we estimate that the Pareto-worst equilibrium is not identical to the Pareto-best equilibrium, we find that the two equilibria are very similar to each other in their implied ACH adoption decisions. Because the bank fixed costs are so low, the equilibrium bank ACH adoption is very close to the first best adoption level. In contrast, the first-best consumer adoption level implies about 16% more ACH transactions than the observed equilibrium. Policies that subsidize a portion of consumer fixed costs can unambiguously increase welfare.

References

Angrist, Joshua D. and Guido Imbens W. (1994). "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62: 467-76.

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

Gowrisankaran, Gautam and Joanna Stavins (2003), "Network Externalities and Technology Adoption: Lessons From Electronic Payments," forthcoming, *RAND Journal of Economics*.

Heckman, James J. (1978), "Dummy Endogenous Variables in a Simultaneous Equation System," *Econometrica* 46, 931-959.

Heckman and Robb (1985). “Alternative Methods for Evaluating the Impact of Interventions.” In Longitudinal Analysis of Labor Market Data ed. by James J. Heckman and Burton Singer. New York: Cambridge University Press.

Keane, Michael P. and Kenneth I. Wolpin (1997), “The Career Decisions of Young Men,” *Journal of Political Economy* 105, 473-522.

Keane, Michael and Kenneth Wolpin (2000). ‘Estimating the Effects of Welfare on the Education, Employment, Fertility and Marriage Decisions of Women.” Mimeo, Yale University.

Manski, Charles (1993). “Identification of Endogenous Social Effects: The Reflection Problem,” *Review of Economic Studies* 60: 531-42.

Milgrom, Paul and Chris Shannon (1994). “Monotone Comparative Statics,” *Econometrica* 62: 157-180.

Moro, Andrea (2000). “The Effect of Statistical Discrimination on Black - White Wage Inequality: Estimating a Model With Multiple Equilibria.” Mimeo, University of Minnesota.

Rust, John (1987). “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher.” Econometrica 55: 999-1033.

Rysman, Marc (2000). “Competition Between Networks: A Study of the Market for Yellow Pages.” Mimeo, Boston University.

Topa, Giorgio (2001). “Social Interactions, Local Spillovers and Unemployment.” *Review of Economic Studies* 68: 261-296.

Table 1: Characteristics of Banks in Network

Number of banks based in network	Number of networks/time periods	Mean deposits	Mean percent of banks using ACH	Mean ACH transactions by bank
1	2730	$45.8 Mil.	64.3%	457.7
2	1310	$49.5 Mil.	64.5%	452.0
3	367	$59.4 Mil.	67.8%	1217
4	172	$73.0 Mil.	74.4%	1348
5	83	$50.1 Mil.	74.2%	912.5
6	51	$125 Mil.	70.3%	3485
7	31	$139 Mil.	73.7%	2155
8	41	$57.5 Mil.	66.2%	991.5
9	39	$79.9 Mil.	69.5%	897.9
10	25	$81.6 Mil.	57.6%	732.2

Table 2: Characteristics of Branches of Non-Local Banks

Number of banks based in network	Mean number of non-local banks	Std. dev. of number of non-local banks	Mean deposits within network by non-local banks	Mean total deposits by non-local banks	Percent of non-local banks using ACH
1	3.43	2.74	$59.8 Mil.	$10.4 Bil.	88.5%
2	2.50	2.38	$60.2 Mil.	$6.8 Bil.	85.8%
3	4.05	3.21	$96.0 Mil.	$9.2 Bil.	89.0%
4	4.34	3.32	$92.0 Mil.	$8.6 Bil.	88.5%
5	6.15	5.16	$187 Mil.	$4.8 Bil.	83.3%
6	5.67	5.20	$96.9 Mil.	$8.5 Bil.	84.1%
7	9.13	4.26	$78.6 Mil.	$5.0 Bil.	91.9%
8	6.80	4.65	$95.2 Mil.	$7.9 Bil.	86.0%
9	8.72	5.80	$104 Mil.	$6.9 Bil.	87.4%
10	6.56	3.80	$120 Mil.	$4.9 Bil.	81.1%

Note: Table based on observations kept in sample.

Table 3: Usage Over Time by Banks in Network

Time Period	# of networks with no firm using ACH	# of networks with some, but not all, firms using ACH	# of networks with all firms using ACH
1995: Q2	14.3%	57.1%	28.6%
1995: Q3	16.8%	57.4%	25.7%
1995: Q4	17.3%	55.8%	26.9%
1996: Q1	14.3%	55.6%	30.1%
1996: Q2	10.9%	51.6%	37.5%
1996: Q3	12.5%	51.0%	36.5%
1996: Q4	8.4%	50.3%	41.4%
1997: Q1	7.3%	46.1%	46.6%
1997: Q2	5.8%	41.3%	52.9%
1997: Q3	7.1%	42.9%	50.0%
1997: Q4	6.1%	42.2%	51.7%

Note: Table includes networks with 2 or more banks kept in sample.

Table 4: Parameter Estimates

Parameter	Value	Standard Error
λ (transactions coefficient)	20.43***	0.3609
β_0 (consumer fixed benefit)	-2.174***	0.0597
β_1 (consumer marginal benefit)	0.556***	0.0946
β_2 (price coefficient)	-0.235***	0.0603
β_3 (time coefficient)	0.063***	0.0033
Markup	338.9***	0.0888
$\overline{FC}$ (bank fixed costs)	7.942***	0.2732
Ω (equilibrium selection parameter)	-0.461***	0.0242
σ_{α^A} (std. dev. of random effect α^A_{jmt})	-0.0642***	0.0218
σ_{α^B} (std. dev. of random effect α^B_{jm})	1.939***	0.0757
σ_{α^C} (std. dev. of random effect α^C_{m})	0.485***	0.2000
σ_{α^D} (std. dev. of random effect α^D_{mt})	0.0096	0.0159
σ_e (std. dev. of measurement error e_{jmt})	0.270***	0.0136

*** Significant at 1% level.

Table 5: Goodness of Fit

Moment	Data	Model (with measurement error)
% of banks adopting	0.665	0.678
Mean # of transactions	826	1,172
Standard deviation of # of transactions	3.95	3.02
Mean # of transactions, Q2:95	517	727
Mean # of transactions, Q4:97	1,253	1,627
Correlation between deposits and bank adoption	0.182	0.150
Correlation between deposits and # of transactions	0.424	0.375
Correlation between bank adoption decisions for a given bank at Q2:95 and Q4:97	0.427	0.310
Correlation between # of transactions for a given bank at Q2:95 and Q4:97	0.652	0.782
Correlation between bank adoption decisions within a network / quarter	0.083	0.301
Correlation between # of transactions within a network / quarter	0.112	0.209

Sample includes local banks only.

Table 6: Economic Significance of Parameters

Change	% of banks adopting	% of consumers adopting	% of transactions completed with ACH
Estimates	79.4%	24.6%	8.3%
No mean bank fixed costs	96.5%	24.6%	8.3%
No mean consumer fixed costs	98.2%	54.5%	33.1%
Always in bad equilibrium	79.3%	24.5%	8.3%
Always in good equilibrium	79.5%	24.5%	8.3%
Local banks internalize externality	81.0%	24.6%	8.3%
All consumers internalize externality	80.3%	26.4%	9.6%

Table 7: Policy Experiments

Policy	None	Subsidize Consumer Mean FC	Subsidize Bank Mean FC	Very Large Consumer Subsidy	Very Large Bank Subsidy	Subsidize 0.25 Cons. Mean FC
% of Local Banks Adopting	79.56	97.94	96.44	99.98	100	87.53
% of Consumers Adopting	23.64	53.66	23.73	93.83	23.73	30.41
% ACH transactions	7.69	32.13	7.70	89.57	7.70	11.84
Firm Profits	7.25e6	30.66e6	7.31e6	85.76e6	10.81e6	11.22e6
Consumer Utility	2.56e5	5.09e5	2.63e5	286.07e5	2.64e5	3.03e5
Cost to Govt. (in profit units)	0	0	0.07e6	0	3.5e6	0
Cost to Govt. (in utils)	0	3.30e5	0	288.72e5	0	0.47e5
Total Profits	7.25e6	30.66e6	7.24e6	85.76e6	7.23e6	11.22e6
Total Utility	2.56e5	1.78e5	2.63e5	-2.65e5	2.64e5	2.56e5

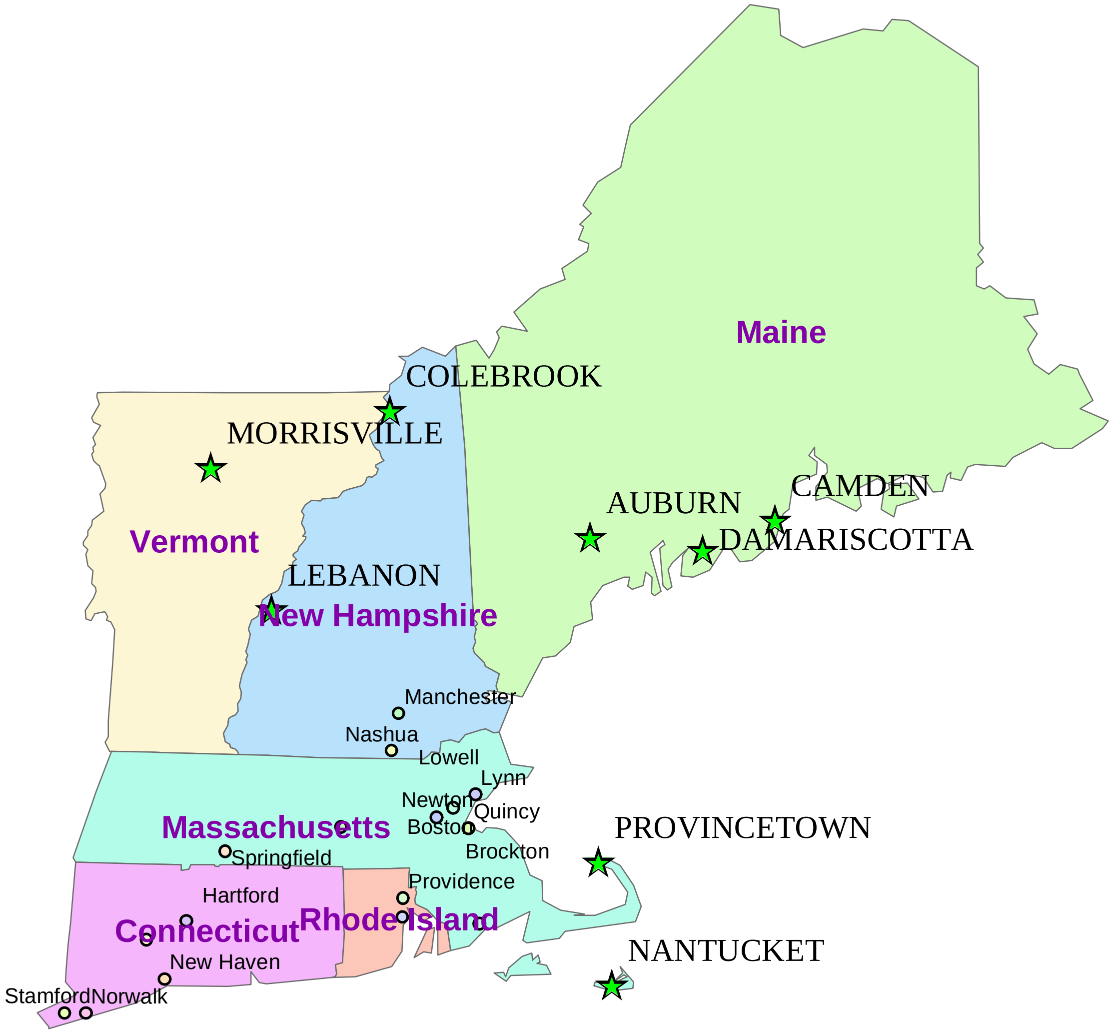
Maine
COLEBROOK
MORRISVILLE
Vermont
LEBANON
New Hampshire
AUBURN
CAMDEN
DAMARISCOTTA
Manchester
Nashua
Lowell
Lynn
Newton
Quincy
Boston
Massachusetts
Springfield
Brockton
PROVINCETOWN
Providence
Hartford
Rhode Island
Connecticut
New Haven
NANTUCKET
Stamford
Norwalk

Figure 2: Per-item origination fees for Federal Reserve ACH Processing

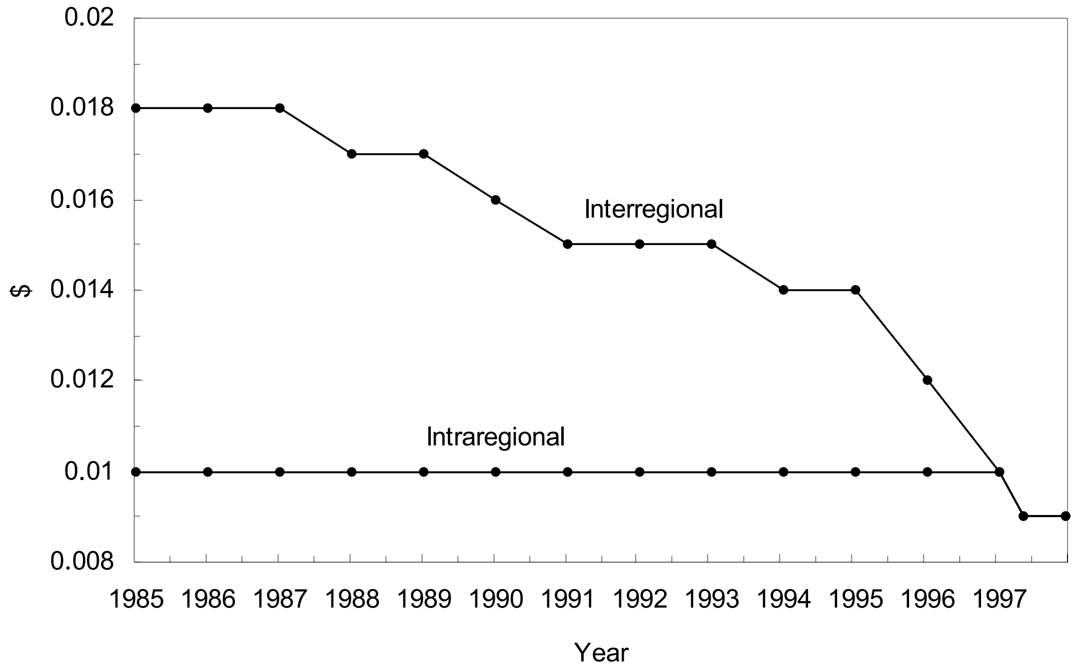

Note: In May 1997, volume-based pricing was introduced, with price set to 0.9 cents per item for files with less than 2500 items and 0.7 cents per item for files with 2500 or more items.