

Quantifying equilibrium network externalities in the ACH banking industry

Daniel A. Akerberg*
Gautam Gowrisankaran**

Abstract:

We seek to determine the causes and magnitudes of network externalities for the automated clearinghouse (ACH) electronic payments system. We construct an equilibrium model of customer and bank adoption of ACH. We structurally estimate the parameters of the model using an indirect inference procedure and panel data. The parameters are identified from exogenous variation in the adoption decisions of banks based outside the network and other factors. We find that most of the impediment to ACH adoption is from large customer fixed costs of adoption. Policies to provide moderate subsidies to customers and larger subsidies to banks for ACH adoption could increase welfare significantly.

* Department of Economics, University of California Los Angeles, Box 951477, Los Angeles, CA 90095 and National Bureau of Economic Research, ackerber@econ.ucla.edu.

** John M. Olin School of Business, Washington University in St. Louis, 1 Brookings Drive, Campus Box 1133, St. Louis, MO 63130 and National Bureau of Economic Research, gowrisankaran@wustl.edu.

Acknowledgements:

We acknowledge funding from the NET Institute and the National Science Foundation (Grant SES-0318170), and thank the Editor, two anonymous referees, Steve Berry, Liran Einav, Jinyong Hahn, Fumiko Hayashi, Brian McManus, Andrea Moro, Joanna Stavins, Klaas van't Veld, Bill Vogt and seminar participants at numerous institutions for helpful comments and Anita Todd for editorial assistance.

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 and adoption. ACH is an electronic payment mechanism developed by the Federal Reserve and used by banks and customers. It is essentially an electronic alternative to paper checks, typically used for recurring transactions such as direct deposit paychecks and automated utility bill payments. Since banks on both sides of a transaction must adopt ACH for an ACH transaction to occur, ACH is a two-sided market.

As a two-sided market, ACH is characterized by network effects. A bank will be more likely to adopt ACH as other banks adopt because it will be able to originate more transactions with ACH, thus justifying the fixed costs of adoption. The importance of the network effect depends on the incremental profits to the bank from each ACH transaction relative to the fixed cost of adoption. As banks likely cannot compensate each other for ACH adoption, these network effects are network externalities and typically cause underutilization of the network good. The underutilization is particularly relevant for ACH – in an age when computers and technology have become prevalent, many more payments are performed with checks than with ACH.

Network effects may also exist at the level of the banks' ACH customers, which include employers, utility companies, and small businesses. These customers must also bear a fixed cost of adoption (e.g., updating their disbursement systems), and similarly may be more likely to adopt ACH if more of the customers with whom they transact also adopt. In contrast to banks, while it is clear that customers must adopt ACH to *originate* transactions, the extent to which a customer *receiving* an ACH transaction must actively adopt ACH is unknown. Starting direct deposit payroll, for example, takes very little effort on the part of the employee, the recipient in this case. This extent to which receiving customers must adopt, along with the magnitudes of customer and bank fixed costs of adoption and the customer and bank incremental benefits from ACH transactions, will all impact equilibrium adoption and may play a role in the low observed rates of adoption.

To understand the causes and extent of the network externalities, we specify a simple static two-sided market model of ACH technology adoption in local markets. Each market contains a set of banks, each with a given set of customers. Each customer must make a fixed number of transactions to other customers using either checks or ACH. While all banks and customers accept checks, some may not have

adopted ACH. Some banks are locally based while others are branches of big banks based outside the network. Local banks decide whether to adopt ACH based on whether the variable profits from ACH transactions conditional on adoption are greater than the fixed costs of adoption; the decisions of nonlocal banks are made exogenously and are known to the local banks. Following bank adoption, customers at banks that have adopted ACH choose whether or not to adopt ACH. We model two types of transactions: one-way transactions, which are passively accepted by the receiving customers, and two-way transactions, which require the receiving customers to adopt ACH actively.

The implications of the model depend on parameters that specify bank and customer costs and benefits, the proportion of two-way transactions, and other details including scales and time trends. We structurally estimate these parameters by applying an indirect inference estimator (a variant of the method of simulated moments) to bank-level panel data on ACH adoption and the number of ACH transactions. The idea of the estimator is to simulate data from the model (which requires solving for the equilibrium of the model conditional on structural parameters and unobservables) to find the parameters for which the simulated data most closely match the observed data.

This work builds on a recent literature on empirically estimating the extent of network effects for different industries (see Goolsbee and Klenow, 2002; Gowrisankaran and Stavins, 2004; Ohashi, 2003; Park, 2002; and Rysman, 2004a, 2004b, among others). Network externalities imply an interdependence of preferences, leading to simultaneity in equilibrium adoption decisions. This makes identification of the network externalities potentially difficult. These papers have all tried to find evidence of network effects by estimating correlations among usage decisions or by estimating reaction functions with techniques such as instrumental variables. Our work differs from this literature in that we fully specify an equilibrium model of interactions among banks and their customers, and estimate the structural parameters of our model by computing and matching equilibrium predictions of the model to data.¹

Our use of a fully specified structural model has a number of advantages. Relative to examining correlation among usage decisions (e.g., Gowrisankaran and Stavins, 2004), our structural model also allows us to estimate the magnitudes of the network effects, rather than just being able to test for their presence. In addition, we are able to allow for time-varying local shocks that otherwise might be incorrectly interpreted as network effects. Also, the methods that we develop here are novel and contribute to the literature on structural estimation of network games with potential multiple equilibria.

In comparison to papers that structurally estimate reaction functions using linear or log-linear specifications (e.g., Rysman, 2004a), our fully specified equilibrium model has a number of advantages. First, it allows us to specify a strategic model of interactions with simultaneous decisions that is realistic given the discrete nature of adoption decisions in our industry. Second, our model allows us to understand the sources of the network effects. This is crucial for ACH policy analysis—for example, subsidies will have very different impacts on ACH usage depending on whether fixed costs are primarily at the bank level or at the customer level. Last, our structural model also allows us to perform a wider set of policy experiments, since our model predicts the outcomes that will result from any policy intervention. This is particularly important to the extent that there are multiple equilibria: in this case reaction functions are consistent with more than one outcome, implying that estimated reaction functions are not sufficient to simulate counterfactual equilibria.

Since it uses the same data as the present study, the work of Gowrisankaran and Stavins (2004) deserves particular attention. As mentioned above, their study takes a reduced-form approach to examining network externalities, modeling ACH adoption as a simple function of the percent of other banks that had adopted ACH in a local area. In contrast to this paper, the central point of the earlier work was to confirm the *presence* of network externalities statistically, not to measure the *magnitude* and sources (i.e., banks or customers) of the externalities.² Most relevant for this paper is that Gowrisankaran and Stavins (2004) proposed a number of different tests that would identify network effects separately from confounding factors. While our methods and questions are different, this study builds on theirs in that our identification of the network effects is from some of the same sources. In addition, the earlier paper guides us in our choice of indirect inference moments—our indirect inference estimator matches the coefficients from regressions that are similar to theirs evaluated on simulated equilibrium data. As an example, Gowrisankaran and Stavins (2004) propose a regression to exploit the quasi-experimental variation from the adoption decisions of small, remote branches of banks. Our analysis models this variation with more structure, allowing for exogenous nonlocal bank decisions and endogenous adoption decisions for their customers. We then choose structural parameters that match (among other things) the coefficients from regressions on real data that are similar to regressions from Gowrisankaran and Stavins (2004) to coefficients from the same regressions on simulated equilibrium data.

The remainder of this paper is divided as follows. Section 2 describes the model. Section 3 describes the data. Section 4 details the estimation procedure, including the identification of the parameters. Section 5 provides results including policy experiments, and Section 6 concludes.

2. Model

We propose a simple, static, two-sided market model of ACH adoption at a geographically local level. Market m consists of J banks at time t , each with a given measure of customers. Some of the J banks are local to the market, and others are branches of large banks based elsewhere in the country. Adoption decisions of these nonlocal bank branches are assumed to be made exogenously before the start of the game and known to all players.³

In the first stage of the game, local banks simultaneously decide whether or not to adopt ACH technology. Let $A_{mt} = (A_{1mt}, \dots, A_{Jmt})$ be a vector of indicator functions representing bank adoption decisions. In the second stage, customers, typically small businesses, decide whether to adopt ACH for their individual transactions. Importantly, a customer can adopt ACH only if its bank has adopted ACH. Some customers will have accounts at branches of nonlocal banks. Although their banks' adoption decisions are exogenous to our model, these customers decide whether or not to adopt ACH in the same way as do customers at local banks.

Banks in our model each have a captive set of customers. We assume that all of these customers are equal in the sense that they each have the same amount of deposits and each originate the same total number of transactions (N).⁴ As a result, the measure of customers at a given bank is proportional to x_{jmt} , the total deposits under that bank's control. In addition, the number of total (both ACH and check) transactions originated by customers of bank j will also be proportional to x_{jmt} , i.e.,

$$T_{jmt} = \lambda x_{jmt} \tag{1}$$

where λ represents the number of transactions per dollar of assets.⁵ Note that this formulation implicitly assumes that customers' demands for total transactions are perfectly inelastic. We feel that this is

reasonable because the demand for transactions is, in fact, likely to be fairly inelastic and ACH is a small proportion of total transactions.

Each of customer i 's N originating transactions can be made either using ACH or with a paper check. Whether or not this transaction can be made using ACH depends on the adoption decisions of four economic agents: (1) customer i , (2) customer i 's bank, (3) the customer at the receiving end of i 's transaction, and (4) the bank of this receiving customer. For any transaction to be performed with ACH, *both* customers' *banks* must have adopted ACH. Otherwise, it will be completed by paper check. In contrast, we model two types of transactions that vary based on whether the receiving customer must also have adopted ACH actively for it to be performed with ACH. A "two-way" transaction requires *both* originating and receiving customers to have adopted ACH. A "one-way" transaction requires only the *originating* customer of the transaction to have adopted ACH. The latter represents transactions such as payroll direct deposit or mortgage payments, where the recipients have likely not formally adopted ACH in the sense that they do not originate transactions.⁶ We represent the proportion of two-way transactions by δ , a parameter to be estimated in the model.⁷

Note that we model banks and customers asymmetrically in that the customer recipient does not necessarily need to adopt. This is due to the specifics of the industry. In particular, the design of the ACH system is such that banks on both sides of the transaction have to have adopted for an ACH transaction to occur. Although we lack formal data on the necessity of customer adoption, we discussed the nature of ACH transactions with a number of industry professionals. Most believe that recurring payments with passive customers, which we might interpret as one-way transactions, form the bulk of ACH transactions.

Consider now the adoption decision of customer i at some bank j that has adopted ACH. The customer obtains net utility

$$V_{ACH} - V_{CHK} = \tilde{\beta}_1 + \tilde{\beta}_2 (p_i^{ACH} - p_i^{CHK}) \quad (2)$$

from making an ACH (versus check) transaction, where p_i^{ACH} and p_i^{CHK} represent the prices for each type of transaction that customers face.⁸ If the net utility in (2) is positive, then, conditional on both their banks adopting, any pair of customers who have both adopted ACH will use ACH for their two-way transactions, and any customer who has adopted ACH will use ACH to originate one-way transactions with customers at banks that have adopted ACH.

Let P_{jmt} denote the proportion of customers adopting ACH at bank j . We make two additional assumptions about the nature of customers and transactions: (1) that customers are small enough to be considered atomless for aggregation purposes, and (2) that the recipient of any transaction is uniformly distributed over all customers in the market. Then the expected equilibrium proportion of a customer's transactions that will be made through ACH if he adopts is

$$u_{mt} = \delta \left[\frac{\sum_j A_{jmt} P_{jmt} x_{jmt}}{\sum_j x_{jmt}} \right] + (1 - \delta) \left[\frac{\sum_j A_{jmt} x_{jmt}}{\sum_j x_{jmt}} \right]. \quad (3)$$

The first term in brackets is the proportion of customers in the market that have adopted, while the second term is the proportion of customers that are at banks which have adopted ACH. Since transactions are spread evenly over the population, these proportions are also the probability that any particular transaction is with another customer who has adopted (the first term), or with another customer whose bank has adopted (the second term).

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

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

where F_{ijmt} denotes the negative of the fixed costs of adopting. These fixed costs may not only be monetary; F_{ijmt} will also capture any fixed utility or disutility that a customer has from using ACH. For instance, a customer with a strong disutility from using electronic payment technologies would have a very negative value of F_{ijmt} . Note also that we model customers as obtaining utility solely from the N transactions that they originate. An alternate specification would allow utility to be generated on both sides of the transaction.⁹

Econometrically, we specify negative customer fixed costs as

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

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 that is common knowledge to industry players, and ε_{ijmt} is an i.i.d. customer-level logit error. In order to identify the network benefits of ACH (vs. within market correlation in behavior) reliably, we allow for a very general correlation structure across markets, firms, and time.¹⁰ Specifically, we let

$$\alpha_{jmt} = \alpha_{jmt}^A + \alpha_{jm}^B + \alpha_m^C + \alpha_{mt}^D, \quad (6)$$

where the four terms in (6) are all i.i.d. and normally distributed with mean zero and have variances σ_A^2 , σ_B^2 , σ_C^2 , and σ_D^2 , respectively. There are a number of other places where one might include a similarly flexible unobservable structure, including customers' marginal benefits, banks' marginal profits, and banks' fixed costs. However, we felt that it would be difficult to identify more than one set of flexible unobservables credibly. We put the unobservables in customer fixed costs because this is the specification that appears to fit the data best.

Substituting (5) and (2) into (4), we obtain

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

where β_1 and β_2 are defined by $\beta_1 = N \cdot (\tilde{\beta}_1 - \tilde{\beta}_2 p_t^{CHK})$ and $\beta_2 = N \cdot \tilde{\beta}_2$.¹¹ Integrating out over ε_{ijmt} gives us the proportion of customers at bank j that adopt ACH:

$$P_{jmt} = A_{jmt} \frac{\exp(\beta_0 + (\beta_1 + \beta_2 p_t^{ACH}) u_{mt} (P_{1mt}, \dots, P_{Jmt}, A_{1mt}, \dots, A_{Jmt}) + \beta_3 t + \alpha_{jmt})}{1 + \exp(\beta_0 + (\beta_1 + \beta_2 p_t^{ACH}) u_{mt} (P_{1mt}, \dots, P_{Jmt}, A_{1mt}, \dots, A_{Jmt}) + \beta_3 t + \alpha_{jmt})}. \tag{8}$$

Note that we have explicitly written u_{mt} as depending on customer adoption proportions and bank adoption decisions. In equilibrium, customer adoption proportions $(P_{1mt}, \dots, P_{Jmt})$ must satisfy the set of J equations defined by (8) conditional on bank adoption decisions $(A_{1mt}, \dots, A_{Jmt})$.

Turning to the adoption decisions for local banks, let the marginal cost to the bank of ACH and check transactions be denoted as mc_t^{ACH} and mc_t^{CHK} respectively. Assume that bank fixed costs FC_{jmt} from adopting ACH have both a common component and an idiosyncratic component, i.e. $FC_{jmt} = \overline{FC} + \alpha_{jmt}^E$, where the realizations of α_{jmt}^E are common knowledge in the industry. As with the customer fixed cost ε_{ijmt} , we assume α_{jmt}^E has a standard logistic distribution. Importantly, FC_{jmt} is a *per-period* fixed cost, not a one-time sunk cost of adoption. As such, there are no dynamic optimization issues, and firms maximize per-period profits.¹²

Given that bank j adopts ACH, the number of ACH transactions originated through bank j is equal to the total number of transactions originated at that bank, times the proportion of customer adopters, times the proportion of their transactions that are made through ACH; i.e.,

$$T_{jmt}^{ACH} = T_{jmt} P_{jmt} u_{mt}. \tag{9}$$

The increment in profits from adopting ACH (conditional on other players' strategies) is T_{jmt}^{ACH} times the difference in bank margins between ACH transactions and paper checks, minus the fixed cost of adoption:

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

Given our lack of data on prices, we treat “markup” as a parameter to be estimated.

Bank j will adopt ACH at time t if and only if $\Pi_{jmt}(T_{jmt}^{ACH}) > 0$. Equations (8) - (10) illustrate that bank j 's adoption will depend on other banks' decisions through T_{jmt}^{ACH} . An equilibrium $(P_{1mt}, \dots, P_{Jmt}, A_{1mt}, \dots, A_{Jmt})$ requires that all banks' adoption decisions are optimal conditional on all other banks' adoption decisions, i.e.,

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

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

There can be multiple equilibria in this game if the magnitude of the network externalities are sufficiently large. 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 not to 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 itself is adopting ACH, the adoption game is supermodular. Several properties follow from supermodularity (see Milgrom and Shannon, 1994, and Athey, Milgrom, and Roberts, 1996). These properties can easily be proved directly (see Gowrisankaran and Stavins, 2004), 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 exists 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 Pareto-worst subgame perfect equilibria. This last property is particularly important for estimation purposes.

To estimate our model and also to examine counterfactuals, we need to specify a process that determines which equilibrium occurs. The process that we assume allows for the possibility that some

networks are in a good equilibrium while others are in a bad equilibrium. Specifically, 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 this frequency as a parameter. Formally, let $\omega_m \sim iid U(0,1)$. We assume that a market will be in the Pareto-best equilibrium if and only if $\omega_m < \Omega$, where Ω , the probability of being in the Pareto-best equilibrium, is a parameter that we estimate.¹³ Note that we do not allow for the equilibrium to vary within a network across time. Clearly, the restriction to only the two extreme equilibria is a simplification, but our hope is that a mix of our two equilibria will reasonably approximate any alternative equilibria that are being played.

3. Data

Our principal data set is the Federal Reserve's ACH billing data. This provides information on individual financial institutions that processed their ACH payments through Federal Reserve Banks.¹⁴ We use quarterly data on the number of transaction originations by each 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.

We link these data with three other publicly available data sources. First, we use the Call Reports database, which provides quarterly information on deposits and zip code information for federally registered banks at the headquarters level. Second, we use the Summary of Deposits database, which provides annual information on zip code and deposits for banks at the branch level. This allows us to find small branches of nonlocal banks. Last, we use information from the census that provides the latitude and longitude of zip code centroids. Gowrisankaran and Stavins (2004) provide more details of the data sources and linking. The resulting data set contains approximately 11,000 banks over 11 quarters.

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 equilibria of the adoption game, we also

need to include all the banks that are within 30 kilometers of the banks that are within 30 kilometers, and all the banks that are near those banks, etc. We perform 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 limit our sample to banks that are in small markets. Specifically, we keep all networks with 10 or fewer bank headquarters total during every time period of our sample. From this set, we also excluded networks where any one bank had more than 20 percent of its deposits in branches *outside* the network or where, in aggregate, 10 percent of deposits for local banks were in branches *outside* the network. This leaves us with a sample of 456 mutually exclusive networks comprising 878 local banks, observed over 11 time periods. Note that even though we only study this rather small subsample of the banking industry, we are estimating *structural* utility and profit functions. As such, at least in theory we can use our estimates to study larger markets by simply solving for equilibria in larger markets given these primitives. However, one might worry that customers and banks in larger markets are fundamentally different, e.g., they might face different costs of adoption or different relative markups for ACH transactions.

Figure 1 displays a map of New England with the networks from this region marked with asterisks and the major population centers marked with circles. This gives some idea about typical networks in our final data set, which draws from the U.S. as a whole. There are eight networks in New England in our sample, all of which are small, isolated towns, such as Colebrook, NH, and Nantucket, MA. These small, isolated towns are exactly the types of markets our model is intended to represent. The major population centers, such as Boston, MA, and Hartford, CT, are all far away from these networks.

As described in Section 2, we model banks with branches in the network but headquarters outside the network. Recall, however, that these nonlocal banks' adoption decisions are assumed to be exogenous to the local market. Our sample includes 661 of these branches of banks outside the network.

Table 1 gives some specifics on local banks, broken down by the number of local banks in the network. Approximately half of the network time-periods in our sample—2,730 in all—have only one

local bank. Another quarter of the network time-periods have two local banks. However, there are a large number 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 local banks using ACH appears to be quite consistent across network size, although local banks in smaller networks have fewer ACH transactions.

Table 2 examines the nonlocal banks in these networks. Of note is the large number of outside banks, and the variance in this number across markets. For instance, in markets with one local bank, the average number of outside banks is 2.02 and the standard deviation across markets is 2.69. Although the sizes of nonlocal banks and local banks are similar in terms of *local* deposits, nonlocal banks adopt ACH much more frequently. This is because nonlocal banks are on average much larger than the local banks in terms of total deposits. Relatedly, note that the proportion of deposits of nonlocal banks that are in local markets is small, about 1 percent, as can be seen by dividing column 4 by column 5. This motivates treating the adoption decisions of these nonlocal banks as exogenous to the local markets.

Table 3 gives some specifics on the changes in local bank ACH usage over our sample period. The fraction of banks using ACH increased during our sample period. There also appears to be a large fraction of networks where *every* bank uses ACH relative to the marginal probability of banks adopting (from Table 1). While this suggests that there are correlations in usage decisions within networks, it does not necessarily imply network externalities, as the correlations could simply be driven by correlated unobservables within markets.

One factor that can affect adoption and usage of ACH is its price. Prices that the Federal Reserve charges banks for ACH processing were set at a fixed rate and adjusted periodically. The interregional per-item prices (for items exchanged between banks in different Federal Reserve Districts) fell from \$.014 in 1995 to \$.012 in 1996 to \$.01 in 1997, with the intraregional price fixed at \$.01 during this time period. In May 1997, both prices were lowered to \$.009, with a discounted price of \$.007 for banks that processed over 2,500 transactions at once. We use the prices of \$.014 (for 1995), \$.012 (for 1996), \$.01 (for 1997:Q1 and 1997:Q2) and \$.009 (for 1997:Q3 and 1997:Q4) for all transactions. Because these prices are set by fiat and do not respond to changes in local demand, they may be viewed as exogenous. We do not have systematic data on the prices that banks charge to their customers.

In addition to per-transaction costs, banks must pay a variety of fixed costs to adopt ACH, e.g., a small participation fee and the costs of maintaining a Fedline connection. As detailed in Gowrisankaran

and Stavins (2004), accounting data suggest that these are moderate in magnitude, likely significantly lower than \$500 per quarter.

4. Estimation

Our model is based on a vector of unknown parameters $\theta \equiv (\lambda, \beta, \Omega, \overline{FC}, markup, \sigma)$ and econometric unobservables $v = (\alpha, \omega)$.¹⁵ Our estimation algorithm seeks to recover the true parameters θ_0 from the data. In this section, we describe this algorithm and explain how the parameters of the model are identified.

Estimation algorithm

An observation in our model is a market observed over time. Our data set contains M distinct markets. We start by defining the data for market m over time. For ease of notation, suppose that market m has J banks at every time period, that banks $1, \dots, \hat{j}$ are local, and that banks $\hat{j} + 1, \dots, J$ are branches of nonlocal banks. For each bank, our data contain observed predetermined variables, namely its local deposits x_{jmt} , the Fed ACH price p_t , time t , and its local/nonlocal status. Recall that the adoption decisions of the branches of nonlocal banks are also considered exogenous data. Thus, the exogenous data for the market are $z_m \equiv \{x_{1t}, \dots, x_{Jt}, t, p_t, A_{\hat{j}+1,t}, \dots, A_{J,t}\}_{t=1}^T$. The endogenous variables, which we denote y_m , are the local bank transactions $\{T_{1mt}^{ACH}, \dots, T_{\hat{j}mt}^{ACH}\}_{t=1}^T$; note that $T_{jmt}^{ACH} > 0$ implies that $A_{jmt} = 1$.

We estimate our model using an indirect inference (II) estimator, as suggested by Gouriéroux, Monfort, and Renault (1993). II is a variant of the method of simulated moments (MSM) estimator proposed by McFadden (1989) and Pakes and Pollard (1989).

To construct the estimator, we first perform an easily computable and informative estimation procedure on the true data (that we detail in the subsection on choice of indirect inference moments below), resulting in a parameter vector $\mu(y, z)$. The estimator is indirect in the sense that the estimated parameters in the inference do not correspond to any structural parameters; an example similar to one we use is a regression of bank adoption decisions on their competitors' adoption decisions. Although the

regression coefficients in this example will capture causation (i.e., network effects), they also will likely capture within-market correlation in unobservables. As such, they are clearly not unbiased estimates of the network effect parameters.

We next simulate data according to our model at various parameter vectors. This involves taking S simulation draws $v_{sm}, s = 1, \dots, S$ from the distribution of unobservables for each market m . We compute the appropriate Nash equilibrium (either Pareto-best or -worst, depending on ω) for each market given v_{sm} and θ . Call this Nash equilibrium vector $\hat{y}_{sm}(v_{sm}, \theta, z_m)$, and group the equilibria across markets for one simulation draw together as $\hat{y}_s(v_s, \theta, z) \equiv (\hat{y}_{s1}, \dots, \hat{y}_{sM})$. We then perform the same indirect inference on these computed Nash equilibrium data. The II estimator is constructed from the difference between μ evaluated at the true data and the simulated data:

$$\hat{\theta}_{II} \equiv \arg \min_{\theta} G^{II}(\theta) \equiv \arg \min_{\theta} \left\| \mu(y, z) - \frac{1}{S} \sum_{s=1}^S \mu(\hat{y}_s(v_s, \theta, z), z) \right\|. \quad (12)$$

Thus, $\hat{\theta}_{II}$ is the structural parameter vector for which the coefficients of μ estimated from the simulated data most closely match the coefficients estimated from the actual data, where closeness is defined by the norm $\|\cdot\|$. We find $\hat{\theta}_{II}$ with a nonlinear search.¹⁶

A special case of $G^{II}(\theta)$, which can be called $G^{MSM}(\theta)$, occurs when the functional form of $\mu(y, z)$ is a sum across observations. The resulting estimator, defined by

$$\hat{\theta}_{MSM} \equiv \arg \min_{\theta} G^{MSM}(\theta) \equiv \arg \min_{\theta} \left\| \frac{1}{M} \sum_{m=1}^M \left[\left(f(y_m) - \frac{1}{S} \sum_{s=1}^S f(\hat{y}_{sm}(v_{sm}, \theta, z_{sm})) \right) \otimes g(z_m) \right] \right\|, \quad (13)$$

for some arbitrary functions f and g , is a standard MSM estimator. This estimator is well-known to be consistent as the number of observations goes to infinity, even with a fixed number of simulation draws (see McFadden, 1989, and Pakes and Pollard, 1989). Similarly, Gouriéroux, Monfort and Renault (1993) show that the more general II estimator is also consistent for a fixed number of simulation draws. This is because, at the true parameter vector, the simulated criterion function for each simulation draw

$\mu(\hat{y}_s(v_s, \theta_0, z), z)$ converges in probability to $\mu(y, z)$ as the number of observations becomes large. We use II instead of MSM because the functional form is convenient for summarizing important aspects of the data with relatively few moments, a point that we detail below.¹⁷

In order to compute the II estimator, we need to specify the norm $\|\cdot\|$. A norm is defined by a weight matrix—any positive-definite matrix with dimension equal to the number of moment conditions. Gouriéroux, Monfort and Renault (1993) show that the efficient weight matrix is the inverse of the variance of the moment conditions in (12), evaluated at θ_0 . An approximation to this efficient weight matrix is particularly easy to compute in our case. According to the model, each $\mu(\hat{y}_s(v_s, \theta_0, z), z)$ has the exact same distribution as $\mu(y, z)$ and the two terms are i.i.d. given the construction of the simulation draws. Combining the above, we get

$$Var(G''(\theta_0)) = Var\left(\mu(y, z) - \frac{1}{S} \sum_{s=1}^S \mu(\hat{y}_s(v_s, \theta_0, z), z)\right) = \left(1 + \frac{1}{S}\right) Var(\mu(y, z)), \quad (14)$$

and thus we use an estimate of $A = \left[\left(1 + \frac{1}{S}\right) Var(\mu(y, z))\right]^{-1}$ as our weight matrix. Since $\mu(y, z)$ is a function of the observed data but not the structural parameters, we can estimate A using a bootstrap. Specifically, we resample the data with replacement at the market level (i.e., we draw entire networks across time), and calculate $\mu(y, z)$ for 3,000 resampled data sets in order to approximate its variance numerically.¹⁸

Equation (14) also plays a role in computing asymptotic standard errors of our estimated structural parameters. Following Gouriéroux, Monfort and Renault (1993), the asymptotic variance matrix is given by

$$Var(\theta) = (\Gamma' A \Gamma)^{-1}, \quad (15)$$

where $\Gamma = \frac{dG''(\theta_0)}{d\theta'}$. We calculate these derivatives numerically.

One important detail is the computation of $\hat{y}_{sm}(v_{sm}, \theta, z_m)$. This involves solving for either the Pareto-best or Pareto-worst subgame perfect equilibrium of the model conditional on values for z_m and v_{sm} (recall that the ω component of v determines which equilibrium a market is in). We solve for the equilibria by sequentially iterating on reaction functions of individual customers and banks. Specifically, given a vector of bank adoption and customer usage decisions $(A_{1mt}, \dots, A_{Jmt}, P_{1mt}, \dots, P_{Jmt})$, we solve for the new optimal customer usage decisions and then the new optimal bank adoption decisions. To compute the usage decision for a customer at bank j , we assume that bank j has adopted, consistent with our assumption that customers of a bank observe this decision.

The difference between the computation of the two equilibria is only in the starting values. To solve for the Pareto-best subgame perfect equilibrium, we start the iterative procedure at a point where *all* customers and local banks adopt ACH (with nonlocal banks following their observed choices from the data). In contrast, to find the Pareto-worst subgame perfect equilibrium we start at a point where no customers and local banks use ACH. As noted in Section 2, the game is supermodular, which ensures that these iterative processes will converge to the Pareto-best and Pareto-worst subgame perfect equilibrium, respectively.¹⁹ In our policy experiments, we use variants of this algorithm to solve for equilibrium outcomes when either local banks internalize the network externality or when customers internalize the externality.

Choice of indirect inference moments

Our goal is to choose moments (i.e., a set of estimation procedures $\mu(y, z)$) that are informative on the existence and magnitude of network externalities. In particular, we want these moments to be able to separate network externalities from confounding factors such as correlated unobservables. One option would be to use moments of the dependent variables. These would include adoption and usage decisions, second- and higher-order cross moments in these variables (e.g., defining Q to be the number of transactions at bank 7 times the adoption of bank 10, then a moment would be Q minus its expected value), similar second- and higher-order moments across time (e.g., adoption of bank 1 in time period 1 times adoption of bank 1 in time 5) and interactions of these variables with exogenous variables. This approach is problematic because the sheer number of these moments is larger than the number of markets.

Moreover, the literature suggests the possibility of non-negligible finite-sample biases even if we only used a small fraction of them (see e.g., Staiger and Stock, 1997).

To try to summarize the nature of network externalities in the data with an informative but concise set of moments, we make use of Gowrisankaran and Stavins (2004), whose reduced-form analyses attempted to accomplish exactly this purpose. In particular, that paper provides three sets of reduced-form analyses that are useful for characterizing network externalities in this market and for separating them from potentially confounding effects such as local correlations in usage. We therefore choose our II estimator moments $\mu(y, z)$ to be essentially the same regressions as in the earlier paper, although there are a couple of differences. First, we use linearized versions of their non-linear (e.g., logit and probit) specifications. The reason for this is that for our II procedure, we need to perform the $\mu(y, z)$ estimation procedure a very large number times—once for each simulated data set at every structural parameter vector ever evaluated in the non-linear search for θ . This would be computationally impractical given a non-analytical $\mu(y, z)$ procedure. Note that even though these linearized versions may not be correct specifications, our II estimator is still consistent, as we are performing the *same* incorrect procedure on both the simulated data and the actual data. Second, we also cannot use every coefficient from their regressions as moments, since they include thousands of coefficients (mainly because of bank fixed effects). Hence, we use as moments the coefficients that the earlier paper found to be indicative of network externalities, and other coefficients/functions of coefficients that we think would be particularly informative about our structural parameters.

We now go through the exact specifics of the components of μ . Our first set of regressions is intended to capture clustering of bank adoption decisions within a network, controlling for regional variation in usage decisions with bank fixed effects. We regress the adoption decisions of banks on the deposit-weighted adoption decisions of competitors' banks, deposits, squared deposits, and bank and time fixed effects. We include the 3 main regression coefficients and 11 time dummies (including the constant term) as components of μ . We also include the estimated standard error of the coefficient on competitor adoption. This should measure the variance of the unobservable in the regression and thus should help identify the variance of the bank-time random effects (α_{jmt}^A) in our structural model. We also want to try to match the extent to which adoption is explained by unobserved factors that are correlated across banks within a market, either within time periods or across them. These will serve to identify the variances of

the other random effects in the model $(\alpha_{jm}^B, \alpha_m^C, \alpha_{mt}^D)$. We do this by decomposing the estimated bank fixed effects into market fixed effects and residuals and including the standard deviation of both in μ , and by also decomposing the estimated residuals into market-time fixed effects and residuals, and including the standard deviation of both in μ . Note that we do not include the R^2 for this regression in μ —this would likely be redundant given that we include all of these standard deviations.

To capture correlation in customers' adoption decisions (in contrast to banks' adoption decisions), we then run similar regressions using ACH *volume* as the adoption variable; i.e., we regress ACH volume per deposits on competitors' volume per deposits and deposit-weighted adoption decisions, using the same deposit and fixed effect controls. We take three coefficients from this regression as moment—one of the time dummies (representing the constant term) and the coefficients on competitors' volume and adoption decisions. Note that for this regression and the ones that follow we do not decompose the error variance as in the first regression—presumably the variances of the random effects should be pinned down by the first regression.

Our second set of regressions examines the relationship between market concentration and adoption by regressing bank adoption decisions on the market level Hirschmann-Herfindahl Index (HHI), time dummies, and controls for deposits. This regression should help identify network externalities because exogenous differences in market concentration will affect adoption decisions. For example, network externalities will be more internalized in more concentrated markets, implying higher ACH adoption and usage, all else equal. We use two coefficients—one time dummy (i.e., the constant term) and the HHI coefficient—as moments. We also include the R^2 as a moment, in order to match the overall variance of the residuals.

Our third set of regressions examines the relationship between local bank and nonlocal bank adoption decisions. The intuition here is that local bank response to exogenous, nonlocal bank adoption decisions should be directly proportional to the level of network effects. Specifically, we regress local bank adoption decisions on the adoption decisions of nonlocal banks using a variable that indicates the deposit-weighted fraction of nonlocal banks that adopt. Again, we control for deposits and include time dummies. For II moments we use the coefficient on the nonlocal adoption measure, one of the time dummies (i.e., a constant term), and the R^2 from the regression. Markets with no nonlocal banks were dropped from the base version of this regression, since the nonlocal adoption measure is undefined in this

case. However, in an effort to be robust to including these markets, we also utilize an alternative regression where *all* markets are included but where we define the nonlocal adoption measure to be equal to one in the case when there are no nonlocal banks. We take the same three moments for this new regression.²⁰ These three groups of regression coefficients give us a grand total of 31 II moments, all listed in Table 5, as described in Section 5 below.

As with any method of moments estimator, one can never be certain that another choice of moments would not yield substantially different parameter estimates. While our moments are designed to match some of the facets of the data that relate to the magnitude of the network benefits and other parameters, it is possible that there are other unmatched facets of the data that would yield substantially different estimated parameters. To investigate this concern, we examine the robustness of our parameter estimates to an alternate choice of moments. The idea here is to add additional moments that are substantially different from those already in the procedure, but that are hopefully also informative on the extent of network effects. One can then investigate whether these new moments suggest economically different parameter values from our original moments.

Thus, we estimate a second specification with the same 31 base moments plus an additional 11 moments, all simple moments of the data. These are substantially different from our base moments, should be informative in pinning down parameters of the model, and also, because of their simplicity, provide an intuitive sense of the fit of the model. We include the mean adoption probability at local banks (the mean implies the standard deviation with a dichotomous adoption decision), the mean and standard deviation of the number of transactions at local banks, the mean number of transactions in the first and last periods of the sample, the correlation between deposits and adoption probability at local banks, the correlation between deposits and number of transactions at local banks, the correlation between local banks' adoption decisions in the first and last periods of the sample, the correlation between local banks' number of transactions in the first and last periods, and the average local bank within-bank (across-time) standard deviations of ACH adoption and number of transactions.

Identification

In our structural model network effects are captured by five parameters: customer and bank fixed costs of adoption, customer and bank per-transaction (marginal) benefits from adoption, and the proportion of two-way transactions. Understanding the separate identification of each of these parameters

is somewhat difficult because our model is more structural than a typical model of network externalities. For example, unlike prior models, we have to identify network effects *both* at the bank and at the customer level. To start, note that if we knew the relative markup received by a bank from ACH transactions, then the observed proportion of banks adopting ACH would identify the bank fixed cost of adoption. Similarly, if we knew the utility customers obtain from an individual ACH transaction and the proportion of two-way transactions, then the overall level of customer adoption as implied by the number of ACH transactions would identify the customer fixed cost of adoption.²¹

Our data should identify the remaining three parameters using variation in local bank adoption and usage levels across different exogenous industry configurations. We have such exogenous variation in the market structure of local banks (e.g., the sizes of local competitors and the local HHI) and in the adoption decisions and market structure of nonlocal banks. One nice aspect of our structural model is that it allows us to combine these multiple sources of variation into one estimation procedure.

To illustrate the separate identification of these three parameters more formally, we now focus on one of these sources of identification, although our discussion would apply to other sources as well. Specifically, consider exogenous variation in the number of nonlocal banks adopting. The extent to which changes in exogenous nonlocal bank adoption affect local bank adoption should identify the bank relative markup. As the relative markup increases, local banks will be more sensitive to nonlocal bank adoption decisions.

Turning to identification of the customer marginal benefit parameter (β_1) and the two-way transactions parameter (δ), we simplify matters for this discussion by assuming that the market has one small, adopting, local bank (small enough that we can ignore transactions between two customers who both use the local bank) and a number of identical nonlocal banks. Let τ_a be the proportion of outside banks that adopt and $P(\tau_a)$ be the proportion of customers that adopt at each of the adopting nonlocal banks. We can then use (3) to write the gross benefit that a local bank customer obtains from adopting as a function of the exogenous shifter τ_a :

$$f(\tau_a) \equiv \beta_1 u_{mr}(\tau_a) = \beta_1 \delta \tau_a (P(\tau_a) - 1) + \beta_1 \tau_a. \quad (16)$$

While we do not directly observe the gross benefit $f(\tau_a)$, we do observe the number of ACH transactions at the local bank, which depends directly on it. From (16), an increase in τ_a has different implications for $f(\tau_a)$ depending on β_1 . This is easy to see for the simple case of all one-way transactions, $\delta = 0$. In this case, $f(\tau_a) = \beta_1 \tau_a$ implying a higher slope (with respect to τ_a) when β_1 is higher. The *shape* of $f(\tau_a)$ will depend on δ . For example, with all one-way transactions, $f(\tau_a)$ is linear in τ_a . When $\delta > 0$, $f(\tau_a)$ increases nonlinearly in τ_a because it now additionally depends on nonlocal customer adoption $P(\tau_a)$, which increases with τ_a . This suggests that β_1 and δ will be separately identified by the shape (both first and second derivative) of local bank ACH transactions in response to exogenous nonlocal bank adoption.

Although this identification is based on nonlinearities, we believe that these nonlinearities are economically meaningful. Nonlocal banks switching from non-adoption to adoption generate additional potential transaction partners in two ways: (1) through customers at the newly adopted bank (at a linear rate), and (2) through stimulating more adoption by customers at already adopted nonlocal banks (at a higher-than-linear rate). This second effect is only relevant with two-way transactions because with one-way transactions it does not matter to the local bank customers whether customers at other already adopted banks adopt more, since they can be transacted with regardless. This is the underlying intuition for why a more convex $f(\tau_a)$ implies more two-way transactions.

A couple of other notes about identification are in order. First, an important aspect of our model is the flexible unobservable structure, which allows for unobserved shocks that are time-varying and correlated between banks in a market. This implies that, unlike many prior studies, our model will not identify network effects from within-market correlations in adoption decisions across time. We believe that this is an advantage: that source of identification is often thought to be problematic because it is hard to rule out the presence of contemporaneous unobserved local shocks.

Another interesting identification issue concerns the equilibrium selection parameter Ω . As the number of firms increases, the increasing externality should result in many more cases with a Pareto-worst equilibrium that is distinct from the Pareto-best equilibrium. Thus, the nature of the increased unexplained variance in usage levels for networks with more banks should identify Ω . Note that high variance in the usage levels in *all* markets would not necessarily be evidence of multiple equilibria but

could instead simply be attributed to a large variance of α . Also, this logic suggests that the polar cases of “always in a good equilibrium” and “always in a bad equilibrium” will be hard to identify separately.

Last, note that all the terms in our model are identified only in terms of *utility units* (for customers) or *profits units* (for banks). Equivalently, the bank or customer benefit from an ACH transaction is identified only relative to its fixed cost. There is no obvious way in our model to scale either of these values in real terms: neither are expressed in dollar terms, nor are the two values in comparable units. The underlying reason for this is that we do not observe any meaningful price variation in the data.

Even though we cannot get a dollar measure of the size of the externalities, we can measure the sizes of the externalities in terms of utility units and profit units. This limits, but does not eliminate, the policy experiments that we can perform. For instance, we cannot examine the implication of a dollar subsidy to customers for adoption. However, we can find the implications of a policy where customers are subsidized, e.g. 20 percent of average fixed costs. To enact such a policy, one would need the policymaker to come up with a reasonable guess as to the dollar amount of customer fixed costs.

5. Results and implications

We estimate structural parameters of our model using the indirect inference procedure developed in Section 4. We first discuss the results, then turn to analyzing the implications of various counterfactual policies.

Results

Table 4 provides two sets of estimates for the structural parameters. As detailed in Section 4, both estimates pertain to the same model; they differ in the number and type of indirect inference moments that we use to estimate them. The first column contains our base estimates, the second contains the estimates with our additional II moments. The estimated parameters are similar across the two specifications, with a few exceptions. For instance, the time trend is estimated to be .0848 in the base estimates but .0961 in the specification with additional moments. The customer marginal benefit from an ACH transaction is estimated to be .8149 in the base specification and 1.239 in the alternate specification. One parameter that

is quite different is σ_c , which varies between .1279 and 1.908. However, given that this is just a variance, this difference does not significantly affect the general conclusions of our model.

Two other parameters that are fairly different across the specifications are λ , the scale parameter that relates assets to the number of total transactions (both ACH and non-ACH), which is 3.444 in the base model and 9.161 in the alternate estimates, and β_0 , the customer fixed benefit coefficient, which moves from -3.048 to -4.998 between the two specifications. These differences can be explained by the fact that our data only directly tell us the number of ACH transactions, *not* the total number of transactions or the proportion of ACH transactions. The number of ACH transactions is determined by both λ , which affects the total (ACH and check) number of transactions, and β_0 , which affects customer adoption and hence the proportion of transactions completed using ACH. Given we only observe this total number of ACH transactions, λ and β_0 are likely not *separately* identified by any intuitive variation in the data and are more likely separately identified by particular distributional and functional form assumptions that cannot be easily verified.²² In contrast, we do believe that the combined implications of the two parameters are well identified by our data on the total number of ACH transactions. Moreover, we have found that the economic implications of the model (e.g., the policy experiments reported later in Table 7) are similar across the two parameter estimates, likely because they depend largely on the combined implications of these two parameters on the total number of ACH transactions.

Table 5 provides the values of the moment conditions for both specifications. A comparison of the differences in the simulated moments and the data moments (normalized by their bootstrapped standard errors) suggests that our model fits reasonably well. For the base model, only 5 of the 31 differences are significantly different from zero at the 5 percent level; the figure is 10 out of the 42 for the specification with additional moments. Nonetheless, it is worth noting that the point estimates for some of the important moments are substantially different between the real and simulated data. For instance, the coefficient relating adoption to competitors' adoption is estimated at .1893 or .1607, depending on the specification, while it is .0894 in the data. While this particular comparison suggests that our estimated model may overstate the level of network externalities, this pattern is not prevalent. For instance, in both specifications the coefficients relating local adoption to nonlocal adoption are smaller in our estimated model than in the data.

The value of our objective function at the estimated parameters is 51.70 for the base model and 84.79 for the specification with additional moments. As our choice of weighting matrix standardizes the moments, this represents the sum of 31 or 42 squared i.i.d. $N(0,1)$ random variables. While a chi-squared test for either model would fairly easily reject the joint hypothesis that all of the differences are zero, it is not unusual for a structural model to be rejected by the data.

Table 5 also shows that the two models give similar values for the 31 moments that are matched across both specifications. Although the other 11 moments are more closely matched by the specification that attempts to match these moments, even the base model does reasonably well here. Given the similarity of the parameter estimates across the two specifications, we focus on the base specification in the discussion that follows.

Turning to the actual estimates in Table 4, the parameters seem reasonable. For instance, the coefficient on time trend (β_3) is positive, suggesting that there is increased acceptance of technological goods over time. While the ACH price coefficient (β_2) is the anticipated sign, it is hard to know how to interpret its magnitude since it only represents the impact of the wholesale price charged by the Federal Reserve.

On the customer side, both mean customer fixed costs and marginal benefits are positive (recall that $-\beta_0$ is the mean customer fixed cost). The ratio between them indicates that an average customer would make up its fixed costs of operation over a quarter with about 3,700 transactions. While large employers, for example, could clearly cover these fixed costs, the level is high enough to explain a general lack of customer adoption. For banks, the estimated ratio between the relative markup of an ACH transaction and the bank fixed costs of adoption is much different: the average bank would recoup its fixed costs with only about 30 transactions. This high markup relative to bank fixed costs is consistent with the anecdotal evidence of rather low bank fixed costs detailed in Gowrisankaran and Stavins (2004). Evidence from websites of big banks is also supportive of our estimated high markups relative to fixed costs. In 2004, at a time when the Fed per-item fee was \$.003, we found most large banks' posted per-item ACH fees to be between \$.10 and \$.15, numbers that are 30 to 50 times the wholesale price. If we assume that the in-house costs to banks from processing an ACH transaction are roughly equal to the Fed per-item fee, this still suggests a markup of 15 to 25 times. This contrasts with checks, for which banks often set a zero price, suggesting a negative margin.

At the bottom of the table is our estimated value of δ , the proportion of two-way transactions. We estimate $\delta = 0$, suggesting that *all* transactions are one-way transactions. While the standard error of this parameter is relatively small (e.g., we can easily reject that all transactions are two-way, i.e., $\delta = 1$), the estimated parameter is on the boundary of parameter space, making its estimated standard error not completely reliable.²³ Nonetheless, this is an interesting result, suggesting that ACH is mainly being used by big companies (employers, mortgage companies) that have the ability to make ACH transactions with smaller customers who have not formally adopted originating technology.

It is worth examining further what aspects of the data may be generating our estimate of $\delta = 0$. Similar to the discussion on identification in Section 4, information on δ should come from the shape of the relation between bank adoption and transactions. If $\delta = 1$, all transactions have to have both sides of the market adopting. This implies a quadratic relation between customer adoption and transactions: if 1/2 of customers have adopted, 1/4 of transactions would be made with ACH, if 1/3 of customers have adopted, 1/9 of transactions would be made with ACH, etc. In contrast, with $\delta = 0$, the relation can be more linear, depending on the level of bank adoption. While we do not have direct data on customer adoption, we do have data on bank adoption. With $\delta = 1$, the relation between *bank* adoption and transactions will likely be even more convex than quadratic, as more bank adoption causes more customer adoption. We examine this relation by performing a simple regression of ACH transactions per dollar of deposits on the deposit-weighted fraction of banks adopting. Instead of a quadratic relation, we found a roughly linear relation between the two variables, which does not seem consistent with two-way transactions.²⁴ Note that we should be cautious in interpreting this regression too broadly because its variation is not exogenous.

The estimates of the four random effects' standard deviations are interesting in that the estimated standard deviations of α_{jmt}^A and α_{jm}^B are considerably higher than those of α_m^C and α_{mt}^D . This suggests that firm-specific effects are quantitatively more important than the market-specific effects (although in the alternate specification, the standard deviation on α_m^C is much higher). Overall, these standard deviations are reasonably large, with the standard deviation of the sum α_{jmt} roughly equal in magnitude to the mean customer fixed cost of adoption, suggesting a large variation in customer fixed costs of adoption.

Our equilibrium selection parameter Ω is estimated to be .7106, suggesting that approximately 70 percent of these markets are in the good equilibrium. However, the parameter is not estimated very precisely: we could not reject always being in a good equilibrium or always being in a bad equilibrium. Moreover, as we will see in Table 6 below, our estimates of the other parameters suggest that network effects are simply not strong enough to generate economically significant multiple equilibria. As such, we should not interpret our results as being particularly indicative of which equilibria these markets are in. Note also that the lack of economic significance across values of Ω implies that we need not worry much about the lack of statistical significance in the sense that policy conclusions will not vary much based on the value of Ω .²⁵

Table 6 further explores the economic magnitudes of the base parameters from Table 4 by examining their impact on the estimated equilibrium. The table examines three statistics of the estimated equilibrium – the percentage of local banks adopting ACH, the percentage of customers adopting, and the percentage of overall transactions done through ACH. At the estimated parameters, reported in row 1, 68.7 percent of local banks are adopting, 17.7 percent of customers are adopting, and 16.1 percent of all transactions are made through ACH. The latter two figures imply that the potential market for ACH transactions (e.g., the set of recurring payments) is about five to six times its actual size during our data set, and that only about one-fifth of customers who might benefit from originating ACH transactions were actually originating them.

Row 2 of Table 6 examines what would happen if we eliminated mean bank fixed costs of adoption for local banks. The difference between this and row 1 is indicative of the level of the network externalities at the bank level. Although many more banks adopt ACH, the differences in customer adoption rates and in transactions processed with ACH are small. This is due to our small estimated bank fixed cost of adoption, which implies that it is not bank non-adoption that is preventing customers from using ACH. On the other hand, when we eliminate the customer mean fixed cost of adoption in row 3, there are big changes in the equilibrium proportion of customers that adopt ACH: customer adoption increases to 55.3 percent. In response to this expected adoption by customers, local banks also increase adoption to 96.8 percent, and in this equilibrium, 52.6 percent of all transactions are made using ACH. These estimates suggest that local customer fixed costs are the primary impediments to ACH usage.

Anecdotal evidence is consistent with these findings. For instance, by the year 2000, almost all banks had adopted ACH, further suggesting that bank non-adopters were reasonably close to the margin

of adopting during our sample. Yet, ACH market share was still low, increasing only by 50 percent over 1997 according to National ACH Association (NACHA) estimates.²⁶ The fact that this increased bank adoption did not lead to massive increases in the number of transactions is consistent with the hold-up for adoption being customers, not banks, and with the presence of large customer fixed costs.

Rows 4 and 5 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. While not identical, the results across the two equilibria are very similar. This suggests that, at our estimated parameters, multiple equilibria are simply not a significant issue. Again, this implies that our estimate of Ω is not particularly meaningful.

Last, we investigate what would happen if some of these externalities could be internalized. There is no natural way to compare customer utility to firm profits in our model. As such, we cannot solve a social planner's problem. However, we can investigate what happens if all the local banks coordinated decisions to maximize joint profits, or if all customers coordinated to maximize joint utility. We find (from row 6) that joint profit maximization of all the local banks raises adoption, but not by much. Row 7 reports that when all customers coordinate to maximize joint utility, nothing changes. This result follows directly from the fact that all transactions are one-way and that one-way ACH transactions can be performed without both customers adopting. Because of this, customers simply do not exert externalities on each other.²⁷ However, customers are exerting externalities on banks. In particular, by not adopting, customers are lowering the benefit to their bank and other banks from adopting. Our policy experiments below reveal this to be the primary source of externalities in this market, in the sense that customer subsidies for adoption primarily help banks.

Policy experiments

The above results suggest that it is customer fixed costs that are preventing widespread adoption of ACH technology. In contrast, bank fixed costs are small and do not significantly limit ACH use. This suggests that government policy, particularly at the customer level, might increase welfare. We analyze two types of fixed cost subsidies: one where the government pays a fixed amount to customers who adopt ACH, and another where the government pays a fixed amount to local banks that adopt ACH. It is important to note that customer utility is measured in "utils" and profits are measured in "profit units" and the two measures are not comparable. Thus, we cannot compare a government subsidy received by

customers in utils to a government subsidy received by banks in profit units without further normalization assumptions.

Table 7 provides the results of these policy experiments using the base parameters from Table 4.²⁸ Column 1 of Table 7 reports properties of the estimated equilibrium as a baseline measure. In addition to statistics on customer and bank adoption, we compute six welfare measures: bank profits,²⁹ customer utility, the costs to the government of the potential subsidies (in both profit and utility units), and the resulting net profits and net utility.

As a first cut at useful subsidies, columns 2 and 3 report the welfare consequences of two subsidies similar to the experiments from the prior section. We use government subsidies to eliminate either customer mean fixed costs (column 2) or local bank mean fixed costs (column 3). The bank subsidy unambiguously increases welfare, as total profits and utility increase. However, its total effect on usage and welfare is small, because bank fixed costs are not the central cause of non-adoption. The customer subsidy is far more effective at increasing ACH usage; it roughly triples both local bank profits and customer utility. However, the subsidy is inefficient in terms of utils, as the gain in customer utility due to the policy is more than offset by the cost of the subsidy to the government in utils. This suggests that the subsidy might be generating inefficiently high levels of adoption.

Another limitation of these two subsidies is that they are not large enough to force everyone to adopt. Thus, we tried alternate subsidies that are large enough to force virtually everyone to adopt. A large bank subsidy generates only marginally different results than column 3 (with slightly higher total utility and total profit numbers), so we do not report the results. The fact that the large bank subsidy results in outcomes that are virtually the same as the mean fixed cost bank subsidy is not surprising given that 94.8 percent of banks were already adopting in this earlier subsidy. In contrast, introducing a large customer subsidy (column 4) does change things significantly, generating even more ACH transactions and firm profits than the customer mean fixed cost subsidy, but at an even larger utility cost.

Given that we have no way of relating the increase in profits to the decrease in utility, we cannot determine whether policies that raise one and lower the other (such as column 2 and column 4) are welfare-improving. However, we can attempt to construct customer subsidies that unambiguously raise welfare. In particular, the overadoption in column 2 suggests that a smaller, more efficient customer subsidy might leave customer utility unchanged or higher while increasing profits. Column 5 exhibits results from the largest customer subsidy (approximately) that does this, 1.7 percent of customer mean

fixed cost. With this subsidy, total utility is unchanged, but total profits increase by approximately 3 percent. Column 6 adds a local bank subsidy to the policy, since column 3 shows that local bank subsidies, while not increasing adoption by much, do unambiguously increase welfare. We add a very large subsidy to local banks, enough to ensure that they all adopt. With this concurrent bank subsidy, we can increase the customer subsidy to 7.4 percent of customer mean fixed costs and still keep customer utility at its baseline level. This ends up increasing local bank profits by about 18 percent. This suggests that the government can substantially increase welfare in this market with a simple fixed cost subsidization policy. Note that these profit increases are in some sense a lower bound, as we are forcing ourselves to keep total utility constant in both of these subsidy schemes. If we had a way to relate utility to profit units, we could likely use subsidies to increase profits or overall welfare by more than 18 percent.³⁰

The above results also suggest that banks might be able to increase their profits by subsidizing their own customers. However, there are potential problems with doing this. First, banks would need to devise mechanisms to expropriate the excess customer surplus generated by the subsidies. Second, there would be an issue regarding externalities between banks. Specifically, possible subsidies to a bank's own customers make it more likely that the bank will find it profitable to adopt. This in turn creates profits for other banks, i.e. an externality, and as a result, these subsidies might be underprovided.³¹ The ability to switch banks implies that customers might be able to take adoption subsidies (e.g., software for ACH processing and accompanying employee training) from one bank and then switch to the bank with the lowest per-transaction fees. If this is in fact the case, there would be considerable returns to banks from developing methods for reducing the expropriation problem, perhaps through contracts or proprietary software.

6. Conclusions

In this paper, we have estimated a static 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 reasonable and generally precisely estimated, and they fit the data reasonably well. However, our model is fairly simple and static. A key area for future research would be to extend the model to allow for sunk costs and dynamic optimization.

We find that bank fixed costs from ACH adoption are low and do not explain why ACH is not more widely used. In contrast, customer fixed costs of ACH adoption are substantial and are a major explanation for the lack of ACH usage. Most ACH transactions appear to be one-way in the sense that recipient customers do not need to adopt formally for them to be realized. 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. Changes that lower the customer fixed cost of ACH adoption will encourage adoption and usage of ACH. Because of this, policies that subsidize a small portion of customer fixed costs can unambiguously increase total surplus. Adding large bank subsidies for ACH adoption in combination with the customer subsidies can increase welfare even more.

References

- Athey, Susan, Paul Milgrom and John Roberts (1996). "Robust Comparative Statics." Mimeo, Stanford University.
- 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.
- Bjorn, P.A. and Q. Vuong (1985). "Simultaneous Equations Models for Dummy Endogenous Variables: A Game Theoretic Formulation with an Application to Labor Force Participation." Working Paper No. 538, California Institute of Technology.
- Goffe, W.L., G.D. Ferrier and J. Rogers (1994). "Global Optimization of Statistical Functions with Simulated Annealing." *Journal of Econometrics* 60: 65-99.
- Goolsbee, A. and P. Klenow (2002). "Evidence on Network and Learning Externalities in the Diffusion of Home Computers." *Journal of Law and Economics* 45: 317-44.
- Gouriéroux, C., A. Monfort and E. Renault (1993). "Indirect Inference." *Journal of Applied Econometrics* 8: 85-118.
- Gowrisankaran, Gautam and Joanna Stavins (2004), "Network Externalities and Technology Adoption: Lessons From Electronic Payments." *RAND Journal of Economics* 35: 260-276.

- Heckman, James J. (1978). "Dummy Endogenous Variables in a Simultaneous Equation System." *Econometrica* 46, 931-959.
- Heckman, James J. and R. 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 (1994). "A Computationally Practical Simulator for Panel Data." *Econometrica*, 62: 95-116
- 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.
- Kooreman, P. (1994) "Estimation of Econometric Models of Some Discrete Games." *Journal of Applied Econometrics*, 9, 255-268
- Manski, Charles (1993). "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies* 60: 531-42.
- McFadden, Daniel L. (1989). "A Method of Simulated Moments for Estimation of Discrete Response Models without Numerical Integration." *Econometrica* 57: 995-1026.
- Milgrom, Paul and Chris Shannon (1994). "Monotone Comparative Statics." *Econometrica* 62: 157-180.
- Moro, Andrea (2003). "The Effect of Statistical Discrimination on Black - White Wage Inequality: Estimating a Model With Multiple Equilibria." *International Economic Review* 44: 467-500.
- Ohashi, Hiroshi (2003). "The Role of Network Effects in the U.S. VCR Market." *Journal of Economics and Management Strategy* 12(4): 447-494.
- Pakes, Ariel and David Pollard (1989). "Simulation and the Asymptotics of Optimization Estimators." *Econometrica* 57: 1027-1057.
- Park, Sangin (2002). "Semiparametric Instrumental Variables Estimation." *Journal of Econometrics* 112: 381-99.
- Rust, John (1987). "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica* 55: 999-1033.
- Rysman, Marc (2004a). "Competition Between Networks: A Study of the Market for Yellow Pages." *Review of Economic Studies* 71: 483-512.

Rysman, Marc (2004b). "Empirical Analysis of Payment Card Usage." Mimeo, Boston University.

Staiger, Douglas O. and James H. Stock (1997). "Instrumental Variables Regression with Weak Instruments." *Econometrica* 65: 557-86.

Tamer, Elie (2003). "Incomplete Simultaneous Discrete Response Models with Multiple Equilibria." *Review of Economic Studies* 70: 147-65.

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

Table 1: Characteristics of Local Banks

Number of local banks	Number of networks/time periods	Mean deposits (millions)	Mean percent of banks using ACH	Mean ACH transactions by bank
1	2730	\$45.8	64.3	457.7
2	1310	\$49.5	64.5	452.0
3	367	\$59.4	67.8	1,217
4	172	\$73.0	74.4	1,348
5	83	\$50.1	74.2	912.5
6	51	\$125	70.3	3,485
7	31	\$139	73.7	2,155
8	41	\$57.5	66.2	991.5
9	39	\$79.9	69.5	897.9
10	25	\$81.6	57.6	732.2

Table 2: Characteristics of Nonlocal Banks

Number of local banks	Mean number of nonlocal banks	Std. dev. of number of nonlocal banks	Mean deposits within network by nonlocal banks (millions)	Mean total deposits by nonlocal banks (millions)	Percent of nonlocal banks using ACH
1	2.02	2.69	\$59.8	\$10,400	88.5
2	1.73	2.29	\$60.2	\$6,800	85.8
3	3.25	3.30	\$96.0	\$9,200	89.0
4	4.06	3.38	\$92.0	\$8,600	88.5
5	5.27	5.24	\$187	\$4,800	83.3
6	5.67	5.20	\$96.9	\$8,500	84.1
7	9.13	4.26	\$78.6	\$5,000	91.9
8	6.80	4.65	\$95.2	\$7,900	86.0
9	8.72	5.80	\$104	\$6,900	87.4
10	6.56	3.80	\$120	\$4,900	81.1

Note: Table based on observations kept in sample.

Table 3: Usage Over Time by Local Banks in Network

Time period	Percent of networks with no local bank using ACH	Percent of networks with some, but not all, local banks using ACH	Percent of networks with all local banks 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 two or more local banks kept in sample.

Table 4: Parameter Estimates

Parameter	Base estimation	Estimation with additional moments
λ (transactions coefficient) (unit: trans. per \$100,000)	3.444 (1.832)	9.161 (2.218)
β_0 (customer fixed benefit) (unit: utils)	-3.048 (.2407)	-4.998 (.1938)
β_1 (customer marginal benefit) (unit: 1,000 transactions)	.8149 (.2714)	1.239 (.3309)
β_2 (price coefficient) (unit: \$.01)	-.5444 (.1325)	-.8720 (.2544)
β_3 (time coefficient) (unit: quarter)	.0848 (.013)	.0961 (.0200)
Markup (unit: profit units per 1,000 transactions)	190.4 (92.78)	313.7 (80.08)
\overline{FC} (bank fixed costs) (unit: profit units)	6.256 (1.083)	3.395 (.3233)
Ω (Probability of good equilibrium)	.7106 (.4458)	.7546 (1.523)
σ_A (std. dev. of random effect α_{jm}^A) (unit: utils)	.8276 (.0576)	.8675 (.1272)
σ_B (std. dev. of random effect α_{jm}^B) (unit: utils)	1.953 (.1556)	2.028 (.1580)
σ_C (std. dev. of random effect α_m^C) (unit: utils)	.1279 (.0661)	1.908 (.2900)
σ_D (std. dev. of random effect α_{jm}^D) (unit: utils)	.0383 (.0181)	.0044 (.1597)
δ (proportion of two way transactions)	0 (.1847)	0 (.1002)
Moment condition at estimated parameters	51.70	84.79
Number of moments	31	42

Note: Standard errors in parentheses.

Table 5: Indirect Inference Moments at Estimated Parameters

Description	Moment in data	Moment in model		T-stat. for diff.	Moment in model	T-stat. for diff.
		Base specification	Specification with additional moments			
Constant	.6516	.5695	-2.536	.5967	-1.767	
Competitor adoption	.0894	.1893	2.9707	.1607	2.145	
Time dummy	.0040	.0088	.8923	.0085	.8349	
Time dummy	.0173	.0284	1.4209	.0298	1.663	
Time dummy	.0358	.0457	1.193	.0469	1.368	
Time dummy	.0468	.0517	.5477	.0500	.3623	
Time dummy	.0642	.0683	.3847	.0675	.3157	
Time dummy	.0832	.0774	-.5497	.0718	-1.085	
Time dummy	.0900	.0842	-.5023	.0786	-.9956	
Time dummy	.1113	.0963	-1.237	.0889	-1.846	
Time dummy	.1125	.1057	-.5834	.1003	-1.032	
Time dummy	.1142	.1098	-.3679	.1045	-.8147	
Deposits	.1480	.3104	.8904	.2768	.7094	
Squared deposits	-.0879	-.1517	-.2635	-.1245	-.1431	
SE competitor adoption	.0122	.0120	-.6687	.0122	-.1720	
Std. dev. of residual	.1978	.1702	-2.644	.1878	-.8939	
Std. dev. of residual	.2826	.2889	.9644	.2802	-.3671	
Std. dev. of residual	.1232	.1129	-1.994	.1132	-1.856	
Std. dev. of residual	.2168	.2169	.0286	.2133	-.6749	

Adoption on competitor adoption

ACH volume per deposits on comp. adoption	Constant	.0009	.0002	-.7860	.0005	-.4207
	Competitor volume per transaction	.0021	.0587	1.542	.0998	2.482
	Competitor adoption	.000017	.00025	2.178	.00014	1.203
Adoption on HHI	Constant	.1784	-.0947	-1.150	.1507	-.1208
	HHI	.0949	.0653	-.5728	.0255	-1.437
	R ²	.1360	.1120	-1.570	.0694	-4.341
Adoption on nonlocal adoption (Model 1)	Constant	.2570	.3877	1.445	.4251	1.754
	Nonlocal adoption	.2142	.0732	-1.589	.0893	-1.319
	R ²	.0850	.0795	-4.260	.0484	-2.817
Adoption on nonlocal adoption (Model 2)	Constant	.2591	.4062	1.646	.4348	1.830
	Nonlocal adoption	.2220	.0830	-1.557	.1025	-1.253
	R ²	.0922	.0804	-.7179	.0500	-2.539
Simple data moments for robustness check	Mean adoption	.6651	.6867		.6864	1.529
	Mean transactions	.8259	.3392		.5333	-2.329
	Std. dev. of local transactions	3.955	.8074		1.750	-2.321
	Mean transactions 1995:Q2	.5168	.2186		.3348	-2.323
	Mean transactions 1997:Q4	1.253	.4836		.7978	-2.108
	Correlation of deposits and adoption	.1819	.1962		.1398	-1.817
	Correlation of deposits and transactions	.4241	.5258		.3576	-2.037
	Within-bank correlation in adoption	.3426	.4065		.4292	1.683
	Within-bank correlation in trans.	.6191	.7412		.7660	1.275
	Within-bank std. dev. of adoption	.1950	.1979		.1872	-1.285
	Within-bank std. dev. of transactions	.3469	.1784		.2807	-1.484

Table 6: Economic Significance of Parameters

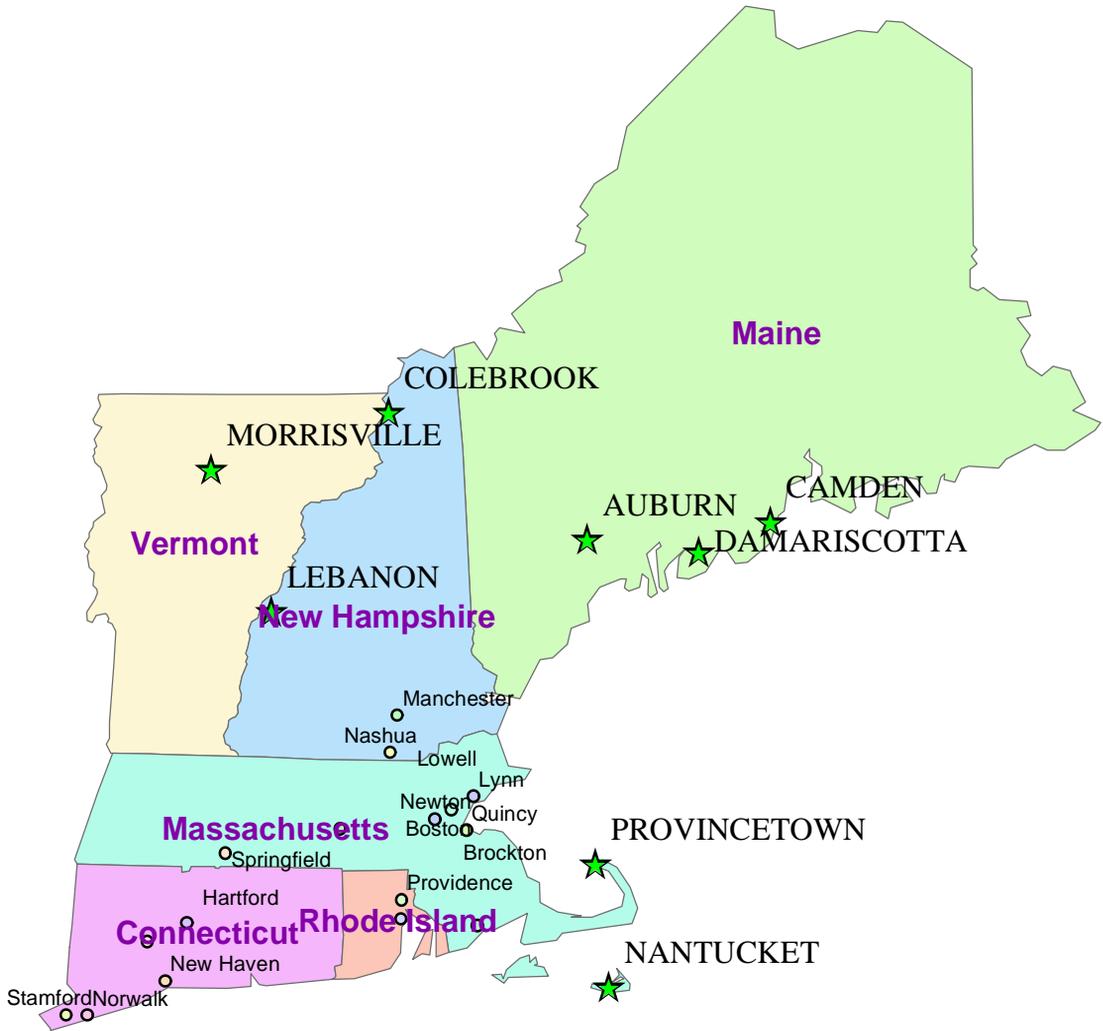
Change	% of local banks that adopt	% of customers that adopt	% of transactions completed with ACH
Estimates (1)	68.7	17.7	16.1
No mean bank fixed costs (2)	94.8	17.8	16.9
No mean customer fixed costs (3)	96.8	55.3	52.6
Always in bad equilibrium (4)	68.3	17.7	16.0
Always in good equilibrium (5)	68.9	17.7	16.1
Local banks internalize externality (6)	83.3	17.8	16.8
All customers internalize externality (7)	68.7	17.7	16.1

Table 7: Policy Experiments

Policy:	No subsidy	Subsidy to cust.: customer mean fixed cost	Subsidy to local banks: bank mean fixed cost	Large customer subsidy	Subsidy to cust.: .017 customer mean fixed cost	Large bank subsidy .074 cust. mean fixed cost
Statistic:	(1)	(2)	(3)	(4)	(5)	(6)
% local banks that adopt (1)	68.7	96.8	94.8	100	69.5	100
% cust. that adopt (2)	17.7	55.3	17.8	93.4	18.2	20.0
% trans. made with ACH (3)	16.1	52.6	16.9	89.2	16.5	19.1
Bank profits (profit units) (4)	1.42 million	4.73 million	1.54 million	8.04 million	1.46 million	2.13 million
Customer utility (utils) (5)	13,264	64,957	13,370	339,181	13,713	15,416
Cost to govt. (profit units) (6)	0	0	53,425	0	0	450,300
Cost to govt. (utils) (7)	0	80,431	0	445,646	449	2,150
Net profits (profit units) (8) = (4) – (6)	1.42 million	4.73 million	1.48 million	8.04 million	1.46 million	1.68 million
Net utility (utils) (9) = (5) – (7)	13,264	-15474	13,370	-106,465	13,264	13,266

Figure 1: Map of New England indicating large cities (with circles) and networks in sample (with stars) in New England

(See next page.)



¹ A related literature considers the empirical implications of social interaction models (see Brock and Durlauf, 2001). One paper in this literature (Topa, 2001) also identifies structural parameters from the steady state of the system, although not in the context of an underlying profit-maximizing model of decision-making.

² While Gowrisankaran and Stavins (2004) did recover magnitudes for some of the specifications, these specifications were limited and somewhat problematic. For instance, they estimated the discrete ACH adoption decision using a linear instrumental variables specification, which is problematic in this context (see Heckman and Robb, 1985, and Angrist and Imbens, 1994).

³ The idea behind this assumption is that the ACH adoption decision appears to be made at the *bank* level, not the *branch* level. Since these nonlocal banks are very large relative to the small markets we examine, we believe that the exogeneity assumption is reasonable.

⁴ We assume a constant N because it would be difficult to identify any sort of distribution on N credibly without customer-level data on the number of transaction originations.

⁵ As well as originating their own N transactions, customers are also the recipients of transactions originated by other customers in the market, a point we explore below.

⁶ A direct deposit for payroll is a credit transaction originated by the employer, while a mortgage payment is typically a debit transaction initiated by the lender.

⁷ Note that we are assuming that this proportion is the same for all customers, i.e., δ percent of each customer's N transactions are two-way. An alternative model might specify two different types of customers, one who only makes two-way transactions, and one who can make one-way transactions. Again, given our lack of customer-level data, it would likely not be very fruitful to try to distinguish these alternative models.

⁸ We specify prices as varying only across time, because of a lack of better data.

⁹ We assume that utility only accrues to the originator of the transaction because we do not have the customer-level data that would likely be necessary to estimate two different utilities.

¹⁰ If one incorrectly disallowed these sorts of correlations, the result would likely be biased estimates of network effects, as within-market correlations in adoption would be incorrectly picked up by the network effects parameters.

¹¹ We fold p_i^{CHK} into β_i in (7) because we do not have data on the price of checks.

¹² 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 (2004) for details.

¹³ Our method of estimating models with multiple equilibria is a generalization of the method used by Moro (2003), who estimates the equilibrium as a parameter in a model of labor discrimination. The difference is that we estimate the frequency of being in either equilibrium as a parameter, since we observe several regional markets, while Moro (2003) has only one market per year. It is also similar to methodologies used by Bjorn and Vuong (1985), Kooreman (1994), and Tamer (2003).

¹⁴ We thank the Federal Reserve's Retail Payments Product Office for making this data set available to us.

¹⁵ Note that since there is a continuum of customers, the customer level unobservables ε are aggregated out of the model at the level of the data.

¹⁶ We use a combination of simulated annealing (see Goffe, Ferrier, and Rogers, 1994) and a simplex procedure with randomness.

¹⁷ Another alternative would be to estimate our model with simulated maximum likelihood (SML), e.g., Keane and Wolpin (2000). One problem with SML in our context is that we have continuous dependent variables. This makes the model prone to generating probability zero events. Keane and Wolpin (2000) deal with this by adding analytically integrable measurement error to the model, which generates a positive likelihood of any event.

¹⁸ One advantage is that this is a one-step procedure. This alleviates the ambiguity that occurs with standard two-step MSM estimators where the choice of the initial weight matrix affects the initial estimates, which in turn affect the second-stage weight matrix and estimates.

¹⁹ The proof is essentially the same as in Gowrisankaran and Stavins (2004).

²⁰ The logic behind this is that banks without any nonlocal competitors are more similar to banks where all nonlocal competitors have adopted than to ones where none have adopted.

²¹ Note that there is a selection issue here, since we do not observe the proportion of customers adopting for banks that do not adopt. Our structural model implicitly accounts for this selection.

²² From the standard errors, one can see that the confidence intervals for these two parameters are also very different across the two specifications. Note that standard errors only measure variance due to the data-generating process, not variance due to uncertainty about model specification. It is likely that by highlighting different aspects of the functional form and distributional assumptions generating identification, the two specifications are generating different statistical predictions.

²³ To compute the variance/covariance matrix in (15), we use a one-sided derivative for the column of the derivative matrix Γ that pertains to δ . An alternative would be to fix $\delta = 0$ and calculate the variance/covariance matrix for the other parameters with this assumption. We tried this and found similar, though generally smaller, standard errors than the ones that we reported, suggesting that our reported precisions may be conservative.

²⁴ Using the fraction of banks adopting and its square as explanatory variables, we obtain a t-statistic of 3.04 on the linear term and 0.18 on the squared term, with the constant term also being insignificant.

²⁵ A caveat is that this statement depends on the nature of the policy experiment. Our policy experiments use subsidies to decrease the externality. This likely decreases the possibility of multiple equilibria, so the estimate of Ω is even less important. In contrast, policies that seek to increase the level of the externality might generate economically significant multiple equilibria, making the estimate of Ω important. Fortunately, policies that increase the externality are generally not that interesting or useful.

²⁶ See http://www.nacha.org/news/Stats/ACH_Statistics_Fact_Sheet_2001.pdf.

²⁷ If we instead estimated a specification where recipients of transactions obtain utility (as in footnote 9), there would be externalities between customers, even though the proportion of two-way transactions would be zero.

²⁸ It would be conceptually straightforward to obtain confidence intervals around these numbers by bootstrapping from the estimated variance/covariance matrix for the parameters.

²⁹ Note that our bank profit measure includes the increment in profits to nonlocal banks.

³⁰ For example, one way to relate profit units to utility units is to make an assumption about how surplus is split in the equilibrium observed in the data. If we assume that banks and customers split the original equilibrium surplus 50/50, then (from Table 7 column 1) we get the equality that 1.42 million profit units = 13,264 utility units, or that 1 utility unit = 107 profit units. Using this welfare conversion mechanism, a customer subsidy that is 55 percent of mean fixed costs plus a large bank subsidy can increase firm profits by almost 60 percent.

³¹ Note that our model is consistent with banks already subsidizing their customers, in which case our estimated fixed costs and benefits can simply be interpreted as their post-subsidy values. However, the point still remains that further subsidies would be profitable.