

Heterogeneity and the Dynamics of Technology Adoption*

Stephen Ryan[†] and Catherine Tucker[‡]

January 13, 2007

Abstract

This paper analyzes the role of heterogeneity and forward-looking expectations in the diffusion of network technologies. Using a detailed dataset on the adoption of a new videoconferencing technology within a firm, we estimate a structural model of technology adoption and communications choice. We allow for heterogeneity in network benefits and adoption costs across agents. We develop a new “simulated sequence estimator” to measure the extent to which agents seek diversity in their calling behavior, and characterize the patterns of communication as a function of geography, job function, and rank within the firm. We find that agents differ significantly in their adoption costs and network benefits. We find that agents have significant welfare gains from having access to a diverse network, and that a policy of strategically targeting the right subtype for initial adoption can lead to a faster-growing and larger network than a policy of uncoordinated or diffuse adoption.

*Financial support from the NET Institute, www.NETinst.org, is gratefully acknowledged. We would also like to thank Dan Akerberg, Amy Finkelstein, Shane Greenstein, Andrew Sweeting and seminar participants at the 2006 International Industrial Organization Conference, Boston University, the University of Mannheim, Central European University, the 2006 Minnesota Applied Microeconomics Conference, the Wharton School of the University of Pennsylvania, the 2006 Quantitative Marketing and Economics Conference, and MIT.

[†]Department of Economics, MIT, Cambridge, MA and NBER.

[‡]Marketing Group, MIT Sloan School of Business, Cambridge, MA.

1 Introduction

Technological improvements lie at the heart of economic growth, and understanding the diffusion of innovation is a centrally important question in economics. In his pioneering work on the diffusion of hybrid corn technology, Griliches (1957) poses three questions which still resonate today: What factors influence the timing of adoption of new technologies? What determines their rates of diffusion? And finally, what factors govern the long-run level of adoption? Griliches, along with other early empirical and theoretical work such as Mansfield (1961) and Rogers (1962), attempted to answer these questions by explaining differences in diffusion curves as arising from heterogeneity in user characteristics, such as profitability, cost, and competitive pressure. Foundational work by Katz and Shapiro (1985) and Farrell and Saloner (1985) greatly extended this literature by identifying a completely different mechanism driving the diffusion of a broad class of technologies. In these “network technologies,” canonical examples of which are telephones, fax machines, and the Internet, an agent’s payoff from the technology explicitly depends on having other agents adopt the technology as well. For these technologies, equilibrium expectations over how the network will evolve in the future are crucial to understanding Griliches’ three questions.

In this paper, we bridge and extend these two literatures. We build on earlier work by constructing a fully dynamic, utility-based model of technology adoption and use. We examine how heterogeneity, as expressed by differences in adoption costs, network effects, and tastes for a diverse network, affects network technology diffusion and use. We apply our model of forward-looking heterogeneous agents to detailed data on the introduction of a videoconferencing technology in a large multinational bank. Our approach allows us to quantify the effects of three dimensions of individual heterogeneity on network evolution and use, and permits analysis of two common policies for jump-starting network technology diffusion. Our research strategy consists of three sequential steps.

First, we construct a fully dynamic model of network technology adoption and use. The model addresses two interrelated technological questions: how the network evolves over time, and how agents in the network use it. Agents vary in their fixed costs of adopting the network technology, and weigh the expected present value of joining the network today against the opportunity costs of not joining today. This naturally leads to formulating the adoption decision as an optimal waiting problem. After an agent has adopted the technology, they then can choose how to use it. We model the sequence of network interactions as a function of two forces: differences in utility each agent receives from interacting with others; and a taste for “dynamic diversification,” or the desire to interact with different agent types in sequence. The latter is motivated by the idea that the utility of making a connection to an agent may depend on whom I have interacted with previously. For example, if the agent is collecting information to solve a problem, she may value a diverse set of resources to draw upon. Our model allows us to provide a rich description of how diversity in the characteristics of network subscribers affects agents’ motivations for adoption.

Second, we apply our model to an extensive data set on the diffusion and use of a videoconferencing technology within a large multinational investment bank. We have detailed data on all 2,169 potential adopters in the firm, from the time that the technology was first offered for installation up to the network’s steady state three and half years later. Our data also encompasses all of the 463,806 videoconferencing calls made during that time, which allows us to estimate a rich model of calling preferences for 64 different types of individuals in the firm. The technology deployment was unusually clean from a modeling standpoint, as the bank took a laissez-faire approach to spreading the technology throughout the firm. Employees were able to get the technology installed upon request at no cost to themselves, but were not otherwise compelled to adopt it. This process falls naturally within the confines of our modeling framework, as otherwise we would have to model the firm’s adoption policies. We use recently-developed techniques from the literature on the estimation of dynamic

games to recover primitives that are consistent with the evolution and use of the network technology within this firm.

Third, we use our parameters to simulate how two different technology adoption policies focused on initial adoption could affect the evolution and use of the network over time. These policies represent potential deployment approaches that a firm or network operator can use to avoid sub-optimal diffusion for their technology. Under the first policy, the firm targets one type of agent as the initial set of technology adopters. The rationale for this policy experiment is that firms commonly roll out a new technology in a specific workgroup, for example among all the IT staff, before allowing wider adoption throughout the organization. In the second policy the firm adopts a uniform adoption strategy, where the technology is spread equally across various types in the initial period. This type of policy can be more effective when agents value being able to communicate with a wide variety of other agents. Comparing these two policies to the baseline case of decentralized adoption will allow us to evaluate the extent to which heterogeneity in agent behavior and characteristics must be accounted for in crafting an optimal policy for jump-starting the diffusion of a network technology.

The structural empirical literature on network effects has been static in nature. For example, Rysman (2004)'s work on two-sided markets evaluates cross-sectional yellow pages data, while Akerberg and Gowrisankaran (2006)'s assumption of free exit enables them to analyze the diffusion of electronic payments as a repeated static game. This orientation towards static models has been driven by three practical challenges. First, in technology adoption models with network effects, the researcher must confront the issue of multiple equilibria. Both Akerberg and Gowrisankaran (2006) and Rysman (2004) tackle this by computing which out of a limited set of equilibria are selected. It is also theoretically possible to not limit the set of potential equilibria, and explicitly model the equilibrium selection process, as in Bajari, Hong, and Ryan (2006). However, this approach requires

the computation of all equilibria to a system, which can take a prohibitive amount of time. This is due to the second difficulty, which is the size of the state space. In the present application, for example, the state space consists of an indicator function for each agent denoting their adoption status. The number of possible combinations of these variables is 2^{2169} , or approximately 10^{602} . It is clear that is this an impossibly large set of points to enumerate, let alone compute equilibria over. However, by using the two-step techniques described by Bajari, Benkard, and Levin (2006), we circumvent the problem of multiple equilibria and the curse of dimensionality which beset estimation of dynamic technology adoption games.¹

The last difficulty is that in any research on network effects, identification is a key challenge. This means that earlier empirical work has focused on documenting causal network effects, see for example Gowrisankaran and Stavins (2004). In this paper we take a different approach. Rather than trying to explicitly estimate a causal network effect, instead we structurally model the entire system of inter-related demand over time. This means that our estimates encompass all drivers of inter-dependent demand. These drivers include informational spillovers, employee coordination and herding as well as causal network effects. This agnosticism resembles modeling approaches such as Bass (1969), which allows for multiple mechanisms by which users' influence each others' adoption.

Another question we answer that has not been tackled by the previous network literature is how to model network usage after adoption. Existing discrete choice models are not appropriate for modeling an employee's sequential and interrelated choices governing which other employee to call over a given period. We propose a new "simulated sequence estimator" to deal with the twin challenges of predicting how many calls an agent will make and whom

¹Our results also contribute to a new literature which explicitly addresses issues of dynamics in technology adoption. One example is Schmidt-Dengler (2005)'s research on dynamic technology adoption timing in the presence of pre-emption effects. Einav (2004) also studies the introduction of new products from the firm's perspective and shows that dynamic estimation can reveal inefficiencies in timing.

they will call.

Our primary finding is that heterogeneity is important at all three levels that we specify. Employees in the firm have very different tastes for using the system, depending on their location, job function, and rank. We find the pattern that, all else equal, each given subtype in the firm is more likely to call someone similar in the firm. However, allowing for dynamic diversification in tastes implies that this taste decreases in the number of times a call is made. Employees therefore have significant positive welfare gains from having access to a diverse network where there are employees of many types for them to call.

Using our estimates, we compared two commonly used technology management policies. Reflecting the complex interplay between heterogeneity in network effects among employees in the firm and heterogeneity in adoption costs, we find that the policy with targeted interventions dominates the uniform adoption policy. The network that is seeded with one subtype grows faster and stays larger, by almost 20 percent, in the long run. Targeting should be used towards a subtype of employee that has high adoption costs, but also large network effects on the adoption decisions of others. By inducing them to enter early, a targeting policy changes other employees' expectations about how the network will evolve. This leads to slightly more calls per adopter, and significantly higher overall welfare.

Our results also shed substantial light on how communication in the firm operates across geography, job function, and rank. There is a burgeoning literature examining the role of hierarchies and communication in firms, e.g. Garicano and Hubbard (2003) and Garicano (2000). While we find evidence that communication in the hierarchy is more likely between similar ranks in the firm, we observe communication across all regions, functions, and ranks. The complexity of the system of communication we uncover suggests that the highly stylized models of communication networks prevalent in the theoretical literature need extension to be capable of reproducing our results.

The paper is organized as follows. Section 2 describes the technology and data used in

this study. Section 3 lays out a dynamic model of technology adoption choice and subsequent interaction choice. Section 4 discusses our estimation strategy. Section 5 discusses the results of our estimation. Section 6 reports results from a policy experiment to test two alternative technology adoption policies. Section 7 concludes and discusses directions for future work.

2 Data and Technology

2.1 Technology

Installing videoconferencing can improve the effectiveness of internal firm communication, by adding visual communication cues to the audio communication cues provided by telephones.² Older videoconferencing systems failed because they were based on rarely-used videoconferencing rooms. This research studies a new videoconferencing technology attached to an employee's workstation. The end-point technology consists of three elements: videoconferencing software; a media compressor; and a camera fixed on top of the computer's monitor. Using the language of Farrell and Saloner (1985), the videoconferencing technology has a "network use" and a "stand-alone use." The network use is television-quality videoconferencing calls. The stand-alone use is watching TV on a desktop computer. In this paper we explicitly abstract from the stand-alone use and focus on the network use, since this is of more general interest.

The videoconferencing technology can only be used for internal communication within the firm. This makes it attractive for empirical studies, because there are comprehensive data on all potential adopters.

²The advantages of visual communication cues are documented in technical literature such as Marlow (1992).

2.2 Firm Setting

We study adoption within a single multinational bank. After the bank chose this technological standard to conduct internal videoconferencing, it invested in an extensive network architecture. We study this particular bank because an existing relationship with the videoconferencing technology manufacturer meant that they adopted a laissez-faire policy towards the distribution of this technology within the firm. The bank publicized the technology to employees and each employee decided if and when to order a videoconferencing unit from an external sales representative. The videoconferencing firm had excess capacity in this period, and through our conversations with them and the bank, we uncovered no evidence of supply constraints restricting adoption. Though such explicitly decentralized adoption is unusual, it is not unusual for companies to install software or ICT equipment for employees and then leave it to the employee's discretion whether or not they use it.

The bank made employees eligible to adopt the technology if they held a position of Associate or higher (85 percent of full time employees). The videoconferencing supplier had excess capacity, so capacity constraints did not affect the timing of employee installation decisions.

2.3 Data

There are complete personnel records for each employee in the bank in March 2004. Throughout our data there were 2,169 employees employed. Entry and exit was around 300 employees, and we exclude employees who left the firm from our data. Data are available for both those employees who adopted videoconferencing and those who did not. The bank was divided up into different divisions. To reduce the number of "types" of employees in our study, we focus on the largest division of the bank and exclude observations on the Credit Analysis and Finance divisions. Employees were divided up into a hierarchy of Associates, Vice-

Presidents, Directors and Managing Directors. Employees are also divided by the function they performed in the firm: Administration, Research, Sales, and Trading. Last, employees are located in 4 broad geographical locations: the US, the UK, Europe and the Rest of the World (mainly Asia). This 4 region x 4 function x 4 title structure gives us a set of 64 broad categories of employees for our empirical analysis. Figure 1 shows the distribution of employees across the firm organized by the 64 broad categories of types. Figure 2 shows the percentage of each of these groups that adopted, where adoption is defined as whether or not that employee ever used the technology for any purpose.

A call database recorded each of 2.4 million calls made using videoconferencing technology from January 2001 to August 2004, within the bank. The call database has two types of call data. For two-way videoconferencing calls, the database records who made the call, to whom they made it, when they made it and how long it lasted. For one-way TV calls, the database records who made the call, to which TV channel, when and for how long. We excluded from our call data: TV-watching calls, calls which involved the Finance/Credit-Analysis division, calls which had multiple participants (5 percent of calls), calls made by employees who left the firm and calls that did not go through or ended in error. Of the original 2.4 million calls, we used 463,806. Figures 3, 4 and 5 illustrate that though on average calls were made most frequently between employees of similar types, there was a great deal of cross-type calling. Figure 6 provides further evidence of heterogeneity at the micro-level by comparing the adoption patterns and number of calls for US researchers in different management positions.

3 Theoretical Model

In this section, we construct a theoretical model of the initial adoption decision and subsequent calling decisions of an agent. This model consists of three elements: the state space, the transition rules over this state space, and agents' per-period payoff given the state space.

3.1 State Space and Timing

Our state space consists of the set of agents in the model, their characteristics, and their adoption decisions at a given time. Each element s_{it} of the state space s_t is a vector encoding the adoption decision and observable characteristics of each agent in the firm. In contrast to the previous literature on dynamic games in technology adoption, in which agents are identical, here agents are endowed with a geographic region, a job function, and a title which describes that agent's relative rank in the firm. We assume that these characteristics are exogenous and do not vary over time. In the exposition below, we slightly abuse notation and indicate that agent i has adopted the network at time t by $s_{i,t} = 1$.

The state space evolves in discrete time, and there is no bound to number of periods that the network can be active. We further simplify by assuming that all agents share the same discount factor, $0 < \beta < 1$, when evaluating future payoffs. We assume that the relevant period of time is one month since much of the rhythm of work is based on monthly cycles in investment banks although we relax this assumption in our specification tests below. We also assume that agents are able to fully use the network in the period in which they adopt the technology.

3.2 Per-Period Payoffs: A Model of Communications Choice

Agents adopt videoconferencing technology to communicate with other agents in the network. The per-period payoff for each agent in the model is the payoff they receive from the calls they make to other agents, relative to the outside communication option. To capture the benefits of these interactions, we propose a model which generalizes the standard discrete-choice utility maximization framework from a single choice to a sequence of interdependent choices. The objective of each agent in the network is to find the sequence of calls which maximizes overall utility.

We denote the ordered sequence of calls of agent i by Ω , where the k th call in the sequence is Ω_k . We use $\Omega_{1:k}$ to refer to the first k calls in the sequence. We will refer to the number of calls in the sequence by $K = |\Omega|$. We suppress the dependence of Ω on i and t wherever possible for expositional clarity. Conditional on being the k th call in agent i 's sequence, the utility of calling agent j is given by:

$$U_{ijk} = f(x_i, x_j, x_i x_j) + g(\Omega_{1:k}) + \epsilon_{ijk}. \quad (1)$$

As in Jackson and Wolinsky (1996), the utility for each call depends on not only the caller's and receiver's characteristics, but also the interaction of these characteristics (the link "synergy"). This allows us to evaluate the extent to which callers seek diversity by calling agents with different characteristics than their own.

The second term in the utility specification, $g(\Omega_{1:k})$, reflects the second source of benefits from heterogeneity in usage. An agent may value the ability to make calls to people with a range of characteristics, as opposed to just repeatedly interacting with the agent who gave the highest initial call utility. For example, the second call an agent makes may be motivated by the need to acquire information that augments the information gathered in the first call. If an agent uses communication for information gathering, they will not necessarily benefit from speaking to the same person twice. Alternatively, the agent may have satisfied their information-gathering needs with the first call, and has moved on in the second call to processing another task with different informational requirements. We call this desire for diversity within a calling sequence "dynamic diversification". The term $g(\Omega_{1:k})$ captures these effects by allowing the marginal utility of calling agent j to depend arbitrarily on all previous calls. This step builds on a growing literature on the estimation of state-dependent discrete choice models in the marketing literature. Typically the state-dependence is expressed as habit formation or variety-seeking, and the current choice is a

probabilistic function of the purchase history. For example, Chintagunta (1999) presents an empirical framework based on the hazard model for dealing with variety-seeking in customer shopping behavior in scanner panel data.

Therefore, we evaluate two potential ways a taste for heterogeneity can affect interactions across a network. The first is a baseline effect. From the beginning, agents may receive utility from interacting with people who have different characteristics to themselves. The second effect is dynamic in origin. After interacting with one type of agent, it may become attractive to interact with another agent who has different characteristics to add diversity to information received.

The agent’s optimization problem is to find the sequence of calls which maximizes overall utility:

$$\max_{\Omega} \sum_{k=1}^{K=|\Omega|} U_{ijk}. \quad (2)$$

Each agent makes calls until the best marginal call has a negative utility. We assume that the $g(\cdot)$ function is invariant to the order of previous calls. This assumption rules out time-specific nonlinearities between any two (or more) calls. This assumption simplifies the optimization problem in Equation 2, since only the composition, and not the specific ordering, of a calling sequence matters in evaluating the utility function. If this assumption is not made, then agents may be strategically forward-looking in their choice of when to time certain calls. This assumption of time invariance is crucial for our empirical strategy.

3.3 Transition between States: A Dynamic Model of Technology Adoption

The second component of our model is the adoption decision of agents currently outside the network. This governs the transition between adoption states in our model. In each period, an agent can choose to adopt the technology or not, which we denote by $adopt_{it} \in \{0, 1\}$.

Agents can use the technology immediately upon adoption. Once agents are in the network, they cannot divest themselves of the technology. Therefore, if an agent adopts the technology in one period, they adopt the technology forever. This seems reasonable, given that the option value of holding the technology is always positive in our model.

When deciding whether to adopt, each agent weighs the costs and benefits. If an agent adopts, she can expect to use that technology to communicate with others in the network, both today and in the future. Her payoff function is a function of the state vector, her adoption decision, her expected communication decisions, her own characteristics, and the characteristics of everyone else in the network. As in Farrell and Saloner (1985), the benefits of adopting a network technology consist of both the network benefit, a stream of expected discounted calling sequence utilities, and the stand-alone benefit, which we denote by Γ . In our empirical application, the stand-alone benefit is the ability to use the videoconferencing technology to watch TV.³ Each agent discounts future benefits according to the common discount factor, β .

The costs to installing this technology for the agent consist of the time spent setting it up and learning how to use it. The firm bears all monetary costs. To reflect this installation cost, we assume that adopters have to pay a one-time up-front fixed cost of F_i . We assume that this cost F_i does not change after the agent has made their initial draw and is private information to the agent. This private information reflects persistent agent-specific heterogeneity in both learning costs and the technology's stand-alone benefits.

The stand-alone benefits and adoption costs are not separately identified in the model. To see this, suppose that there were no network benefits but only stand-alone use benefits. Then agents will be indifferent to adoption if and only if:

$$-F_i = \sum_{t=0}^{\infty} \beta^t \Gamma_i = \frac{1}{1-\beta} \Gamma_i. \quad (3)$$

³We also observe agents that “call” themselves; we leave the interpretation of such behavior to the reader.

For any Γ_i , we can find a F_i such that the agent is indifferent to adoption. Therefore, without loss of generality, we will assume that $\Gamma = 0$.

Given beliefs about the evolution of the network, we can write out the technology adoption decision as an optimal waiting problem. Intuitively, the agent weighs the benefits of adoption now against the opportunity costs of doing so. The opportunity costs here encompass both the outside option and waiting to adopt in a future period. An agent will adopt today if the following inequality is satisfied:

$$U_{it}(s_t) - F_i + \beta EV_{i,t}(s_t; s_{i,t+1} = 1) \geq \beta EV_{i,t}(s_t; s_{i,t+1} = 0). \quad (4)$$

We have explicitly written the continuation values as a function of s_t to emphasize that the agent is making predictions about the future evolution of the network. This expectation raises three important dynamic considerations. First, the agent may have a high draw on F_i , which gives an incentive to wait for the installed base s_t to be larger to cover the fixed costs. A second countervailing effect is that agents anticipate that their adoption now may spur other agents to adopt in future periods. Such forward-looking sequential behavior may help reduce the coordination failure in technology adoption, as pointed out by Farrell and Saloner (1985). This second effect has a wide range of potential outcomes, from nudging inframarginal non-adopters a little bit closer towards adoption without visible effect, to generating an entire cascade of adoptions in future periods. The agents in our model balance these two effects against each other and the payoff of the outside option, when making an optimal choice about whether to enter in the current period. Third, there is an option value in not adopting in this period.⁴ Even though all agents are playing the same equilibria, and have rational expectations about the expected evolution of the network, there is still variance in who is going to join each period, and so there is potential value in waiting to see how the

⁴We are grateful to Dan Akerberg for bringing this to our attention.

network has evolved.

4 Estimation

Computational limitations imposed by the burden of explicitly computing the equilibrium to the theoretical model prevent a straight likelihood approach. Therefore, our empirical strategy follows the approach of Bajari, Benkard, and Levin (2006), who advocate a two-step approach for estimating dynamic games. In the first step, we recover reduced-form policy functions which describe the equilibrium strategies followed by each agent as a function of the state vector. In the second step, we project these functions onto an underlying dynamic model of technology adoption choice and usage. In this manner, we recover consistent estimates of the underlying parameters which govern the process of network evolution and utilization.

There are two separate policy functions in the first stage. The first reduced form addresses the question of how the network will be used by agents who have already adopted the technology. This function describes how many videoconferencing calls these agents will make, and to whom, as a function of the network’s characteristics. We propose a new “simulated sequence estimator,” to capture the relevant aspects of the calling decision, explicitly accounting for both the length and composition of the sequence of videoconferencing calls. The second reduced form estimates the factors that measure the propensity to join the network, given the number and composition of current users.

Throughout our estimation, we focus on region-function-title subtypes rather than individuals. The reasoning for this is two-fold: first, estimating pair-specific connection parameters will quickly exhaust the degrees of freedom in our data set. Second, since we never observe divestment and there is insufficient variation in the calling patterns among agents in the network at the individual level, we will at best be able to estimate one-sided parameter

boundaries. While recent econometric work has shown how to estimate these unbounded parameters, we cannot calculate counterfactuals. For these reasons, we focus on subtype-specific policy functions instead of individual-level functions. We denote the set of these subtypes by M . Given our four regions, four functions, and four titles, we have a total of 64 subtypes.

4.1 Simulated Sequence Estimator

We estimate a reduced-form policy function to capture how agents use the network once they adopt the videoconferencing technology, using our simulated sequence estimator. For a given calling sequence, Ω , of length K , the simulated sequence estimator splits the calling sequence problem into two parts by exploiting the following identity:

$$Pr(\Omega, K) = Pr(\Omega|K)Pr(K). \quad (5)$$

We can then separate the estimation of the composition of the calling sequence from the length of the sequence.⁵ After taking logarithms in Equation 5, we obtain:

$$\ln Pr(\Omega, K) = \ln Pr(\Omega|K) + \ln Pr(K). \quad (6)$$

The simulated sequence estimator first estimates the composition of the call and then estimates the parameters which determine the number of calls. While the composition of the calling sequence is nonparametrically identified, for computational ease we assume that the utility of agent i making the k -th call in the sequence to agent j is determined by the following equation:

$$U_{ijk} = \hat{\theta}_1 x_j + \hat{\theta}_2 x_i x_j - g_j(\Omega_{1:k}; \hat{\theta}_3) + \hat{\theta}_4 x_i + \epsilon_{ijk}, \quad (7)$$

⁵For a related idea in a dynamic optimization context, see Hendel and Nevo (2005).

where ϵ is distributed Type-I extreme value. The $g(\cdot)$ functions capture the change in utility across various subtypes as a function of previous calls in the sequence $\Omega_{1:k}$. We index this function with a subscript j to emphasize that penalties arising from previous calls to subtype j only enter the utility of that same subtype. We assume that $g(\cdot)$ takes the following functional form:

$$g_j(\Omega_{1:k}; \hat{\theta}_3) = \left\{ \sum_{m=1}^M 1(m=j) \hat{\theta}_{3m} \eta_m \right\} + k, \quad (8)$$

where η_m is the count of the number of previous calls in the current sequence to type m . In this specification, previous calls to each type m generate disutility of calling that same type again. To see this, note that the marginal disutility of calling a certain type enters in differentially across receiver types. Differences in θ_{3m} shift the marginal utility of calling a type the agent has called in the past. This marginal difference grows linearly in the number of previous calls. It is this variation in marginal utility across the calls within a sequence that generates dynamic diversification. The normalizations we make are that the error term has unit variance and that the value of not making a call is zero. The assumption that the agent will make calls until the best marginal call gives negative utility serves as the location normalization.

The parametric assumption on the error term generates the usual logit probability of observing a call from agent i to agent j as the k th call of a sequence:

$$Pr(\Omega_{ijk}; s_t, \hat{\theta}) = \frac{\exp(U_{ijk}(\hat{\theta}))}{\sum_{j' \in s_t} \exp(U_{ij'k}(\hat{\theta}))}. \quad (9)$$

Note that the outside option does not enter the probability of a call as it usually does in discrete choice models, as we are conditioning on the length of the sequence. Computationally, the estimation proceeds by finding parameters to maximize the probability of observing each call in the sequence in that order. The ordering of the sequence is valuable in identifying the parameters of the decay functions, as the conditional probability of each call in the sequence

depends on the order of the calls made before it. Specifically, the relative frequency with which we observe two calls to the same subtype in a given sequence contains statistical information about the magnitude of the decay function for that subtype. So while the overall utility of a given call sequence is invariant to the ordering of the calls in that sequence, we can exploit variation in relative frequencies of given runs of calls to more precisely identify the decay functions.

The second step in the simulated sequence estimator is to find parameters which govern the length of the sequences. To solve for these parameters, we use a simulated method of moments approach. As mentioned above, the own-type utility parameters drop out of the conditional utility function. This is what allows us to separate the estimation into two steps. Though theoretically it is straightforward to jointly estimate these two steps, the computational burden of doing so means that we estimate the two steps separately: the second step requires thousands of simulated sequence draws for each guess of the parameter vector. We correct the standard errors through a two-step bootstrap. Given the calling parameters from step one, we generate a large number of calling sequences. Each sequence is computed for a given network by assigning simulated utilities to each potential receiver in the network. The agent then compares among these utilities and calls the receiver with the highest utility. This process stops when the highest utility is dominated by the outside option, which is normalized to be equal to zero. This process is inherently stochastic, however, as the vector of shocks in the receiver utilities can generate variance in the length of the simulated calling sequence. As a first step in recovering the parameters entering an agent of subtype m 's own-characteristics, $\hat{\theta}_4$, we compute the expected sequence length:

$$\hat{K}_m = \frac{1}{N_s} \sum_j^{N_s} K_{jm}(\hat{\theta}_4), \quad (10)$$

where N_s is the number of simulations used to form the expectation. We then form the

following estimator:

$$\hat{Q}(\hat{\theta}_4) = \frac{1}{T} \sum_{t=1}^T \left(\frac{1}{M_t} \sum_{m=1}^{M_t} \left(\frac{1}{N_{mt}} \sum_{i_{mt}=1}^{N_{mt}} \left(K_m(\hat{\theta}_4) - \Omega_{imt} \right)^2 \right) \right). \quad (11)$$

To be clear, N_{mt} is the number of agents of subtype m who are in the network at time t , M_t is the set of subtypes at time t , and T is the time of the final observation. Intuitively, we are matching the expected length of each subtype against the expected length for that subtype, which is a function of the unknown parameters and the composition of the network in each period. Identification follows from the fact that these parameters increase the utility level of each call uniformly across all calls. Therefore we can identify type-specific own-characteristic parameters by matching the variation in the expected sequence length for each subtype against the predicted sequence length.

4.2 Estimating the Adoption Decision

The second policy function that we recover from the data governs the choice of videoconferencing technology adoption. Our goal is simply to estimate the proportion of agents of a given subtype that will adopt the technology in a given period, conditional on the current composition of the network. One natural interpretation of an agent’s reaction to the composition of the network is that it results from a network effect. A wide variety of behaviors, however, could explain a positive reaction to others’ adoption decisions. This includes information diffusion, improved technical support, or emulation of other employees. Therefore, though we refer to the reaction of an agent as a network effect, it is important to be clear that this is a convenient term for a wide variety of behaviors and we do not explicitly identify a causal mechanism. Causal identification of network effects for this data has been studied in detail in other research. Tucker (2006a) and Tucker (2006b) use the videoconferencing technology’s stand-alone use of TV-watching as an exogenous shifter for measuring causal

network effects in adoption. Though the focus of these papers are different they suggest that it is reasonable to think that just under half of the correlation in adoption should be thought of as a strictly causal network effect, where users adopt because another users' adoption allows communication. Our model encompasses many reasons employee's demand may be interrelated because we do not interpret any of our parameters causally.

Our estimation goal is amenable to completely nonparametric estimation. The ideal approach is to simply count the number of times that an agent of a subtype adopts the technology given every combination of the network. The data, however, preclude such an approach, as the majority of network states are never observed. Instead we adapt a simple parametric framework that follows our modeling assumptions.

We have assumed that the fixed cost of adoption is agent-specific private information that does not change over time. The distribution of fixed costs varies across subtypes and is uniform within subtypes. These two assumptions imply that the proportion of people within a subtype who adopt is nondecreasing in the size of the network. Our approach is to estimate the proportion of a given subtype who has joined the network at any point as a function of the current and lagged state:⁶

$$Proportion(adopt_i = 1|s_t, s_{t-1}) = \hat{\lambda}_1 x_i + \hat{\lambda}'_2 x_i \nu_t + \hat{\lambda}_3 x_i \nu_{t-1}, \quad (12)$$

where ν_t is a vector of counts of 12 subtypes (four regions, four titles, four functions) currently in the network. This specification is similar to how we estimate the calling functions. Due to limited observations on entry decisions, we cannot estimate this policy function by interacting all 64 subtypes on counts of all 64 subtypes in the network. Instead, we impose that each component of an agent's subtype has the same marginal effects in adoption probabilities

⁶Although we estimate the proportion of a given subtype, it is possible to directly convert these changes into adoption probabilities at the individual level. For example, if the model predicts that the proportion increases from 10 percent to 20 percent, and there are 100 total adopters, the probability any given individual joins is $(0.2-0.1)*100/(1-0.1)*100$, or 11.11 percent.

across all other subtypes which share that characteristic. For example, all functions and titles in Asia share the same region-specific heterogeneity in the manner they react to a change in the state space of total adoption in Europe. The interpretation of this approach is that the region, function, and title characteristics are additively separable causal determinants of behavior.

The inclusion of lagged polynomial terms of ν_{t-1} is to control for selection. In later periods, if an agent has not adopted, the model implies that they received a high fixed cost draw. The inclusion of polynomials of ν_{t-1} means that higher values of the expected benefit are required for an agent to join the network, all else equal. As Equation 12 is simply a reduced-form that provides a description of how the system evolves, we can successfully control for selection when we project these policy functions onto the underlying model, which we describe next. The lagged terms are sufficient for determining the degree of this selection, as the value to network is always weakly increasing in the number of adopters, regardless of type. Additionally, we impose the constraint the sign of the coefficients on current state variables has to be positive, and the sign on lagged state variables has to be negative. The reasoning for this is that, all else equal, increasing the number of people in the network has to make it weakly more attractive to a potential adopter. This also implies that, all else equal, an agent who decided not to join a larger network in the last period must have higher fixed costs, and thus a lower probability of adopting, than an agent who did not join a smaller network in the previous period.⁷

One concern with our approach is that we are not controlling for unobserved heterogeneity in the form of shocks which are common across groups of individuals, e.g. all agents who are in administration and also located in Asia. Bajari and Hong (2006) have shown that identification of causal effects, as identified by region, function, and title, is still possible in this setting. The intuition for this argument is that time-subset fixed effects can control for

⁷The implementation of such constraints is easy with standard optimization packages such as Matlab.

time-varying shocks to a particular group, as there are shifters that change the propensity of one agent in that subset to join the network separately from other agents in the network. For example, administrators in Asia still vary within that group by title. Variation in the network composition shifts the propensity of agents in this group differently depending on title. This variation is enough to control for group-time specific shocks to the adoption decision. The drawback to this approach is that it soaks up almost all of our identifying variation in the data, but in principle our approach is robust to a wide range of unobserved heterogeneity.

4.3 Estimating the Fixed Costs of Adoption

Once we have policy functions governing the adoption decisions of agents outside the network and the calling decisions of agents on the inside, we have enough information to simulate the evolution of the network and assign payoffs to any given agent who has adopted the technology.

In addition, for each state of the network we can simulate the volume and pattern of all calls made in a given period using our call policy parameters. We simulate the process described in Equation 7 and compute the expected value of network use. By simulating this sequence many times for a given agent, we form a consistent estimate of the expected utility of using the network in that period. Performing this calculation for a large number of periods in the future, discounting properly, will give us a consistent estimate of the discounted present value of using the network to an agent who adopts in the current period.

To recover the fixed costs of adoption for each agent subtype, we exploit the structure of the optimal adoption choice given by Equation 4. The optimal adoption decision is a function of today's expected values and the continuation value of not adopting. The continuation value of not adopting is a complex object: it is derived from the optimal decision rules for all possible states of the network in the future. In principle, this is a computable object.

One solves out a nested sequence of dynamic programming problems, starting at a distant point in the future where the expected value of the network is relatively unchanged and working backward. However, this is exceedingly burdensome, as each step in the computation requires solving out the expected value of a nest of dynamic programs. These nested dynamic programs capture the three dynamic effects discussed above: expected changes in the network outside my own adoption decision, my influence on other agents to adopt the network, and the option value of waiting to see how the network evolves. Instead of solving out the explicit continuation value of not adopting, we consider a related approximation to that problem.

Suppose that the option value of waiting is zero beyond a finite number of periods, denoted by T . This will hold if there is no uncertainty about the evolution of the network in the next period beyond T , e.g. the network is already hit the steady state. Under this assumption, the agent's decision boils down to adopting in the current period, one of the periods up the point where the option value is zero, or never adopting at all. At time zero, we can write this as:

$$EV_0 \geq \max \{0, \beta EV_1, \beta^2 EV_2, \dots, \beta^T EV_T\}. \quad (13)$$

In the case where adopting immediately dominates adopting in the future, agent i of subtype m would not adopt is if the draw on fixed costs outweighs the benefits of adopting. The probability of this happening is a function of the distribution of the fixed costs:

$$Pr(adopt_{im} = 1 | s_t) = Pr(EV_0 - F_i > 0) \quad (14)$$

$$= Pr(EV_0 > F_i) \quad (15)$$

$$= CDF(EV_0; \mu_m, \sigma_m^2). \quad (16)$$

Although this distribution is nonparametrically identified, for simplicity we assume that this

distribution is normally distributed with subtype-specific mean μ_m and variance σ_m^2 . These parameters are the unknowns that we wish to infer from the agents' behavior. Note that we already have an estimate of the term on the left-hand side: $Pr(adopt_{im} = 1)$. This is exactly the adoption policy function found in the first step. Also, we have computed EV_0 by using our policy functions to simulate the evolution of the network and to assign discounted present values derived from the usage of that network. The only unknowns are the parameters of the CDF. To obtain estimates of those unknowns we simply match empirical probabilities from the policy function to their computer counterparts:

$$\hat{Q}_n(\mu_m, \sigma_m^2) = \left(\sum_{i=1}^{N_s} Pr(adopt_{im} = 1) - \Phi(EV_0; \mu_m, \sigma_m^2) \right)^2. \quad (17)$$

Changes in the way that the probability of joining the network changes traces out the shape of the underlying distribution of fixed costs. By matching the curvature of this change in the probability of adoption against the implied model parameters, we uniquely identify the normal distribution that best fits the data.

In future periods, the distribution of fixed costs in the agents who have not adopted is not the same as the unconditional distribution in Equation 15. Since adoption is an absorbing state, it is sufficient to rewrite the probability of adoption in terms of the following conditional probabilities:

$$Pr(adopt_{im} = 1|s_t) = Pr(EV_t - F_i > 0|F_i > \overline{EV}) \quad (18)$$

$$= Pr(EV_t > F_i|F_i > \overline{EV}) \quad (19)$$

$$= CDF(EV_t|F_i > \overline{EV}; \mu_m, \sigma_m^2) \quad (20)$$

$$= CDF(EV_t; \mu_m, \sigma_m^2) - CDF(\overline{EV}; \mu_m, \sigma_m^2). \quad (21)$$

Here we have denoted the highest previous expected value of joining the network by \overline{EV} .

Intuitively, we know that the agents who are still in the network have fixed costs at least as high as this value, since otherwise they would have joined in a previous period. The modification of Equation 17 is straightforward:

$$\hat{Q}_n(\mu_m, \sigma_m^2) = \left(\sum_{i=1}^{N_s} Pr(adopt_{im} = 1) - \Phi(EV_0; \mu_m, \sigma_m^2) - \Phi(\overline{EV}; \mu_m, \sigma_m^2) \right)^2. \quad (22)$$

It is also possible that the discounted expected network benefits of the future are greater than the current expected benefits. In this case, we would observe an agent i of subtype m joining the network if:

$$EV_0 - F_i \geq \beta^{t^*} (EV_{t^*} - F_i), \quad (23)$$

where t^* solves the right-hand side in Equation 13. Rearranging terms, we obtain:

$$\frac{EV_0 - \beta^{t^*} EV_{t^*}}{1 - \beta^{t^*}} \geq F_i. \quad (24)$$

Writing this relationship in terms of the CDF of F_i , we obtain:

$$Pr(adopt_{im} = 1) = \Phi\left(\frac{EV_0 - \beta^{t^*} EV_{t^*}}{1 - \beta^{t^*}}; \mu_m, \sigma_m^2\right). \quad (25)$$

For cases where the expected value of waiting in the future is positive, we set up and solve an analogous estimator like that in Equation 17:

$$\hat{Q}_n(\mu_m, \sigma_m^2) = \sum_{i=1}^{N_s} \left(Pr(adopt_{im} = 1) - \Phi\left(\frac{EV_0 - \beta^{t^*} EV_{t^*}}{1 - \beta^{t^*}}; \mu_m, \sigma_m^2\right) \right)^2. \quad (26)$$

In periods beyond the first, we need to correct for selection. This time we keep track of the maximal values on the left-hand side of Equation 24, which denote by \widehat{EV} . The appropriate

estimator for $t > 0$ is then:

$$\hat{Q}_n(\mu_m, \sigma_m^2) = \sum_{i=1}^{N_s} \left(Pr(adopt_{im} = 1) - \Phi \left(\frac{EV_0 - \beta^{t^*} EV_{t^*}}{1 - \beta^{t^*}}; \mu_m, \sigma_m^2 \right) - \Phi \left(\widehat{EV} \right) \right)^2. \quad (27)$$

In our estimation, we simply solve out the expected values of waiting to adopt in future periods until the value of waiting is dominated by earlier choices or the outside option. This happens quite quickly given the discount factor and the relative speed with which the network stabilizes. Once we have these expected values in hand, we solve the right-hand side of Equation 13. If the current period dominates waiting for future periods, we apply the estimator defined in Equation 17. Otherwise, we match the empirical moments in the data using the estimator defined in Equation 26. We repeat this process for each subtype to obtain estimates of the distribution of fixed costs for each different type of agent in the data.

We use a discount value β equal to 0.9. This high discount factor reflects that risk of exit for employees in this industry was high - there was annual employee turnover of around 8% in the three and a half year period that we study.

Our assumption about the option value of waiting to be zero after some period T is clearly an approximation. However, given that in our application most of the uncertainty about the evolution of the system has been resolved by period 10, this seems a reasonable approximation to the proper infinite-horizon solution.

4.4 Multiple Equilibria

One of the concerns of the network effects literature has been dealing with the potential for multiple equilibria in outcomes. One advantage of our empirical approach is that we recover the equilibrium actually played in the data. Furthermore, since there is only one network, we can be assured that the equilibrium that we estimate from the data is the only equilibrium

being played. To our knowledge, this is unique among applications of the BBL framework, as we do not have to confront the possibility of multiple equilibria across markets, as in Ryan (2006).

4.5 Monte Carlo Evidence

To evaluate the efficacy of our estimation approach, we ran a simple Monte Carlo experiment. The results, along with the true parameters, are shown in Tables 1 and 2. The Monte Carlo evidence suggests that our estimator precisely estimates the calling parameters, even including the decay rates.

We discovered that the performance of the single-step estimator was poor even for large samples in the Monte Carlo. The intuition behind this is clear when considering the identification of the two-step approach. In the first step we estimate connection utilities and decay rates. The connection utilities are identified off even just one call, as the probability of making a call between types reveals the magnitude of the connection synergy once we have normalized the error term. The decay rates are then identified from within-sequence variation in the ordering of the calls—having called a certain type in the past, the conditional probability of calling that type in the future, holding connection utilities constant, depends on the decay rate. By comparing those conditional probabilities across a large sample of calls, we can precisely identify the decay rates jointly with the connection utilities.

Once these connection parameters are in hand, we can estimate the parameters which govern the length of calling sequences. The identification of the intercept and decay parameter governing how fast utility decreases as a function of the number of calls is a bit subtle. In infinite samples, the intercept is identified off the frequency of calling sequences with no calls. The reason for this once we know the calling parameters is it possible to put a probability that no one makes a call as a function of the intercept. All else equal, a higher intercept leads to a lower set of sequences with no calls in a period. Once the intercept is

identified from this probability, it is straightforward to match the decay slope parameter to the average number of calls that a type makes. We find that two-step process works well in the Monte Carlos we have run.

Since our two-step method is inefficient, we are interested in exploring using a one-step method, where the probability of not making a call is incorporated in the likelihood of observing a given sequence of calls. The issue with this approach is not econometric—in infinite samples the same identification arguments as above can be made in the simultaneous setting. Rather, the problem is practical. In a finite sample the intercept is going to be poorly identified off of sequences with zero calls if the model suggests this happens infrequently. Given that the intercept is poorly identified, doubling both the intercept and decay slope parameters give very similar empirical predictions, as both parameterizations give the same number of average calls in simulations. The difference is that the variance of calls around that average is estimated incorrectly; with enough noise in the data, the estimator cannot discern between the truth and linear transformations of the true intercept and slope parameters. In Monte Carlos we performed, we obtained suites of estimates which were all biased upward roughly by a factor of two.⁸ Therefore, despite the loss of efficiency, we estimate the model in two steps.

One additional drawback of the two-step method is that the second step involves matching sequence lengths against predicted lengths. With a finite number of agents in the model, this introduces flat spots into the objective function, a problem well-known from the discrete choice literature. A simple fix to this problem is the use the Laplace-type estimator of Chernozhukov and Hong (2003), a method using Markov Chain Monte Carlo techniques to help perform inference on objective functions with local minima and flat spots.

⁸Results available upon request from the authors.

5 Results

This section reports the results of our estimation. There are two main sets of results. The first are the calling parameters which capture the per-period payoffs from adoption. The second are the fixed costs, which determine adoption decisions and the transition between states in our model.

5.1 Call Utilities

We use observations on 463,806 calls from February 2001 to August 2004 to estimate the utility parameters in Equation 7. Tables 4 through 6 display the results of our estimates. In general, all else equal, agents prefer to call other agents with similar characteristics. For example, employees in Asia prefer to call other employees in Asia, and employees in the US prefer to call other US-based employees. By contrast, in the UK employees prefer calling employees in Europe to calling other UK-based employees. We speculate, therefore, that the propensity to call within-regions could be influenced by time zones. Employees' work hours in the US and Asia barely overlap, but the work hours of British and European employees overlap greatly.

Employees on average exhibit a preference for calling employees in similar functions to themselves. In all cases an employee prefers to communicate with someone within their own functional group than outside it. Given the perception that the research, sales and trading functions should support each other in a banking environment, it is striking that all such employees prefer to call administrators rather than anyone in one of their sister functions. This might reflect the fact that the videoconferencing is an internal firm technology, and that employee compensation is based on the ability to sell, research and trade financial products for outside clients, rather than communicating information to each other.

The estimates on preferences for calling across the hierarchy suggest that this technology

is being used to pass information within a rung of the hierarchy rather than transmitting information up or down it. Associates are most likely to call each other and become decreasingly likely to call with the number of rungs the receiver is above them. Managing Directors are similarly most likely to call each other and less likely to call employees further down the ladder of command. These results augur against the technology being used successfully for monitoring.

The results for the parameter $\hat{\theta}_3$ which captures the role of the dynamic decay rates are displayed in Table 3. The taste for dynamic diversification is strong—the decay rates are large enough to have a significant effect at the margin of calling the same group twice in a row, especially with respect to workers at the associate level, workers in research, and workers in Asia.

5.2 Fixed Costs of Adoption

Tables 7 through 10 display the results of our fixed cost estimates from estimating Equations 17 and 26. There are a few patterns to highlight. First, out of the four regions, US-based employees have on average the highest fixed costs of adoption. Second, fixed costs seem to be declining with position in the hierarchy - Managing Directors seem to have higher net stand-alone benefits or lower net costs than associates. Finally, there are within-group differences in fixed costs. For example, administrators have a lower fixed cost of adoption than employees in other functions in Asia and the Europe but not in the US or UK.

If we compare these results with the results for calling choices in Tables 4 through 6, we see that it is not the case that the employees whom most employees preferred to call had the lowest fixed costs of adoption. Instead, for example in the case of Managing Directors of Research in the Europe, this was on average a group whom callers received high utilities from calling. However, they also had some of the relatively highest fixed adoption costs.

5.3 Alternative Specifications

To examine the sensitivity of our results to timing assumptions we have made, we run several robustness checks. Our decision to make the relevant calling period one month, starting at the first day of each month, is somewhat arbitrary. Although several processes in the firm are scheduled on a monthly basis (e.g. payroll, monthly market forecasts, discussion of economic indicators), it is reasonable to assume that the relevance of previous calls could be shorter or longer than a month, and could be unrelated to the first of the month. To address these issues, we group calls into two additional periods: two weeks and two months. In the two week samples, calls are split within each month depending on whether they take place before or after the 15th day. By comparing estimates across the two week periods we can examine the sensitivity of our estimates to our assumptions about when calling sequences “reset” and start over. We estimated the calling model on each of the three period lengths. We then evaluated the conditional log likelihood at each of those three parameter vectors on one month of data. We perform a likelihood ratio test, which is asymptotically distributed χ^2 with 57 degrees of freedom. The critical values are equal to 75.62, 84.73, and 95.75 at the ten percent, one percent, and one-tenth of a percent levels, respectively. The test statistics when restricting the parameters to be from the two-week data and the two-month data are 21.0 and 35.0, respectively. Therefore, we fail to reject the hypothesis that the parameters are equal at all three levels of significance.

We also estimated an adoption model without any heterogeneity in the adoption policies across subtypes. Equation 12 was replaced with:

$$Proportion(adopt_i = 1 | s_t, s_{t-1}) = \hat{\lambda}'b(x_i), \quad (28)$$

where $b(x_i)$ is a set of b-spline basis function. This flexible series estimator does very well at matching the overall adoption levels, but obviously does quite poorly in explaining the

cross-subtype variation in adoption rates evidenced in Figure 6.

6 Policy Experiment

Carr (2003) documents that the typical company spends 3.7 percent of its revenues on IT. A challenge for managers is to ensure that their employees actually use the firm’s technology investment to its full advantage. The videoconferencing context that we study is unusual because adoption decisions were decentralized to employees. A far more common challenge facing IT managers is how to get employees to start using a costly technology which has already been installed for them. The focus of our policy experiments, therefore, is how best to encourage actual interactions using a new IT technology. Consequently, in our discussion, we interpret “adoption” in our data as the equivalent of the more general idea of “activation”, the active usage of a new technology by an employee.

As discussed by Liebowitz and Margolis (1994), network owners can prevent coordination failure if they offer targeted incentives to reflect the network benefits to network participants brought by new adopters. In the presence of network effects which are heterogenous in interactions, however, the optimal policy is more complex, because each potential network entrant should be compensated for the varying positive network effects they have for a large set of different users. Since firms rarely engage in personalized subsidies and the information burden of an optimal policy would be large, we evaluate two possible “rule of thumb” technology management policies: a targeted policy where a single subtype joins the network, and a uniform policy where a few agents from every subtype join the network. The intuition here is that the firm will install the physical hardware and provide whatever training is necessary to overcome the fixed costs of adoption for a selected set of agents under each policy.

The first policy we consider is where the firm picks one subtype to adopt/test the technol-

ogy first. This resembles the way that many firms roll out new IT technologies. IT managers usually pick this initial seed from employees who are similar by virtue of their operational similarity and location. Therefore we conduct a policy experiment where the starting network is seeded with all 112 research associates located in the United States. This group constitutes the single largest subtype within the firm, and may be considered a natural place to seed the network, as agents in the United States generally have high adoption costs.

The second policy takes a diffuse approach to adoption. Here the firm spreads 112 installations across the entire set of subtypes. The idea here is that diversity increases the value of the network, and seeding the initial network with a broad range of types may most efficiently jump-start the growth of the network. Given there are 64 subtypes, there are 16 groups which start with only one agent. We choose the last 16 types, which correspond to all subtypes located in the United States.

In each counterfactual simulation, we start by seeding the initial network in accordance with the desired policy. Starting at time zero, the network is then simulated forward for fifty months. This amount of time is sufficient to allow the network to achieve the steady state where it is no longer growing at a significant rate. Also, the discounted present value of utility of months more than 50 periods from now is essentially zero for the discount rate of 0.9 that we use. To simulate the evolution of the network, we draw uniform random variables for each potential adopter, and check these against each agent's corresponding subtype-specific policy function. If the policy function indicates that the agent will join the network, we draw a sunk entry cost from the associated truncated normal distribution. After determining the evolution of the network in that period, we then calculate the sum of expected utilities for all agents in the network. This calculation is greatly simplified by the fact that it is possible to do this on a subtype basis, rather than agent by agent. The results of the two policy experiments and a baseline comparison against the empty starting network are shown in Table 11. Figures 7, 8 and 9 contrast the results graphically for the total adoption, calls and

average utility.

The first result concerns the average number of phone calls. Across each specification, the undiscounted average number of calls in each month is roughly similar, with slightly higher amounts in the baseline and targeted policy than in the uniform policy.

The maximum number of adopters is considerably higher in the targeted case than in the baseline or uniform cases. This occurs because the adoption of that group is considered particularly valuable to the overall network. Conversely, the maximum number of adopters is slightly lower in the uniform case than the baseline. This difference reflects the difference in network evolution paths that the three policies take. The results for the uniform policy suggest that a broad-based adoption process may be highly inefficient, as some agent subtypes are inefficiently forced to join the network.

We calculate the expected discounted monthly utility for each subtype across the three policies. We report the mean utility for the population of agents, and also report utilities by quartiles. The uniform policy improves over the baseline case, but only marginally. We have assumed the agents who start out in the initial seeding network have paid no fixed costs, while those in the baseline case have. If utility were monetized, it is possible that the costs to the firm are dominated by the alternative of just allowing agents to pay their own adoption costs. Accounting for these costs reinforces the idea that decentralized adoption allows for the efficient agents to join the network through a process of self-selection.

The case is significantly different for the targeted policy. In this case, there is an increase of over 6% in present discounted utility (discounted at $\beta = 0.9$) for the mean type. This increase is also reflected across the other quartiles of the utility distribution. If the objectives of the firm are positively related to the utility of the agents, then this policy can have a significantly positive effect from the firm's perspective. In addition the utility gains appear to shift the utilities equally across subtypes in the firm, even in the targeted case. This is not necessarily intuitive, as it could have been the case that a targeted case would lead to

more dispersion in the utility levels across subtypes, whereas the simulations actually show a small decrease in relative dispersion.

The last two panels in the table illustrate inter-temporal differences in adoption rates and network usage. We assume that, everything else being equal, the firm would prefer to have a given number of phone calls or agents in the network sooner rather than later. We report the discounted sums of users who have adopted the network in a month and the number of calls they have made, using two contrasting potential monthly discount rates for the firm. The differences are quite stark: the uniform policy makes marginal improvements over the baseline case, while the targeted policy dominates along both dimensions. When $\beta = 0.9$, user counts increase by 20% and calls increase by 23%. In an investment bank where the opportunity cost of time is high, these results suggest that the dominating policy is to target a specific group for initial adoption.

7 Conclusion

This paper explores how heterogeneity in users' adoption costs, network effects and tastes for a diverse network affects how network technologies diffuse. We estimate the first empirical model to combine dynamic technology adoption choice with adopters' subsequent interactions. We obtain parameter estimates from a detailed dataset on the adoption and subsequent usage of a videoconferencing technology in a large investment bank. This dataset has the advantage that it allows us to study adoption at the micro as opposed to the aggregate level this allows greater understanding of heterogeneity. We find substantial heterogeneity in adoption costs, benefits from the network, and a strong willingness-to-pay for a diverse network with many types. Our estimates of heterogeneity in adoption costs, network effects and interactions provide guidelines about rules of thumb that network operators should use when trying to jump-start growth of their network technology. In general, to have the biggest

impact on the evolution of the network, firms should jump-start the diffusion of the network by targeting individuals who have high adoption costs and with whom other users want to interact.

References

- Akerberg, D. A. and G. Gowrisankaran (2006). Quantifying equilibrium network externalities in the ACH banking industry. Forthcoming in *RAND Journal of Economics*.
- Bajari, P., C. L. Benkard, and J. Levin (2006). Estimating Dynamic Models of Imperfect Competition. Forthcoming in *Econometrica*.
- Bajari, P. and H. Hong (2006). Semiparametric estimation of a dynamic game of incomplete information. University of Minnesota Working Paper.
- Bajari, P., H. Hong, and S. Ryan (2006). Identification and estimation of a normal form game. University of Minnesota Working Paper.
- Bass, F. (1969, January). A new product growth model for consumer durables. *Management Science* 15, 215–227.
- Carr, N. G. (2003). It doesn't matter. *Harvard Business Review* 81(5), 41–49.
- Chernozhukov, V. and H. Hong (2003). A MCMC approach to classical estimation. *Journal of Econometrics* 115(2), 293–346.
- Chintagunta, P. K. (1999). Variety seeking, purchase timing, and the “lightning bolt” brand choice model. *Management Science* 45(3), 486–98.
- Einav, L. (2004, August). Not all rivals look alike: An empirical model for discrete games with asymmetric rivals. Econometric Society 2004 North American Winter Meetings 626, Econometric Society.

- Farrell, J. and G. Saloner (1985). Standardization, compatibility, and innovation. *RAND Journal of Economics* 16, 70–83.
- Garicano, L. (2000). Hierarchies and the organization of knowledge in production. *Journal of Political Economy* 108(5), 874–904.
- Garicano, L. and T. N. Hubbard (2003). Hierarchies and the organization of specialization. Mimeo, University of Chicago GSB.
- Gowrisankaran, G. and J. Stavins (2004). Network externalities and technology adoption: lessons from electronic payments. *RAND Journal of Economics* 35(2), 260–276.
- Griliches, Z. (1957, October). Hybrid corn: An exploration in the economics of technological change. *Econometrica* 25.
- Hendel, I. and A. Nevo (2005, March). Measuring the implications of sales and consumer inventory behavior. Northwestern University working paper.
- Jackson, M. and A. Wolinsky (1996). A Strategic Model of Social and Economic Networks. *Journal of Economic Theory* 71, 44–74.
- Katz, M. L. and C. Shapiro (1985). Network externalities, competition, and compatibility. *American Economic Review* 75(3), 424–40.
- Liebowitz, S. and S. E. Margolis (1994, Spring). Network externality: An uncommon tragedy. *Journal of Economic Perspectives* 8(2), 133–150.
- Mansfield, E. (1961, October). Technical change and the rate of imitation. *Econometrica* 29(4), 741–66.
- Marlow, E. (1992). *Corporate television programming: applications and technologies*. Knowledge Industry Publications, Inc.
- Rogers, E. M. (1962, February). *Diffusion of Innovations*. Free Press.

- Ryan, S. (2006). "the costs of environmental regulation in a concentrated industry". MIT Working Paper.
- Rysman, M. (2004, April). Competition between networks: A study of the market for yellow pages. *Review of Economic Studies* 71(2), 483–512.
- Schmidt-Dengler, P. (2005, January). The timing of new technology adoption: The case of MRI. Mimeo, LSE.
- Tucker, C. (2006a, January). Interactive, option-value and domino network effects in technology. Mimeo, MIT.
- Tucker, C. (2006b, January). The role of formal and informal influence in technology adoption. Mimeo, MIT.

Table 1: Monte Carlo Results: Two-Step Estimator—64 Agents

Variable	Truth	Mean	Median	MeanAD	MedianAD	RMSE	StdDev	Left	Right
Intercept	2.0000	2.5314	2.4445	1.4742	1.1537	1.8710	1.8121	-0.4093	6.3934
N	-1.0000	-1.0671	-1.0203	0.3597	0.2913	0.4558	0.4554	-2.1243	-0.3128
decay Asia	1.5000	1.5970	1.5564	0.3008	0.2756	0.3804	0.3716	0.9511	2.4193
decay UK	1.6000	1.7813	1.7940	0.3042	0.3801	0.3380	0.3850	1.2079	2.5032
decay Europe	1.7000	1.8421	1.8244	0.3383	0.2910	0.4045	0.3825	1.1030	2.5464
decay USA	1.8000	1.9877	1.9839	0.3685	0.2889	0.4593	0.4234	1.0459	2.9241
decay Admin	1.9000	2.0797	2.0913	0.4026	0.3725	0.4890	0.4594	1.3027	2.9727
decay Research	2.0000	2.1781	2.1967	0.3417	0.2669	0.4223	0.3868	1.4918	2.9735
decay Sales	2.1000	2.2363	2.1620	0.2536	0.2000	0.3357	0.3099	1.6563	2.9183
decay Trading	2.2000	2.3491	2.3120	0.2913	0.3800	0.3531	0.3531	1.7246	3.0132
decay Associate	2.3000	2.4131	2.3690	0.3409	0.2639	0.4244	0.4132	1.6926	3.3167
decay Vice President	2.4000	2.5955	2.6035	0.3488	0.2648	0.4466	0.4056	1.8311	3.5413
decay Director	2.5000	2.7232	2.7276	0.3961	0.3546	0.4971	0.4487	1.9452	3.7289
decay Managing Director	2.6000	2.7332	2.7729	0.3530	0.2959	0.4176	0.3998	2.0743	3.4533
Asia to UK	0.3400	0.2642	0.2524	0.3461	0.3066	0.4107	0.4077	-0.3811	1.0053
Asia to Europe	0.3600	0.3430	0.3114	0.2908	0.2334	0.3592	0.3625	-0.2397	1.0128
Asia to USA	0.3800	0.3734	0.3664	0.2862	0.2592	0.3616	0.3652	-0.4006	0.9956
UK to Asia	0.4000	0.4080	0.3979	0.1925	0.1184	0.2537	0.2561	-0.0844	0.9054
UK to UK	0.4200	0.4793	0.4414	0.2414	0.1942	0.2877	0.2844	0.0254	1.0257
UK to Europe	0.4400	0.4084	0.3448	0.1851	0.1942	0.2531	0.2536	0.0254	1.0110
UK to USA	0.4600	0.4199	0.4096	0.1759	0.1164	0.2308	0.2296	-0.0068	0.9034
Europe to Asia	0.4800	0.4478	0.4338	0.1732	0.1340	0.2139	0.2136	0.1099	0.8429
Europe to UK	0.5000	0.4870	0.4880	0.1878	0.1591	0.2368	0.2389	0.0805	0.9376
Europe to Europe	0.5200	0.5798	0.6083	0.2244	0.1747	0.2846	0.2811	-0.0725	1.0290
Europe to USA	0.5400	0.5209	0.5132	0.2245	0.1841	0.2761	0.2782	0.0910	1.0413
USA to Asia	0.5600	0.5638	0.5457	0.1676	0.1266	0.2163	0.2185	0.2307	0.9824
USA to UK	0.5800	0.5671	0.5439	0.1986	0.1584	0.2432	0.2453	0.1047	0.9719
USA to Europe	0.6000	0.6058	0.6199	0.1884	0.1320	0.2386	0.2409	0.1578	0.9820
USA to USA	0.6200	0.6189	0.5661	0.2379	0.1802	0.2951	0.2981	0.0442	1.1583
Administration to Research	0.6400	0.7271	0.6005	0.3222	0.2492	0.4303	0.4257	0.0933	1.7478
Administration to Sales	0.6600	0.6730	0.6478	0.3102	0.2628	0.3954	0.3992	-0.0240	1.5041
Administration to Trading	0.6800	0.7350	0.7089	0.3583	0.3078	0.4581	0.4594	-0.1755	1.7047
Research to Administration	0.7000	0.7170	0.7071	0.2252	0.2052	0.2762	0.2785	0.2581	1.2856
Research to Research	0.7200	0.7462	0.7924	0.1773	0.1522	0.2081	0.2085	0.3884	1.1591
Research to Sales	0.7400	0.7328	0.7301	0.1851	0.1520	0.2281	0.2303	0.2132	1.1318
Research to Trading	0.7600	0.7198	0.7002	0.1983	0.1812	0.2370	0.2360	0.2682	1.0918
Sales to Administration	0.7800	0.7936	0.7707	0.2319	0.1848	0.2907	0.2934	0.3237	1.3810
Sales to Research	0.8000	0.7809	0.7813	0.1931	0.1565	0.2471	0.2488	0.1414	1.2452
Sales to Sales	0.8200	0.7992	0.7942	0.2418	0.1820	0.2921	0.2943	0.2152	1.3403
Sales to Trading	0.8400	0.8619	0.8479	0.1792	0.1318	0.2359	0.2373	0.4049	1.4342
Trading to Administration	0.8600	0.8796	0.8514	0.2266	0.1749	0.2756	0.2777	0.3779	1.4018
Trading to Research	0.8800	0.9364	0.9366	0.2077	0.1913	0.2518	0.2479	0.4722	1.3977
Trading to Sales	0.9000	0.8629	0.8447	0.2054	0.1574	0.2562	0.2560	0.3693	1.3237
Trading to Trading	0.9200	0.8765	0.8452	0.2473	0.2291	0.2920	0.2917	0.4222	1.4577
Associate to VP	0.9400	0.9360	0.8754	0.2885	0.2300	0.3625	0.3661	0.3479	1.6143
Associate to Director	0.9600	1.0128	1.0056	0.3282	0.2903	0.4152	0.4160	0.0960	1.8018
Associate to Managing Director	0.9800	1.0467	1.0170	0.3098	0.2706	0.3668	0.3643	0.4899	1.6523
Vice President to Associate	1.0000	0.8799	0.8648	0.2235	0.1713	0.2707	0.2451	0.4559	1.4004
Vice President to VP	1.0200	1.0366	1.0070	0.1744	0.1741	0.2216	0.2233	0.7043	1.5335
Vice President to Director	1.0400	1.1013	1.0477	0.2018	0.1487	0.2626	0.2579	0.7123	1.6207
Vice President to Managing Director	1.0600	1.0977	1.0567	0.2182	0.1585	0.2969	0.2975	0.4910	1.7878
Director to Associate	1.0800	1.0654	0.9879	0.2109	0.2050	0.2428	0.2449	0.6535	1.5356
Director to VP	1.1000	1.1743	1.1482	0.2007	0.1528	0.2563	0.2477	0.6988	1.6108
Director to Director	1.1200	1.0477	1.0645	0.1935	0.1826	0.2437	0.2534	0.5534	1.4706
Director to Managing Director	1.1400	1.1479	1.1454	0.2015	0.1670	0.2504	0.2528	0.6729	1.6556
Managing Director to Associate	1.1600	1.1801	1.1981	0.1714	0.1536	0.2118	0.2130	0.7766	1.5662
Managing Director to VP	1.1800	1.1409	1.1503	0.1988	0.1322	0.2694	0.2692	0.5320	1.5971
Managing Director to Director	1.2000	1.2297	1.1868	0.1763	0.1364	0.2290	0.2294	0.8328	1.7970
Managing Director to Managing Director	1.2200	1.2044	1.2215	0.2067	0.1544	0.2693	0.2716	0.6449	1.6588

Table 2: Monte Carlo Results: Two-Step Estimator—128 Agents

Variable	Truth	Mean	Median	MeanAD	MedianAD	RMSE	StdDev	Left	Right
Intercept	2.0000	2.0078	1.8336	1.0690	0.8351	1.4170	1.4314	-0.7418	5.6943
N	-1.0000	-0.9667	-0.9303	0.2752	0.1985	0.3470	0.3489	-1.7645	-0.3652
decay Asia	1.5000	1.5868	1.5835	0.1935	0.1327	0.2489	0.2356	1.1629	2.1148
decay UK	1.6000	1.7027	1.6383	0.1929	0.1385	0.2530	0.2335	1.3107	2.1876
decay Europe	1.7000	1.8044	1.7701	0.2433	0.2050	0.3131	0.2982	1.2787	2.4115
decay USA	1.8000	1.8815	1.8686	0.2206	0.1748	0.2863	0.2773	1.3856	2.4849
decay Admin	1.9000	2.0211	1.9947	0.2183	0.1732	0.2712	0.2452	1.4876	2.4698
decay Research	2.0000	2.0977	2.0639	0.2440	0.1994	0.3097	0.2969	1.5443	2.7553
decay Sales	2.1000	2.1525	2.1506	0.1967	0.1632	0.2466	0.2465	1.7731	2.6178
decay Trading	2.2000	2.2756	2.2861	0.2483	0.2183	0.2989	0.2921	1.8371	2.8561
decay Associate	2.3000	2.3598	2.3949	0.2000	0.1654	0.2388	0.2335	1.8588	2.7453
decay Vice President	2.4000	2.4655	2.4414	0.1795	0.1353	0.2404	0.2337	2.0360	2.9395
decay Director	2.5000	2.5889	2.5776	0.2169	0.1878	0.2692	0.2567	2.1291	3.1761
decay Managing Director	2.6000	2.7106	2.6862	0.2466	0.1910	0.3184	0.3016	2.1655	3.2562
Asia to UK	0.3400	0.3320	0.3252	0.2038	0.1776	0.2645	0.2671	-0.1433	0.9442
Asia to Europe	0.3600	0.2925	0.2859	0.2060	0.1712	0.2581	0.2517	-0.2049	0.7610
Asia to USA	0.3800	0.3868	0.4015	0.2387	0.2179	0.2901	0.2930	-0.1032	0.8721
UK to Asia	0.4000	0.3937	0.3712	0.1856	0.1337	0.1808	0.1826	0.0512	0.7276
UK to UK	0.4200	0.3909	0.3891	0.1209	0.1019	0.1578	0.1567	0.1075	0.7475
UK to Europe	0.4400	0.4824	0.4805	0.1244	0.0987	0.1593	0.1551	0.2020	0.7935
UK to USA	0.4600	0.4487	0.4695	0.1695	0.2032	0.2032	0.2049	0.1305	0.8449
Europe to Asia	0.4800	0.4543	0.4595	0.1231	0.1000	0.1505	0.1498	0.1877	0.8165
Europe to UK	0.5000	0.5436	0.5361	0.1346	0.1172	0.1657	0.1615	0.2979	0.8587
Europe to Europe	0.5200	0.4991	0.5187	0.1175	0.0988	0.1492	0.1492	0.1911	0.7830
Europe to USA	0.5400	0.5388	0.5393	0.1039	0.0912	0.1270	0.1282	0.3197	0.7766
USA to Asia	0.5600	0.5447	0.5749	0.1617	0.1164	0.2160	0.2176	0.0324	0.8482
USA to UK	0.5800	0.6231	0.6037	0.1342	0.1298	0.1636	0.1676	0.3247	0.8686
USA to Europe	0.6000	0.6148	0.6227	0.1460	0.1222	0.1797	0.1809	0.3351	0.9836
USA to USA	0.6200	0.5731	0.5644	0.1473	0.1230	0.1874	0.1833	0.2480	0.9970
Administration to Research	0.6400	0.6419	0.6053	0.2446	0.2209	0.2861	0.2890	0.1582	1.1567
Administration to Sales	0.6600	0.6581	0.6828	0.1996	0.1314	0.2631	0.2657	0.1973	1.1047
Administration to Trading	0.6800	0.6520	0.6612	0.2354	0.1935	0.3038	0.3056	0.0620	1.1951
Research to Administration	0.7000	0.7368	0.7556	0.1464	0.1120	0.1881	0.1863	0.2192	1.0699
Research to Research	0.7200	0.6994	0.6751	0.1366	0.1162	0.1657	0.1661	0.3814	1.0271
Research to Sales	0.7400	0.7196	0.6999	0.1432	0.1175	0.1790	0.1797	0.4270	1.1091
Research to Trading	0.7600	0.7601	0.7555	0.1517	0.1345	0.2071	0.2092	0.3296	1.1410
Sales to Administration	0.7800	0.8046	0.8197	0.1414	0.1197	0.1723	0.1722	0.4593	1.1294
Sales to Research	0.8000	0.7784	0.7512	0.1309	0.1257	0.1550	0.1550	0.5273	1.1361
Sales to Sales	0.8200	0.7842	0.7956	0.1131	0.1117	0.1379	0.1345	0.5672	1.0508
Sales to Trading	0.8400	0.8686	0.8297	0.1245	0.0967	0.1566	0.1555	0.5720	1.1797
Trading to Administration	0.8600	0.8732	0.8632	0.1324	0.1018	0.1678	0.1690	0.5215	1.1900
Trading to Research	0.8800	0.8567	0.8626	0.1419	0.1081	0.1837	0.1840	0.3987	1.1500
Trading to Sales	0.9000	0.8648	0.8576	0.1300	0.1120	0.1630	0.1608	0.5586	1.2092
Trading to Trading	0.9200	0.9611	0.9563	0.1093	0.0886	0.1358	0.1308	0.7057	1.1993
Associate to VP	0.9400	0.9824	0.9636	0.2274	0.1685	0.2847	0.2844	0.4493	1.5052
Associate to Director	0.9600	1.0073	0.9868	0.1929	0.1616	0.2444	0.2423	0.5961	1.4226
Associate to Managing Director	0.9800	1.0419	0.9499	0.2889	0.2330	0.3659	0.3643	0.5396	1.8456
Vice President to Associate	1.0000	1.0144	0.9966	0.1177	0.0924	0.1509	0.1517	0.7422	1.3226
Vice President to VP	1.0200	0.9876	0.9601	0.1319	0.1116	0.1575	0.1557	0.7319	1.2529
Vice President to Director	1.0400	1.0510	1.0606	0.1208	0.1024	0.1481	0.1492	0.7755	1.2777
Vice President to Managing Director	1.0600	1.0629	1.0636	0.1266	0.1053	0.1639	0.1655	0.7425	1.3266
Director to Associate	1.0800	1.0948	1.0928	0.1271	0.0986	0.1600	0.1610	0.7192	1.3476
Director to VP	1.1000	1.1412	1.1535	0.1449	0.1051	0.1764	0.1732	0.8342	1.4226
Director to Director	1.1200	1.1097	1.1039	0.0907	0.0801	0.1141	0.1133	0.9115	1.3578
Director to Managing Director	1.1400	1.1009	1.1189	0.1348	0.1100	0.1673	0.1643	0.8179	1.4338
Managing Director to Associate	1.1600	1.1178	1.1250	0.1371	0.0990	0.1708	0.1672	0.8172	1.4146
Managing Director to VP	1.1800	1.2073	1.1937	0.1290	0.1098	0.1675	0.1670	0.8576	1.5328
Managing Director to Director	1.2000	1.1833	1.2079	0.1287	0.1215	0.1586	0.1593	0.8657	1.4637
Managing Director to Managing Director	1.2200	1.2469	1.2353	0.1263	0.0884	0.1600	0.1593	1.0002	1.5548

Figure 1: Distribution of Employees by Type



Figure 2: Distribution of Adoption Rates by Type



Figure 3: Calls Across Regions




Figure 4: Calls Across Functions



Figure 5: Calls Across Titles


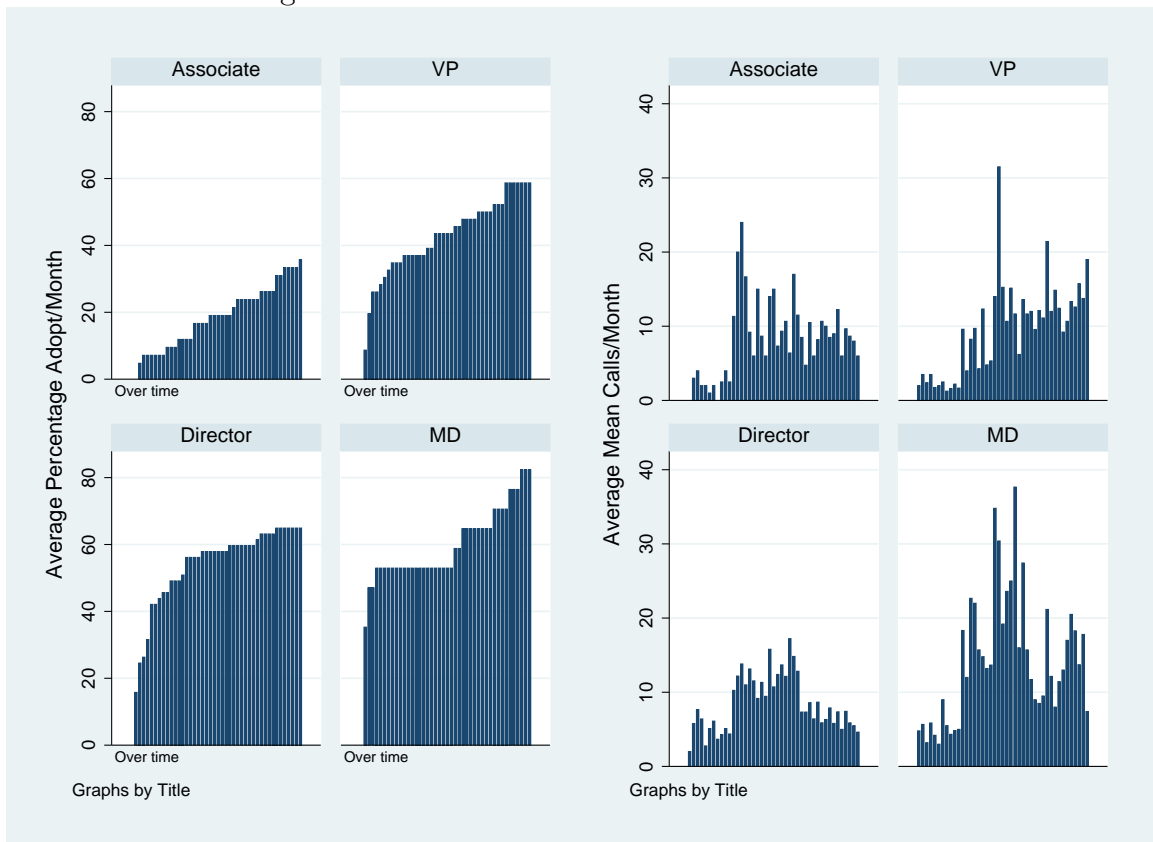


Figure 6: Calls Across Titles for US Researchers



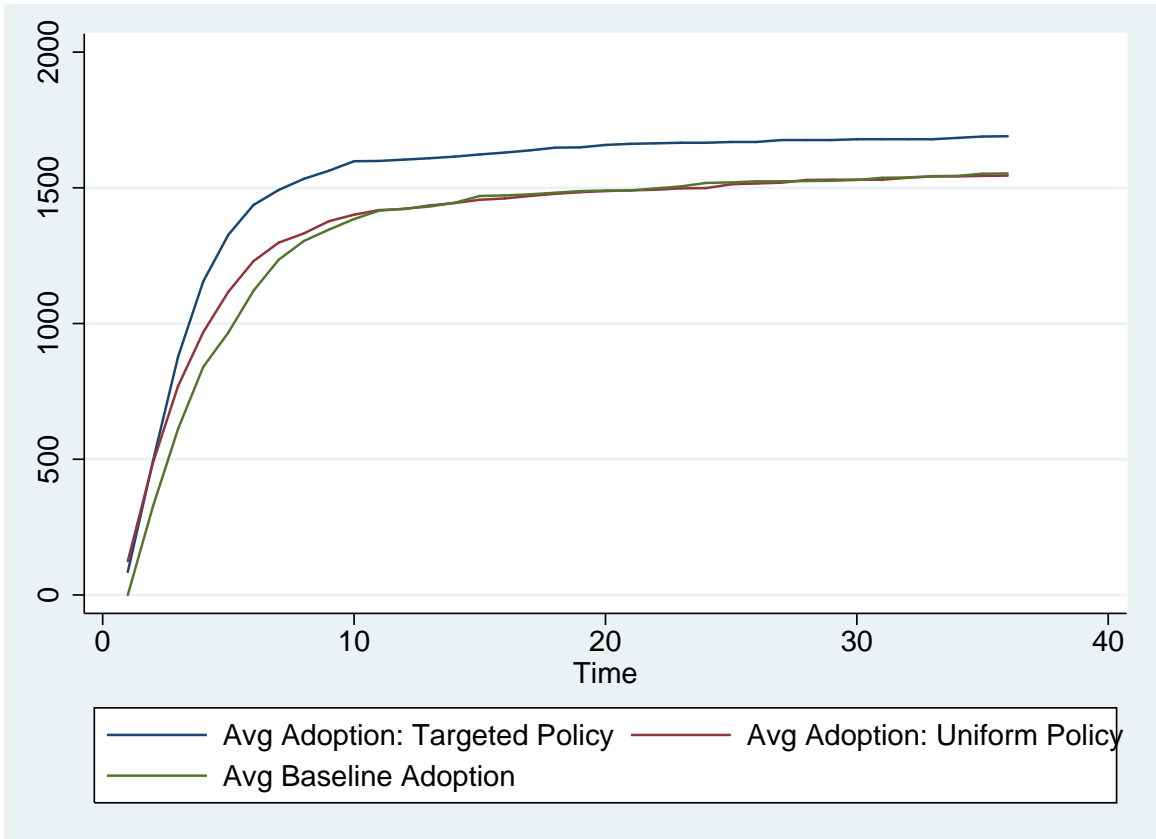


Figure 7: Adoption: Targeted vs Uniform

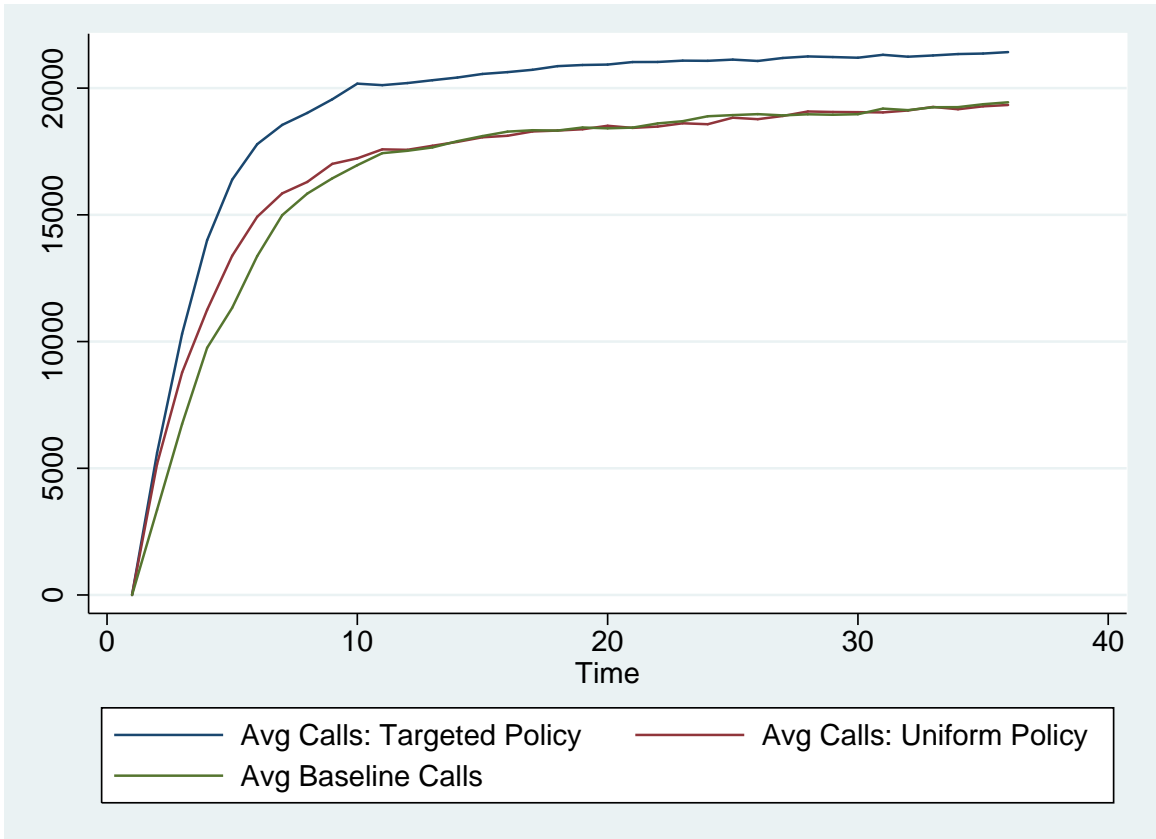


Figure 8: Calls: Targeted vs Uniform




Figure 9: Utility: Targeted vs Uniform

Table 3: Decay Rates by Receiver Characteristic

Variable	Mean	StdDev
Intercept	-0.6862	0.0151
N	-0.6735	0.0010
decay Asia	-0.1685	0.0074
decay UK	-0.0674	0.0016
decay Europe	-0.0478	0.0015
decay USA	-0.0702	0.0017
decay Admin	-0.0569	0.0020
decay Research	-0.1210	0.0028
decay Sales	-0.0520	0.0021
decay Trading	-0.0446	0.0016
decay Associate	-0.1001	0.0024
decay Vice President	-0.0521	0.0011
decay Director	-0.0396	0.0012
decay Managing Director	-0.0546	0.0021

Table 4: Static Interactions of Caller and Receiver Regions on Calling Choice

Variable	Mean	StdDev
Asia to UK	-0.6600	0.0597
Asia to Europe	-1.0436	0.0942
Asia to USA	-1.9795	0.1309
UK to Asia	0.6670	0.0909
UK to UK	0.9514	0.0746
UK to Europe	1.5223	0.0699
UK to USA	0.9829	0.0733
Europe to Asia	0.5919	0.2800
Europe to UK	1.6874	0.2647
Europe to Europe	2.7498	0.2664
Europe to USA	0.1695	0.2769
USA to Asia	-0.6244	0.1519
USA to UK	0.9069	0.0979
USA to Europe	0.2601	0.1060
USA to USA	1.5474	0.0879

Table 5: Static Interactions of Caller and Receiver Functions on Calling Choice

Variable	Mean	StdDev
Administration to Research	-2.0443	0.0496
Administration to Sales	-1.4193	0.0472
Administration to Trading	-1.3955	0.0459
Research to Administration	2.6370	0.2032
Research to Research	2.4206	0.2023
Research to Sales	1.9498	0.2049
Research to Trading	1.7574	0.2013
Sales to Administration	0.2744	0.0841
Sales to Research	-0.6013	0.0707
Sales to Sales	0.3052	0.0819
Sales to Trading	-0.1223	0.0846
Trading to Administration	-0.2484	0.0678
Trading to Research	-1.5532	0.0749
Trading to Sales	-0.9859	0.0741
Trading to Trading	0.0832	0.0731

Table 6: Static Interactions of Caller and Receiver Titles on Calling Choice

Variable	Mean	StdDev
Associate to VP	0.0246	0.0436
Associate to Director	-0.3665	0.0516
Associate to Managing Director	-0.4850	0.0599
Vice President to Associate	-0.6165	0.0809
Vice President to VP	-0.4307	0.0801
Vice President to Director	-0.5650	0.0819
Vice President to Managing Director	-0.7572	0.0742
Director to Associate	-1.6287	0.1106
Director to VP	-1.1712	0.1006
Director to Director	-1.0022	0.0988
Director to Managing Director	-0.9161	0.1114
Managing Director to Associate	0.4116	0.1405
Managing Director to VP	0.6804	0.1395
Managing Director to Director	1.1702	0.1444
Managing Director to Managing Director	1.8718	0.1393

Table 7: Fixed Costs by Function and Title for Asia

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	0.907	0.031	1.024	0.008
Vice President	-0.091	0.037	1.220	0.010
Director	1.551	0.039	1.094	0.034
Managing Director	0.616	0.034	1.106	0.016
Research				
Associate	2.450	0.014	0.086	0.001
Vice President	2.465	0.016	0.086	0.001
Director	1.979	0.032	0.707	0.032
Managing Director	0.873	0.031	1.040	0.010
Sales				
Associate	2.158	0.034	0.527	0.041
Vice President	2.030	0.036	0.681	0.037
Director	1.279	0.033	0.948	0.006
Managing Director	0.394	0.030	1.273	0.007
Trading				
Associate	2.202	0.038	0.484	0.043
Vice President	1.887	0.039	0.833	0.033
Director	1.613	0.028	1.055	0.022
Managing Director	0.638	0.032	1.107	0.011

Table 8: Fixed Costs by Function and Title for United Kingdom

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	2.015	0.039	0.666	0.042
Vice President	1.417	0.088	0.999	0.090
Director	1.616	0.033	1.043	0.024
Managing Director	0.199	0.028	1.303	0.006
Research				
Associate	1.992	0.043	0.644	0.047
Vice President	0.900	0.037	1.031	0.012
Director	0.613	0.032	1.120	0.016
Managing Director	0.138	0.036	1.326	0.009
Sales				
Associate	1.749	0.034	0.923	0.026
Vice President	0.363	0.031	1.283	0.007
Director	0.783	0.027	1.062	0.009
Managing Director	0.274	0.034	1.307	0.007
Trading				
Associate	1.583	0.035	1.062	0.026
Vice President	0.783	0.030	1.062	0.010
Director	0.771	0.029	1.061	0.008
Managing Director	0.050	0.038	1.292	0.007

Table 9: Fixed Costs by Function and Title for Europe

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	1.693	0.035	0.963	0.027
Vice President	1.687	0.039	0.974	0.030
Director	1.223	0.040	0.953	0.007
Managing Director	2.403	0.015	0.082	0.001
Research				
Associate	1.827	0.036	0.804	0.037
Vice President	1.496	0.070	1.048	0.062
Director	0.665	0.038	1.107	0.015
Managing Director	0.363	0.044	1.323	0.011
Sales				
Associate	2.086	0.046	0.548	0.054
Vice President	1.375	0.083	0.968	0.088
Director	0.893	0.040	1.030	0.011
Managing Director	0.146	0.040	1.321	0.006
Trading				
Associate	1.904	0.036	0.769	0.036
Vice President	1.537	0.046	1.076	0.039
Director	1.255	0.040	0.942	0.008
Managing Director	0.070	0.032	1.310	0.008

Table 10: Fixed Costs by Function and Title for USA

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	2.093	0.031	0.617	0.034
Vice President	1.622	0.034	1.040	0.026
Director	1.198	0.034	0.959	0.006
Managing Director	0.359	0.040	1.300	0.010
Research				
Associate	2.216	0.067	0.361	0.086
Vice President	1.420	0.081	0.993	0.079
Director	0.747	0.037	1.077	0.013
Managing Director	0.205	0.040	1.316	0.007
Sales				
Associate	1.904	0.030	0.779	0.030
Vice President	1.541	0.051	1.074	0.041
Director	1.179	0.041	0.961	0.007
Managing Director	0.719	0.030	1.087	0.010
Trading				
Associate	2.204	0.041	0.441	0.050
Vice President	2.092	0.031	0.581	0.033
Director	1.935	0.032	0.777	0.027
Managing Director	0.919	0.035	1.024	0.011

Table 11: Policy Experiment Results

Variable	Baseline	Targeted	Uniform
Average Number of Calls	12.237	12.496	11.93
Maximum number of Adopters	1553	1690	1545
Present Value utility (mean)	403.1	426.6	418.9
Present Value utility (median type)	371.3	392.2	382.7
Present Value utility (25% type)	272.8	286.7	283.3
Present Value utility (75% type)	516.8	542.8	538.5
Discounted Value to Firm with $\beta = 0.9$			
Present Discounted Monthly Users	8904.8	10761.5	9542.9
Present Discounted Calls	107603.5	132763.1	114451.1
Discounted Value to Firm with $\beta = 0.99$			
Present Discounted Monthly Users	39371.1	44892.9	40218.8
Present Discounted Calls	484299.6	563224.0	494145.2