Heterogeneity and the Dynamics of Technology Adoption

Stephen Ryan* and Catherine Tucker†

October 6, 2006

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 find that ignoring heterogeneity in the interplay between adoption costs and network effects will underpredict the size of the steady-state network size by almost 50 percent. 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 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.

*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 drive 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. In doing so, we are able to provide answers to Griliches' three questions from a utility-based framework. 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 activities: 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 either never adopting or of waiting to adopt in future periods. We characterize their adoption timing 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 types of people over time. This heterogeneity allows us to explicitly model choices to interact with different users both initially and over time. This model allows us to capture how diversity in the characteristics of network subscribers affects agents' motivations for adoption.

Second, we use extensive data on the diffusion and use of a videoconferencing technology within a large multinational investment bank to estimate the parameters of our model. We have detailed data on all 2,112 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.

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 Ackerberg 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 Ackerberg 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^{2112}, 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

[1] Our 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.

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 challenge we face 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 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. This interaction is critical to understanding the evolution of the network, as our agents perceive equally-sized homogenous and diverse networks very differently. Our specification without heterogeneity underpredicts the extent of adoption of the technology by 50 percent and obviously does a very poor job of matching adoption rates across different types. These findings echo Tucker (2006b)'s previous empirical research using this data, which documents how network effects vary in size with both formal and informal influence in a firm.

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 elaboration to be capable of reproducing our results.

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 marketing 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.

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.

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.

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

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

data on all potential adopters.

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 video-messaging 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 video-messaging 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.

This decentralization focuses analysis on the private benefits to installation for employees, as opposed to firm-level productivity benefits. Studies such as Lazear (2000) discuss how firms find it hard to monitor and reward improved communication. Information asymmetries mean that employees' installation benefits may be small relative to firm-level benefits from the videoconferencing system. We cannot quantify these firm-level productivity benefits.

2.3 Data

There are complete personnel records for each employee in the bank in March 2004. Throughout our data there were around 2000 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 or 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, 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.

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.

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, 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 video-messages 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 kth 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 kth call in agent i's sequence, the utility of calling agent j is given by:

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

As in Jackson and Wolinsky (1996)'s model of network formation, 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:

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

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

[3] We also observe agents that "call" themselves; we leave the interpretation of such behavior to the reader.

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:

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

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.

Denote the expected discounted present value of using the network after adopting t periods in the future by $EV_t(s)$, where all expectations are taken with respect to the current period. An agent will adopt today if and only if the following inequality is satisfied:

$$EV_0(s_0) - F_i \geq \max\left\{0, \max_{t>0} \beta^t \left(EV_t(s_t) - F_i\right)\right\}. \tag{3}$$

The agent compares the benefits to adopting today against the best alternative, which is either the outside option or waiting to adopt until a period in the future. We have explicitly written the best future expected value, $EV_t(s_t)$, as a function of s_t to emphasize that the agent is making predictions about the future evolution of the network. This expectation raises two 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.

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:

$$Pr(\Omega, K) = Pr(\Omega|K)Pr(K). \tag{4}$$

We can then separate the estimation of the composition of the calling sequence from the length of the sequence.[4] After taking logarithms in Equation 4, we obtain:

$$\ln Pr(\Omega, K) = \ln Pr(\Omega|K) + \ln Pr(K). \tag{5}$$

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

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:

$$U_{ijk} = \hat{\theta}_1 x_j + \hat{\theta}_2 x_i x_j + g(\Omega_{1:k}; \hat{\theta}_3) + \epsilon_{ijk}, \tag{6}$$

ϵ is a 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 assume that $g(\cdot)$ takes the following functional form:

$$g(\Omega_{1:k}; \hat{\theta}_3) = \sum_{m=1}^{M} \exp(\hat{\theta}_{3m}) \eta_m - k, \tag{7}$$

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 generates constant 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. We normalize the coefficient on k, the linear utility penalty of making multiple calls, to be equal to -1. This is a standard normalization in discrete choice models, as the parameters are only identified up to scale and location. 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 logit probability of observing a call from agent i to agent j as the kth call of a sequence:

$$Pr(\Omega_{ijk}; s_t, \hat{\theta}) = \frac{\exp(U_{ijk}(\hat{\theta}))}{\sum_{j' \in s_t} \exp(U_{ij'k}(\hat{\theta}))}. \tag{8}$$

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.

Equation 6 does not contain own-characteristics. They drop out of the conditional calling probability, as they enter equally into all the calls that agent i could make. Since there is no variation in these parameters within the calling sequence of a given length, we cannot identify these characteristics. However, as discussed below, these parameters will be central to the second step of the simulated sequence estimator, where we use them to recover the type-specific average number of calls.

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:

$$\hat{K}_m = \frac{1}{N_s} \sum_{j}^{N_s} K_{jm}(\hat{\theta}_4),$$

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

$$\hat{Q}(\hat{\theta}_4) = \frac{1}{T} \sum_{t=1}^{T} \left(\frac{1}{M_t} \sum_{m=1}^{M_t} \left(\frac{1}{N_{mt}} \sum_{i_{mt}=1}^{N_{mt}} \left(\widehat{K_m(\theta_4)} - \Omega_{imt} \right)^2 \right) \right). \tag{9}$$

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 probability that an agent of a given subtype will adopt the technology in a given period, conditional on the extant 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.

The optimal approach in such a setting 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. We recover the adoption policy for agent i of subtype m as a function of the network at time t using a probit function:

$$Pr(adopt_{im} = 1|s_t) = \Phi(\hat{\lambda}_m, s_t),$$

where $\Phi(\cdot)$ is the standard normal cumulative distribution function. Note that the $\hat{\lambda}_m$ varies for each m, so we will recover different policy functions for each agent subtype. We represent the state of the network at time t by including counts of every subtype currently in the network as explanatory variables. We assume that these counts enter in linearly with the same marginal parameters for all subtypes. It would be desirable to not have to assume that the addition of a given subtype to the network has the same marginal effect on adoption, holding utility constant across subtypes. There is, however, not enough variation in the data to allow us to relax this assumption. We take $adopt_{it} = 1$ as the first time an employee either makes or receives a videoconferencing call. Since many employee actions are measured and rewarded on a monthly basis, we use a month as our time-frame for adoption. We include the stand-alone use of TV-watching as an explanatory variable in our adoption policy functions as well. This helps to control for heterogeneity in the stand-alone benefit. We assume that the utility from TV-watching is identical across agents in the same region, and varies across time. The variable TV_{rt} is measured by the percentage of adopters who watch TV in that region in that period. Controlling for TV in this manner allows us to focus our analysis in the rest of the paper on the video-conferencing use of this technology, and abstract away from a highly idiosyncratic feature of this technology suite.

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.

As an illustration of how this works, consider the empty network. For each potential adopter of the videoconferencing technology, we draw a uniform random number on the interval $[0, 1]$, check to see if this is lower than the value returned by the probit function governing that adoption decision, and update the network in the next period accordingly. Looping over all the potential adopters in each period, we mimic the evolution of the network as it occurred in the data.

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 6, drawing shocks appropriately, and generate calling sequences. 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 first period. This is exactly the logic that we will use to form estimates of the fixed costs of adoption.

An agent weighs the expected discounted present value of joining the network in the present period against the best alternative. The outside option gives a payoff of zero in this period, but this may be the best choice at a given time since the network continues to grow, which in turn increases future utility. The agent computes the expected value of waiting to adopt across all future periods, and finds the period that yields the highest expected utility. If the utility of waiting is higher than the utility of joining at once, then the agent finds it optimal to wait.

To illustrate these incentives at work, suppose that there is one agent who is given the option to join the empty network one period before her peers. If this agent joins immediately, she will have no one to talk to, although she may find some benefit from watching TV and may induce additional agents to join in the next period, increasing her continuation value. However, it is also possible that these benefits are swamped by the fixed costs of joining in the current period, relative to the alternative of waiting one period to join with everyone else.

To recover the fixed costs of adoption for each agent subtype, we exploit the structure of this optimal waiting problem. The empirical component has two parts. First, discounting means that the current value of adoption is higher than any other period:

$$EV_0 \geq \max\left\{\beta EV_1, \beta^2 EV_2, \ldots, \right\}.$$

In this case, the only reason that 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:

$$\begin{aligned} Pr(adopt_{im} = 1) &= Pr(EV_0 - F_i > 0) \\ &= Pr(EV_0 > F_i) \\ &= CDF(EV_0; \mu, \sigma^2). \end{aligned}$$

We use a cumulative distribution function of the fixed costs to capture the cost condition and assume that this distribution is normally distributed with mean μ and variance σ^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 unknown is the parameters of the CDF. To obtain estimates of those unknowns we simply match empirical probabilities from the policy function to their computer counterparts:

$$\hat{Q}_n(\mu_m, \sigma_m^2) = \sum_{i=1}^{N_s} \left(Pr(adopt_{im} = 1) - \Phi\left(EV_0; \mu_m, \sigma_m^2\right)\right)^2. \tag{10}$$

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. The same idea can be used to trace out the entire CDF nonparametrically, if the policy function is estimated flexibly.

The second case that can occur is that a period in the future dominates the expected value of joining the network immediately. In this case, the only reason that we would observe an agent i of subtype m joining the network is if the following is true:

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

where t^* solves the right-hand side in Equation 4.3. Rearranging, we find it is optimal to enter in this period even when the usage value is higher in the future at a time when:

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

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

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

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 10:

$$\hat{Q}_n(\mu_m, \sigma_m^2) = \sum_{i=1}^{N_s} \left(Pr(adopt_{im} = 1) - \Phi\left(\frac{EV_0 - \beta^{t^*} EV_{t^*}}{1 - \beta^{t^*}}; \mu_m, \sigma_m^2 \right) \right)^2. \tag{11}$$

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 4.3. If the current period dominates waiting for future periods, we apply the estimator defined in Equation 10. Otherwise, we match the empirical moments in the data using the estimator defined in Equation 11. 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.

4.4 Selection

One of the attractions of the two-step approach we use is that the first-stage reduced form policy function, which maps adoption decisions to the state space, automatically reflects selection. When estimating models over time, selection is an issue. In later periods, if an agent has not adopted, it seems likely that they received a high fixed cost draw. Since we do not observe divestment of the technology, standard panel data techniques do not allow estimation of individual-level fixed effects. By contrast, our inclusion of state space polynomials in our reduced form policy function explicitly controls for selection.

4.5 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). In the counterfactual policy evaluations below we have to assume that the same equilibrium is played in those simulations, since we have no facility for calculating even one equilibrium to our game, much less the entire set of all equilibria.

4.6 Monte Carlo Evidence

To see how good our estimation approach is at recovering the underlying parameters of our model, we ran a simple Monte Carlo experiment. The results, along with the true parameters, are shown in Tables 1 to 4. The Monte Carlo evidence suggests that our estimator precisely estimates the calling parameters, even including the decay rates. The constant is less well-identified, but in all case the true parameter is contained in statistically-significant 95% confidence intervals, with only the constant exhibiting bias in the small sample.

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 6. Tables 5 through 7 display the results of our estimates. In general, in a static setting, agents prefer to call other agents with similar characteristics. 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 8. Our estimates of $\hat{\theta}_3$ in Equation (6) suggest that dynamic preferences for heterogeneity in communication vary across types but are on average relatively small. The parameters in Table 8 are estimated as negatively-constrained functions, $-exp(\cdot)$, so small numbers represent large decay rates and large numbers represent smaller decay rates. People seem satisfied to have fewer repeated interactions with employees in Asia and employees in sales and trading. The decay rates on titles suggest that employees have similar needs for repeated interactions with associates, vice presidents and managing directors, but have less of a need to repeatedly interact with directors. We speculate that this is because a director is often the highest in the chain of command for a support function in the bank - for example "Director of Marketing" or "Director of Human Resources". Regular employees may not benefit from repeated interactions with the employees who occupy these roles; instead a single conversation may be sufficient to provide the needed information.

5.2 Fixed Costs of Adoption

Tables 11 through 14 display the results of our fixed cost estimates from estimating Equations 10 and 11. 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, administrators have on average a lower fixed cost of adoption than employees in other functions. Finally, there are within-group differences in fixed costs. For example, Managing Directors appear to have the lowest fixed costs of adoption in Europe but the highest fixed costs of adoption in the US.

If we compare these results with the results for calling choices in Tables 5 through 7, 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 US, this was on average a group whom callers received high utilities from calling. However, they also had some of the highest fixed adoption costs.

5.3 Robustness Checks

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 first estimated the calling model on each of the three period lengths. We then evaluated the log likelihood at each of those three parameter vectors on one month of data. The likelihoods are reported in Table 9. 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.

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 technology 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 10. Figures 6 and 7 contrast the results graphically for the total adoption 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. This is somewhat surprising, given that the starting network for the uniform policy is much larger than the baseline case.

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 in the baseline case is just 4% lower than the uniform case. 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 8.5% in present discounted utility 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 57% and calls increase by 64%. 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.

6.1 Homogeneity and Heterogeneity Comparison

A central emphasis of this paper is that it is important to allow for heterogeneity in network effects when modeling technology adoption, especially in a dynamic setting. To demonstrate this, we also repeated the estimation procedure described in Section 4, but this time with a first stage policy function which allowed only homogeneity in network effects. In other words, adoption was constrained to be a non-linear function of the number of other adopters in the network, as opposed to being a non-linear function of the number and types of the other adopters in the network.

Figure 8 displays the results when we force network effect to be homogenous, an assumption which has been prevalent in the previous literature. There are two things that are clear when comparing Figure 8 with Figure 6. Unsurprisingly, if you model network effects as homogenous there is very little variation in the predicted adoption outcomes of using a targeted or a uniform adoption seeding policy. Second, an assumption of homogenous network effects leads predicted adoption over time to be far lower than predicted adoption under a heterogenous network effects model. These patterns in predicted adoption are echoed for the results for utility from calling displayed in Figure 9. A comparison of the original predictions of calling utility with the heterogenous network effects in Figure 7 shows that the predicted calling utility is smaller if the model assumes homogenous network effects.

This under-prediction is due to the failure of a homogenous model to reflect the interplay between heterogeneity in adoption costs and heterogeneity in network effects. Put simply, in the data, the subtypes of employees who have the lowest fixed adoption costs also have the largest network effects on the adoption decisions of others. By modeling just the average network effect and the average fixed cost, a model which assumes homogeneity misses this crucial kick start to adoption in the network, and the effect this has on later adoption decisions.

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. 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. We also show that by disregarding the interplay of heterogeneity in network effects and fixed costs in dynamic models of adoption, adoption would be under-predicted by 50 percent.

References

Ackerberg, 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., 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.

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.

Lazear, E. (2000, May). Power of incentives. *American Economic Review 90*(2), 410–414.

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.

Figure 1: Distribution of Employees Across Firm

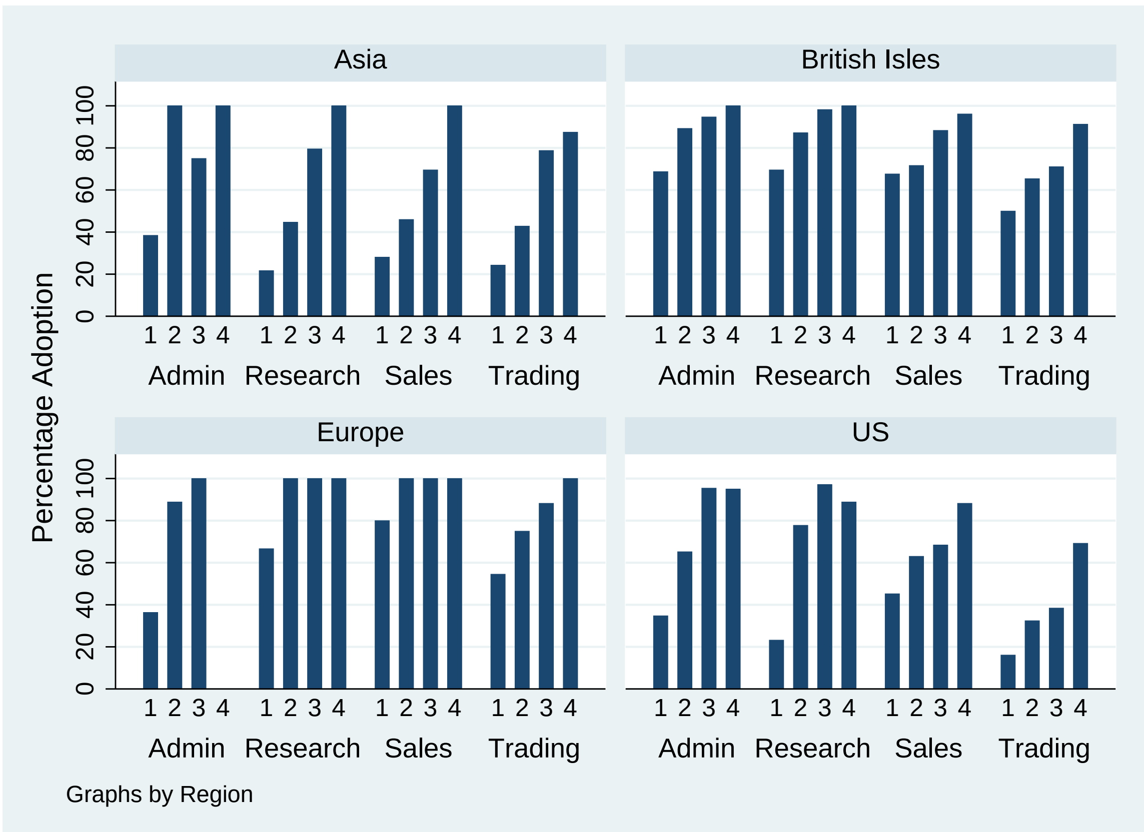

Figure 2: Distribution of Adoption Across Firm

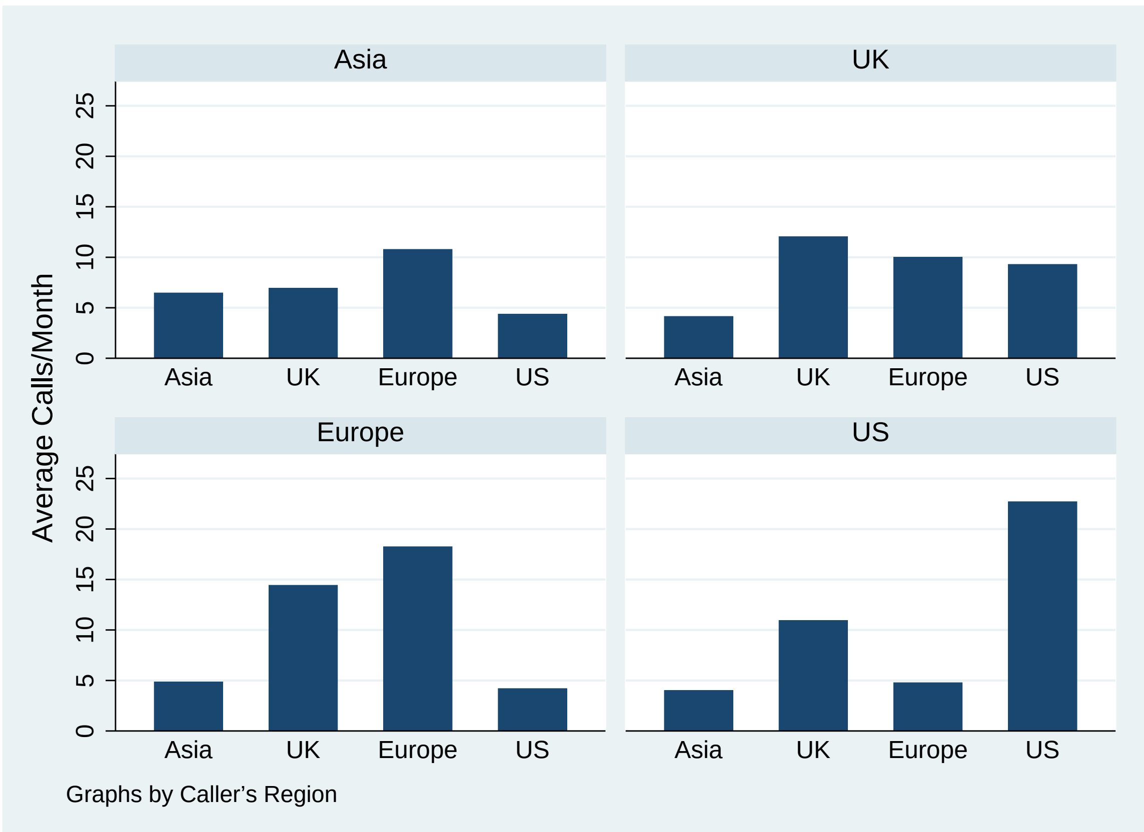

Figure 3: Calls Across Regions

Figure 4: Calls Across Functions

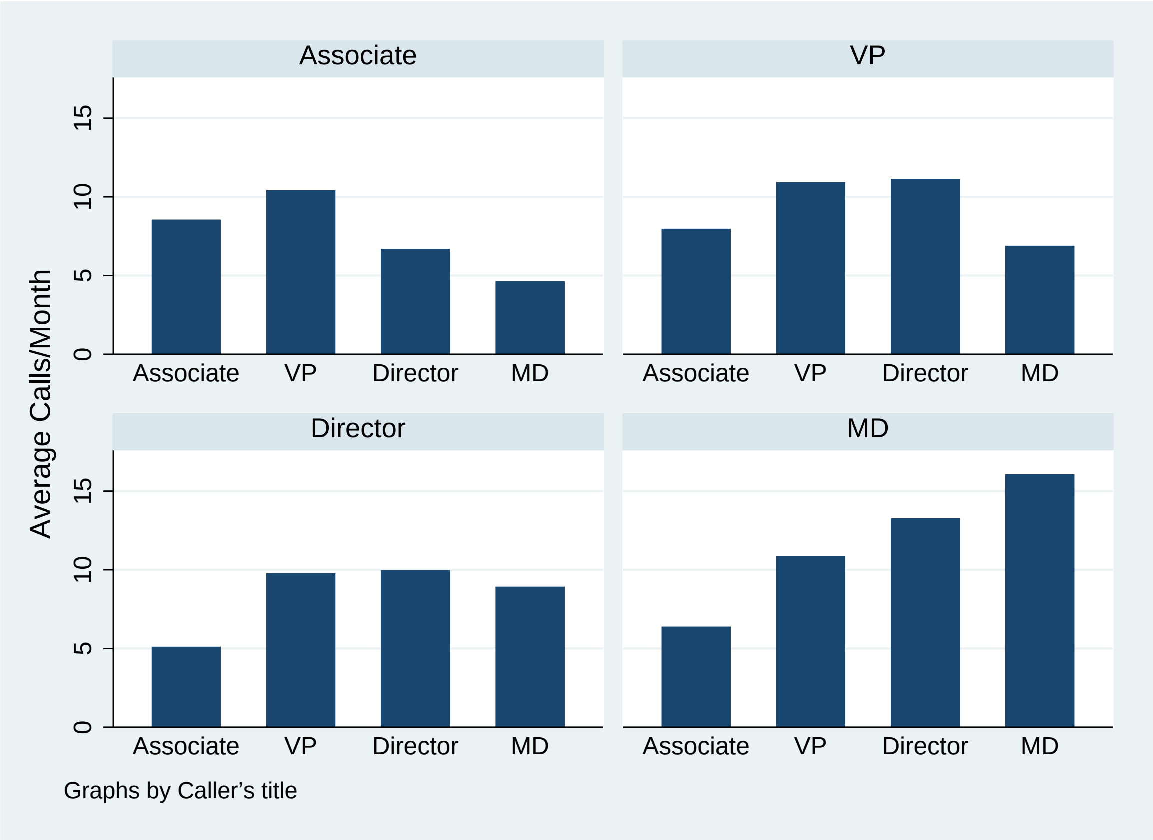

Figure 5: Calls Across Titles

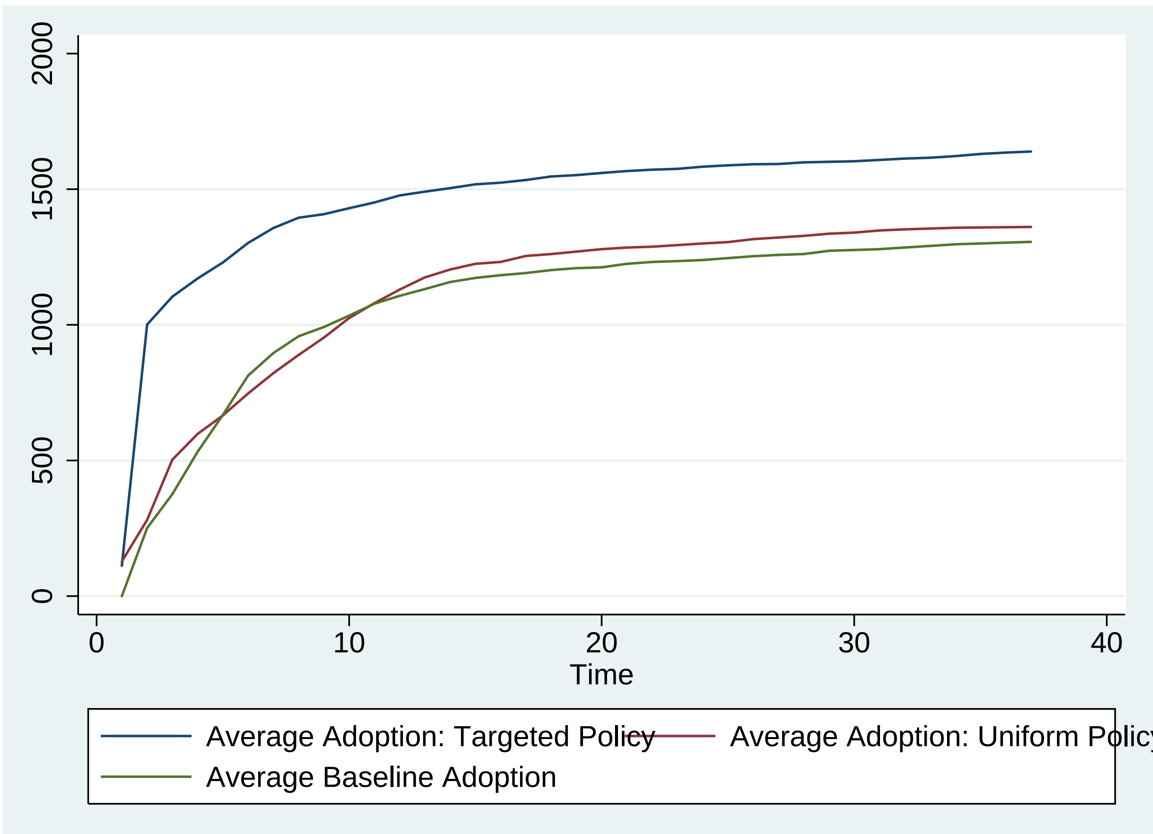

Figure 6: Adoption: Targeted vs Uniform

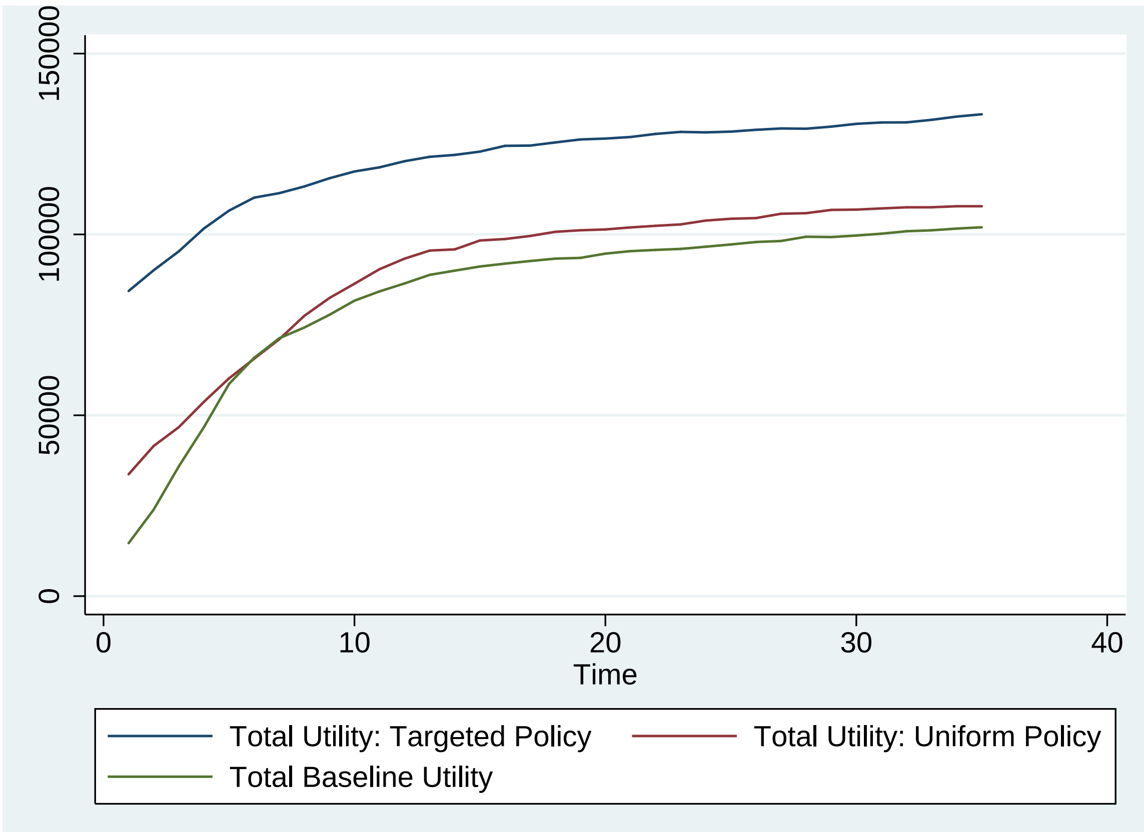

Figure 7: Utility: Targeted vs Uniform

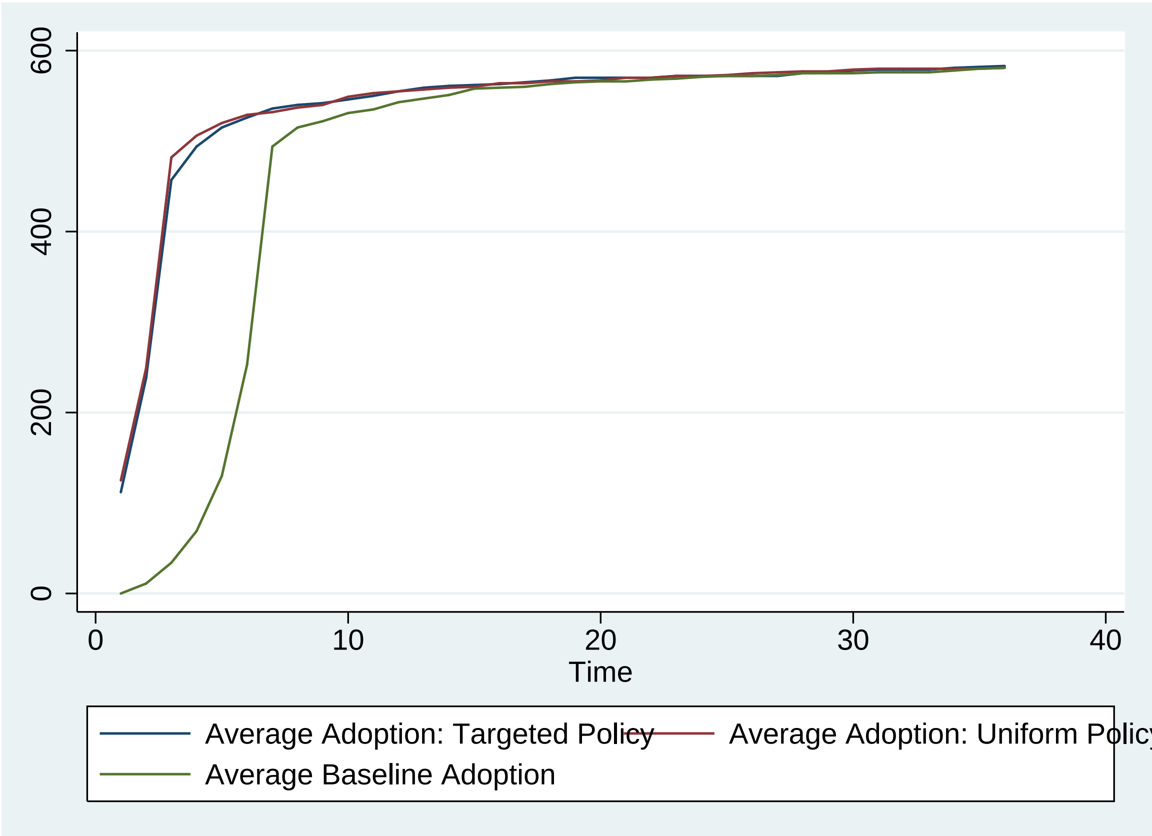

Figure 8: Adoption: Targeted vs Uniform under Homogeneity assumption

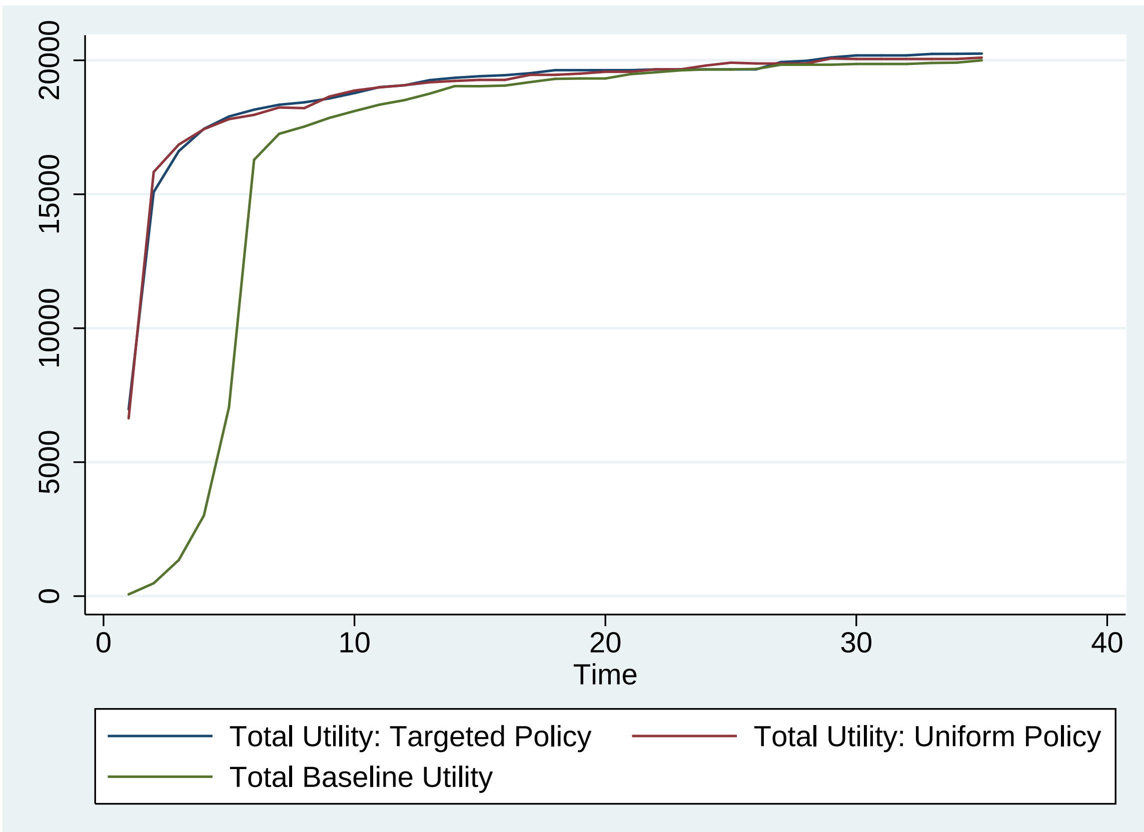

Figure 9: Utility: Targeted vs Uniform under Homogeneity assumption

Table 1: Monte Carlo Results—Caller and Receiver Regions

Variable	**Actual**	**Mean**	**Median**	**StdDev**	$\alpha = 0.025$	$\alpha = 0.975$
Constant	3.0000	4.0133	4.6322	1.1846	1.0094	4.9342
AsiatoUK	2.3000	2.4506	2.4019	0.4501	1.5372	3.2915
AsiatoEurope	2.4000	2.5452	2.6340	0.5281	1.2618	3.3280
AsiatoUS	2.5000	2.7437	2.7134	0.4629	2.0550	3.6327
UKtoAsia	2.6000	2.5623	2.5774	0.2516	1.9433	3.1785
UKtoUK	2.7000	2.6263	2.6166	0.2662	2.0501	3.0453
UKtoEurope	2.8000	2.8034	2.8061	0.2071	2.4494	3.1001
UKtoUS	2.9000	2.8634	2.8184	0.3336	2.1472	3.5383
EuropetoAsia	3.0000	2.9251	2.8600	0.2537	2.5498	3.5025
EuropetoUK	3.1000	3.0564	3.0635	0.2740	2.3765	3.6628
EuropetoEurope	3.2000	3.1311	3.1263	0.2801	2.6017	3.5899
EuropetoUS	3.3000	3.3143	3.3325	0.2389	2.6932	3.7455
UStoAsia	3.4000	3.3771	3.3734	0.2564	2.8271	3.8506
UStoUK	3.5000	3.5101	3.4904	0.3040	3.0433	4.1931
UStoEurope	3.6000	3.5612	3.5765	0.2377	3.1232	4.0486
UStoUS	3.7000	3.5527	3.5277	0.2504	3.0080	4.1036

Table 2: Monte Carlo Results—Caller and Receiver Functions

Variable	**Actual**	**Mean**	**Median**	**StdDev**	$\alpha = 0.025$	$\alpha = 0.975$
AdmintoResearch	3.8000	4.6602	4.1578	1.7681	2.8313	10.2658
AdmintoSales	3.9000	4.6963	4.3894	1.7028	3.1709	10.0637
AdmintoTrading	4.0000	4.9913	4.5655	1.7035	3.2356	10.3027
ResearchtoAdmin	4.1000	4.1057	4.0544	0.2813	3.5606	5.1116
ResearchtoResearch	4.2000	4.2620	4.2427	0.2601	3.8010	4.7657
ResearchtoSales	4.3000	4.0882	4.1378	0.2316	3.4896	4.4939
ResearchtoTrading	4.4000	4.3330	4.3282	0.2365	3.8786	4.8233
SalestoAdmin	4.5000	4.5268	4.5353	0.2502	4.1398	4.9848
SalestoResearch	4.6000	4.4569	4.4223	0.2745	4.0923	5.0474
SalestoSales	4.7000	4.6487	4.5828	0.2965	4.1265	5.1283
SalestoTrading	4.8000	4.7266	4.7060	0.3364	4.0224	5.2424
TradingtoAdmin	4.9000	4.8834	4.8786	0.2126	4.4984	5.2424
TradingtoResearch	5.0000	4.8672	4.8779	0.1879	4.5173	5.1781
TradingtoSales	5.1000	5.0267	5.0382	0.2415	4.5510	5.5406
TradingtoTrading	5.2000	5.1458	5.1685	0.2379	4.5488	5.5492

Table 3: Monte Carlo Results—Caller and Receivers Titles

Variable	**Actual**	**Mean**	**Median**	**StdDev**	$\alpha = 0.025$	$\alpha = 0.975$
AssociatetoVP	5.3000	6.5366	7.4483	1.8900	3.4330	9.5059
AssociatetoDirector	5.4000	6.6047	7.5642	1.9292	3.4032	9.4869
AssociatetoMD	5.5000	6.7104	7.5770	1.9174	3.1688	8.9927
VPtoAssociate	5.6000	5.5212	5.4968	0.2640	4.9405	6.0473
VPtoVP	5.7000	5.6720	5.7062	0.3407	5.0849	6.4654
VPtoDirector	5.8000	5.6670	5.6799	0.3132	4.9597	6.2673
VPtoMD	5.9000	5.8639	5.8672	0.2821	5.4253	6.5790
DirectortoAssociate	6.0000	5.9145	5.9341	0.2355	5.3027	6.3437
DirectortoVP	6.1000	5.9923	5.9466	0.2839	5.5386	6.6651
DirectortoDirector	6.2000	6.2086	6.1701	0.2466	5.7096	6.6947
DirectortoMD	6.3000	6.1771	6.1603	0.2787	5.3511	6.6585
MDtoAssociate	6.4000	6.2242	6.1520	0.3208	5.2353	6.6983
MDtoVP	6.5000	6.4336	6.3814	0.2371	6.0233	6.9568
MDtoDirector	6.6000	6.6028	6.6263	0.2392	6.1291	7.1032
MDtoMD	6.7000	6.5975	6.5379	0.2477	6.0032	7.1056

Table 4: Monte Carlo Results—Decay Rates on Receiver Characteristics

Variable	**Actual**	**Mean**	**Median**	**StdDev**	$\alpha = 0.025$	$\alpha = 0.975$
decayAsia	1.1000	1.1527	1.1244	0.1143	0.9802	1.3788
decayUK	1.2000	1.2615	1.2636	0.1143	1.0221	1.5192
decayEurope	1.3000	1.3490	1.3627	0.0971	1.0627	1.4777
decayUS	1.4000	1.4774	1.4852	0.1253	1.2271	1.7182
decayAdmin	1.5000	1.6543	1.5870	0.2632	1.3689	2.3487
decayResearch	1.6000	1.7177	1.6422	0.2449	1.4570	2.3589
decaySales	1.7000	1.8323	1.7668	0.2307	1.6127	2.4686
decayTrading	1.8000	1.9198	1.8415	0.2130	1.6913	2.4547
decayAssociate	1.9000	2.0576	2.1790	0.2703	1.5371	2.3426
decayVP	2.0000	2.1610	2.2316	0.2208	1.6604	2.3926
decayDirector	2.1000	2.2636	2.3333	0.2202	1.8128	2.5143
decayMD	2.2000	2.3888	2.4321	0.2293	1.9137	3.0199

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

Variable	**Mean**	**Median**	**StdDev**	$\alpha = 0.025$	$\alpha = 0.975$
Intercept	-1.8559	-1.8497	0.0264	-1.9103	-1.8056
AsiatoUK	-1.1954	-1.1960	0.0667	-1.3173	-1.0233
AsiatoEurope	-1.1385	-1.1655	0.0930	-1.3015	-0.9363
AsiatoUS	-2.8845	-2.8561	0.1008	-3.1261	-2.7558
UKtoAsia	-0.3446	-0.3364	0.0685	-0.4543	-0.1915
UKtoUK	0.1446	0.1458	0.0661	0.0188	0.2199
UKtoEurope	0.5669	0.5916	0.0690	0.4307	0.6566
UKtoUS	0.0965	0.1092	0.0604	-0.0152	0.1932
EuropetoAsia	-0.9972	-0.9712	0.1316	-1.3195	-0.7980
EuropetoUK	0.5087	0.5095	0.0526	0.3813	0.6323
EuropetoEurope	1.5593	1.5512	0.0459	1.5010	1.6601
EuropetoUS	-1.6116	-1.5525	0.1573	-1.9485	-1.4412
UStoAsia	-0.8045	-0.7659	0.1226	-1.0548	-0.5599
UStoUK	1.1171	1.1128	0.0452	1.0500	1.2111
UStoEurope	0.1054	0.1129	0.0690	-0.0439	0.2510
UStoUS	2.0642	2.0752	0.0305	2.0217	2.1812

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

Variable	**Mean**	**Median**	**StdDev**	$\alpha = 0.025$	$\alpha = 0.975$
AdmintoResearch	-2.5160	-2.3658	0.3835	-3.9523	-2.3240
AdmintoSales	-2.1164	-1.9995	0.2553	-2.8293	-1.9808
AdmintoTrading	-2.1659	-2.0696	0.2784	-3.2949	-1.9801
ResearchtoAdmin	0.2430	0.2513	0.0612	0.0838	0.3372
ResearchtoResearch	0.3893	0.3889	0.0730	0.2570	0.4981
ResearchtoSales	-0.5220	-0.5148	0.0858	-0.6822	-0.3983
ResearchtoTrading	-0.7048	-0.6986	0.0892	-0.8789	-0.5657
SalestoAdmin	0.2218	0.2233	0.0492	0.0989	0.2892
SalestoResearch	-0.5510	-0.5476	0.0715	-0.6907	-0.4227
SalestoSales	0.4916	0.4977	0.0465	0.3977	0.5988
SalestoTrading	-0.2848	-0.3013	0.0513	-0.4002	-0.1869
TradingtoAdmin	1.1713	1.1851	0.0579	1.0661	1.2655
TradingtoResearch	-0.2052	-0.1908	0.0876	-0.4367	-0.0901
TradingtoSales	0.4621	0.4746	0.0716	0.3040	0.5700
TradingtoTrading	1.8229	1.8275	0.0611	1.7178	1.9151

Table 7: Static Interactions of Caller and Receiver Titles on Calling Choice

Variable	**Mean**	**Median**	**StdDev**	$\alpha = 0.025$	$\alpha = 0.975$
AssocctoVP	-0.0147	-0.0190	0.0440	-0.1292	0.0825
AssoctoDirector	-0.6394	-0.6266	0.0497	-0.7464	-0.5488
AssoctoMD	-0.8381	-0.8325	0.0782	-0.9865	-0.6911
VPtoAssoc	0.3235	0.3191	0.0291	0.2810	0.3792
VPtoVP	0.3361	0.3384	0.0265	0.2793	0.3748
VPtoDirector	0.2590	0.2558	0.0294	0.2106	0.3093
VPtoMD	-0.1437	-0.1421	0.0561	-0.2352	-0.0228
DirectortoAssoc	-0.1544	-0.1542	0.0529	-0.2409	-0.0292
DirectortoVP	0.4232	0.4356	0.0726	0.3203	0.5289
DirectortoDirector	0.3854	0.3850	0.0606	0.2586	0.4843
DirectortoMD	0.4429	0.4465	0.0496	0.3180	0.5311
MDtoAssoc	-0.8562	-0.8619	0.0633	-0.9875	-0.7270
MDtoVP	-0.5670	-0.5856	0.0700	-0.6862	-0.4318
MDtoDirector	0.0094	-0.0041	0.0679	-0.1109	0.1407
MDtoMD	0.9614	0.9593	0.0515	0.8683	1.0332

Table 8: Dynamic Decay Rates by Receiver Characteristic

Variable	**Mean**	**Median**	**StdDev**	$\alpha = 0.025$	$\alpha = 0.975$
decayAsia	-7.7561	-8.3017	1.6317	-9.8410	-4.9448
decayUK	-9.9379	-11.8607	3.5149	-14.0998	-4.0073
decayEurope	-11.7070	-13.1628	4.4734	-17.6831	-3.8583
decayUS	-8.7108	-9.0317	1.5681	-11.0208	-4.8461
decayAdmin	-7.7248	-8.1032	2.0515	-10.2060	-3.1607
decayResearch	-9.1281	-9.2071	1.4415	-11.1483	-5.2281
decaySales	-10.1359	-10.7995	3.1513	-14.3523	-5.4567
decayTrading	-11.2336	-11.5360	2.4664	-16.0255	-5.6909
decayAssoc	-10.4469	-10.4877	3.1911	-14.9434	-4.4089
decayVP	-8.3577	-8.5584	2.6715	-13.6759	-4.1194
decayDirector	-9.4891	-10.6652	3.0824	-14.4404	-4.2738
decayMD	-9.6324	-9.7235	1.9136	-13.1718	-5.3363

Table 9: Specification Test: Calling Period Length

Model	Negative LLH
One month data with one month parameters	149,575
One month data with two week parameters	149,586
One month data with two month parameters	149,593

The χ^2 critical values with 57 degrees of freedom are 75.62, 84.73, and 95.75 at the ten percent, one percent, and one-tenth of a percent levels, respectively.

Table 10: Policy Experiments: Summary of Results

Variable	**Baseline**	**Targeted**	**Uniform**
Average Number of Calls	11.6	11.8	11.5
Maximum number of Adopters	1,303	1,635	1,363
Present Value utility (mean)	589.4	643.4	599.9
Present Value utility (median type)	600.0	652.0	606.0
Present Value utility (25% type)	508.6	546.2	511.5
Present Value utility (75% type)	701.0	760.0	709.8
Discounted Value to Firm with 0.9 Monthly β			
Present Discounted Monthly Users	7,485.9	11,805.5	7,686.5
Present Discounted Calls	87,855.5	143,041.6	90,581.8
Discounted Value to Firm with 0.99 Monthly β			
Present Discounted Monthly Users	32,478.4	44,479.2	33,643.4
Present Discounted Calls	387,450.8	543,134.7	403,738.6

These counterfactual simulations were conducted by varying the starting configuration of the network. In the baseline case, the network is simulated from the starting condition of no adopters. In the targeted case, the starting network contains one entire subtype. In the uniform case, the starting network contains the same number of initial adopters as in the targeted case spread uniformly over all subtypes. The last two panels in the table give the present discounted monthly counts of users and calls. Discounting reflects the difference in the time-value to the firm of having a given number of adopters or calls today versus some point in the future. Each policy simulation consists of 50 months.

Table 11: Fixed Costs by Function and Title for Asia

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	2.604	0.002	0.567	0.001
Vice President	1.373	0.086	1.189	0.647
Director	2.517	0.002	0.671	0.002
Managing Director	3.164	1.024	3.284	3.220
Research				
Associate	3.125	0.001	0.418	0.001
Vice President	3.043	0.002	0.536	0.002
Director	2.965	0.002	0.598	0.001
Managing Director	2.658	0.004	0.801	0.004
Sales				
Associate	2.946	0.002	0.540	0.001
Vice President	2.924	0.002	0.597	0.001
Director	2.782	0.002	0.710	0.002
Managing Director	2.445	0.190	1.103	0.408
Trading				
Associate	3.031	0.001	0.505	0.001
Vice President	2.979	0.002	0.584	0.001
Director	2.782	0.004	0.732	0.002
Managing Director	2.761	0.199	0.966	0.279

Table 12: Fixed Costs by Function and Title for United Kingdom

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	2.699	0.001	0.614	0.002
Vice President	2.642	0.005	0.694	0.005
Director	2.739	0.455	1.146	0.798
Managing Director	2.791	0.562	1.802	1.245
Research				
Associate	3.047	0.001	0.655	0.001
Vice President	2.964	0.002	0.762	0.002
Director	2.698	0.008	0.914	0.010
Managing Director	2.719	0.004	0.888	0.004
Sales				
Associate	3.006	0.002	0.652	0.001
Vice President	2.995	0.001	0.703	0.003
Director	2.867	0.003	0.791	0.001
Managing Director	2.482	0.003	0.981	0.003
Trading				
Associate	3.114	0.003	0.608	0.001
Vice President	3.081	0.002	0.673	0.002
Director	3.002	0.002	0.727	0.001
Managing Director	2.665	0.004	0.907	0.006

Table 13: Fixed Costs by Function and Title for Europe

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	2.877	0.002	0.603	0.001
Vice President	3.831	1.225	3.397	3.466
Director	2.586	0.408	2.767	1.833
Managing Director	1.593	0.045	1.195	0.301
Research				
Associate	3.105	0.004	0.819	0.002
Vice President	3.567	1.003	2.454	2.412
Director	2.672	0.317	2.483	1.482
Managing Director	1.737	0.026	0.433	0.110
Sales				
Associate	3.125	0.001	0.781	0.002
Vice President	2.902	0.019	0.956	0.028
Director	3.329	1.484	2.783	5.820
Managing Director	1.963	0.376	3.221	6.479
Trading				
Associate	3.190	0.003	0.764	0.002
Vice President	3.410	0.003	0.692	0.002
Director	3.063	0.289	1.075	0.484
Managing Director	1.920	0.079	1.519	0.609

Table 14: Fixed Costs by Function and Title for United States

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	4.382	1.432	2.176	1.815
Vice President	3.277	0.916	1.227	1.467
Director	3.005	1.004	1.581	2.465
Managing Director	4.888	1.350	7.495	5.114
Research				
Associate	3.454	0.003	0.562	0.002
Vice President	3.249	0.003	0.762	0.002
Director	3.051	0.030	0.895	0.036
Managing Director	4.904	1.612	6.325	5.169
Sales				
Associate	3.282	0.003	0.623	0.002
Vice President	3.215	0.005	0.730	0.002
Director	3.161	0.002	0.767	0.002
Managing Director	3.079	0.003	0.810	0.003
Trading				
Associate	3.471	0.003	0.534	0.002
Vice President	3.438	0.003	0.630	0.001
Director	3.395	0.004	0.670	0.001
Managing Director	3.244	0.002	0.771	0.001