Heterogeneity and the Dynamics of Technology Adoption*

Stephen Ryan† and Catherine Tucker‡

March 16, 2007

Abstract

This paper analyzes the role of heterogeneity and forward-looking expectations in the diffusion of network technologies. Using a detailed data set 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 employees. We develop a new "simulated sequence estimator" to measure the extent to which employees 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 employees differ significantly in their adoption costs and network benefits. We find that employees 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 Ackerberg, Amy Finkelstein, Shane Greenstein, Thierry Magnac, Chris Snyder, Marc Rysman, 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, the Tuck Winter IO Conference, the Fourth IDEI Conference on the Economics of the Software and Internet Industries, MIT, and the 2007 NBER Winter IO Meetings.

†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 an alternative 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 employee's payoff from the technology explicitly depends on having other employees adopt the technology as well. For these technologies, variation in equilibrium beliefs can lead to differences in rates and depths of diffusion, even for identical users.

In this paper, we bridge and extend these two literatures by constructing a fully dynamic, utility-based model of technology adoption and use which allows for both individual heterogeneity and network effects. 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 employees to detailed data to 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 employees in the network use it. Employees 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 employee 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 employee receives from interacting with others; and a taste for "dynamic diversification," or the desire to interact with different employee types in sequence. The latter is motivated by the idea that the utility of making a connection to an employee may depend on whom I have interacted with previously. For example, if the employee 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 employees' 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. Using data on the universe of 463,806 videoconferencing calls made using this technology, we 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 parameters of our model which are consistent with the adoption patterns and use of the network technology within this firm. Our approach to identification is novel, as our structural model allows us to identify network effects through ex-post calling behavior. The alternative approach using instrumental variables would require finding a valid instrument for each type of employee in the firm, which is infeasible given the rich set of types in our model.

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 employee 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 work group, 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 employees value being able to communicate with a wide variety of other employees. Comparing these two policies to the baseline case of decentralized adoption will allow us to evaluate the extent to which heterogeneity in employee behavior and characteristics must be accounted for in crafting an optimal policy for jump-starting the diffusion of a network technology.

Our paper makes several contributions to the existing literature on technology adoption and network effects. Despite the fact that technology adoption is fundamentally a dynamic process, the extant 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 estimating which equilibrium out of a limited set is 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 employee 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. 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.

The last difficulty is that in any research on network effects identification is a key challenge. Much of the early empirical work 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 employee 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.

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 Technology and Data

We study adoption of a desktop-based videoconferencing technology within a single multinational bank. The primary benefit of videoconferencing is that it can improve the effectiveness of oral and written communication by adding visual cues.[1] While older videoconferencing systems failed in part because they were based on expensive and inconvenient videoconferencing rooms, the videoconferencing technology studied in this paper was 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" of making high-quality videoconferencing calls to other adopters within the

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

firm, and a "stand-alone use" of watching television on your desktop computer.

There are two institutional details which are central to our analysis. The first is that the videoconferencing component could only be used for internal communication within the firm, which means that we have comprehensive data on both the set of all potential adopters and how the technology was used after adoption. The second fact is that the bank pursued an unusually laissez-faire approach to promotion and adoption of the technology in the firm. After the bank chose this technological standard to conduct internal videoconferencing, it invested in the basic network components which would form the backbone of the network infrastructure. The bank then publicized the availability of the technology to employees and each employee independently decided if and when to order a videoconferencing unit from an external sales representative. The firm paid for all costs associated with the adoption, and through our conversations with outside videoconferencing vendor and the bank we verified that there were no supply constraints which may have restricted adoption. Though such explicitly decentralized adoption is unusual, it is not uncommon for companies to install software or IT equipment for employees and then leave it to the employee's discretion whether or not they use it.

Our data consists of two databases: personnel information for all employees as of March 2004, and complete records of videoconferencing adoption and use from the first call in January 2001 through August 2004. There were 2,169 employees who qualified as potential adopters, of which slightly over 1,600 eventually adopted the videoconferencing technology.[2] Based on data on each employee's job description, rank, and location, we classified all potential adopters into broad types. Employees were sorted into a hierarchy of Associates, Vice-Presidents, Directors, and Managing Directors. Depending on their job title within the firm employees were also assigned to one of four functions: Administration, Research, Sales, and Trading. Finally, employees were sorted into four broad geographical locations: Asia,

[2] We exclude 300 employees who left the firm during our sample period from our analysis.

Britain, Europe, and the United States.[3] The combination of these classifications results in all employees being sorted into one of 64 categories.

Figure 1 shows the count of employees by type in the firm. With the exception of Asian and European workers in Administration, there are a significant number of employees in each of our types. This firm has a diamond-shaped hierarchy, with most employees holding the second-lowest rank of Vice President, followed by a smaller number of Associates, Directors, and finally a very small number of Managing Directors. The lowest rank in the firm is concentrated among Associates, primarily working in Research in Asia and the US. It is interesting to contrast these counts to Figure 2, which enumerates the adoption rates of each of these groups. The pattern of adoption rates in the firm by August 2004 is highly regular: the higher the employee's rank in the firm, the higher the adoption rate. This graph suggests that there are significant differences in the benefit to adoption across different groups in the firm.

The call database recorded each of the 2.4 million uses of the videoconferencing technology from January 2001 to August 2004. 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 watched 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 more than two participants, which comprised roughly five percent of all calls; calls made by employees who left the firm; and calls that did not go through or ended in error.[4] Of the original 2.4 million recorded uses of the technology, this left us with a data set of 463,806 person-to-person calls. Figures 3, 4 and 5 illustrate the distribution of calls between regions, functions, and titles. These summary statistics suggest that while calls within type are very common, there is significant

[3] The Asia region also includes a small number of isolated offices in other locales.

[4] We drop the Finance/Credit Analysis division from our data due to a lack of observations.

cross-calling between all different regions, functions, and titles. Without further information on numbers of callers within each reason, it is not possible to infer if the connection utility across types varies, but it does suggest that our model of technology use should account for this possibility. We also note that our data is inconsistent with theoretical models of intrafirm communication which do not predict calling across hierarchical levels, as there are significant amounts of calls between even Associates and Managing Directors.

Figure 6 provides evidence of heterogeneity within our types by graphing the adoption patterns and number of calls for US researchers across different ranks. The left panel shows the cumulative adoption levels for each of the four ranks. It is interesting to note two things: first, cumulative adoption varies across titles, with Managing Directors adopting the technology at a much higher level than Associates; and second, not all employees within each type perceive the value of the technology equally. If they did all share the same expected value of adoption, we would expect the adoption graphs to be a single stepwise function: at some point the net benefit of adopting the technology becomes large enough, and all the employees within that type would adopt at once. Clearly this pattern is not supported in the data, and these graphs imply that any model of adoption in this setting must account for differences in adoption rates both across and within types of employees.

The right panel of Figure 6 shows the average volume of monthly calls by rank over time. It is evident that call volumes differ across types, e.g. Vice Presidents make more calls on average than Directors. Also, while the data is noisy, it appears that call volume per month is growing over time, which is intuitive given that there are more potential receivers for an employee to call later in the data. We interpret this as weak evidence of a "network effect," in the sense that intensity of use of the technology is growing the number of adopters.

3 Theoretical Model

We construct a theoretical model with the goal of providing a utility-based foundation that can rationalize variation in adoption rates and calling patterns across and within types of employees. We characterize the adoption decision as an optimal waiting problem, with each employee joining the videoconferencing network when the expected benefits of adoption exceeds the opportunity cost of not doing so. Our model is inherently dynamic, as each employee computes the expected benefit of joining the network as a sum of utility flows accruing from future network use, which critically depends on expectations about growth of the installed base. There are three basic ingredients to the model: a set of state variables describing the adoption decisions of all employees at a given time, payoffs which accrue to each employee as a function of the state variables and their actions, and the process governing changes in the state space over time. We discuss each of these components in turn.

3.1 State Space and Timing

The state space, denoted by s, consists of the adoption decisions of all employees in the firm. Each element of the state space, s_{it}, is an indicator function representing the adoption decision of employee i at time t. Each period in our model is one month, and the model has an infinite time horizon. We assume that all employees share the same discount factor, $\beta = 0.9$, when evaluating future payoffs.

As discussed above, each employee is also endowed with a set of characteristics which describe the employee's position in the firm. We assume that these characteristics are exogenous and do not vary over time. We denote the vector of characteristics of employee i by:

$$x_i = \{\hat{e}^r, \hat{e}^f, \hat{e}^t\}, \tag{1}$$

where each $\hat{e}$ is a 1×4 unit vector representing the region, function, and title of each employee, respectively. We order the regions as Asia, Britain, Europe, and the United States; the functions as Administration, Research, Sales, and Trading; and titles as Associate, Vice President, Director, and Managing Director. For example, a Vice President of Administration in Europe would be represented as $x_i = \{(0,0,1,0),(1,0,0,0),(0,1,0,0)\}$.

3.2 Per-Period Payoffs: Communications Choice

Employees derive utility from using the videoconferencing technology to make calls to other employees in the network and to watch television. We first consider the payoffs from making videoconferencing calls before discussing the payoffs from television use.

The basic building block of our calling model is the utility that an employee receives from making a call to another employee. Since we are interested in understanding the patterns of calling across types and within sequences, the utility employee i obtains from making the k-th call in a calling sequence to employee j is a function of both caller and receiver characteristics and the set of previous calls already made this month. To address the first interaction, define $\Gamma = (\gamma^r_{ij}, \gamma^f_{ij}, \gamma^t_{ij})$, where each γ_{ij} is a vector defined by the interaction of each e_i in Equation 1 with each corresponding e_j of the receiver:

$$\gamma_{ij} = (\hat{e}_{i1}\hat{e}_{j1}, \ldots, \hat{e}_{i4}\hat{e}_{j1}, \hat{e}_{i1}\hat{e}_{j2}, \ldots, \hat{e}_{i4}\hat{e}_{j2}, \hat{e}_{i1}\hat{e}_{j3}, \ldots, \hat{e}_{i4}\hat{e}_{j3}, \hat{e}_{i1}\hat{e}_{j4}, \ldots, \hat{e}_{i4}\hat{e}_{j4}). \tag{2}$$

Intuitively, the Γ function zeros out all the interaction terms which are not relevant for the connection between two given employees. Define η_j to be a 12×1 vector counting the number of times employee i has made calls to each of the characteristics listed in Equation 1. We then define connection utility as:

$$U_{ijk} = \underbrace{\theta_1 + \theta_2'\Gamma}_{\delta_1} - \underbrace{\theta_3'\eta_j + \theta_4(k-1)}_{\delta_2} + \epsilon_{ijk}. \tag{3}$$

This function is composed of three parts: a static connection utility, δ_1; a dynamic component, δ_2, which depends on the set of previous calls made in the current period; and an idiosyncratic error term, ϵ_{ijk}, which we assume is distributed Type-I extreme value with unit variance. This specification allows us to evaluate the extent to which callers value network diversity through both calls to employees with characteristics different than their own, and by calling different types of employees within a calling sequence. The utility of not making a call is normalized to zero.

The static component, δ_1, is composed of a constant, θ_1, and the Γ function which interacts the characteristics the caller and receiver. In the terminology of Jackson and Wolinsky (1996), this second component is a measure of the "link synergy" between two types of employees. This is our first measure for a taste for diversity in the network composition, as connection utilities between groups may vary considerably. For example, the connection utility between two Managing Directors in Administration may be much higher than the utility between a Managing Director in Administration and an Associate in Trading. The constant determines a baseline utility, which influences how many calls a given employee will make in any period.

The second component of the utility function, δ_2, which we term "decay functions," reflects the dynamic changes to the utility of a connection as the employee makes additional calls. The first component, $\theta_3'\eta_j$, reflects the intuition that employees may value the ability to make calls to people with a range of characteristics within a sequence. For example, a Managing Director in Administration may find it valuable to call a number of Directors in different geographic areas before making a change in policy. Alternatively, the employee 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 $\theta_3'\eta_j$ captures these effects by allowing the marginal utility of calling employee j to depend on the number of previous calls to other employees with similar characteristics. The second component of δ_2, $\theta_4(k-1)$, shifts the marginal utility of making any calls linearly in the number of calls previously made in the current month. This captures the idea that the opportunity cost of using the videoconferencing technology is increasing due to the need to attend to other work-related activities.

This model of calling utility generates three of sources of variation within and across sequences for different types of employees. First, relative differences in connection utilities between employees depends on the type of the caller and receiver, implying different groups of employees in the firm can have very different calling patterns in who they call. Second, the dynamic diversification terms in δ_2 generate variation within a calling sequence, as the marginal utility of calling the same employee type is linearly decreasing in the number of times an employee has previously called other employees with similar characteristics. Third, differences in the levels of connection utilities across employee types implies that call volumes may differ across groups.

Our decision to model utility as accruing only to the sender and not the receiver is motivated by how we think this technology is being used in the firm. We assume that the decision to have a two-way call is the outcome of a joint decision. That is, we do not model any negative calling externality which might drive a wedge between receivers and senders of calls, such as are discussed by Hermalin and Katz (2004). This assumption is backed up by anecdotal evidence from the firm that video-messaging calls were often set up in advance as a result of e-mail negotiation between the two parties. In earlier work using this data, Tucker (2006) found no evidence that there was a difference in network effects when directionality of calls were explicitly modeled in social relationships. We consider only an individual employee's payoffs from a call in our model, because the decentralized adoption policy means that we do not have data to identify the firm's payoff. Given that this is an investment bank, however, where communication is crucial for profits, it seems reasonable to assume a positive correlation between the payoffs of a video-messaging call for a pair of employees and the payoffs of the firm.

In each period, each employee makes as many videoconferencing calls as he or she desires to any other employees in the network. 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 employee in the network is to find the sequence of calls which maximizes overall utility:

$$\max_{\Omega} \sum_{k=1}^{K=|\Omega|} U_{ijk}, \tag{4}$$

where Ω denotes the set of calls made in this month. Each employee makes calls until the best marginal call has a negative utility.[5] Due to our linearity assumption on how a call's marginal utility shifts as a function of its slot in the calling sequence, δ_2 in Equation 3, the utility of a sequence is not a function of the order in which the calls were made. That is, we can reorder the sequence of calls in Ω and still obtain the same level of utility. This assumption rules out time-specific nonlinearities between any subset of calls, which implies that when solving the optimization problem in Equation 4 we need only consider the composition, and not the specific ordering, of a calling sequence. If this assumption did not hold, then employees may be strategically forward-looking in their choice of when to time certain calls.

Equation 4 highlights the important role that the error term, ϵ_{ijk}, plays in our model. First, it helps rationalize why employees do not make the same number of calls each period to the same set of receivers. The reason is that the sequence of calls generated by the

[5]We note that this assumption may drive some users to stop making phone calls even if they expected positive utility in the next set of draws. The reason is that we cannot distinguish a model where employees are forward-looking with respect to expected utility in future periods from our current model without making explicit assumptions about the timing of potential calls. If we are willing to assume that employees receive a new set of error terms once a day, we could solve for parameters which are consistent with that notion; however, imposing timing assumptions is not palatable given the nature of the technology, and that the benefits of doing so are minor with regard to predicting the sequence of observed calls.

optimization program in Equation 4 is driven by selecting the highest utility call until the best marginal call gives negative utility. Due to the presence of the error term, this process generates call sequences of random length, for example if the employee receives a particularly low set of draws on all calls.

Second, the error term plays a critical role as a foundational primitive of a network effect when a new employee joins the network. The addition of a marginal adopter is important to the installed base for two reasons: first, that new adopter may be of a different type than exists in the network, which means there are new possibilities for connection synergies between that employee and the installed base, and second, there is now one more draw from the set of stochastic connection utilities. This is important because the calling sequence that results from the optimization problem in Equation 4 is driven by order statistics: the expected value of the maximum over random utilities is increasing in the number of potential receivers. Therefore, the more employees there are in a network, even if they all share the same type, the higher is the number of expected calls. This is the underlying causal foundation of the standard reduced-form network effects model, where the usefulness of a network good is found to increase in the number of adopters, without having to assume that the number of adopters directly enters the utility function.

We do not model the length of calls. The reason is that we do not have a good *a priori* model that suggests how utility flows should change with the length of communication; one can imagine low-intensity calls of long length being as equally useful as high-intensity short calls. Secondly, if we interpret the utility of a call as the total surplus from a pre-coordinated action, then the length of call does not reveal any information about the utility of that call.

3.3 Transitions Between States: Technology Adoption

The second component of our model concerns the adoption decision of employees outside the network. At the beginning of each period, every employee who has not already adopted the videoconferencing technology can do so. Adoption is instantaneous and the employee is able to make calls immediately. We assume that it is not possible to divest the technology. This seems reasonable, given that the option value of holding the technology is always positive in our model, and we did not observe any divestitures in our data.

If an employee adopts, she can expect to use that technology to communicate with others in the network, both today and in the future. We can write the value function for each potential adopter as:

$$V_i(s_t) = \max\left\{E\left[U(\Omega_{it}) + \tau - F_i + \beta V_i(s_{t+1}; s_{i,t+1} = 1)\right], \beta E\left[V_i(s_{t+1}; s_{i,t+1} = 0)\right]\right\}, \quad (5)$$

where expectations are taken with respect to that employee's beliefs about how the network is going to evolve in the all future periods. As in Farrell and Saloner (1985), the benefits of adopting the videoconferencing technology consist of both the network benefit derived from the stream of expected discounted calling utilities, $E\left[U(\Omega_{it}) + \beta V_i(s_{t+1}; s_{i,t+1} = 1)\right]$, and the stand-alone benefit of watching television, denoted by τ_i. The stand-alone benefit varies across employees and is private information. If the employee does not adopt the technology, she receives the expected discounted continuation value. The employee solves an optimal waiting problem, adopting the technology when the benefits exceed the opportunity cost, which may include adopting in a future period.

The cost of adopting the technology consists of the time spent setting it up and learning how to use it, with the firm bearing 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, which is drawn from a distribution that is known to all employees in the firm. We assume that F_i does not change after the employee has made their initial draw and is private information to the employee.

Without loss of generality, we set $\tau = 0$, since 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 employees will be indifferent to adoption if and only if:

$$F_i = \sum_{t=0}^{\infty} \beta^t \tau_i = \frac{1}{1-\beta} \tau_i. \tag{6}$$

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

The employee makes predictions about the future evolution of the network. This expectation raises three important dynamic considerations. First, the employee 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 employees anticipate that their adoption now may spur other employees 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. Third, there is an option value in not adopting in this period. Even though all employees have rational expectations about the expected evolution of the network, the presence of private information in the fixed costs of joining the network implies there is variance in who actually joins in any period. The resolution of this uncertainty over time creates the option value of delaying adoption.

3.4 Equilibrium and Network Evolution

The adoption decision in Equation 5 depends critically on each employee's beliefs about how the network is going to evolve in all future periods. The standard notion of Bayesian Nash equilibrium obtains when all employees have beliefs which ensure that no employee has an incentive to change their action or beliefs in response. Without formally deriving any properties of the equilibria of our model, we note that such models typically a large set of admissible equilibrium beliefs. As is discussed in the next section, we do not need to solve for the equilibrium of our model even once in order to estimate its underlying structural primitives.

4 Estimation

The most direct approach to estimating the unknown parameters of our model is to take a simulated method of moments approach. After specifying functional forms for the distribution of adoption costs and the calling utility function, we would search for parameters which made simulated moments from the model match their empirical counterparts as closely as possible. Such an approach requires solving the dynamic model of Section 3 for each iteration of a nonlinear optimization program. Unfortunately, the computational burden of this approach is astronomical, as we would have to solve for the fixed point of a system of 2^N interrelated equations, where $N = 2,169$ is the number of employees in our sample. The sheer size of the state space makes it impossible to enumerate all the possible states of the network, let alone compute even just one equilibrium.

To circumvent this problem, our empirical strategy follows the approach of Bajari, Benkard, and Levin (2006), who advocate a two-step approach for estimating dynamic games. The intuition of their approach is that we can let the employees in the firm solve that dynamic program for us. Under the assumption that the employees optimize their adoption decision as in Equation 5, we find parameters such that their observed behavior is optimal. In the first step, we recover reduced-form policy functions which describe the equilibrium strategies followed by each employee as a function of the state vector. In the second step, we project these functions onto our 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 policy function addresses the question of how the network will be used by employees who have adopted the technology. We propose a "simulated sequence estimator" to estimate the calling utility parameters embedded in Equation 4, which defines the program which determines how employees use the videoconferencing technology. The second reduced form estimates the factors that measure the propensity to join the network, given the number and composition of current users.

4.1 Simulated Sequence Estimator

Our goal is to estimate the utility calling parameters which govern how employees use the network once they adopt the videoconferencing technology. 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:

$$
\begin{aligned}
Pr(\Omega, K) &= Pr(\Omega|K)Pr(K) & (7)\\
\ln Pr(\Omega, K) &= \ln Pr(\Omega|K) + \ln Pr(K). & (8)
\end{aligned}
$$

The simulated sequence estimator first estimates the composition of the call and then estimates the parameters which determine the number of calls. There are two reasons for separately estimating the parameters which govern which calls we make from the parameters which govern how many calls we make. The first reason is computational: the conditional probability of a calling sequence has a closed-form solution under our parametric assumptions, while the probability of a call length does not, and requires numerically costly simulation to calculate. The second reason is we have found that the estimator which estimates both sets of parameters jointly is badly biased in small samples. We provide Monte Carlo evidence that our simulated sequence estimator performs well in small samples in the

Appendix.

The assumption that the error term in Equation 3 is distributed type-I extreme value generates a logit probability of observing a call from employee i to employee j as the k-th call of a sequence:

$$Pr(\Omega_{ijk}; s_t, \theta_2, \theta_3) = \frac{\exp(U_{ijk}(\theta_2, \theta_3))}{\sum_{j' \in s_t} \exp(U_{ij'k}(\theta_2, \theta_3))}. \tag{9}$$

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 identifies the magnitude of the decay function for that subtype. We apply this estimator to all of the videoconferencing calls made by employees in this firm during the last three months of our data. Formally, the first step of our estimator is defined by the following maximization:

$$\max_{\{\theta_2, \theta_3\}} \sum_t \sum_i \sum_k \ln Pr(\Omega_{ijk}; s_t, \theta_2, \theta_3). \tag{10}$$

We assume that all employees of a given subtype share the same utility parameters. We could adopt a random coefficients framework, where the synergy parameters between two types is individual-specific, but this poses a selection problem. Employees with low fixed costs of adoption and employees with higher draws on their connection utilities are both likely to enter the network. Since we only observe selection for one network over a relatively short time, it is highly unlikely that we would be able to separately identify the distribution of low fixed cost types from the distribution of higher fixed cost types with high connection synergies. For this reason, we fix the connection parameters and allow the fixed costs of adoption to vary across individuals.

The second step in the simulated sequence estimator recovers the parameters which govern the length of the sequences. To solve for these parameters, we use a simulated method of moments approach. Given the extant installed base at any time, and using the calling parameters found in step one, we generate N_s independent calling sequences by repeatedly simulating the process defined in Equation 4 for each employee in each month. We then compute the expected sequence length by averaging over these simulated sequences:

$$\hat{K}_{it}(s_t; \theta_1, \theta_4) = \frac{1}{N_s} \sum_{j}^{N_s} |\Omega_{itm}(s_t; \theta_1, \theta_4)|, \tag{11}$$

where $|\Omega_{itm}(s_t; \theta_1, \theta_4)|$ is the length of the m-th simulated calling sequence for employee i in month t. We then perform the following minimization program:

$$\min_{\{\theta_1, \theta_4\}} \hat{K}_{it}(s_t; \theta_1, \theta_4) - |\Omega_{it}|. \tag{12}$$

Intuitively, we find parameters such that we match the length of observed calling sequences against the sequence lengths predicted by the process in Equation 4. Identification of these parameters is straightforward. Since θ_4 does not enter the utility function for the first call, the intercept, θ_1, is identified by the proportion of agents do not make any calls in that period. Once the intercept is known, identification of θ_4 follows directly from variation in the length of calling sequences across different types of employees, since they generically have different expected utilities of making a marginal call.

4.2 Estimating the Adoption Decision

The second policy function that we recover from the data governs the choice of videoconferencing technology adoption. We estimate the proportion of adopters of employee type m as a function of current and lagged state variables:

$$Proportion(adopt_m = 1; s_t, s_{t-1}, \lambda) = \lambda_1' x_m + \lambda_2'(x_m \otimes \nu_t) + \lambda_3'(x_m \otimes \nu_{t-1}), \qquad (13)$$

where x_m is defined in Equation 1, ν_t is a 12×1 vector enumerating the counts of employee characteristics currently present in the installed base, and the operator $\otimes$ represents element-wise multiplication.

The functional form of this policy function is guided by our model of adoption. First, we allow for the possibility that different employee groups in the firm have different propensities to join the network, as captured by λ_1. Second, we have assumed that the fixed cost of adoption is employee-specific private information that does not change over time. In conjunction with the fact that expected calling utility is weakly increasing in the size of the network, this assumption implies that the proportion of people within an employee type who adopt is nondecreasing in the size of the network. Therefore, we restrict the coefficients on these state variables, λ_2, to be positive. Third, we include lagged state variables to correct for selection. The intuition is that, all else equal, the larger the network was in the last period, the higher the fixed costs employees who are still outside the network in this period must have. We can imagine a thought experiment where we exogenously increase the size of the network in the last period. Employees who are on the margin of adoption will be induced into joining the network, while those with higher fixed costs will not. This implies that agents in future periods have systematically higher draws on fixed costs than other employees from their group who decided to adopt in previous periods. We capture this effect by including one-month lagged state variables, and restrict the coefficients on these state variables, λ_3, to be negative.

4.3 Estimating the Fixed Costs of Adoption

Once we have estimated the policy functions governing adopt and use the videoconferencing technology it is possible to estimate the fixed costs of adoption. Rearranging Equation 5, the necessary and sufficient condition for adoption is:

$$F_i \leq E\left[U(\Omega_{it}) + \beta\left(V_i(s_{t+1}; s_{i,t+1} = 1) - V_i(s_{t+1}; s_{i,t+1} = 0)\right)\right] \tag{14}$$

We assume that F_i is normally distributed with mean μ and variance σ^2, with associated cumulative distribution function $\Phi(x; \mu, \sigma^2)$. It follows directly that the probability that employee has a draw of F_i low enough to induce adoption is:

$$Pr(adopt_i) = \Phi\left(E\left[U(\Omega_{it}) + \beta\left(V_i(s_{t+1}; s_{i,t+1} = 1) - V_i(s_{t+1}; s_{i,t+1} = 0)\right)\right]; \mu, \sigma^2\right). \tag{15}$$

With exception of the mean and variance parameters of $\Phi(\cdot)$, the terms in Equation 15 are either known or computable. We can calculate $Pr(adopt_i)$, the empirical probability of adoption, from Equation 13. The first set of policy functions gives an estimate of $U(\Omega_{it})$, the expected calling utility, for any configuration of the network. The second set of policy functions describe how that network evolves over time as a function of current and lagged state variables. In combination, these policy functions allow us to simulate the evolution of the network and compute $EV_i(s_{t+1}; s_{i,t+1} = 1)$, the expected present discounted utility of joining the network in this period.

Computing the expected value of not joining the network in this period is a little more involved. In principle, one needs to solve out an infinite series of nested dynamic programming problems, starting at a time infinitely far in the future and working backward, solving Equation 5 at each point in time. However, suppose there is a time T at which the network has stopped growing. We can compute $EV_i(s_T)$ by noting that the adoption decision is a simple comparison of whether or not the expected benefits of joining the network exceed F_i. Since the network is not growing, there is no reason to delay adoption to a future period; an employee either adopts now or never adopts. Once we have this terminal value, we can solve the value function backwards to the current time period, thus approximating $EV_i(s_t)$. We think this is a reasonable approach in our application, since most of the uncertainty about the evolution of the system has been resolved by period 10; we use this time horizon in computing the continuation value of not joining in the current period.

To recover the parameters underlying the distributions of fixed costs, we estimate Equation 15 by forming the following moment:

$$\min_{\{\mu,\sigma^2\}} \sum_M \sum_T \sum_i (Pr(adopt_{im}) - \Phi\left(\cdot; \mu_m, \sigma_m^2\right)). \tag{16}$$

We index μ and σ^2 by m to emphasize that we estimate the fixed cost distributions separately for each type of employee in the firm. The identification of these parameters is straightforward: μ is identified off of the percentage of employees who join the network at $t = 0$, and σ is identified by the rate at which subsequent employees join as the network grows.

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).

5 Results

This section reports and discussed the implications of our two sets of results. The first set of estimates are for the calling parameters which capture the per-period payoffs from adoption. The second set of estimates are for 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 calling utility parameters in Equation 3. Tables 3 through 6 display the results of our estimates.

Table 3 illustrates that with the exception of UK-based employees, who prefer to call other employees from Europe, employees prefer to call other employees within their region. Given that this within-region propensity is larger for employees in the US and Asia we speculate 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.

Table 4 illustrates that employees on average, exhibit a preference for calling employees in similar functions to themselves. On average 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 also 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.

Most of the theoretical literature on hierarchies and firm organization pose abstract models of why the need to process information may lead a firm to organize itself into a hierarchy. See for example Garicano (2000), Radner (1992), Radner (1993), Van Zandt (1999), and Wernerfelt (2004). These theories predict that communication in a firm will be predominantly directed up and down a hierarchy. By contrast, our results on calling preferences across the hierarchy in 5 suggest a more nuanced pattern of communication. Managing Directors are most likely to call each other and less likely to call employees further down the ladder of command. Other employees appear to have similar preferences for calling other employees in similar positions in the hierarchy or one step above them. However, they are less likely to call employees either lower in the hierarchy than they are or a step removed above them in the hierarchy. These results augur against the technology being used successfully for monitoring, but instead suggest that it is being used to exchange information about tasks assigned to one layer of the hierarchy or occasionally gathering information from a superior one rung up in the hierarchy.

One of the auxiliary aims of this paper is to provide some empirical evidence on communication patterns within a firm. Lack of data has meant that most of the literature on hierarchies and firm organization is theoretical. The estimates presented in tables 6 through 5 have the advantage over existing empirical research on the organization of firms such as Rajan and Wulf (2006) and Garicano and Hubbard (2003) that we study and model actual communication flows. This means that we are able to provide evidence on whether hierarchies are fulfilling the communications role assigned to them by theory. The obvious caveat of our study is that to get this level of detail in data we have to follow the example of researchers such as Baker, Gibbs, and Holmstrom (1994) and study the internal communications of only one firm.

The results for the parameter $\hat{\theta}_3$ which captures the role of the dynamic decay rates are displayed in Table 6. 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. We speculate that this reflects the fact that these workers are more on the periphery of the firm and that their roles are more to provide one-time information than engaging in consistent exchange of information.

5.2 Fixed Costs of Adoption

Tables 7 through 10 display the results of our fixed cost estimates. There are a few patterns to highlight. What is striking is how costs of adoption vary across the hierarchy and function in the four regions. For example, generally we see declining fixed costs of adoption as we move up the hierarchy, with managing directors having the lowest costs of adoption. However, in Europe we see that managing directors actually had the highest fixed costs of adoption. We speculate that this is because the managing directors with the greatest operational responsibilities tended to be located in Europe and as a result the value of their time was high. On average, we see that employees in the US and Asia had the highest fixed costs of adoption, but that in these countries administrators have lower net costs of adoption. We speculate that this reflects the fact that these administrators occupied positions which tended to be more peripheral to the central working of the bank and that their time costs were lower as a result.

If we compare these results with the results for calling choices in Tables 3 through 5, 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, in the case of Managing Directors of Research in the Europe, while callers received high utilities from calling them, they also had some of the relatively highest fixed adoption costs [6].

[6] These results on decay rates rely particularly heavily on the assumption that the appropriate calling period to examine is a calendar month. In earlier versions of this paper we contrasted the results obtained from using a month with the results of using longer and shorter periods. We found no evidence that the

One powerful benefit of the reduced form approach is that we implicitly capture and model correlations in adoption in clusters which are the result of ad-hoc managerial coordination. However we do not control for unobserved heterogeneity in the form of shocks which are common across groups of individuals, e.g. all employees 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 time-varying shocks to a particular group, as there are shifters that change the propensity of one employee in that subset to join the network separately from other employees in the network. For example, administrators in Asia still vary within that group by title. Variation in the network composition shifts the propensity of employees in this group differently depending on title. This variation is enough to control for group-time specific shocks to the adoption decision. Due to limited variation in data we take a simpler approach, but in principle our approach is robust to a wide range of unobserved heterogeneity.

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",

choice of a month influenced our results.

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 employees 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 employees 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 employees 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 that 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 employee. 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 employee's corresponding subtype-specific policy function. If the policy function indicates that the employee 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 employees in the network. This calculation is greatly simplified by the fact that it is possible to do this on a subtype basis, rather than employee by employee. 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 which have identical outcomes. This occurs because the adoption of that group is considered particularly valuable to the overall network, and their high fixed costs of adoption made them unlikely to adopt without the initial policy. As a result targeting this group has the largest impact on employees; adoption decisions and their expectations about how the network will evolve. The results for the uniform policy suggest that a broad-based adoption process may be highly inefficient, since it doesn't target employees with high fixed costs of adoption and consequently does not alter expectations of how the network will evolve.

We calculate the expected discounted monthly utility for each subtype across the three policies. We report the mean utility for the population of employees, and also report utilities by quartiles. The uniform policy improves over the baseline case. This reflects the taste for diversity in the network which is apparent in the decay rates reported in table 6.

However, the utility for employees which results from the targeted policy is higher yet. In this case, there is an increase of over 9% 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 employees, 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 illustrates that in this setting the utility benefits of changing the number of people in the network by targeting those with high fixed costs outweigh trying to encourage diversity in the network.

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 employees 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 25% 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.

It is important to note, however, that we have assumed throughout this discussion that it is costless for the firm to place employees in the installed base in the first period, and that such policies can be accomplished by fiat without any compensation towards the employees concerned. If employees had to be compensated the targeted policy would prove more expensive to the firm relative to more diffuse policies merely because of the higher fixed costs involved in the targeted policy.

7 Conclusion

This paper prevents empirical evidence on the importance of understanding the role heterogeneity in the dynamics of technology adoption. In doing so the paper combines an older literature which sought to explain s-curve diffusion patterns by user heterogeneity and a newer literature on network effects, which emphasizes that employees' utilities from technology adoption are often interdependent. We bridge and extend these two literatures by estimating a model of technology and usage for forward-looking employees and that explicitly models heterogeneity over adoption costs, network effects and usage behavior. We estimate this model using unusually detailed data on 463,806 calls made after the introduction of a videoconferencing technology in a large investment bank. We quantify how different types of heterogeneity affect network evolution and use, and analyze two common policies which are used to jump-start network technology diffusion. We find evidence that in this case it was better to cater towards employees' taste for diversity by explicitly targeting a group with high fixed costs of adoption which had a large impact on expectations about network size and evolution, as opposed to policies more explicitly geared to ensuring diversity in the network.

A 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.[7] 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 simulated sequences our objective function has flat spots, a well-known problem 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.

References

Ackerberg, D. A. and G. Gowrisankaran (2006). Quantifying equilibrium network externalities in the ACH banking industry. Forthcoming in *RAND Journal of Economics*.

[7]Results available upon request from the authors.

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.

Baker, G., M. Gibbs, and B. Holmstrom (1994). The internal economics of the firm: Evidence from personnel data. *The Quarterly Journal of Economics 109*(4), 881–919.

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.

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*.

Hermalin, B. E. and M. L. Katz (2004, Autumn). Sender or receiver: Who should pay to exchange an electronic message? *RAND Journal of Economics 35*(3), 423–447. available at http://ideas.repec.org/a/rje/randje/v35y20043p423-447.html.

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.

Radner, R. (1992, September). Hierarchy: The economics of managing. *Journal of Economic Literature 30*, 1382–1415.

Radner, R. (1993, September). The organization of decentralized information processing. *Econometrica 61*, 1109–1146.

Rajan, R. and J. Wulf (2006). The flattening firm: Evidence from panel data on the changing nature of corporate hierarchies. Forthcoming in Review of Economic Statistics.

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. (2006, January). Interactive, option-value and domino network effects in technology. Mimeo, MIT.

Van Zandt, T. (1999). Real-time decentralized information processing as a model of organizations with boundedly rational agents. *Review of Economic Studies 66*(3), 633–58.

Wernerfelt, B. (2004). Organizational languages. *Journal of Economics & Management Strategy 13*(3), 461–472.

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	0.2801	0.3805	0.3380	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	0.2899	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	0.2294	0.3800	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	0.2077	0.1851	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	0.2351	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	0.2496	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	0.1556	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	0.1652	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	0.1676	0.1636	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	1.1990
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	1.0987	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

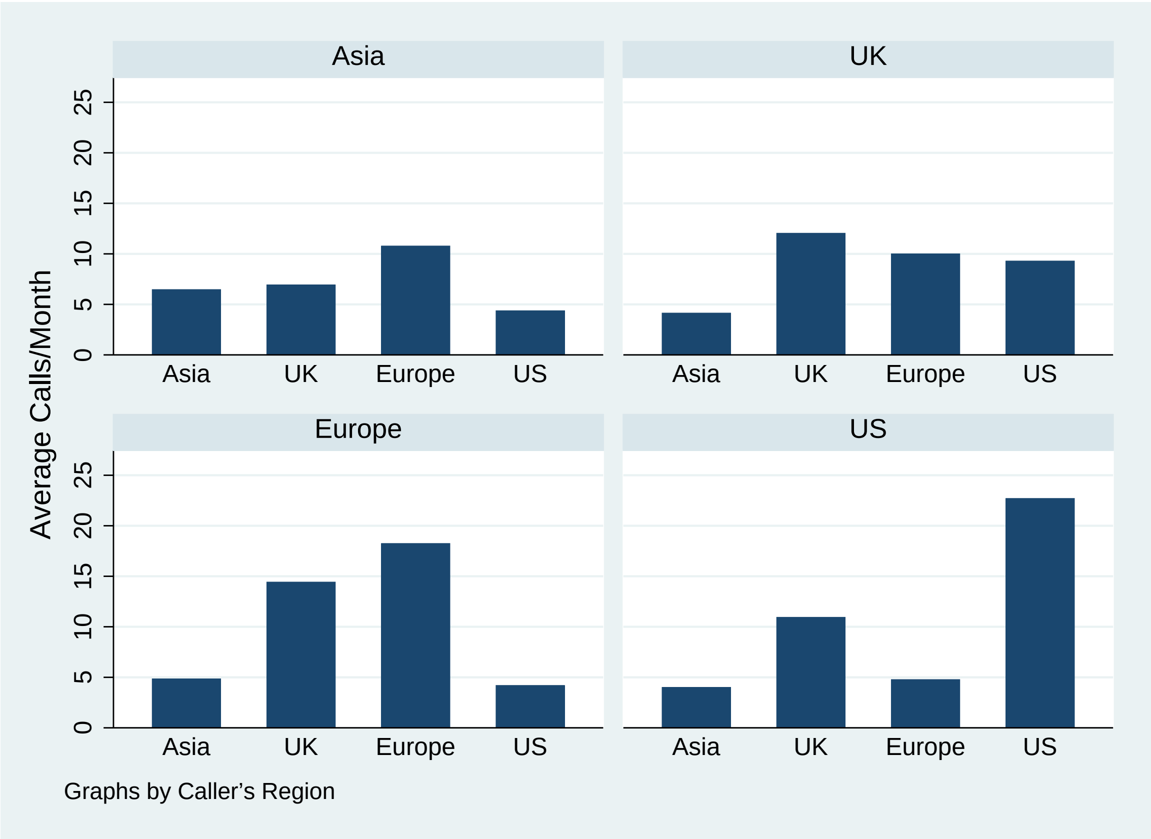

Figure 4: Calls Across Functions

Figure 5: Calls Across Titles

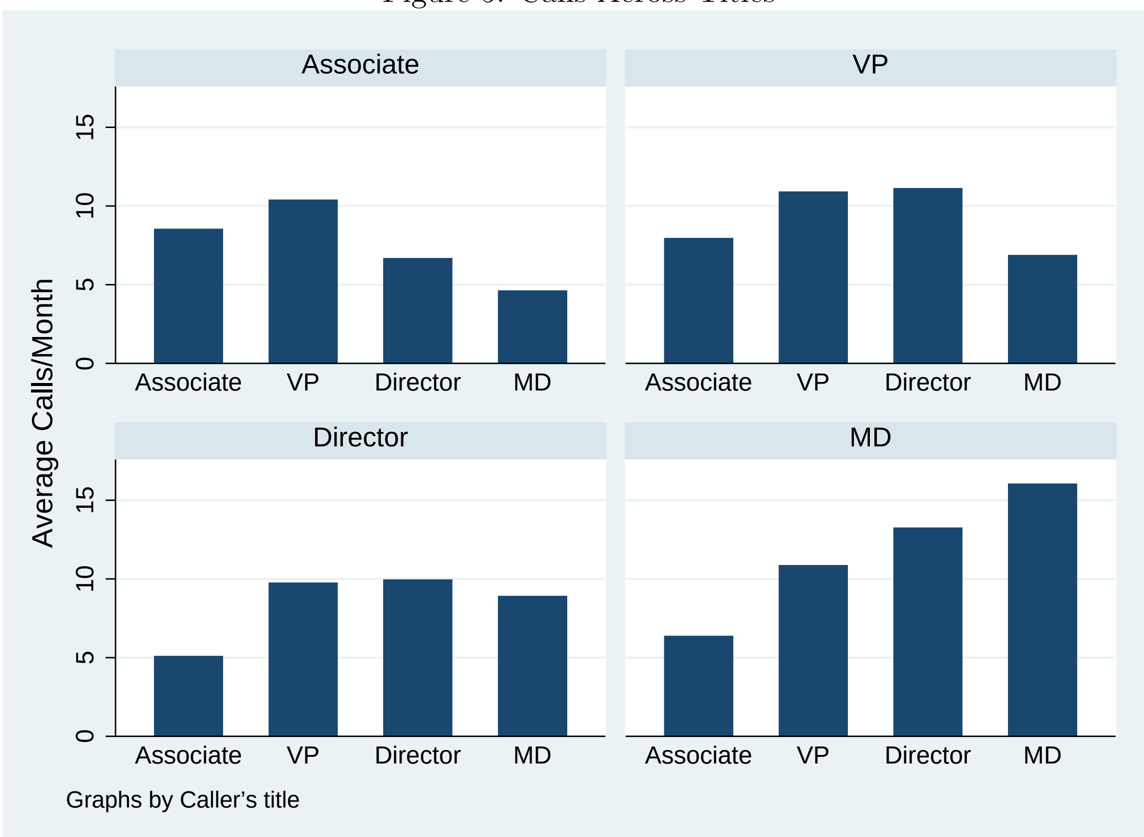

Figure 6: Cumulative Adoption and Call Volume for US Researchers

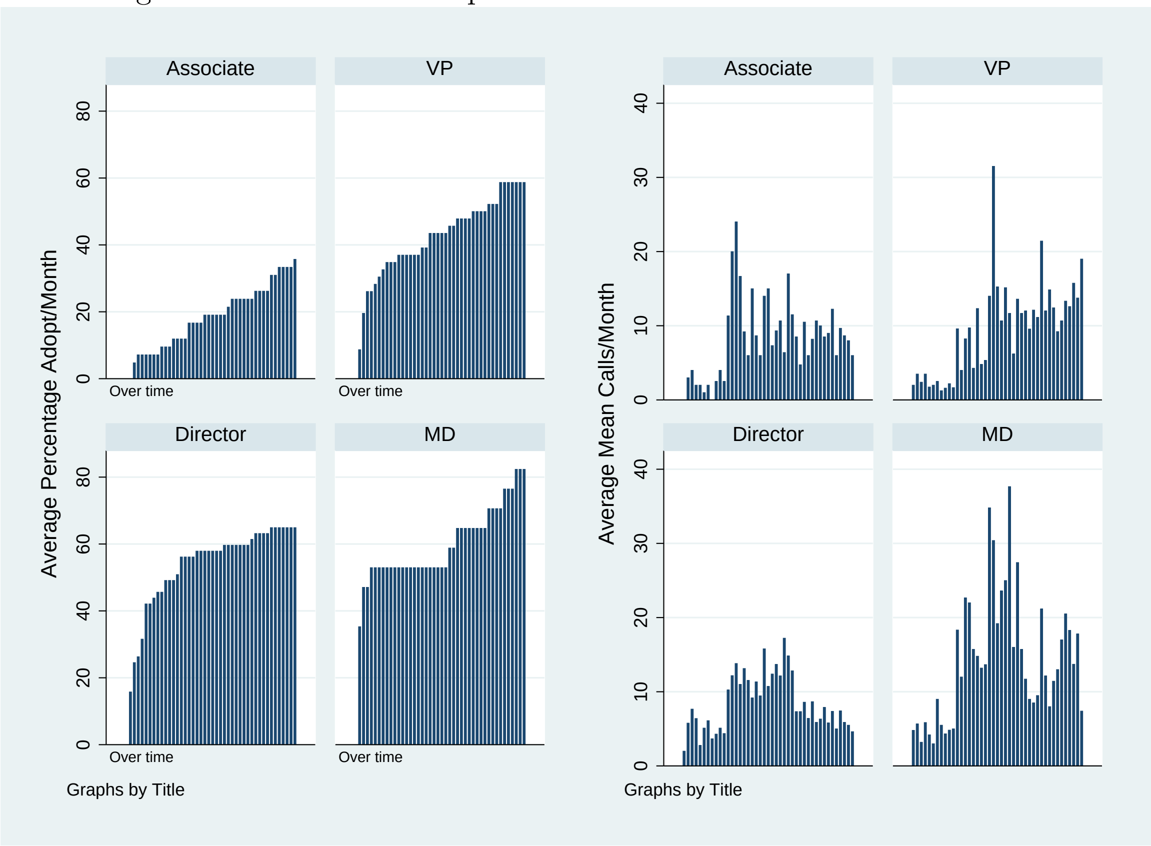

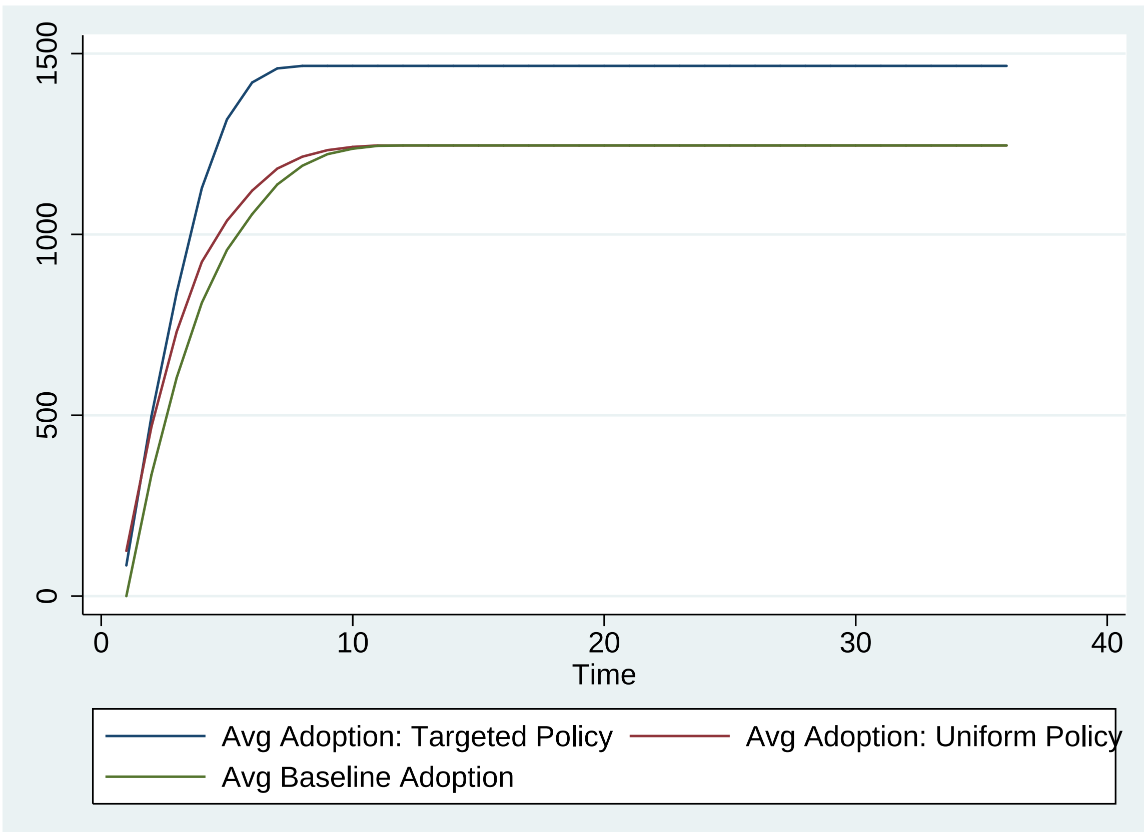

Figure 7: Adoption: Targeted vs Uniform

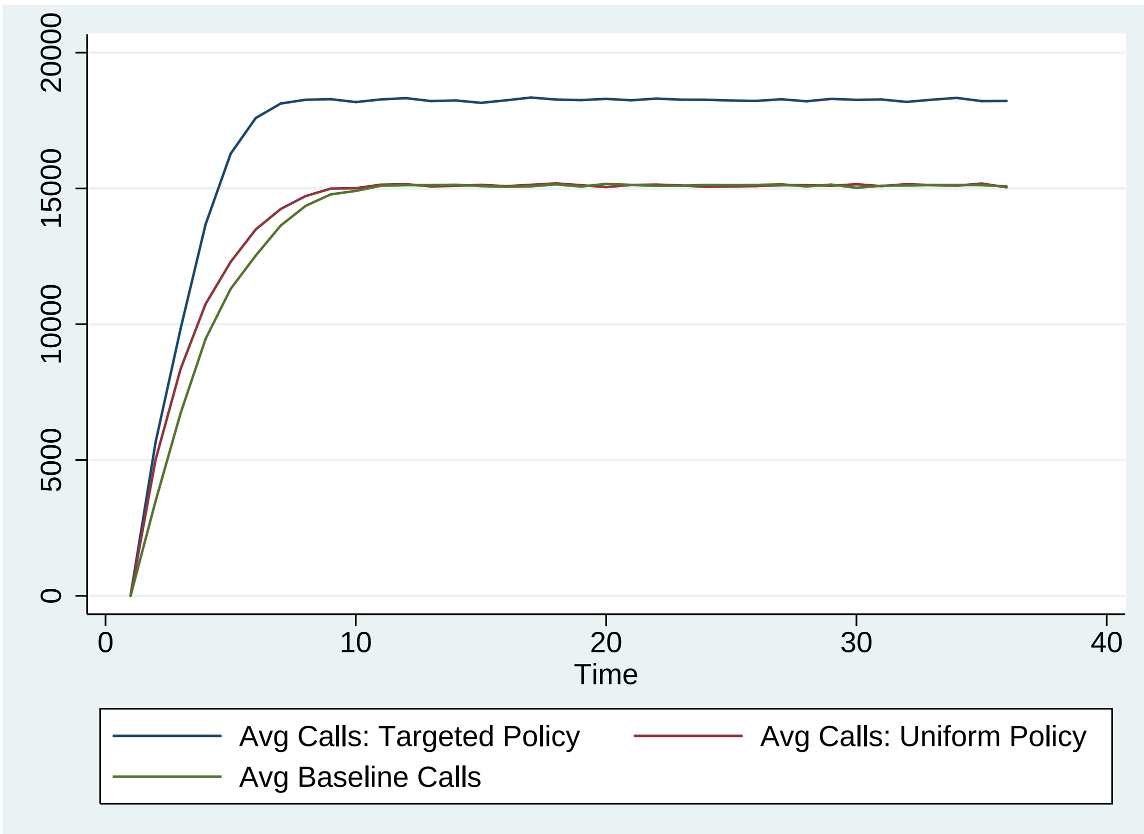

Figure 8: Calls: Targeted vs Uniform

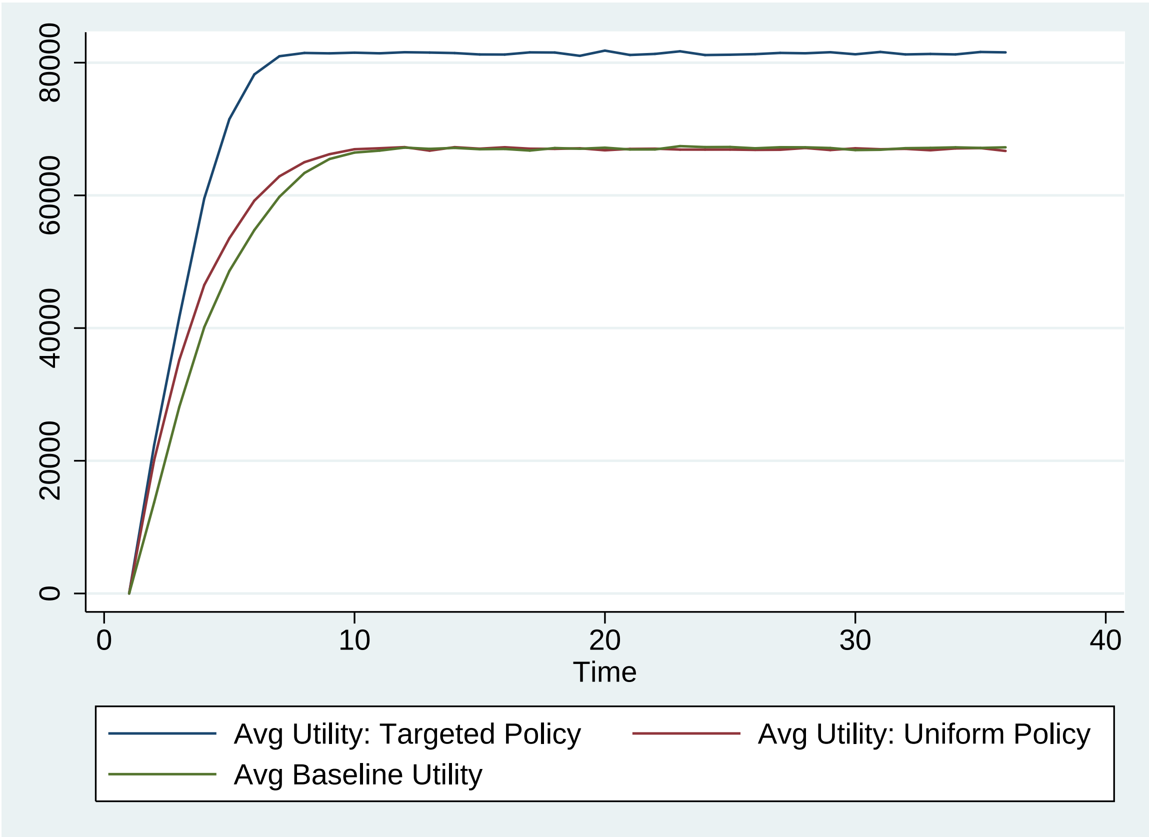

Figure 9: Utility: Targeted vs Uniform

Table 3: 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 4: 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 5: 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 6: 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 7: Fixed Costs by Function and Title for Asia

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	0.310	0.043	1.255	0.002
Vice President	-0.677	0.044	0.970	0.023
Director	0.693	0.038	1.074	0.009
Managing Director	-0.006	0.042	1.218	0.011
Research				
Associate	2.221	0.017	0.535	0.021
Vice President	2.193	0.016	0.558	0.020
Director	1.576	0.024	1.072	0.016
Managing Director	0.727	0.015	1.061	0.003
Sales				
Associate	1.868	0.021	0.873	0.017
Vice President	1.737	0.019	0.964	0.014
Director	0.963	0.016	1.009	0.004
Managing Director	0.190	0.017	1.248	0.003
Trading				
Associate	2.079	0.019	0.680	0.020
Vice President	1.533	0.019	1.098	0.010
Director	1.188	0.018	0.967	0.003
Managing Director	0.550	0.014	1.132	0.049

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

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	1.306	0.054	0.953	0.024
Vice President	0.629	0.041	1.093	0.024
Director	0.739	0.038	1.062	0.009
Managing Director	-0.400	0.046	1.094	0.021
Research				
Associate	1.772	0.020	0.945	0.016
Vice President	0.699	0.013	1.070	0.003
Director	0.387	0.017	1.248	0.002
Managing Director	-0.003	0.017	1.203	0.007
Sales				
Associate	1.332	0.015	0.943	0.003
Vice President	0.208	0.016	1.250	0.003
Director	0.530	0.016	1.191	0.057
Managing Director	0.151	0.015	1.245	0.004
Trading				
Associate	1.221	0.016	0.963	0.003
Vice President	0.585	0.013	1.100	0.004
Director	0.537	0.018	1.171	0.061
Managing Director	-0.137	0.017	1.176	0.007

Table 9: Fixed Costs by Function and Title for Europe

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	0.810	0.036	1.043	0.008
Vice President	0.862	0.042	1.032	0.010
Director	0.480	0.047	1.226	0.041
Managing Director	2.540	0.007	0.128	0.001
Research				
Associate	1.543	0.020	1.092	0.012
Vice President	1.113	0.019	0.980	0.003
Director	0.459	0.021	1.239	0.003
Managing Director	0.242	0.019	1.243	0.004
Sales				
Associate	1.824	0.023	0.901	0.018
Vice President	1.080	0.015	0.984	0.003
Director	0.591	0.015	1.097	0.003
Managing Director	-0.031	0.017	1.205	0.007
Trading				
Associate	1.612	0.020	1.047	0.014
Vice President	1.119	0.018	0.978	0.003
Director	0.959	0.016	1.010	0.004
Managing Director	-0.088	0.019	1.194	0.008

Table 10: Fixed Costs by Function and Title for USA

Subtype	Mean	StdDev	Variance	StdDev
Administration				
Associate	1.376	0.056	0.964	0.055
Vice President	0.969	0.034	1.007	0.008
Director	0.692	0.033	1.074	0.008
Managing Director	-0.042	0.037	1.208	0.011
Research				
Associate	2.348	0.014	0.394	0.017
Vice President	1.270	0.011	0.958	0.002
Director	0.648	0.011	1.084	0.003
Managing Director	0.169	0.012	1.237	0.003
Sales				
Associate	1.724	0.011	0.975	0.008
Vice President	1.334	0.009	0.944	0.002
Director	1.014	0.011	0.997	0.002
Managing Director	0.693	0.007	1.072	0.002
Trading				
Associate	2.191	0.013	0.553	0.017
Vice President	2.002	0.015	0.759	0.015
Director	1.708	0.016	0.984	0.011
Managing Director	0.853	0.009	1.033	0.002

Table 11: Policy Experiment Results

Variable	Baseline	Targeted	Uniform
Average Number of Calls	12.011	12.352	11.701
Maximum number of Adopters	1246	1466	1246
Present Value utility (mean)	392.766	420.042	409.786
Present Value utility (median type)	353.783	386.349	371.797
Present Value utility (25% type)	270.43	285.985	277.501
Present Value utility (75% type)	495.832	533.61	522.935
Discounted Value to Firm with $\beta = 0.9$			
Present Discounted Monthly Users	7989.307	10075.749	8496.759
Present Discounted Calls	95250.675	123249.251	100198.508
Discounted Value to Firm with $\beta = 0.99$			
Present Discounted Monthly Users	33645.408	40621.199	34349.059
Present Discounted Calls	405623.719	503114.059	413060.789