Interaction- and Non-Interaction-Based Network Effects in Technology Adoption

Catherine Tucker*

August 31, 2005

Abstract

Network effects and benefits for communication technologies are often modeled as depending on the total number of subscribers in a network. Since people often only communicate with a small subset of existing users in a network, assuming dependence on the entire network assumes that there are non-interaction based network effects. A potential adopter can receive a non-interaction based network benefit from someone they don't interact with entering the network for two reasons. Either they place an option value on being able to talk to them, or they anticipate that a large subscriber base makes those they actually want to talk to more likely to adopt.

The extent to which users then value the adoption decisions of users they don't regularly talk with is therefore a matter for empirical research. This paper uses extensive data on all potential adopters of a firm's internal video-messaging system and their subsequent video-messaging patterns. The technology can also be used to watch TV. Exogenous shocks to the benefits of watching TV are used to identify the causal (network) effect of changes in the installed base on adoption decisions. I find evidence that network effects are interaction-based, and that potential adopters only react to adoption by people they wish to communicate with. This suggests that the network benefits to adding an additional user for a communications technology are more restricted in scope than has previously been supposed.

*I am very grateful to my dissertation advisors Susan Athey,Timothy Bresnahan, and Liras Erin for their numerous insights and suggestions. My thanks also go to Kenneth Arrow, Shade Greenstick, Rebecca Henderson Jonathan Levin, Roger Nol, Esteban Rossi-Hansberg, Garth Saloner, Kathryn Shaw and Scott Stern for their helpful comments. All mistakes are mine alone. Thank you to John Carlson for the data, and to the video-messaging users who agreed to be interviewed. Financial support from the John D. Olid Foundation and the Korea Foundation is gratefully acknowledged.

1 Introduction

Communication technologies are the classic example of network effects: If only one person has a telephone or video-messaging unit, they receive no network benefits. A common assumption when analyzing network and communication technologies is that all users put equal weight on the probability and value of interacting with others in the future. This is a convenient simplification for complex models of equilibria, because symmetry implies that network effects are only a function of network size. Like any theoretical abstraction, however, it does not wholly describe observed behavior. People may want to use a communication technology to communicate with only a few other people in the network. In this paper I use an extensive dataset on the adoption of an internal video-messaging system within an investment bank, to see whether patterns of interaction are reflected in the structure of network effects.

I find that potential adopters only react to changes in the installed base if the person installing is someone they wish to communicate with. They do not react particularly to the adoption of users with whom they do not have a direct communications link. As not everyone communicates with each other, adding one more user to the network does not make all other users equally more likely to adopt. I take this as empirical evidence that network effects should be thought of as interaction-based phenomena. This empirical work largely supports the contention of researchers such as Rohlfs (2001) that the relative importance to your adoption decision of someone joining the network who you did not communicate with is small.

This paper tests the hypothesis that network effects are only empirically significant for the subset of users whom a potential adopter actually interacts with over the network. I distinguish between "interaction-based" network effects, which depend on actual interactions with the network users, and "non-interaction-based" network effects, which depend on the physical effect of having more people out there to make calls with, but doesn't depend on actual interactions. Both of these call-related phenomena are distinct from indirect network effects, associated with indirect non-physical benefits from increasing the installed base such as an improvement in the quality of technical support. Two motivations are given by the literature for why there might be non-interaction-based direct network effects. First, potential adopters may value the option of talking with a wider group of people than they actually end up talking with. For example, if the emergency services install phones, I may receive utility from having the option to call them in an emergency even if I don't actually do so. Second, potential adopters may value the addition of subscribers who are friends of friends, because they value the positive effect this has on the adoption decisions of their friends. In the face of these two possibilities, the extent to which network effects are confined to users a potential adopter actually interacts with, is an empirical question.

The extent to which network effects are interaction-based is a public policy issue, because network externalities are often taken as a justification for intervention in pricing schemes and government regulations in the communications and technology networks field. Network externalities occur when network effects are not internalized by network participants. These policy interventions necessarily make assumptions about the dispersion of network effects and the potential for network externalities if these network effects are not internalized. In particular the efficiency of such policies is affected by the extent to which network effects are interactions-based. Without significant non-interaction-based network effects, the potential scope for network externalities is reduced. This may augur against policy intervention. If network effects are confined within social networks then social mechanisms may facilitate co-ordination between potential users, and mean that these network effects never become network externalities.

Establishing the extent of non-interaction-based network effects is crucial for many communications and network technology policy concerns. For example one of the key justifications for universal service provisions in land-lines is that there are network externalities involved in adding marginal subscribers to the network (Riordan (2001)).

Similar arguments are used in pricing schemes which cross-subsidize between existing users and new subscribers in cell-phone networks.

Empiricists have long recognized that the linear assumption in theoretical network effects games, while a convenient assumption, is not applicable to real world data where there are thousands of subscribers in the user installed base. It seems natural to assume that as the subscriber base increases, network effects decline since subscribers cannot not take advantage of everyone's adoption. Generally the way empirical papers have sought to deal with this is by either adding non-linearities to the installed base as an independent variable or by assuming a functional form (such as a probit) which implicitly models network effects as declining with the size of the network. On this research, the unusually detailed dataset about individual adoption and subsequent usage enables more nuanced analysis, as opposed to opposing a functional form on an aggregate measure of the subscriber base. For non-linearities to fully capture the complexity of the structure of network externalities, users must exhibit symmetric interaction patterns. For example early research by Griffin (1982) which allows the network effects to vary with N and N^2, may not fully capture the extent of network effects if the proportion of people a user typically interacts with varies across users. The lack of significance of non-interaction-based network effects, however, suggests that the right empirical approach with aggregate data is to disaggregate networks into smaller networks where interactions take place. This supports the modeling assumptions in models of banking technology adoption which break up payment networks into regional phenomena (Saloner and Shepard (1994), Gowrisankaran and Stavins (2004)).

This paper's emphasis on the relationship between network structure and network effects aims to bridge the gap between empirical analysis of aggregate data and the growing theoretical literature on the importance of network structure and topology for modeling network effects. For example Sundararajan (2004) provides one of the first theoretical models of multiple equilibria with local network effects. The social-network-learning literature also commonly makes the assumption that social learning behavior takes place locally and is directly influenced only by those you directly interact with (Gale and Kariv (2003)).

The rest of the introduction summarizes how I obtained these findings. Call data cover each of the 2.4 million video-messages made by employees in a financial firm, from the inception of the technology in January 2001 until August 2004. These calls are matched with personnel data which give the precise function and role in the firm of all the employees who could have adopted the technology. The video-messaging technology is used only for internal firm communication, so the dataset covers all potential adopters and network interactions.

Time and information constraints induced management to devolve installation decisions to employees. This means that the unit of analysis is the private benefits and consequent installation decisions of employees, not productivity benefits at the firm level.

Employee-level contact data allow an employee's installation decisions to be modeled as a function of their contacts' installation decisions. A "contact" is someone an employee would video-message, if both the contact and employee installed the technology. For those employees who did not install the technology, contacts are projected using the video-messaging patterns of similar employees who did adopt. To compare interaction and non-interaction based network effects, I look at the effect of the installed base of direct contacts who the potential adopter does interact with, and the installed base of employees who are only contacts of contacts who the employee does not interact with on adoption decisions.

The ability to pinpoint who video-messages whom is also key to identifying the causal network effect of changes in the installed base for new installation decisions. The installed base is endogenous for an employee's installation decision because it is composed of installation decisions by other employees who face similar unobserved incentives to adopt the technology. A range of external events, such as the 2002 Soccer World Cup, means that the value of TV-watching to employees varies both over time and across regions. Since similar employees have contacts in different regions, each month two employees in the same office receive different exogenous shocks to their individual installed contact base. I used these exogenous shocks to the installed base, calculated at the individual level, to identify the causal effect of an increase in an employee's installed base on their installation decision.

This identification strategy allows measurement of network effects associated two types of employees: those directly connected and those only indirectly connected with each other. I find that only when someone an employee communicates with adopts, there is a positive and significant network effect. If someone adopts who does not have a direct communications link to the adopter, then this does not have a similarly positive and significant effect. This is evidence that network benefits and consequently network effects are interactions-based and restricted to direct contacts for potential adopters.

2 Technology and Data

2.1 Technology

Installing video-messaging can improve the effectiveness of internal firm communication, by adding visual communication cues to the audio communication cues provided by telephones. Marlow (1992) describes the benefits of video-messaging as greater intimacy, geographic reach, flexibility and effectiveness in communications.

Many older video-messaging systems failed to achieve popularity because they were based on dedicated video-conferencing rooms. This research studies a new type of video-messaging technology which is attached to each employee's workstation. The end-point technology consists of three elements: Video-messaging software; a media compressor attached to the employee's computer; and a camera fixed on top of the computer's monitor. Using the language of Farrell and Saloner (1985), the video-messaging technology has a "network use" and a "stand-alone use." The network use is television-quality video-messaging calls. The stand-alone use is watching TV on a desktop computer.

The video-messaging technology is a proprietary network that can only be used for internal communication within the firm. This makes it an attractive technology for empirical research, because it is possible to obtain comprehensive data on all potential adopters.

2.2 Institutional Setting

After this financial firm chose this particular technological standard to conduct internal video-messaging, they invested in an extensive network architecture. The firm felt that informational and time constraints made it too costly to identify which employees should use the technology. So, instead, the firm decided to decentralize installation decisions to the individual employee. The firm publicized the technology to employees and each employee decided if and when to order a video-messaging unit from an external sales representative. The firm made employees eligible to adopt the technology if they held a position of Associate or higher (85% of full time employees). The supplier of the equipment had excess capacity, so capacity constraints did not affect the timing of individual employee installation decisions.

This decentralization means that the unit of analysis is the private benefits to installation for employees, as opposed to firm-level productivity benefits. Net benefits to the firm from an employee using video-messaging may be larger than net benefits to the employee. More effective communications are a part of employee effort that is hard to monitor and reward (Lazear 2000). Information asymmetries therefore mean that employees' installation benefits may be small, relative to how much the firm benefits from their using the video-messaging system. I cannot, unfortunately, quantify these firm-level productivity benefits.

2.3 Data

2.3.1 Personnel Database

There are complete personnel records for each employee in a financial firm. This personnel database contains data on all employees in March 2004. Data are available for both those who adopted video-messaging and those who did not. There are data on what kind of work the employee did in the firm. The six principal functions were business management, sales, trading, research, IT and financing.

Employees were associated with two main products: Equities (with 60% of employees) and derivatives. There were 46 further business groups within the firm, with 8 to 176 employees within them. Examples of business groups are "European equities research for chemical industries" and "Japanese equities sales web database management." The precise street address is known for each employee, in each of the 33 cities the firm operates in. For simplicity, cities are assigned to the British, North American, European or Asian/Sub-Equatorial region.

2.3.2 Call Database

A call database recorded each of 2.4 million calls made using video-messaging technology from January 2001 to August 2004, within the firm. The call database has two types of call data. For two-way video-messaging calls, the database records who made the call, to whom they made it, when they made it and how long it lasted. For one-way TV calls, the database records who made the call, to which TV channel, when and for how long.

For the empirical analysis, the call database is divided up into two parts. The first period, January 2001-July 2003, is used to examine installation decisions. The second period, August 2003-August 2004, is used to construct the communications network within the firm. The data appendix fully documents how the personnel and call data were used to construct the dataset used in estimation.

When Rohlfs (1974) first developed the theory of network effects, he stated that "In any practical problem we could never hope to have a complete empirical list of principal contacts."[1] The value of these two datasets is that they get close to what Rohlfs considered unobtainable, that is, a complete list of potential adopters, their attributes and who their principal contacts may be.

3 Conceptual Model

3.1 Latent Variable Framework

A latent variable setting means that only the installation decision $inst_{i,t}$ for employee i in month t is observed, not the benefits of installation $inst^*_{i,t}$:

$$inst_{i,t} = \begin{cases} 1 & \text{if } inst^*_{i,t} > 0 \\ 0 & \text{if } inst^*_{i,t} \leq 0 \end{cases}$$

Every month, an employee makes the installation decision which maximizes their utility given a vector of coefficients, the installed base, the stand-alone TV benefit, heterogenous benefits captured by controls $X_{i,t}$ and unobserved heterogeneity ($\epsilon_{i,t}$). All employees take other employees' behavior as given. The question of interest for this paper is establishing what effect different measures of the scope of the installed base have on individual adoption decisions.

Though the firm bears the monetary costs of installing video-messaging, the employee still bears non-monetary costs. These non-monetary costs are the time the employee spends away from their regular work to get the desktop video-messaging unit installed and the time it takes to learn how to make video-messaging calls. Since the installation costs born by an employee are predominantly time costs, these costs are

[1] Rohlfs based his pioneering work on his observation of an early video-messaging technology, the AT&T Picturephone, and its failure to achieve critical mass.

sunk. Therefore, installation of video-messaging is treated as irreversible.

3.2 Interaction- and Non-Interaction-Based Network Effects

Employees have a unique set of links with a subset of other employees (their "contacts" in the firm.) I define these contacts as the other employees they would video-message with, if both they and the contact installed the technology. These links with contacts come from pre-existing communications relations within the firm. I compare the extent to which installation decisions are affected by the installation decisions of direct contacts, and those who are only indirect contacts (once removed).

I am interested in examining whether people's installation decisions are only affected by those they talk to (an interaction-based network effect), or whether installation decisions are also affected by people they don't talk to but whose installation they may value for other usage reasons(a non-interaction-cased network effect). The most straightforward way of measuring this would be to compare how the effect on adoption decisions of changes in the entire installed base, and the effect on adoption of changes in the installed base the employee actually interacts with. Unfortunately however, this method would not give much tractability for identification, because changes in the installed base are identical across users and are therefore empirically indistinguishable from the series of time dummies I use to capture monthly firm-level shocks.

Instead I use a measure of the non-interacted-with installed base which will vary from user to user; the number of contacts of contacts the employee doesn't communicate with who have installed the technology. I pick this measure because the logic inherent in the two arguments for non-interaction based network effects, suggests that non-interaction based network effects are greater the less steps removed someone is in the communications network.

The first argument for non-interaction-based network effects is that there is an option value to being able to call someone in a network even if you don't actually end up calling them. Empirically you are more likely to call someone who is close to you in the communications structure of the firm. I am more likely to call a friend of a friend than someone totally unrelated to me. This suggests that the option value of a contact of a contact adopting should be greater than someone who is less closely related to the employee.

The second argument for non-interaction-based network effects is that you may anticipate that others' adoption will make those you want to talk to more likely to adopt. Contacts of contacts are precisely the adopters whose adoption an employee would anticipate having the greatest effect on their actual contacts, adoption decisions.

3.3 TV Stand-Alone Benefits

The extent to which employees value the TV-watching capacity of the video-messaging technology varies both over time and across region. There are two types of television employees can watch: News TV programming on CNN and CNBC which covers stories of international interest to financial markets; and local TV programming (both news and non-news) which is broadcast by country-specific channels. An example of a shock to the local news TV benefit is the coverage of the German Chancellor election on the German channel ZDF in September 2002. An example of a local non-news shock is coverage of the Soccer World Cup on the British channel BBC1 in June 2002.

The percentage of adopters watching TV in employee i's region r in month t is used to capture the TV benefit, $TV_{r,t}$. News shocks are always treated as unpredictable and only known about in the month they happen. Employees may learn about some local TV shocks, such as sporting events, the month before. Therefore $TV_{r,t}$ contains both the percentages of adopters watching "News TV" in month t and the percentage of adopters watching "Local TV" in month t and month $t+1$.

3.4 Controls for heterogeneity

$X_{i,t}$ is a set of controls which captures observable heterogeneity in net installation benefit. $X_{i,t}$ contains location effects to capture variation in net benefits for employees in each of the 33 cities. It also contains business group effects to capture variation in net benefits for employees in each of the 46 different business groups within the firm. It also controls for the position of the employee in the firm's hierarchy.

The non-monetary costs of adoption for employees may vary over time. The time trend is allowed to vary by the firm's two products (equities and derivatives). These product-specific time effects are also included in $X_{i,t}$ alongside the business group and location effects.

3.5 Unobserved Heterogeneity

Not all heterogeneity is observable. Every month, each employee draws $\epsilon_{i,t}$ from the normal distribution, $\epsilon_{i,t} \sim N(0, \sigma^2)$. A possible interpretation of this error term is as a negative cost shock, specific to that employee in that month, which affects how close that employee is to the adoption threshold. For example, a low draw of $\epsilon_{i,t}$ may reflect higher time costs for an employee who is particularly busy that month.

$\epsilon_{i,t}$ may also be correlated with measures of the installed base, since the installed base consists of installation decisions by similar employees. The next section discusses the identification strategy I employ to ensure that a true causal network effect is estimated.

4 Identification of Network Effects

I want to quantify the effect of an employee adopting on the adoption decisions of other employees. Using a positive correlation of installation behavior to measure network effects is problematic, however, if all the employees are subject to similar unobserved shocks to their common incentive to adopt. Take the case of two employees who both are instructed to install the technology by their boss. I need a clear identification strategy to avoid interpreting the subsequent correlation in their adoption decisions as a causal network effect.[2]

For identification I use shocks to the TV-watching benefit. I use these shocks as instrumental variables to identify the causal network effect of the installed base on an individual's adoption. The instrumental variables strategy exploits three kinds of variation in the data: Regional variation in the benefit of watching TV; time variation in the benefit of watching TV; and variation in the regions in which similar employees have contacts. Each month, the installed base of two employees in the same business group and location will receive different shocks, because they have contacts in different regions.[3] The instrument for each employee's installed base is the TV benefit ($TV_{r,t}$) for that employee's manager and worker contacts, weighted to reflect the region of each contact. The instrument is calculated every month for each employee.

The Soccer World Cup in June 2002 illustrates the identification strategy. Figure 3 shows how the percentage of employees who watch local TV programming varies across the US and UK in 2002. While the Soccer World Cup in June 2002 elicited great interest from employees in the UK, it did not interest many employees in the US. The World Cup is associated with a spike in installations in the UK. Figure 4 shows that the spike in installations in the US in June 2002 is dominated by employees in the US reacting to the TV-inspired installation of the technology by their contacts in the UK. This anecdote illustrates the identification strategy. I do not count all earlier adoption by i's contacts as necessarily causing i's installation. Instead I use variation in i's contacts' adoption that can be predicted by variation in the stand-alone (TV) benefit.

[2]This is similar to (Manski (1995), Manski (1993))'s distinction between contextual/correlated effects and endogenous effects in the social interactions literature.

[3]On average, fewer than 20% of employees in a business group had an identical regional composition of contacts. Figure 1 shows the wide dispersion of regional links.

5 Estimation

5.1 Adopters' contacts

To estimate the conceptual model outlined earlier, I need information on contacts and the network topology, that is, whom each employee would use the video-messaging to communicate with if they both adopted the technology. It would be ideal to establish who is a contact for whom using data on existing communication networks such as e-mail records or telephone records. Unfortunately, however, these data are not available. In their absence, I use the data on video-messaging to establish not only adoption timing, but also contact network topology.

Figures 7 and 8 show that installation leveled off in the last 12 months. This stability motivates my use of call data from August 2003-August 2004 to determine who each adopter's contacts were. Using the last 12 months as a representative communications network presumes that video-messaging did not change relations between employees in the firm.[4]

5.2 Predicting non-adopters' contacts

The call database provides data on whom the 1,294 adopters video-messaged with, but not on whom the 824 non-adopters would have video-messaged with if they had adopted. If I were to use only data for those who adopted, my estimates would have selection bias. To avoid this, I predict the links of non-adopters to contacts using the communication patterns of similar adopters.[5]

This is a sparse network. For those who adopted, out of 1.5 million potential links, there were only 23,805 actual links. This sparsity means I use directed links to predict

[4]This would be a problem if time spent using video-messaging increased the number of contacts employees have. Figure 5 compares the distribution of contacts for those who installed in 2001 with those who installed later. The relative distribution of contacts does not indicate that using video-messaging for longer markedly increased how many contacts an employee had. If anything, they decreased.

[5]The methodology used is similar to that developed by computer scientists to predict hypertext links between different web-pages ((Popescul and Ungar 2003) and (Zhou and Scholkopf 2004)).

contacts. A "directed link" is whether employee i initiates a video-messaging call to employee j. Data on whether i called j or j called i are used only for link-prediction, not for estimating installation decisions. None of the coefficients are interpreted.

An indicator variable for whether or not i video-messaged j is regressed on a vector of interaction dummies for i and j's characteristics.[6] These interaction dummies include an indicator variable for every possible combination of caller city and callee city. The interaction variable for a caller in New York and a callee in London, captures the incremental effect on the probability of a link if the caller is based in New York and the callee is based in London. There are also interaction variables for every possible combination of caller and callee title, product, product market, and function in the firm. The interaction variable between a caller who is a managing director in European convertible swaps sales and a callee who is an associate in American automobile equities research captures the likelihood of an outward link between employees with these attributes.

For each caller i, the predicted total number of outward links to all employees in the firm is calculated. This is the predicted number of out-degrees $\hat{D}_i$ for employee i. The predicted probability of each outward link is ranked from most likely (1) to least likely (2118) for each employee. Links to employees who have a rank below $\hat{D}_i$ for a non-adopter are taken as that non-adopter's unobserved outward links. Additionally, if everyone had adopted, adopters may have video-messaged non-adopters. These outward links are identified by taking the difference between $\hat{D}_i$ and the observed out-degrees for each non-adopter. The links to non-adopters whose rank is below this difference for that particular adopter are taken as the unobserved part of the adopter's outward link network.

The projected outward link network is then symmetrized. If an outward link is projected in either direction, then those two employees are considered to be linked

[6] A linear probability model is used, as there over 10,000 variables to be estimated and 1.5 million observations.

and part of each other's M_i or W_i. Figure 6 shows that non-adopters have on average fewer contacts than adopters. Since they had fewer contacts, non-adopters would have received a lower network benefit on average than adopters of the technology.

To evaluate how well this procedure predicts contacts, I redid the above, using data for adopters from August 2001-August 2002 to predict contacts for those who adopted in August 2002-August 2003. The results suggested that contacts are predicted correctly approximately 60% of the time.

There is also the potential problem that non-adopters could have had substantially different contact networks and communication patterns from adopters, if they had adopted. This is a potential source of measurement error which is hard to resolve with the existing data.

5.3 Estimation

The irreversibility of installation decisions means that estimation only includes observations where the employee installs video-messaging in that month, or has not yet installed it. Measures of the installed base, however, include all contacts' installation decisions up to and including month t. Figure 7 shows permanent installation over time by employee function. Figure 8 shows permanent installation over time by region. I use the first 2.5 years of data (January 2001- July 2003) to look at individual employee installation decisions. The dependent variable $inst_{i,t}$ is defined as whether in that month the employee used video-messaging technology for the first time. The sample is all employees who have up until that month not installed the technology. Table 1 gives summary statistics for the dependent and independent variables.

$$inst^*_{i,t} = \Delta_1 DirectContact^{Installed}_{i,t} + \Delta_2 IndirectContact^{Installed}_{i,t} + \lambda TV_{r,t} + \beta X_{i,t} + \epsilon_{i,t} \quad (1)$$

Equation (1) captures the installation decision at time t for employee i. It is estimated using Amemiya Generalized Least Squares (AGLS) estimators for probit with endogenous regressors (Amemiya (1974) and Newey (1987), eq. 5.6.) As described in section 4, both the installed base $DirectContact_{i,t}^{Installed}$ of direct contacts and the installed base of indirect contacts $IndirectContact_{i,t}^{Installed}$ are instrumented. The instrument is the average TV benefit of the employees in that measure of the installed base, $TV_{r,t}$, weighted to reflect which region r each of an employee's contacts and indirect contacts work in. The independent variables are TV shocks in i's region r, and $X_{i,t}$, which includes month-product controls, business group effects and location effects.

6 Results and Interpretation

The estimates for equation (1) is reported in Table 2. For now I will focus on column two, which focuses on the estimates which control for erogeneity. The marginal estimates for installed base on average if an direct contact installs show the technology that this increases the propensity of that employee to install video-messaging that month by 0.012. If on the other hand an employee who is a contact of a contact installs the technology then the point estimate suggests that this only increases the propensity to install by 0.0005, though this point estimate is not significantly different from zero.

The first column of Table 2 displays ordinary probit estimates of the correlation between changes in the installed base and installation decisions. These estimates are larger than the estimates in the second column, where the installed base is instrumented. This suggests that without the instrumentation strategy, correlated effects would wrongly be identified as network effects.

The stand-alone benefits for TV watching are significant for an employee's installation decision. The estimated coefficients suggest that compelling news TV has a greater impact on employee's installation than compelling local TV.

Space constraints mean that the coefficients for time-product, business group and city effects are not reported in Table 1. The estimates for the city effects indicate that employees in Edinburgh were the most likely to install the technology, followed by Amsterdam. The least likely to install were employees in Johannesburg. Estimates for each month's time effect on installation show that employees were most likely to install the technology in the first two months, and least likely to install the technology in December 2002. The product-time interactions for equities and derivatives indicate that on average employees in equities were more likely to install the technology and that they did so earlier. Estimates for business group effects show that employees in European equities sales were the most likely to install, and employees who worked on North American trading floors were the least likely to install.

Other unreported results show that the choice of August 2003-August 2004 as a representative communications period did not drive the results. Repeating the estimation using (March 2003-June 2003) and (March 2003-December 2004) as representative communications periods does not significantly change the estimates.

6.1 Network Diffusion

Most policies for network diffusion are guided by theoretical models which predict S-curve patterns of diffusion for network technologies. After a period of slow expansion, the network technology reaches the number of users in the installed base required for critical mass, and technology adoption increases exponentially. The installation patterns shown in Figures 7 and 8 show no such S-curve. The results suggest that this is because there is no single tipping point for the technology. Instead, there are many smaller individual networks which tipped in succession.

Further evidence that Figures 7 and 8 aggregate the result of a succession of tips in smaller individual networks is provided by Figure 9. Figure 9 shows that the number of other employees each installed employee videomessaged only increased slightly, if at all, from year to year. If, however, network effects increase with the installed base, the number of employees that installed employees videomessaged should also increase with the installed base.

6.2 Policy Inplications

Network externalities occur when network participants do not internalize network effects inherent in adoption decisions. One key policy issue where network externalities have been cited is the provision of universal service in telecommunications. Riordan (2001), for example, argues that because universal service encourages the adoption of the remaining marginal users for landline telecommunication services, universal service regulation can be thought of as internalizing a network externality. Cross-subsidization by existing subscribers to these marginal subscribers therefore is efficient. The findings of this paper suggest, though, that for this to a compelling argument there would have to be evidence that the existing subscribers would interact with the marginal subscribers.

Another recent case has been the competition policy surrounding termination charges for cellphone networks (Armstrong 2004), (Rochet and Tirole 2002). In (Vodafone and T-Mobile (2003)) major wireless networks argued that government regulation should allow high termination charges. One of the many arguments put forth was that these high termination charges allowed cellphone companies to facilitate internalization of network externalities, by enabling cross-subsidization towards new subscribers, for example by offering low-cost handsets. If, however, as suggested by the results in this paper, network effects were interaction-based then existing subscribers may receive no network benefits from these new subsidized network participants, if they don't interact with them.

6.3 Indirect Network Effects

This paper has focused on analyzing interaction and non-interaction based ”direct network effects”. That is the network effects which result from the direct physical effect of having more people to talk to in the network. There is a wide literature, however, which deals with indirect network effects. See for example (Gowrisankaran and Stavins 2004). Indirect network effects occur, for example, when an increase in the user base increases the level of knowledge and familiarity amongst the technical support staff, which means that adoption is more attractive for a potential installer.

One way of estimating such indirect network effects would be to estimate the effect that an increase in the installed base in someone's office would have on the installation decisions of those who work there. Unfortunately, though, there is no obvious identification strategy in this case, since unlike before it is impossible to distinguish between two colleagues in the same location in terms of the TV-watching propensity of their colleagues. The most suggestive results I can produce are straight probit estimates on the effect of installation decisions of changes in the installed office-location base. These estimates are small and not significantly different from zero. This suggests that the installed base in the building/city does not have a large effect on current installation decisions. Without a clear identification strategies these coefficients are hard to interpret. If, however, there is no correlation between an employee's installation decision and the installed base in his work group, it seems likely that there are no large indirect network effects either.

7 Conclusion

I find that the empirical structure of network effects does reflect the interaction patterns within communications networks. In the video-messaging technology studied, adoption cascades were confined to small subsets of individuals who actually interacted with each other.

Often policy is framed under the assumption that network effects for communications and network technologies depend on the total number of subscribers. My estimates suggest that only the smaller subset of individuals who a potential adopter communicates with, play a significant part in the adoption decision. If these results hold for other technologies, then this suggests that the scope for network externalities in communications may be relatively small. This may augur against policy interventions. For example, subsidizing marginal consumers by mandating universal service may not benefit all existing users of the network, as they will not necessarily value the marginal consumers' installation.

A Appendix: Data Documentation

A.1 Call Database

Employees made 1,680,120 two-way user to user video-messaging calls. The dataset includes only the 1,052,110 video-messaging calls where the callee accepted the call. Callees did not accept calls either because the ring time rang out (52%), the caller hung up before the callee answered (25%), the callee refused the call (12%) or the call met an error (9%). The average length of accepted calls was 5 minutes 46 seconds. Calls could be made to more than one employee at a time. Multi-party calls were simplified into their pairwise equivalents. For example, a three-way call is treated as three calls (one call between each two of the participants).

Employees made 754,327 one-way caller to media device calls. 741,926 of these calls were successful and included in the dataset. The empirical work makes the distinction between local and global news channels for users in the four regions. Global news channels were CNN and CNBC. Local channels for Europe were ZDF (German), ARD (German), Kanal (Swedish), ORF (Austria) and Eurosport. Local channels for Britain were ITV, SkySports, Channel 4 and BBC. Local channels for the US were CSPAN, FOX, NBC and CBS. Local channels for Asia were NT (Nippon TV), CATV (Japanese), MTV-Asia, and BBC 24 World Service.

A.2 Personnel database

The second dataset is a complete personnel database which gives the rank, function, product group and street address for both adopters and non-adopters of video-messaging, in March 2004. The dataset unfortunately does not indicate whether these personnel details changed between January 2001 to August 2004. It is more likely that an employee got promoted than changed business group or city, given their geographic immobility and group-specific human capital. Accordingly, estimates based on whether employee i is a manager should more accurately be described as whether employee i was on an upwards career trajectory which meant he would be a manager by 2004. The personnel records do not include data on employees who left the firm before 2004. Observations of these employees' calling patterns are excluded from the dataset.

The dataset excludes personnel records of employees who joined the firm after January 2001. Including them in estimation would be problematic, as they did not actively choose not to install video-messaging in earlier months. The dataset excludes the 127 employees marked as recent recruits or trainees. The unfavorable business climate from 2001-2003 means that the firm made few new appointments.

The firm ruled that only Associates or higher could officially install video-messaging. The dataset excludes all employees without a formal title. A few employees without titles installed the technology. These were mainly executive secretaries to managing directors, so were unlikely to have discretion over the timing of their installation. The dataset excludes 18 employees in Moscow, Bangkok and Athens, since the video-messaging infrastructure did not connect to these cities.

References

Amemiya, T. (1974, July). The nonlinear two-stage least-squares estimator. *Journal of Econometrics 2*, 105–110.

Armstrong, M. (2004). Two-sided markets. Mimeo, University College of London.

Farrell, J. and G. Saloner (1985). Standardization, compatibility, and innovation. *RAND Journal of Economics 16*, 70–83.

Gale, D. and S. Kariv (2003, 11). Bayesian learning in social networks. *Games and Economic Behavior 45*(2), 329–346.

Gowrisankaran, G. and J. Stavins (2004). Network externalities and technology adoption: lessons from electronic payments. *RAND Journal of Economics 35*(2), 260–276.

Griffin, J. M. (1982). The welfare implications of externalities and price elasticities for telecommunications pricing. *The Review of Economics and Statistics 64*(1), 59–66.

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

Manski, C. (1993, July). Identification of endogenous social effects: the reflection problem. *Review of Economic Studies 60*(3), 531–542.

Manski, C. F. (1995). *Identification Problems in the Social Sciences*. Harvard University Press, Cambridge.

Marlow, E. (1992). *Corporate television programming: applications and technologies.* Knowledge Industry Publications, Inc.

Newey, W. (1987). Simultaneous estimation of limited dependent variable models with endogenous explanatory variables. *Journal of Econometrics 36*, 231–250.

Popescul, A. and L. H. Ungar (2003). Statistical relational learning for link prediction. IJCAI03 Workshop on Learning Statistical Models from Relational Data.

Riordan, M. H. (2001). *Handbook of Telecommunications Economics*, Chapter 10: Universal Residential telephone service. Amsterdam: Elsevier Science.

Rochet, J.-C. and J. Tirole (2002, Winter). Cooperation among competitors: Some economics of payment card associations. *RJE 33*(4), 549–570.

Rohlfs (2001, October). MIT Publishing.

Rohlfs, J. (1974, Spring). A theory of interdependent demand for a communications service. *Bell Journal of Economics 5*(1), 16–37.

Saloner, G. and A. Shepard (1994, Autumn). Adoption of technologies with network effects: An empirical examination of the adoption of automated teller machines. *RAND Journal of Economics 26*(3), 479–501.

Sundararajan, A. (2004, July). Local network effects and network structure. Mimeo.

Vodafone, O2, O. and T-Mobile (2003). Vodafone, O2, Orange and T-Mobile: Reports on references under section 13 of the Telecommunications Act 1984 on the charges made by Vodafone, O2, Orange and T-Mobile for terminating calls from fixed and mobile networks. Technical report, Competition Commission. ISBN 0-11- 515503-1.

Zhou, D. and B. Scholkopf (2004). A regularization framework for learning from graph data. In *Workshop on Statistical Relational Learning, International Conference on Machine Learning.*

Figure 1: Number of Links across Different Regions

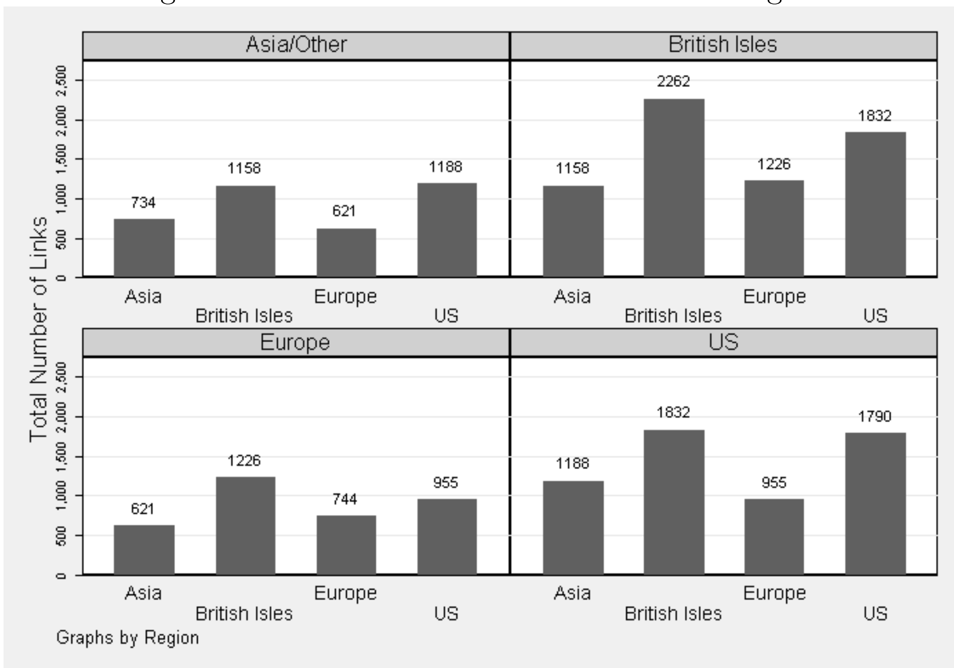

Figure 2: Number of Calls made across Different Regions

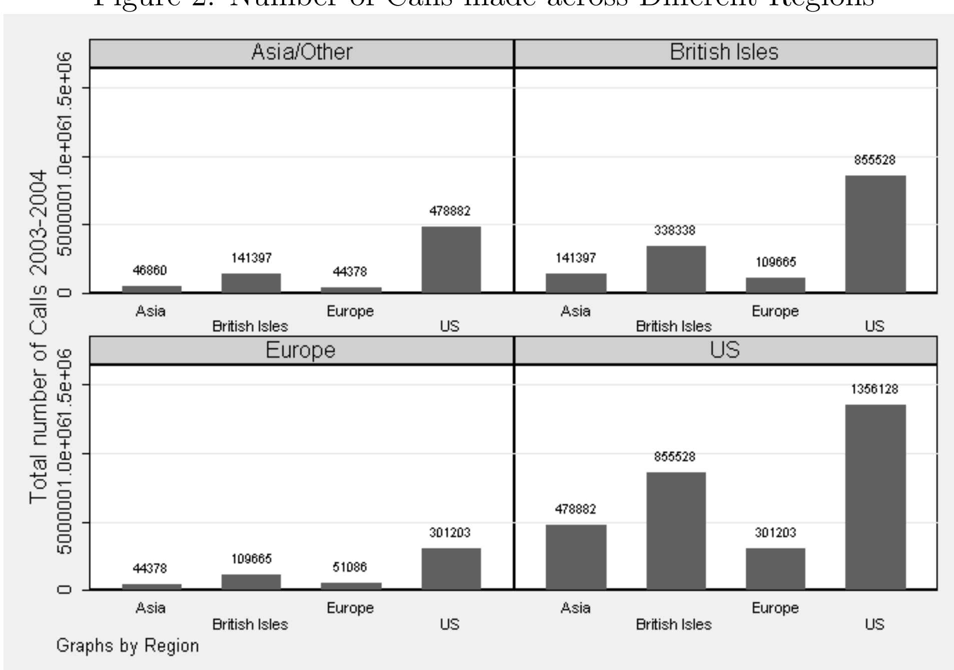

Figure 3: Relationship between New Installations and TV-watching in the US and UK, 2002

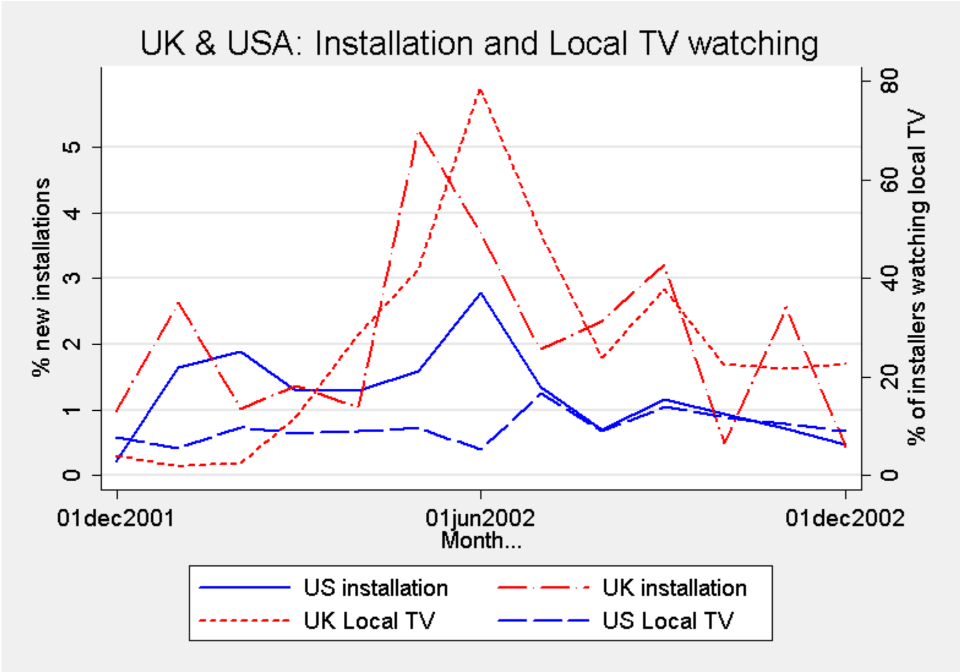

Figure 4: Relationship between New Installations and US employees having any UK contacts

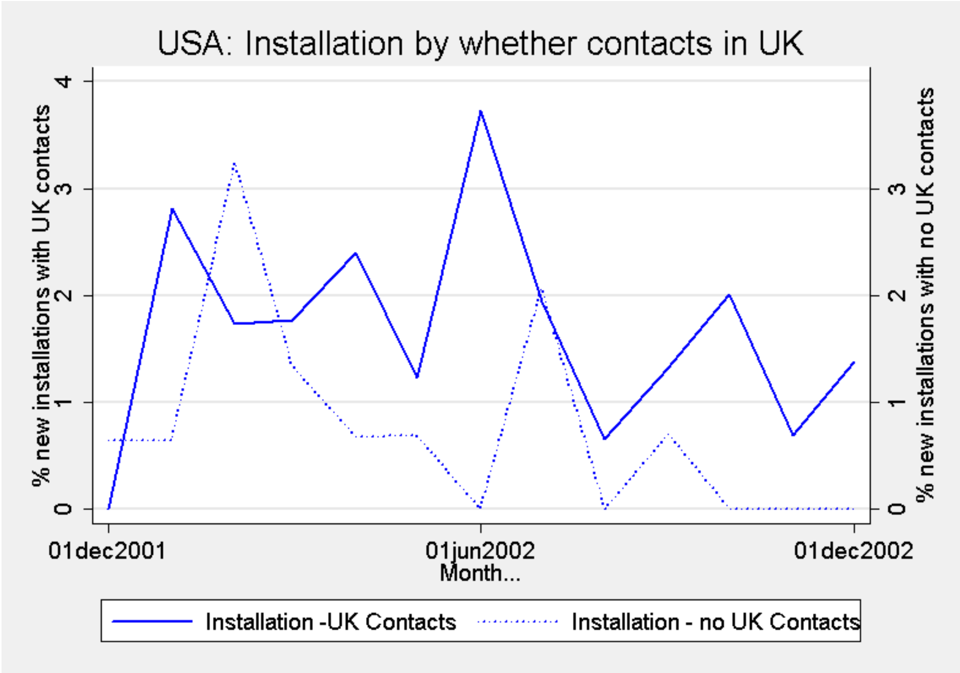

Figure 5: Comparing Distribution of Worker and Manager Contacts for Early and Late Installers

Figure 6: Distribution of Number of Contacts (both manager and worker) by Managerial Status and Adopter Status

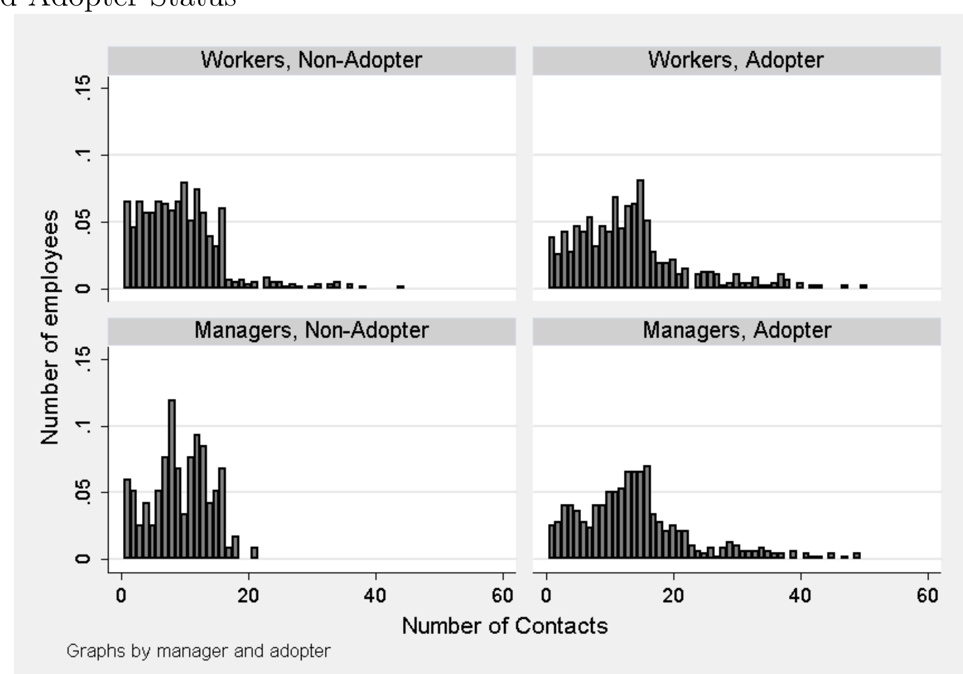

Figure 7: Installation over Time by Function

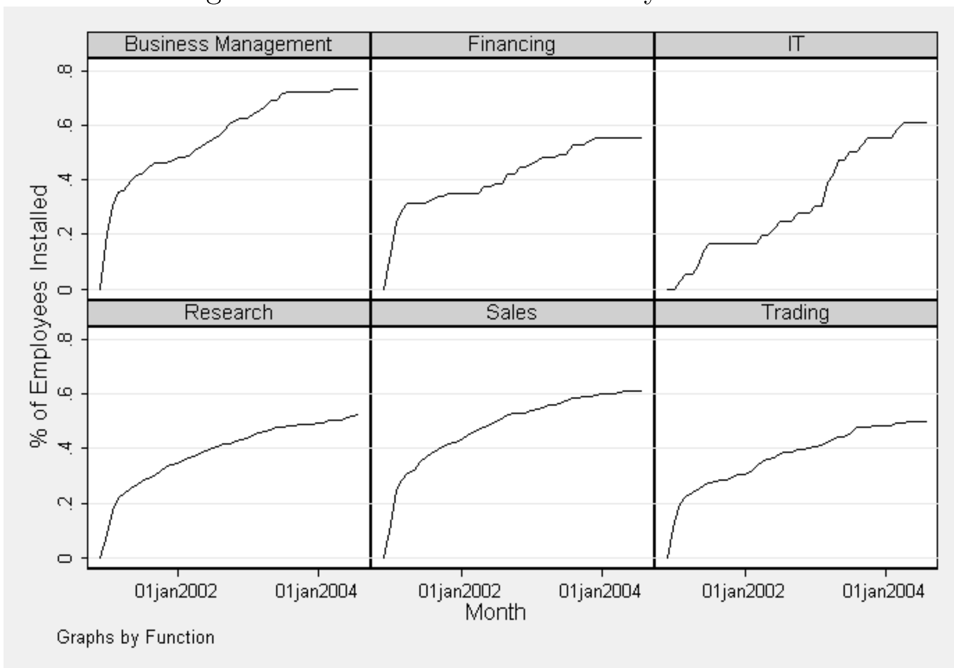

Figure 8: Installation over Time by Region

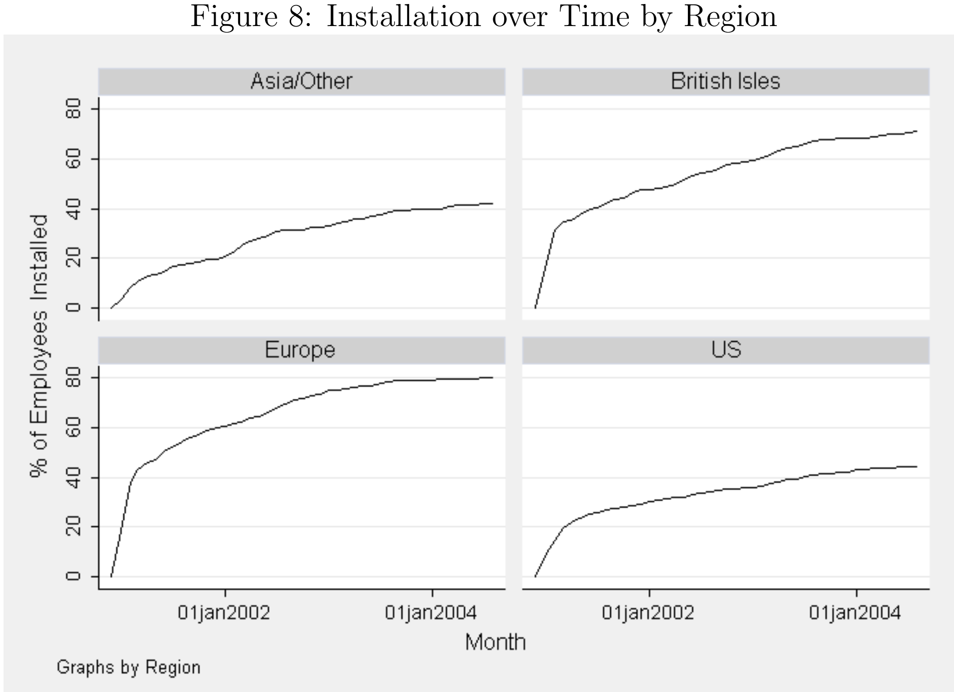

Figure 9: Stability in Video-messaging Behavior over Time

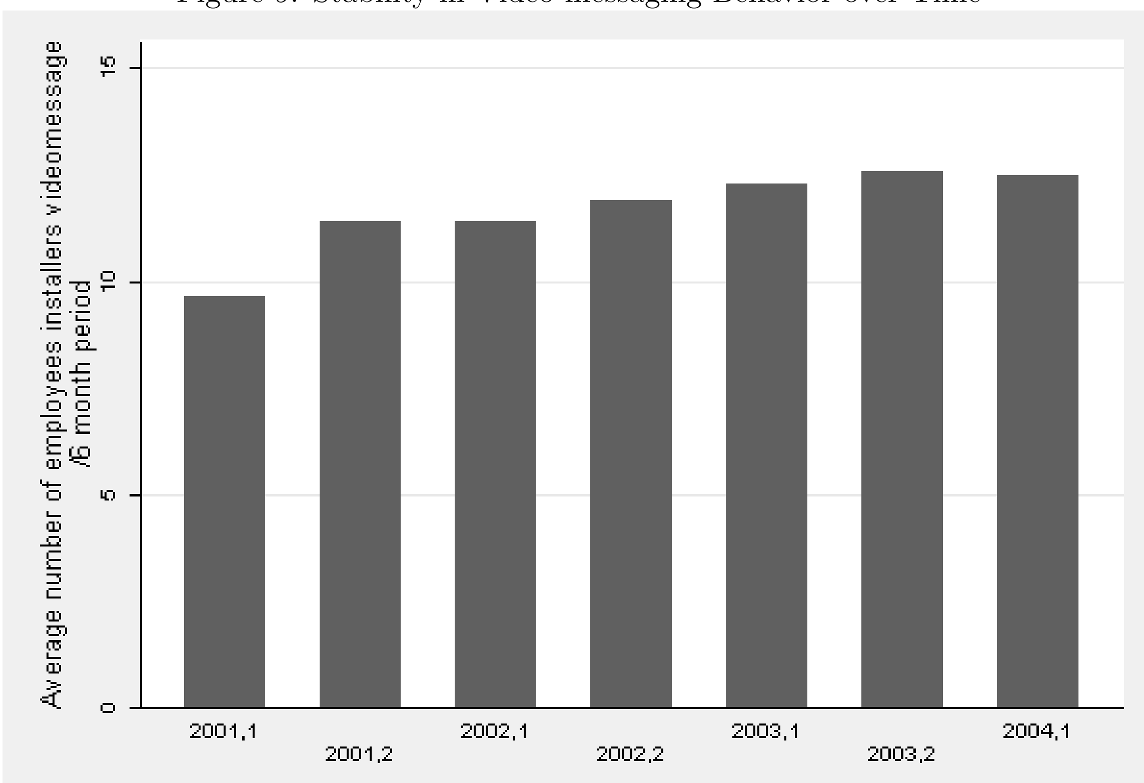

Table 1: Variable Description and Summary Statistics

Label	Variable Description	Mean	S.D.	Range
$inst_{i,t}$	Whether employee i makes a video-messaging call for the first time in month t	0.041	0.169	0-1
$DirectContact_{i,t}^{Installed}$	Number of i's contacts who have installed video-messaging by month t	2.550	3.040	0-22
$IndirectContact_{i,t}^{Installed}$	Number of i's contacts of contacts who have video-messaging by month t	4.015	3.678	0-120
$TV_{r,t}^{News}$	% of installers in i's region watching TV news in month t	0.253	0.211	0-1
$TV_{r,t}^{Local}$	% of installers in i's region watching local TV in month t	0.085	0.132	0-1

Table 2: Direct Network Effects vs Two-Degrees Removed Network Effects

Dependent Variable: Whether employee participated in a video-messaging call for the first time in that month from January 2001-July 2003.

Installation is permanent, so previous installers are excluded as dependent variables.

	Sub-set of observations used in regressions	
Variable Description	Probit	Probit-IV
Δ Confined Network Effects		
Installed Direct contacts (IV)	0.0241 (.0002)***	0.0124 (.0025)*
Non-Confined Network effects		
Installed contacts of contacts(IV)	0.0072 (.0013**)	0.0005 (.0156)
Variables for TV Benefit, Business Group, location, product-time interactions and month included		
Observations Log-Likelihood	35,207 -2640.1789	35,207 -2843.2567

AGLS estimates of **marginal effects** for probit with endogenous regressors.

Instruments are % of installers watching news-TV and local-TV in contacts' regions in that month and the next month, weighted to reflect regional distribution of contacts for each employee.

Standard errors in parenthesis below.

* Significant at the 10% level ** 5% level *** 1% level