

NBER WORKING PAPER SERIES

NETWORK STABILITY, NETWORK EXTERNALITIES AND TECHNOLOGY
ADOPTION

Catherine Tucker

Working Paper 17246
<http://www.nber.org/papers/w17246>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2011

I gratefully acknowledge financial support from the NET Institute (www.netinst.org) and the Kauffman Foundation (www.kauffman.org). All errors are my own. This paper is based on an earlier working paper entitled 'Interactive, Option-Value and Domino Network Externalities in Technology Adoption'. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2011 by Catherine Tucker. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Network Stability, Network Externalities and Technology Adoption
Catherine Tucker
NBER Working Paper No. 17246
July 2011
JEL No. L0,L86,L96

ABSTRACT

This paper investigates how the destabilizing of a social network may increase the scope of network externalities, using data on sales of a video-calling system made to an investment bank's employees and subsequent usage by these customers. The terrorist attacks of 2001 led potential customers in New York to start communicating with a new and less predictable set of people when their work teams were reorganized as a result of the physical displacement that resulted from the attacks. This did not happen in other comparable cities. These destabilized communication patterns were associated with potential adopters in New York being more likely to take into account a wider spectrum of the user base when deciding whether to adopt relative to those in other cities. Empirical analysis suggests that the aggregate effect of network externalities on adoption was doubled by this instability.

Catherine Tucker
MIT Sloan School of Management
100 Main Street, E62-533
Cambridge, MA 02142
and NBER
cetucker@mit.edu

1 Introduction

This paper examines empirically the effect of instability in social networks on network externalities, and therefore on the rate of adoption of technologies whose usage reflects these social networks. The effect of instability is not clear. On the one hand, instability may make potential adopters more responsive to adoption by more users, because they can no longer be reasonably certain they will not talk to them in the future. New products and technologies may therefore be able to spread faster in contexts where communications networks are unstable. On the other hand, instability may make potential adopters less inclined to respond to adoption by anybody, because they are more uncertain about their future communications patterns; this would mean that unstable communications networks hinder the adoption of new technologies. Without knowing the answer to this question, firms who wish to market new technologies that exploit social networks may mishandle product launches through not understanding the true *scope* of their product's network externalities, where 'scope' describes how many other people in the installed base influences a potential adopter (Katz and Shapiro 1985).

To analyze the scope of network externalities under network instability, I use data on the adoption of a video-calling technology sold by a business-to-business technology developer to employees of an international investment bank. It is hard to establish a causal link between communications network instability and the adoption rate. There are many unobserved factors that could lead to both an unstable communications structure and high or low adoption. For example, in a firm setting, poor management could lead to both unstable communication patterns and employees who are unwilling to adopt new technologies. Alternatively, a fast-paced work environment could lead to both unstable communication patterns and employees who embrace cutting-edge technologies. To address these identification problems, this paper uses an exogenous shock to the stability of the communications network (the terrorist attacks of September

11, 2001), that forced the bank's teams in New York to be physically relocated and reorganized, but did not affect the teams in other comparable cities. This shock led New York employees to communicate with different people than they did before the attacks, in a way which was difficult to predict. I study empirically whether this shock to the communications network led to more or fewer people influencing an adoption decision.

Even with this exogenous shock to the communications network, it would still be problematic to conclude that it was network externalities that prompted two employees who communicate to adopt at the same time. Network externalities occur when the performance of the good increases with the number of users. However, these two may have adopted simultaneously because they shared similar external impulses to adopt rather than because the usefulness of the good increased when they both adopted. To identify network externalities, the paper uses the fact that the video-calling system can also be used for watching TV. Some employees adopted the technology to watch one-off regional TV events such as the 2002 Six Nations rugby championship. TV-inspired adoption among employees' possible contacts varied across country and over time, even for employees in a single city such as New York. The network externalities are then measured based on how other employees responded to adoption by these TV-motivated adopters.

Combining these identification techniques enables this research to highlight some surprising findings. Normally, potential adopters react to changes in the installed base only if they communicate with that person in the month after they adopt. In other words, network externalities appear to be highly localized. However, in New York following the attacks, adoption by other related employees with whom adopters do not actually communicate in the month after adoption had a significant positive effect on adoption decisions. This increase in the scope of network externalities means that the importance of network externalities for adoption decisions doubles for those New York customers affected by communications instability.

This result contrasts with the established literature on the role of uncertainty in technology adoption. One plausible explanation for the results is that network externalities become less localized during periods of instability, because employees are unsure about whom they will need to talk with in the future. Models of technology diffusion have usually treated unpredictability about usage as a hindrance to diffusion (Jensen 1982; Kalish 1985; Mahajan et al. 1990); we present new evidence that argues for the opposite perspective when considering interactive technologies whose usefulness depends on social networks. Along similar lines, consumer studies such as Castamo et al. (2008), have treated consumer uncertainty over their future usage of a technology as something that firms need to try to minimize. By contrast, in the case studied here, instability made the adoption decision more responsive to a broader set of adoption decisions. This was beneficial, since it initiated further adoption cascades and nearly wholly compensated for the base-level decrease in adoption associated with the attacks.

Katz and Shapiro (1985) emphasize that the scope of network externalities is just as important as the size of the network externalities themselves, and that this scope may vary. However, most current empirical research (Saloner and Shepard 1994; Brynjolfsson and Kemerer 1996; Gowrisankaran and Stavins 2004; Tucker 2008) has focused on quantifying the size of network externalities and has taken the scope as fixed. Following the analysis of Shankar and Bayus (2003); Chang and Park (2005) who highlight difference in size of network externalities across firms, there have been some attempts to see if the size of network externalities varies across firms individuals. Tucker (2008), which uses the same data as the current paper, shows that for this good network externalities were larger if an adopter was a manager or acted as a gatekeeper to other social networks. However, none of these studies quantifies how uncertainty affects perceptions of network externalities. Mobius (2001) and Puffert (2002) emphasize the importance of localness of networks in understanding network evolution, emphasizing that the growth of a network changes network externalities as the localness of the customer base diminishes. This paper builds on this coordination literature by showing that it is not

just the growth of the network that affects the role of network externalities, but also the stability of the network.

From a managerial perspective, the results suggest that instability in communication networks may make it harder to coordinate the adoption of a new technology whose usefulness is characterized by network interactions. This suggests that firms need to adjust their strategies to reflect network externalities more if networks are unstable. If a firm were trying to encourage the adoption of a web-based video-calling service, such as Skype, they would need to target incentives particularly at groups of people who had less stable communications networks (students rather than seniors, for example). My findings may also help explain why network externalities have proved so important for social networking sites such as Facebook: The younger demographic of such web services is characterized by ever-changing and fluid relationships, which mean that potential users are influenced by broader adoption rates in addition to adoption by people who they are currently friends with.

The result that without network instability network externalities are highly localized is also important for managers. Managers often assume that network externalities for communications and social networking technologies depend on the total number of subscribers. My estimates suggest, by contrast, that if networks are stable, only the smaller subset of people with whom a potential adopter interacts plays a significant part in the adoption decision. If these results hold for other technologies, then network externalities for network technologies with stable communications networks may be limited in scope to predictable communications, reducing the extent to which initial price cuts for such technologies are warranted. This limited scope for network externalities may also help explain findings such as Tellis et al. (2009), that suggests that product quality is more important than network externalities in explaining the success of high tech products. If a firm only has to ‘tip’ smaller communication clusters, there may be less room for a long-run strategic advantage from network size.

2 Description of Technology and Data

2.1 Technology

Video-calling offers network benefits because users can see the person they are talking to rather than just talking to them. This research studies a business-to-business technology manufacturer's efforts to sell a new video-calling technology attached to an employee's workstation to individual employees at an investment bank. The end-point technology consists of three elements: video-calling software; a media compressor allowing TV-quality video; and a camera fixed on top of the computer's monitor. The technology also has a stand-alone use of transforming the employee's desktop computer into a television.

An unusual feature of the technology is that for data security reasons communication was limited to being between employees at the firm. This sets clear limits for the potential network for each user.

2.2 Firm setting

After the bank chose this technology to allow its employees to conduct internal video-calling and invested in the appropriate network architecture, the bank decentralized installation decisions to each employee. The bank publicized the technology to employees and each employee decided whether and when to order a video-calling unit from an external sales representative.¹ The equipment's supplier had excess capacity, meaning that capacity constraints did not affect the timing of individual employee installation decisions. This decentralization means that I focus my analysis on the private benefits of installation for employees, as opposed to firm-wide decisions.

Though the price to employees was zero, there were still non-monetary costs. Interviews with employees at the firm confirmed that the most salient of these non-monetary costs were the time they had to spend with their computer out of service while a tech-

¹Only employees who had a position of Associate or higher were eligible, meaning that staff in secretarial and support functions were not eligible.

nician installed the video-calling unit and the time spent learning how to use it. This cost was generally viewed as high because they could not predict when they would need to respond to rapid changes in the financial markets. The installation required the employee to be present so that the technician could demonstrate the product.

One obvious question is how well informed employees were about the spread of the technology across the firm, since my economic analysis assumes that employees were well-informed about who had adopted. When questioned, employees were most likely to respond that they learned about who else had the technology by asking a colleague who had adopted the technology to look at the computer briefly and examine the electronic list of people who were on the network. The second most popular response was that they gauged who had adopted the technology through emails where someone suggested the use of the technology. Though the sample size for these interviews was low, there were no obvious systematic differences in how New York employees learned about the spread of the technology compared to other employees after the attacks.

2.3 Data

A call database recorded the 2.4 million video-calls made within the bank from January 2001 to August 2004. For two-way video-calls, the database records the caller and callee, the time the call was made, and the length of the call. Employees made 1,768,348 two-way user-to-user video-calls. The data include only the 1,052,110 video-calls where the callee accepted the call and there was a non-zero time spent on the call. Each accepted call lasted on average 5 minutes and 46 seconds. Calls could be made to more than one employee at a time. Multi-party calls (less than 5% of the calls) were simplified into their pairwise equivalents: Three-way calls were treated as three calls between each two of the participants. The results are unchanged whether or not these three-way calls are dropped or included in the sample, suggesting that they do not drive the results. For one-way TV calls, the database records who made the call, which TV channel was viewed, the time, and the length of the call. Employees

made 741,926 successful one-way caller-to-media-device calls. Complete (anonymized) personnel records describe every employee in the investment bank in March 2004. Employees were divided by function (administration, research, trading, and sales), region (Britain, North America, Europe, and Asia/Sub-Equatorial, title ('Associate', 'Vice-President', 'Director', and 'Managing Director'). Managing Directors) and product (equity and derivative) associated with two main products: equities and derivatives.²

These data were used both to evaluate when an employee makes a first call and 'adopts' the technology and also to reconstruct the communications network within the firm. The data only partially describe the communications network, because they describe whom the 1,294 adopters video-called but not whom the 824 non-adopters would have video-called if they had adopted. Therefore, non-adopters were excluded from the regressions. Consequently, the empirical results are representative only of adopters. They should be interpreted as reflecting how network externalities affect the adoption timing of adopters, as opposed to the decision to adopt for all employees.³ Further, it should be noted that these data fall short of an ideal dataset that might include data from regular phone calls to define an employee's communication circle.

Though the data span 3.5 years, my interest in the effect of instability means that I focus on a limited time-frame of data which bridged the terrorist attacks of September 11, 2001. I divide the data into January 2001–August 2001 (the 'stable' period) and October 2001–August 2002 (the 'unstable' period). The choice of the 'unstable' period was based upon interviews with New York employees who said that this was how long the period of upheaval to and uncertainty about their roles within the firm lasted after the attacks.⁴ Using shorter windows such as 2 months before the attacks and 2 months after the attacks provides directionally similar though less precise results. Table 1

²The data is identical to that used in Tucker (2008), which provides a slightly fuller description.

³Tucker (2008) compares predictions from studying adoption only by adopters and a predicted communications network that includes non-adopters. The results were similar.

⁴This is slightly longer than the period described by Bloom (2009) in his study of how multiple New York-based industries recovered from the attacks, probably because this was a Wall Street firm and was more directly affected.

describes the summary statistics for major dependent and independent variables used in the data.

[Table 1 about here.]

2.4 Physical dislocation

Like many firms based in New York, the investment bank had to physically relocate their offices after the terrorist attacks of September 2001, in this case because their office close to the World Trade Center had to be evacuated due to the potential for structural damage. This shift in physical location also led to a change in workgroup dynamics. Management took the opportunity of the new office space to reorganize teams to better reflect the new workspace. Typical of the relocation was that employees who had previously sat in a workgroup that had analyzed ‘small-caps’ or smaller firms, were now reallocated to work with teams that analyzed particular industrial sectors. Similarly, there were shifts in how derivatives traders were organized into teams depending on the particular kind of bonds that they created derivatives for. The reorganization did not lead to drastic changes, such as shifting people across functions from sales to trading, but this did represent a genuine shift to group behavior. This physical relocation was accompanied by employee uncertainty over how their roles would develop and change within the firm and whether their new role would find a place within the physical reorganization. New York employees confirmed in interviews that they now called different people as a result of the reorganization that followed the attacks, but that they could have not initially have predicted who these people would have been.

It is crucial to establish whether there was indeed a shock to the communications network empirically before moving on to the major analysis of paper which presupposes this communications shock. Therefore in this section, I present some descriptive regressions that explore the extent to which this occurred. This focus on repeated correlations to illuminate relative stability, echo the work of Hagedoorn (1995) on firm-networks.

These descriptive regressions compare the calling behavior of people who had adopted the technology before the attacks in New York and elsewhere. This analysis uses a Poisson specification, where the dependent variable is a count of the distinct employees each adopter talks to in each month. The key variable of interest is an interaction term $NY * Unstable$, which captures the effect of being based in New York rather than elsewhere after the attacks. To control for observable and unobservable heterogeneity, controls for the month, region, title, and function of the employee as well as an employee-level random effect were included.⁵

Table 2 displays this initial descriptive analysis. Column (1) shows that there was no statistically significant change in the raw number of other employees each adopter spoke to in each month in New York after the attacks. Relative to other employees in the ‘unstable’ period, previous adopters in New York called the same number of people, so there was no increase in aggregate use. Column (2) of Table 2 shows, however, that there was an increase in the number of new people that an adopter in New York spoke to in the period after the attacks, relative to other employees in other locations in the firm.

Columns 3–5 of Table 2 explore the effect of $NY * Unstable$ on the number of new people that an employee talks to in different workgroups, functions, product groups, and global markets. The insignificance of $NY * Unstable$ in columns 3–5 indicates that the new people that New York employees talked to in the period after the attacks were from similar functions, product groups, and global markets. The fact that these new calls were not placed outside the function, product group and market helps to guide the definition of ‘potential contacts’, the group of employees that an employee may have talked to, in my main empirical analysis of technology adoption.

These results suggest that employees in New York called more new people following the attacks relative to employees in other cities, though these new contacts were in the same function, product group and market. In the remainder of this paper I

⁵Specifications that included employee-level fixed effects had similar results.

investigate whether this change in underlying communication patterns mattered for new technology adoption decisions.

[Table 2 about here.]

3 Modeling Technology Adoption

This analysis suggests that communication patterns in New York were destabilized by the attacks while those in other cities were not. To understand how this affects technology adoption behavior, I model potential customer technology adoption at the individual level. This is a latent variable setting, where only installation decisions $inst_{it}$, not installation benefits $inst_{it}^*$, are observed for customer i in month t :

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

where

$$inst_{it}^* = (ExpectedCalls_{it}, StandAloneBenefits_{it}) \tag{1}$$

Each month, each potential customer chooses whether or not to install the technology. As described in Farrell and Klemperer (2007), adoption depends on both the network benefit that a customer expects to receive from making video-calls and the stand-alone value that the potential customer receives from the technology, independent of whether they use it to make video-calls. This is a net stand-alone benefit that, besides reflecting the TV-watching capacity of the device, also reflects that non-monetary adoption costs are high for that customer. I discuss my conception of each in turn.

‘Network externality,’ is a term used by social scientists to describe a broad class of phenomena. These include ‘bandwagon-effects’ and other types of informational

spillovers, as well as direct technological complementarity where the usefulness of a network good (such as a fax machine or video-calling unit) depends on other users also having it. Since the data used in this study are for a classic network technology, my initial interpretation and analysis concentrates on technological complementarity as the source of network externalities.

The number of calls that a customer expects to be able to benefit from on depends on three factors. First, the number of other customers that potential customer i would call if that person also adopted. There are two potential sets of people that a potential customer i may call $\{C, P\}$. C is the set of customers, whom i thinks they are very likely to call because they are already close-knit contacts. P is the customers that there is some chance that the employee might call, that is, merely potential contacts. The second factor is whether another customer j (either in set C or P) has also adopted the video-calling technology. The last factor is the probability $\alpha_{k \in \{c,p\}}$ that employee i attaches to calling colleagues in their regular contact circle or less closely connected colleagues.

$$ExpectedCalls_{it} = \sum_{j \in C_i} inst_j \times \alpha_c + \sum_{j \in P_i} inst_j \times \alpha_p \quad (2)$$

In some circumstances, customer i may be able to predict very well whom they will talk with. In such stable circumstances, they may be able to ascribe a probability of α_c close to one for a subset of their closest collaborators, and an α_p close to 0 for other potential contacts. If future communication patterns are less certain, that is there is network instability, however, a customer i may ascribe a higher probability α_p to the likelihood of calling one of these potential contacts, as they are less certain that they will talk only with their closest collaborators.

3.1 Applying the model to data

The dependent variable $inst_{it}$ is whether the employee used video-calling technology to make an outward call for the first time in that month. This allows me to pinpoint the actual timing of adoption, because one of the protocols that the installation technician followed was to make an outward test-call as soon as they installed the technology on the employee's laptop. In all instances, this outward call was the first video-calling activity observed for the employee.

To operationalize the split between C people and P people for each employee i with observable customer behavior in the data, I construct two mutually exclusive groups for each employee. The C group, $Contacts_i$, are the other employees that i actually did talk to in the month after their adoption. The P group, $Potentials_i$, are those employees the customer did not talk to in the month following adoption but that there was a chance that they could have talked to as they shared the same product area, function or market specialization.

This definition of $Potentials_i$ is based on the empirical analysis in Table 2, which suggests that existing users contacted more new employees, but that these new contacts were not outside their function, market specialization, or product group. In addition to the evidence from the communication patterns of early adopters, the choice of these particular (reasonably broad) characteristics to define potential contacts was supported by interviews with investment bankers that suggested that it is very unusual for employees to switch from investment banking to equity sales or from researching Asian equities to researching European derivatives, since an employee has built up employee-specific expertise in that area. Therefore, even in view of the dislocation and uncertainty caused by the terrorist attacks, it seems likely that an employee would still perceive their broad product area as being the set of people in the firm they were most likely to communicate with.

Basing my measures for $ExpectedCalls_{it}$ on the subsequent month after adoption

presumes very high discount rates. However, this short time frame is necessary to ensure comparability. It ensures that when I compare the influences of adoption before the attacks and after the attacks, I am using the same measure of the putative installed base.⁶

I model that the extent to which adoption by $Contacts_i$ and $Potentials_i$ will matter depends on the stability of the underlying communications network. If communications are predictable, employees should be able to better anticipate whom their $Contacts_i$ are, or in other words ascribe a high α_c and a low α_p , and respond only to adoption by $Contacts_i$ and not to $Potentials_i$. However, if communications are not predictable, employees may also anticipate that they could talk to $Potentials_i$ whom they do not call, and as a result they will ascribe a higher α_p and a lower α_c relative to a stable network.

The installed base for $Contacts_i$ is captured by $InstalledContacts_{it}$. This installed base is a count of all contacts' installation decisions up to and including month t . The count of adoption by $Potentials_i$ who have adopted the technology, but with whom the employee did not communicate in the month after adoption, is captured by $InstalledPotentials_{it}$. I assume that each employee i takes the adoption choices of other employees as given, that is, that employees do not anticipate that their adoption could affect others' adoption in the future.⁷ Employees did not divest themselves of the technology after it was attached to their desktops. Therefore, I treat adoption decisions as permanent, and include only observations where the employee adopts video-calling in that month or has not yet adopted when estimating adoption decisions. This means an observation is an employee who did not adopt the technology in the previous months but does ultimately adopt.

In order to capture the stand-alone benefit of the technology, I use a rich set of

⁶I cannot, for example, use the final year of data to establish stable communication patterns for all adopters, because as illustrated in Table 2, these communication patterns were not predictable for those who adopted before the attacks.

⁷See Ryan and Tucker (2008); Dube et al. (2010) for a formal equilibrium model.

largely non-parametric controls. It is likely that the net costs of adopting the technology vary across employees. For example, it may be easier for employees in more flexible areas, such as research, to schedule time for their computers to be down than for employees who work in fast-paced areas, such as derivatives trading. Such heterogeneity is captured using a series of controls for each employee X_i . These include indicator variables for each of the different functions, product groups, regions, and titles. A series of dummies was also used for each month that employees could potentially adopt the technology, T_t . Since these time dummies will also pick up selection and the changing baseline hazard rate, they cannot be interpreted, and instead serve to capture the varying underlying hazard rate in a flexible manner (Jenkins 1995).

This technology also had specific benefits that were independent of any network usage. In particular, employees enjoyed being able to watch television on their desktop computer. There were two types of television employees could watch: News TV programming on CNN and CNBC, which covers financial news; and local TV programming (often non-news) broadcast by country-specific channels. There was little variation across regions in the percentage of adopters watching news programming (implying that this is largely captured by the month dummies), but there was large variation in employee interest in local TV programming across regions.⁸ For example, many employees in the UK watched the 2002 Six Nations rugby championship, while employees in the US did not. Empirically, these local broadcast events were correlated with adoption in the month prior to the month they occur. This suggests that employees adopted the technology in advance to ensure they could watch predictable ‘must-see’ television.⁹ These regional shocks to the technology’s stand-alone benefit were captured by the variable TV_{rt} , which contains the percentage of previous adopters

⁸The focus here was on the viewing habits for local television as the instrument. 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 the BBC. Local channels for the US were CSPAN, FOX, NBC, and CBS. Local channels for Asia were NTV (Nippon TV), CATS (Japanese), TV-Asia, and BBC 24 World Service.

⁹In the UK, over 60% of users watched television at the time of the World Cup soccer matches.

watching ‘local TV’ in region r in the month following time t . The video-calling unit’s TV use led to a less systematic pattern of adoption than is common for communication technologies. For example, there is no statistically significant relationship between how many months after the launch date the adopter adopts the technology, and the number of people a new adopter calls or the length of time they spend on video-calling.

Putting these components together, transforms equation (1) into the equation I use for estimation:

$$\begin{aligned} inst_{it}^* = & f(\Delta_1 InstalledContacts_{it} \\ & + \Delta_2 InstalledPotentials_{it} + \lambda TV_{r,t} + \beta X_i + \omega T_t + \epsilon_{it} | Stability_{it}) \end{aligned} \quad (3)$$

To summarize: Installation decisions are modeled as depending on the installed base of $Contacts_i$ and $Potentials_i$, the stand-alone TV benefit, their own idiosyncratic net benefits captured by controls X_i , a series of controls for month-by-month changes in their underlying hazard rate T_t , and unobserved heterogeneity (ϵ_{it}). I estimate this specification with a probit in a discrete hazard model framework (Jenkins 1995; Tellis et al. 2009).¹⁰

4 Correlations In The Data

To investigate the effects of the shock to the communication networks on adoption, I compare adoption behavior in the New York offices to adoption behavior in non-New York offices before and after the attacks. As with any difference-in-differences approach, the identifying assumption is that the New York office, in the absence of the physical relocation, would have had the same change in responsiveness to the installed base as employees in other cities. There was no statistically significant relationship between an employee working in New York and their number of contacts (see Table 2). However, New York employees talk less on average to employees in different cities. Employees

¹⁰The direction and significance of the coefficients in the probit specification are similar to those obtained from a Cox-proportional hazards model.

attributed this to limitations imposed by the time difference between the New World and Old World and did not think that this propensity changed in the period studied.

Table 3 presents the results of running the specification represented by equation (3) separately for before and after the attacks and for New York and non-New York employees.

[Table 3 about here.]

As documented by Manski (1993), measures such as $InstalledContacts_{it}$ and $InstalledPotentials_{it}$ are plagued by the reflection problem, making any causal interpretation of correlations in adoption problematic. Instead, the aim of Table 3 is to compare the extent of correlation between adoption and actual/potential contacts' adoption for New York employees before and after the attacks relative to employees elsewhere. Subsequently, instrumental variables are used to identify actual network externalities.

There are some suggestive shifts in the correlations. Both New York and non-New York employees before 9/11 only exhibit positive and significant correlation in adoption with their 'actual contacts'. However, after 9/11, New York employees no longer have a statistically significant correlation in adoption with their actual contacts, but instead have a statistically significant correlation in adoption with their potential contacts. Non-New-York employees' adoption, by contrast, continues to be correlated only with their actual contacts. These results are robust to whether I use the total number of potential contacts adopting (as I do in Table 3), or the percentage of potential contacts adopting.¹¹ This fact that whether I use percentage or absolute numbers for potential contacts does not matter suggests that these correlations are not driven by the effect of employee group size.

In order to test whether the coefficients on $InstalledContacts_{it}$ and $InstalledPotentials_{it}$ for New York employees are statistically different to those outside New York before

¹¹The proportion of actual contacts that have adopted is always reasonably close to one because a contact necessarily has adopted the technology by month $t + 1$, meaning that it is not very instructive to use it as a regressor.

and after the attacks, I ran a three-way differences-in-differences specification. Following the economics literature, this differences-in-differences specification captures relative changes in the time trend for different groups of employee by the use of three-way interactions and panel data. The results are reported in Table 5. The coefficients on the three-way interactions $InstalledContacts_{it} * NY_i * Unstable_t$ and $InstalledPotentials_{it} * NY_i * Unstable_t$ are significantly different from zero, suggesting that there was a statistically meaningful change ($p < 0.05$) in adoption responsiveness for New York employees after the attacks.

5 Measuring network externalities

Table 3 suggests that employees in New York were more likely to adopt when a wide circle of their acquaintances adopted the technology after the attacks, compared to before and to other comparable employees. Correspondingly, Table 2 suggests that employees called more new people after the attacks. However, while suggestive, these two pieces of evidence do not prove that the changes in correlations in adoption can be ascribed to anticipated changes in calling behavior due to the potential for unobserved heterogeneity (Manski 1993). Consider two employees who are both encouraged to install the technology by their New York boss who is concerned about aviation safety after the attacks; a clear identification strategy is needed to avoid interpreting the subsequent correlation in their adoption decisions as a causal network externality.

I use variation in the TV-viewing benefit of the technology as an instrumental variable to identify how changes in the installed base causally affect an employee’s adoption. This exploits three types 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 employees have contacts or potential contacts. The instrument for actual contacts for employee i is $\sum_{j \in C}^i (TV_{r,t})$, or the average TV benefit ($TV_{r,t}$) for each actual contact j ’s region r at month t . The instrument for potential contacts for employee i is $\sum_{j \in P}^i (TV_{r,t})$, or the average TV benefit ($TV_{r,t}$) for each potential

contact j 's region r at month t .

The value of these instruments varies over time and by the regions these contacts and potential contacts work in. Crucially, the instrument for the New York office for potential contacts also has considerable variation in the data. This is because the variety of different functions, product markets and market specializations means that there were at least 48 different viable sets of potential contacts for any one employee in the New York office. Not all earlier adoptions by i 's actual or potential contacts were counted as necessarily causing i 's installation. Instead, the research uses variation in adoption by i 's contacts or potential contacts that can be predicted by variation in the stand-alone (TV) benefit. In all specifications, the joint F -test of the first-stage regressions was significant at the 0.01 level.

One concern is that employees who had contacts in a region that had heavy TV-watching might be different in unobserved ways from similar employees who did not have contacts in that region. To rule out such challenges to the exclusion restriction, I ran multiple regressions examining whether there were significant differences in the observable correlates of having contacts in a particular region. Except for a relationship with the global market area that an employee worked in, there was no statistically significant difference between employees who had different regional concentrations of contacts in terms of their observable characteristics, TV-watching or network use.

Table 4 displays the estimates for this instrumented version of equation (3). The first two columns display results for the 'stable' pre-period before September 2001. They suggest that for both New York and other financial centers, the effect of actual contacts adopting was significant while the effect of potential contacts was insignificant.

After September 2001, the 'unstable' period, employees in New York reacted differently to adoption by other employees in other locations. For New York employees, the potential installed base now exerts a positive influence on adoption. Potential adopters who are not based in New York, however, still do not react to the potential installed base. For both New York and non-New York employees, the estimated network exter-

nalities for the influence of actual contacts decrease. In the case of New York, they become statistically insignificant at conventional levels with a p -value of 0.11, though the point estimate is still higher than for $InstalledPotentials_i$. This suggests that there was increased variance on the extent to which employees in New York responded to adoption by their actual contacts, which would again fit an explanation whereby New York employees became less sure about whom they would talk to after the attacks. The fact that TV viewing is positive and significant in the unstable period outside of New York reflects the relative degree of interest in the World Cup outside of the US in 2002 (as documented by Tucker (2008)).

[Table 4 about here.]

In order to test the difference in magnitude of the coefficients in Table 4, I repeated the regression using a pooled sample. The main variables of interest are the three-way interactions, $InstalledPotentials_{it} * NY_i * Unstable_t$ and $InstalledContacts_{it} * NY_i * Unstable_t$. These measure the differential effect on the installed base of the relocation of the New York offices. Therefore, the differences-in-differences specification of equation (3) becomes:

$$\begin{aligned}
inst_{it}^* = & \Delta_1 InstalledContacts_{it} + \Delta_2 InstalledPotentials_{it} \\
& + \Delta_3 InstalledContacts_{it} * NY_i * Unstable_t + \Delta_4 InstalledPotentials_{it} * NY_i * Unstable_t \\
& + \gamma(Unstable_t * NY_i + InstalledContacts_{it} * NY_i + InstalledPotentials_{it} * NY_i \\
& + InstalledContacts_{it} * Unstable_t + InstalledPotentials_{it} * Unstable_t) \\
& + \lambda TV_{r,t} + \beta X_i + \omega T_t + \epsilon_{it}
\end{aligned} \tag{4}$$

The level effects of $Unstable_t$ and NY_i are captured by the vector of time dummies and location dummies (X_i).

[Table 5 about here.]

Column (1) of Table 5 presents the results for all employees without IV. Column (2) presents the results using instrumental variables. The instruments are the average TV benefit for an employees' actual and potential contacts and appropriate weights to reflect the endogenous interactions. The key variables of interest are the three-way interactions $InstalledContacts_{it} * NY_i * Unstable_t$ and $InstalledPotentials_{it} * NY_i * Unstable_t$. Though precise estimation is problematic in a regression with eight endogenous variables, the interaction for $InstalledPotentials_{it} * NY_i * Unstable_t$ was positive and significant (p<.10). New York showed no change in responsiveness to the installed base of their contacts, but the adoption of employees in New York became increasingly correlated with the measure of the installed base of potential contacts in the firm after the attacks. The coefficient on $InstalledContacts_{it} * NY_i * Unstable_t$ was not, however, significantly different from zero, suggesting there was no measurable change in how New York employees responded to adoption by employees they called in the month after they adopted. Other variables are much in line with past evidence. There are positive correlations between adoption and the narrow installed base $InstalledContacts$. The negative coefficient on $Unstable * InstalledContacts$ may reflect the fact that later adopters were in general less responsive to adoption by others than earlier adopters.

To address concerns about the reliability of the Probit-IV model when there are multiple dummy variables and interaction terms (Ai and Norton 2003), I also estimated a linear probability model. The estimates are very similar, with a slight increase in significance, especially for the installed base of actual contacts.

5.1 Robustness and discussion

The major finding in Tables 4 and 5 is that New York employees became more likely to adopt if their broader group of potential contacts adopted, relative to employees in other cities. The interpretation that this paper has highlighted so far is that when communication patterns are more unstable, potential adopters take into account adoption by a wider group of network users. However, there are other potential explanations,

which I now discuss in turn.

The first set of alternative explanations are concerned with unobservable heterogeneity. The identifying assumption underlying these results is that New York employees would have had a similar change in their adoption patterns to non-New York employees if it had not been for the attacks. However, there may be concerns that changes in communication in general after 9/11 can provide an alternative explanation for what occurred, and that these changes may have had a larger effect in New York than in cities such as Hong Kong that were far away from the attacks. For example, perhaps concerns over increased security measures in airports led New York employees to anticipate using the technology more broadly. To alleviate such concerns, I reran the regressions from Table 4 for employees in the London office. These were the closest employees to the New York office, both in terms of the economy of the country they operated in and the cultural practices of the firm. London itself was considered at the time to be a reasonably likely location for a terrorist attack, though an attack did not occur there until July 7, 2005. Since both countries participated in the war in Afghanistan, both sets of employees experienced similar changes to their previous air-travel patterns and airport security measures. Both economies' financial sectors experienced similar declines as a result of the attacks. However, the effect of $InstalledPotentials_{it}$ was both insignificant and negative for London employees, suggesting that the broadening of scope of network externalities was exclusive to the New York office which experienced the actual physical reorganization. Furthermore, the communication patterns of existing adopters did not display the shift shown in Table 2.

Another interpretation of the positive coefficient on $InstalledPotentials$ after the attacks is that it reflects a different kind of network effect from those associated with the complementarity of video-calling technology. For example, it could reflect an increase in bandwagon effects or informational spillovers. This is a plausible interpretation if after 9/11 New York employees were more likely to mimic their neighbor's adoption

due to an increase in sensitivity to other people's actions. To study this interpretation further, I examined whether there was any increase in correlations in adoption between one employee and other employees on the same floor as them. If informational spillovers or bandwagon effects were important, it would seem likely that they would be most concentrated amongst people who had daily physical contact. However, I found no statistically significant change in the relationship between New York employees' adoption and their physically close colleagues' adoption before and after the attacks.

Back-of-the-envelope estimates allow a comparison of the overall effect of network externalities for New York employees after the attacks compared with other employees. There is no effect that can be measured accurately from actual contacts for New York employees. An average New York employee after the attacks had 102 potential contacts. The marginal effect of the potential installed base of 0.0004, based on the probit estimates in table 4, therefore suggests that the average New York employee received a boost in their adoption propensity of around 4% from their potential contacts (0.0004×102). To put this boost into context, the average adoption propensity is 8% each month. By contrast, the average non-New York employee, who has nine actual contacts installed after the attacks, receives a boost of 2% from the marginal effect of the coefficient of 0.0021 on actual contacts (0.0021×9). There is no effect that can be measured accurately for potential contacts. Therefore, roughly, network externalities were twice as important for adoption for New York employees after the attacks than for employees elsewhere.

This is a marginal boost to adoption. As suggested by the negative probit coefficient on *Unstable*NY* in Table 5 that translates to a marginal effect of -0.0275 , on average the level of adoption in New York after the terrorist attacks was slightly lower (around 1% in the raw data). This is in line with traditional theories of how instability or uncertainty affects investment. There are of course other explanations too, such as employees having more spare time to spend acquiring the technology as business volume went down after the attacks. What is novel is that the increased scope of network

externalities from the broader installed base nearly compensates for this decrease in the base level of adoption during this time of instability.

Also striking is that in Table 4 is the persistently null result for the effect of $InstalledPotentials_{it}$ for the majority of employees (87.5 % of the sample) who either adopted before the attacks or were not based in New York. This null result is important in two ways. First, it suggests that broader network externalities in the form of informational spillovers and bandwagons were not that important for the adoption of this technology, since it seems likely that if these were present they would be reflected in a correlation in adoption between an employee and other employees similar to him. Second, it suggests that network externalities due to the complementarity of technologies such as video-calling are limited to smaller subsets of network users who actually communicate if network relationships are reasonably stable.

6 Implications

This paper finds that when social networks were stable, adoption cascades for a new network technology were confined to small subsets of people who interacted with each other. However, when the communications network became less stable, potential adopters looked more broadly at the network and were more responsive to adoption by a broader circle of potential contacts. These empirical results are based on the destabilization of a communication network due to the terrorist attacks in 2001 for a subset of potential users.

This finding is important for managers trying to introduce new products. Both managerial practice and the academic literature have assumed that unpredictability in usage for a new product deters consumers from adopting new technology (Jensen 1982; Kalish 1985; Mahajan et al. 1990; Castamo et al. 2008). By contrast, the findings here suggest that if there are network externalities and there is uncertainty over usage patterns, an adoption decision may become responsive to a broader set of adoption decisions. This increase in the scope of network externalities led to a doubling of their

role in adoption decisions.

This research also underscores that for many technologies where communication and social networks are stable, network externalities are extremely local and only depend on the people that a person directly communicates with (Sundararajan 2007). This may undercut marketing strategies for network products that are based on a ‘winner takes all’ mentality. It may also explain research such as Tellis et al. (2009) that has found that network externalities are less important for technology diffusion than other product features such as quality.

From a managerial perspective, my results suggest that instability may make it harder for firms to spread new technologies for which network externalities could theoretically be important. If the network is stable, and interactions are small and localized, managers are unlikely to need to adjust their marketing strategically to take network externalities into account. However, if the network is unstable, and interactions may not be predictably localized, managers should be very active in targeting promotions to kick-start network growth. For example, managers at Skype would ideally target incentives toward groups of people, like students, whose communications networks are relatively unstable. These findings help to explain why the social networking sites that have grown fastest are ones like MySpace and Facebook, which are targeted towards young people who have ever-changing and fluid social relationships. This instability means that potential users are influenced by broader adoption rates as well as only adoption by people they are currently friends with.

There are of course limitations to this research. I study the effect of instability of communications within a specific business-to-business setting. Therefore, I can only speculate how these results translate to direct consumer markets. Further, the nature of the uncertainty in this particular context was unusual in that it was particularly salient. The employees knew that as a result of re-organization that their communication patterns would change - they just did not know whom they would talk to. The effect of uncertainty might be more muted if the fact that interactions may be uncer-

tain was less salient. However, notwithstanding these limitations, my research does suggest that uncertainty can change the influence that network externalities have on adoption decisions, and that recognition of this should be reflected both in academic models and in managerial practice.

References

- Ai, C. and E. C. Norton (2003, July). Interaction terms in logit and probit models. *Economics Letters* 80(1), 123–129.
- Bloom, N. (2009, 05). The impact of uncertainty shocks. *Econometrica* 77(3), 623–685.
- Brynjolfsson, E. and C. Kemerer (1996). Network externalities in microcomputer software: An econometric analysis of the spreadsheet market. *Management Science* 42, 1627–1647.
- Castamo, R., M. Sujan, M. Kacker, and H. Sujan (2008). Managing consumer uncertainty in the adoption of new products: Temporal distance and mental simulation. *Journal of Marketing Research (JMR)* 45(3), 320 – 336.
- Chang, S.-J. and S. Park (2005). Types of firms generating network externalities and mncs’ co-location decisions. *Strategic Management Journal* 26(7), 595–615.
- Dube, J.-P. H., G. J. Hitsch, and P. K. Chintagunta (2010). Tipping and Concentration in Markets with Indirect Network Effects. *Marketing Science* 29(2), 216–249.
- Farrell, J. and P. Klemperer (2007, December). *Coordination and Lock-In: Competition with Switching Costs and Network Effects*, Volume 3 of *Handbook of Industrial Organization*, Chapter 31, pp. 1967–2072. Elsevier.
- Gowrisankaran, G. and J. Stavins (2004). Network externalities and technology adoption: lessons from electronic payments. *RAND* 35(2), 260–276.
- Hagedoorn, J. (1995). A note on international market leaders and networks of strategic technology partnering. *Strategic Management Journal* 16(3), pp. 241–250.
- Jenkins, S. P. (1995, February). Easy estimation methods for discrete-time duration models. *Oxford Bulletin of Economics and Statistics* 57(1), 129–38.
- Jensen, R. (1982, June). Adoption and diffusion of an innovation of uncertain profitability. *Journal of Economic Theory* 27(1), 182–193.
- Kalish, S. (1985). A new product adoption model with price, advertising, and uncertainty. *Management Science* 31(12), 1569–1585.
- Katz, M. L. and C. Shapiro (1985). Network externalities, competition, and compatibility. *American Economic Review* 75(3), 424–40.
- Mahajan, V., E. Muller, and F. M. Bass (1990). New product diffusion models in marketing: A review and directions for research. *The Journal of Marketing* 54(1), 1–26.
- Manski, C. (1993, July). Identification of endogenous social effects: the reflection problem. *Review of Economic Studies* 60(3), 531–542.
- Mobius, M. (2001). Death through success: The rise and fall of local service competition at the turn of the century. Harvard Working Paper.
- Puffert, D. J. (2002). Path dependence in spatial networks: The standardization of railway track gauge. *Explorations in Economic History* 39(3), 282 – 314.

- Ryan, S. and C. Tucker (2008). Heterogeneity and the Dynamics of Technology Adoption. *Mimeo, MIT*.
- Saloner, G. and A. Shepard (1994). Adoption of technologies with network effects: An empirical examination of the adoption of automated teller machines. *RAND* 26(3), 479–501.
- Shankar, V. and B. L. Bayus (2003). Network effects and competition: an empirical analysis of the home video game industry. *Strategic Management Journal* 24(4), 375–384.
- Sundararajan, A. (2007, July). Local network effects and complex network structure. *The B.E. Journal of Theoretical Economics* 7(1).
- Tellis, G. J., E. Yin, and R. Niraj (2009). Does quality win? network effects versus quality in high-tech markets. *Journal of Marketing Research (JMR)* 46(2), 135 – 149.
- Tucker, C. (2008). Identifying Formal and Informal Influence in Technology Adoption with Network Externalities. *Management Science* 54(12), 287–304.

Table 1: Description of all variables used in regressions

Variable	Description	Before Attacks		After Attacks	
		Mean	Std. Dev.	Mean	Std. Dev.
Dependent Variables					
$inst_{it}$	Indicator variable for the first month an employee makes outward video-call	0.112	.3157	.046	0.209
Right-hand Side Variables					
$InstalledContacts_{it}$	Sum of cumulative adoption by employee i 's contacts by month t , where contacts are the employees in the firm that the employee talks to in the month after they adopt.	8.086	9.080	8.523	10.109
$InstalledPotentialContacts_{it}$	Sum of cumulative adoption by employee i 's potential contacts who work in the same function, product, and global market area but whom the employee did not contact in the month after they adopt	68.511	31.316	111.133	45.458
TV_{rt}	Proportion of adopters in the employee's region r who have adopted prior to month t who watch local television channels in month $t + 1$	0.249	0.345	0.393	0.374
Controls for regions	Indicator variables for Europe, Asia, US, and UK				
Controls for month	Indicator variables for each month from February 2001 to August 2002				
Controls for function	Dummies for working in administration, research, trading, and sales				
Total Observations		12,723			

Table 2: Preliminary analysis exploring number of new people that employees who adopted prior to September 2001 talked to after the attacks

Dependent Variable	Number of people spoken to	Number of new people spoken to	Number of new people spoken to in different functions	Number of new people spoken to in different product markets	Number of new people spoken to in different global markets
	(1)	(2)	(3)	(4)	(5)
<i>NY* Unstable</i>	0.005 (0.038)	0.184*** (0.105)	-0.068 (0.102)	0.045 (0.106)	0.035 (0.087)
<i>NY</i>	-0.008 (0.129)	-0.175 (0.137)	0.025 (0.143)	-0.055 (0.138)	-0.154 (0.121)
<i>Unstable</i>	0.063 (0.057)	0.008 (0.160)	0.028 (0.160)	-0.011 (0.172)	0.184 (0.134)
Observations	34,125	34,125	34,125	34,125	34,125
Log-likelihood	-3.35×10^4	-1.26×10^4	-1.18×10^4	-1.18×10^4	-1.63×10^4
Month Dummies	Yes	Yes	Yes	Yes	Yes
Title Dummies	Yes	Yes	Yes	Yes	Yes
Region Dummies	Yes	Yes	Yes	Yes	Yes
Function Dummies	Yes	Yes	Yes	Yes	Yes

Sample: Employees who adopted video-calling prior to September 2001.

Poisson specification. No zeros dependent variables observed in the data. Poisson goodness of fit χ^2 statistic suggested that the hypothesis data were dispersed according to the Poisson distribution could not be rejected at the 5 percent level. Random effects at employee level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Table 3: Correlations between Employee’s Own Adoption and Potential and Actual Contacts’ Adoption

	Stable Period		Unstable Period	
	NYC	Not NYC	NYC	Not NYC
InstalledContacts	0.0419*** (0.0065)	0.0316*** (0.0040)	0.0117 (0.0075)	0.0201*** (0.0040)
InstalledPotentials	0.0010 (0.0036)	-0.0024 (0.0023)	0.0057** (0.0026)	0.0009 (0.0015)
TV-Watching	0.4482 (0.3808)	0.1151 (0.1643)	0.0054 (0.4572)	0.6067*** (0.1828)
Observations	1294	3654	1171	3248
Log-Likelihood	-372.13	-1054.58	-171.10	-586.26
Month Dummies	Yes	Yes	Yes	Yes
Region, Title, Function Dummies	Yes	Yes	Yes	Yes

Dependent Variable: Indicator for when an employee first makes an outward video-call

Sample: Adopters who have not yet made a video-call. 2,506 observations dropped due to multi-collinearity with fixed effects.

Probit Specification. Clustered standard errors at the work-group level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.014$.

Stable period is January 2001–August 2001 and unstable period is October 2001–August 2002

Table 4: Probit specification with instrumental variables

	Stable Period		Unstable Period	
	New York	Non-New York	New York	Non-New York
<i>InstalledContacts</i>	0.0389*** (0.0077)	0.0325*** (0.0040)	0.0162 (0.0100)	0.0174*** (0.0045)
<i>InstalledPotentials</i>	0.0016 (0.0037)	-0.0003 (0.0021)	0.0060* (0.0036)	-0.0004 (0.0016)
TV viewing	-0.3807 (0.2984)	0.0676 (0.1558)	-0.0430 (0.3977)	0.7402*** (0.2126)
Observations	1,383	4,017	1,276	3,541
Month Fixed Effects	Yes	Yes	Yes	Yes
Region, Title, Function Controls	Yes	Yes	Yes	Yes

Dependent Variable: Indicator for whether an employee first makes an outward video-call that month

Sample: Employees who ultimately adopt the technology who have not yet made a video-call. 850 additional observations dropped from sample used in Table 3 due to collinearity.

Probit two-step Newey estimates: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Instruments for the different installed base measures are the TV valuation of each employee's actual and potential contacts.

First-stage regressions significant at 1% level. Regression is exactly identified.

Stable period is January 2001–August 2001 and unstable period is October 2001–August 2002

Table 5: Hypothesis testing using three-way diff-in-diff

	All Employees	All Employees (IV)
<i>Unstable * NY * InstalledContacts</i>	-0.0041 (0.0038)	-0.0037 (0.0097)
<i>Unstable * NY * InstalledPotentials</i>	0.0039*** (0.0006)	0.0040* (0.0023)
<i>InstalledContacts</i>	0.0332*** (0.0028)	0.0324*** (0.0032)
<i>InstalledPotentials</i>	0.0006 (0.0014)	0.0014 (0.0015)
<i>Unstable * NY</i>	-0.4976*** (0.0802)	-0.5371* (0.2796)
<i>NY * InstalledContacts</i>	0.0056** (0.0023)	0.0061 (0.0057)
<i>Unstable * InstalledContacts</i>	-0.0153*** (0.0033)	-0.0142*** (0.0045)
<i>NY * InstalledPotentials</i>	-0.0006 (0.0006)	-0.0004 (0.0011)
<i>Unstable * InstalledPotentials</i>	-0.0001 (0.0012)	-0.0011 (0.0015)
TV viewing	0.2915*** (0.0960)	0.2983*** (0.0961)
Observations	10217	10217
Month Fixed Effects	Yes	Yes
Region, Title, Function Controls	Yes	Yes

Dependent Variable: Indicator for whether an employee first makes an outward video-call that month

Sample: Employees who ultimately adopt the technology who have not yet made a video-call

Column (1) displays Probit estimates. Column (2) displays Probit two-step Newey estimates: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Instruments for column (2) the different installed base measures are the TV valuation of each employee's actual and potential contacts interacted with Unstable and New York Indicator variables. First-stage regressions significant at 1% level. Regression is exactly identified.

Stable period is January 2001–August 2001 and unstable period is October 2001–August 2002