American Economic Review 2008, 98:4, 1578–1590
http://www.aeaweb.org/articles.php?doi=10.1257/aer.98.4.1578

Restructuring Research: Communication Costs and the Democratization of University Innovation

By Ajay Agrawal and Avi Goldfarb*

We examine the effect of a decrease in collaboration costs resulting from the adoption of Bitnet (an early version of the Internet) on university research collaboration in engineering. Our interest in this question stems not from a concern about either Bitnet or engineering research specifically, but rather about the broader question of how changes in collaboration costs may affect the structure of knowledge production. Exploiting the variation in year of adoption and publication output over time in the 270 universities that published in seven top electrical engineering journals from 1981 to 1991, we find that a Bitnet connection did seem to facilitate a general increase in multi-institutional collaboration (by 40 percent, on average). At the same time, not all adopters benefited equally. Overall, Bitnet seems to have facilitated a disproportionate increase in the role of middle-tier universities, particularly those co-located with top-tier institutions.

The non-uniform effect of Bitnet across university pairs offers insight into the nature of collaborative knowledge production. A researcher deciding whether to add a collaborator to a project will do so if the benefit exceeds the cost such that the returns from collaboration are positive for both parties. Due to the way in which knowledge is produced, a technology shock like the introduction of Bitnet might affect the returns to collaboration differently, depending on characteristics of collaborating pairs, such as the quality of the institutions and the geographic distance between them. Indeed, our finding that certain university pair types benefited disproportionately from Bitnet adoption enables us to make inferences about the relative benefits and costs of collaboration across pair types.

For instance, we examine whether the returns to Bitnet adoption were mediated by pair quality. One might expect that pairs comprised of two top-tier universities would benefit most since individually these institutions produced the highest volume of research and thus had the most on which to collaborate. However, we find that top-tier/middle-tier pairs benefited most from adoption. These results suggest that the most salient effect of Bitnet may have been to facilitate gains from trade through the increased use of underutilized research equipment or the heightened specialization of research tasks.[1]

* Agrawal: Rotman School of Management, University of Toronto, 105 St. George St., Toronto, Canada, M5S 3E6, and NBER (ajay.agrawal@rotman.utoronto.ca); Goldfarb: Rotman School of Management, University of Toronto, 105 St. George St., Toronto, Canada, M5S 3E6 (e-mail: agoldfarb@rotman.utoronto.ca). We thank Pierre Azoulay, Iain Cockburn, Wes Cohen, Shane Greenstein, Scott Stern, two anonymous referees, and especially Ig Horstmann, as well as numerous seminar participants, for useful comments. We also thank Raghav Misra, Swapnil Kotecha, Cara Saunders, and Alex Oettl, all of whom provided excellent research assistance. Errors and omissions are our own. This research was funded by the Social Sciences and Humanities Research Council of Canada (grants 410-2004-1770 and 538-02-1013). Their support is gratefully acknowledged.

[1] With respect to the latter, the intuition is similar in spirit to models that examine trade between developed and developing countries (e.g., Avinash K. Dixit and Victor D. Norman 1980; Elhanan Helpman and Paul Krugman 1985). Many of these models show that the type of trade in equilibrium (i.e., developed-developed or developed-developing) will depend on the nature of the specialization and on the size of the economies. While we focus on specialization to explain our results, we acknowledge it is only one possible mechanism for differences of the observed effect of Bitnet across qualities. Other possibilities include monitoring (George P. Baker and Thomas N. Hubbard 2003) and heterogeneity in research interests (Tanya S. Rosenblat and Markus M. Mobius 2004). The aim of this paper is not to identify the particular mechanism, but to empirically measure the impact of Bitnet connection on different types of collaborations.

Why might this be? With respect to increasing the use of underutilized research assets, consider the indivisible nature of capital intensive laboratory equipment (lasers, robots, simulators, etc.). Using Bitnet, faculty at top-tier institutions could more easily coordinate the shared use of their expensive research equipment that would otherwise sit idle between experiments. This would lead to increased multitasking across multiple research projects. Indeed, in a detailed, task-level study of knowledge worker productivity, Sinan Aral, Erik Brynjolfsson, and Marshall Van Alstyne (2007) show an association between network usage, multitasking, and productivity.

With respect to increased specialization, consider universities of two research quality types: top-tier and middle-tier. The former have a stronger orientation toward research, which is reflected in larger resource allocations to research activities and a broad range of doctoral programs. Researchers at top-tier schools may focus on winning grants, supervising the use of specialized equipment, attending international conferences to present results, and other such high-cost activities. Researchers at middle-tier institutions, who may not have the resources necessary for running certain types of experiments entirely on their own, may have the expertise and equipment necessary for certain steps in the research process. Using Bitnet, data could be transferred to these researchers for data analysis and computing. Indeed, this pattern of activity is consistent with prior descriptive findings that characterize early electronic networks as facilitating a division of labor leading to a greater involvement of researchers at "peripheral" institutions (Bradford W. Hesse et al. 1993; John Walsh and Todd Bayma 1996).

We also examine whether the returns to Bitnet adoption were mediated by the distance between pairs. One might expect that since Bitnet substitutes for other communication mechanisms (phone, fax, travel, etc.) and communication costs increase with distance, Bitnet would have disproportionately benefited pairs that were farther apart since such pairs would have enjoyed the greatest cost reduction. However, for top-tier/middle-tier collaborations in particular, our results show that the benefits of Bitnet were greatest for pairs that were close together.

These results suggest that network communication complements other collaborative tools. Since collaboration is predicated on shared ideas, which are often the unplanned output of direct interaction,[2] researchers may benefit significantly from face-to-face communication when they collaborate.[3] Although the cost reduction per collaboration is greater for pairs that are farther apart, pairs that are closer together may interact more and thus create more opportunities for collaboration. Furthermore, electronic communication may be more valuable when paired with face-to-face meetings (Jess Gaspar and Edward L. Glaeser 1998).[4]

Overall, we find that middle-tier schools significantly increased their collaboration rates with co-located top-tier schools after Bitnet connection. These findings imply that the reduction in collaboration costs further accentuated tendencies for research activity to agglomerate rather than disperse; they are also consistent with the notion that the drop in costs facilitated a more efficiently functioning market for inputs into the production of knowledge, thereby broadening the set of institutions that participated—and continue to participate—in the production of high-quality research.

[2] See, for example, Robert K. Merton (1973) and Jacques Mairesse and Laure Turner (2005).

[3] This is one of the arguments advanced to explain empirical evidence of agglomeration, particularly in knowledge-intensive industries (David B. Audrestch and Maryann P. Feldman 1996; Lynne Zucker, Michael Darby, and Marilynn B. Brewer 1998).

[4] A rich theoretical literature has established the ambiguous effect of an improvement in communications technologies on interaction and collaboration across distances (e.g., Gaspar and Glaeser 1998; Rosenblat and Mobius 2004). We draw on this literature to interpret our results.

I. A Brief Description of Bitnet

Bitnet was an early leader in network communications for the research and education community. It allowed communication via e-mail, access to remote file archives, use of Listserv, file transfer protocol (FTP), and compatibility with other operating systems such as UNIX.[5] The first Bitnet adopters were the City University of New York and Yale University in May 1981. By the end of the 1980s, Bitnet had become the largest academic network in the world for computer-based communications.[6] Even still, Bitnet did not have all the capabilities of today's Internet. For example, familiar Internet features such as the World Wide Web and the browser were not invented until the end of our study period.

While other networks (e.g., ARPANET, EDUNET, USENET, CSNET) existed at the same time, Bitnet is most suitable for the purposes of our study for a number of reasons. First, rather than being narrowly focused in such areas as defense or computer science like some of the other networks, Bitnet was made available to all scholars; it was consequently adopted more widely than any other network at the time, allowing us to explore how adoption changed collaboration patterns across a diverse set of institutions. Second, Bitnet adoption was carefully documented; data exist on the exact date of adoption for every institution in the network through 1990. This is not the case for other networks. Third, the ability of Bitnet users to exchange data through FTPs as opposed to certain other networks that allowed only bulletin board postings and text messages offers insight into collaboration in fields that particularly benefit from data sharing, such as electrical engineering.

II. Data

We use a variety of data sources to examine collaboration between institution-pairs across universities in top electrical engineering journals from 1981 to 1991. We describe each of our main data sources below and provide descriptive statistics in Table 1.

Publication Data.—Since we are interested in identifying the effect of Bitnet on collaboration, we use publication data from researchers in technical areas who were likely to be early adopters of this communications technology and who thus closely reflect the time variation in adoption. Specifically, we collect publication data (16,495 papers) from seven electrical engineering journals over the 11-year period 1981–1991.[7] Each of these journals is considered among the top outlets for research in the specified field. Since we focus only on these seven journals, the total number of publications in our analysis does not change systematically over time.[8] This means that we capture an overall change in multi-institution collaboration relative to single-institution collaboration, rather than simply an overall increase in research output.

We extract the unique author-affiliated institution information from each paper and categorize each paper as either single- or multi-institution (i.e., collaborative).[9] We identify 720 unique insti-

[5] http://computing.dcu.ie/~humphrys/net.80s.html, Mark Humphrys, *The Internet in the 1980s* (September 15, 2006).

[6] Vijay Gurbaxani (1990) provides a detailed account of the diffusion of Bitnet.

[7] The journals are: 1) *IEEE Transactions on Aerospace and Electronic Systems,* 2) *IEEE Transactions on Nuclear Science,* 3) *IEEE Transactions on Biomedical Engineering,* 4) *IEEE Journal of Quantum Electronics,* 5) *IEEE Transactions on Electron Devices,* 6) *IEEE Transactions on Communications,* and 7) *IEEE Transactions on Education.*

[8] A total of 1,989 papers were published in the first year of observation (1981) and 1,401 papers in the last year (1991). The total number of publications fluctuates from year to year due to the publication of special issues and occasional conference proceedings. The distribution of article quantity across journals also varies.

[9] Papers with multiple authors are still classified as single-institution if all authors are from the same university.

TABLE 1—DESCRIPTIVE STATISTICS

Variable (by year)	Mean	Standard deviation	Minimum	Maximum	Number of observations
Institution level					
Total papers	2.779	6.453	0	131	2970
Multi-institution papers	1.587	3.369	0	92	2970
Single-institution papers	1.191	3.546	0	39	2970
R&D in electrical engineering (millions of $, lagged)	1.350	4.752	0	67.613	2970
No. of electrical engineering doctorates given (lagged)	2.895	6.864	0	67	2970
No. of electrical engineering post-doctoral students present (lagged)	0.652	2.283	0	50	2970
Average year adopting Bitnet[a]	1985.5	2.018	1981	1990	2970
Has Bitnet	0.400	0.490	0	1	2970
Multi-institution papers if Tier 1	3.819	4.889	0	39	990
Multi-institution papers if Tier 2	0.596	1.107	0	7	990
Multi-institution papers if Tier 3	0.346	1.197	0	13	990
Institution-pair level					
No. of collaborative papers between the pair	0.00165	0.0521	0	6	399,465
Dummy for pair-years where there is collaboration	0.00134	0.0365	0	1	399,465
Dummy for collaboration if at least one has not adopted Bitnet	0.000724	00269	0	1	298,491
Dummy for collaboration if both have adopted Bitnet	0.00315	0.0560	0	1	100,974
Distance	1767.9	1301.2	0	8293.7	399,465
Sum of no. of single-institution papers produced by the pair	2.381	5.025	0	117.0	399,465
Sum of R&D in electrical engineering (millions of $, lagged)	2.700	6.755	0	132.2	399,465
Sum of no. of electrical engineering doctorates given (lagged)	5.789	9.765	0	128.0	399,465
Sum of no. of electrical engineering post-doctoral students present (lagged)	1.304	3.225	0	62.0	399,465
Dummy if at least one of the pair has adopted Bitnet	0.547	0.498	0	1	399,465
Dummy if only the lower-tier university in the pair has adopted Bitnet	0.174	0.379	0	1	399,465
Dummy if both institutions have adopted Bitnet	0.253	0.434	0	1	399,465

[a] Conditional on adopting Bitnet by the end of 1990.

tutions, of which 270 are US universities, our institution type of interest. These form the basis of our unit of analysis.[10] Thus, our primary dataset consists of 36,315 institution-pairs over 11 years, resulting in a balanced panel with 399,465 observations.[11]

Bitnet Connection Data.—We use an online reference, Cyber Geography Research, for a record of Bitnet connections.[12,13] Importantly, there is significant variation in these data. Although only three institutions were connected in 1981, 66, 183, and 225 were connected by 1984, 1987, and 1990, respectively.

[10] We focus on US universities because many of the international institutions and US nonuniversity research labs used networks other than Bitnet. Less than 1 percent of the connected institutions were for-profit (http://computing.dcu.ie/~humphrys/net.80s.html, Mark Humphrys, *The Internet in the 1980s*, September 15, 2006).

[11] For Table 4, we also construct a single-institution dataset that includes the same 11 years of publishing from the specified journals by the 270 institutions of interest. Therefore, this is a balanced panel dataset of 2,970 observations.

[12] http://www.cybergeography.org/atlas/bitnet_topology.txt (September 15, 2006).

[13] We use the year following the technical connection as the first year Bitnet was available at the university. In the journals examined here, six months is a typical publication lag from manuscript submission to publication. All results are robust to using the same year of adoption.

Quality Data.—Since we are interested in the way university research orientation (or "quality") mediates the effect of Bitnet adoption on collaboration propensity, we categorize each university as being Tier 1 (high-research orientation, or top-tier), Tier 2 (medium-research orientation, or middle-tier), or Tier 3 (lower-research orientation, or lower-tier). We define institution quality based on ranking by total university-level National Science Foundation (NSF) funding over four years prior to our sample (1977–1980).[14] Thus, we classify the 270 universities in our data into three tiers, with 90 universities in each. In Agrawal and Goldfarb (2006), we ensure robustness using a number of alternative definitions of institution quality.

Distance Data.—In order to understand how distance between universities mediates the effect of Bitnet adoption on their propensity to collaborate, we calculate the straight-line distance between all possible pairs. We establish the location of each university's primary research campus from its official Web site and collect latitude and longitude data from the US Geological Survey based on city-state information.[15] We determine the distance between each university pair by employing the great circle method.[16]

III. Empirical Strategy and Results

A. *Did Bitnet Facilitate Collaboration across Institutions?*

Our estimation strategy is based on difference-in-differences identification. Using the paired institution data, we examine changes in collaboration between institution-pairs that both adopted Bitnet relative to pairs in which one or both did not adopt. We label the first institution in the pair i, the second j, and the year t.

We run linear regressions on the data using the following equation:

$$Collaboration_{ijt} = \alpha X_{ijt} + \beta Both\ Have\ Bitnet_{ijt} + \mu_t + \phi_{ij} + \varepsilon_{ijt}, \tag{1}$$

where the key explanatory variable, $Both\ Have\ Bitnet_{ijt}$, is a dummy that equals 1 if both institution i and j have connected to Bitnet by year t.[17] In addition, ϕ_{ij} measures institution-pair fixed effects, μ_t measures year fixed effects, and X_{ijt} is a vector of observable institution-pair-year characteristics including number of single-authored publications, number of electrical engineering doctorates awarded,[18] number of electrical engineering postdoctoral students,[19] and R&D expenditure in electrical engineering.[20] We lag the latter three covariates by one year to reflect the time between their input into research and final publication. The fixed effects mean that institution-pair level explanatory variables like collaboration behavior in the 1970s cannot be included as a control. For this linear equation to identify the average effect of Bitnet adoption on collaboration between two given institutions, we implicitly assume that unobserved institution-pair quality can be decomposed into an additively separable fixed component and a time-varying component that is constant across institution-pairs (Susan Athey and Scott Stern 2002).

[14] http://www.nsf.gov/awardsearch/tab.do?dispatch=4 (October 2, 2006).

[15] US Geological Survey: http://geonames.usgs.gov/, Web query application: http://geonames.usgs.gov/pls/gnis/web_query.gnis_web_query_form (September 15, 2006).

[16] $\operatorname{acos}(\cos(lat1)\cos(long1)\cos(lat2)\cos(long2) + \cos(lat1)\sin(long1)\cos(lat2)\sin(long2) + \sin(lat1)\sin(lat2)) \times earthRadius.$

[17] We also examine time since Bitnet adoption, the effect of which is illustrated in Figure 1.

[18] NSF Survey of Earned Doctorates.

[19] Survey of Graduate Students & Postdoctorates in Science and Engineering.

[20] Survey of R&D Expenditures at Universities and Colleges. This variable measures annual spending by electrical engineering departments. We include spending from NSF grant money in the value.

We treat $Collaboration_{ijt}$ as a dummy variable for whether institutions i and j had any collaborations in year t. We estimate equation (1) using a fixed effects linear probability (OLS) regression with the fixed effects differenced out using average values.[21] We treat collaboration as a dummy variable because 78 percent of all institution-pair-years with at least one collaboration had only one. Heteroskedasticity-robust standard errors are clustered by institution pair–Bitnet adoption status.

The first column of Table 2 shows our baseline specification. We regress collaboration on both universities in the pair being connected to Bitnet (*Both have Bitnet*), institution-pair fixed effects, and year fixed effects. Collaborations increased by approximately 50 percent after both universities in the pair were connected. This represents a significant increase in the propensity to collaborate. However, as we will show in the following sections, Bitnet adoption had an even greater effect (more than double) on collaboration between certain types of institutions, namely top-tier/middle-tier pairs that were co-located.

First, though, we provide a combination of statistical and institutional evidence that our main findings are not likely to be a result of omitted variables or Bitnet adoption endogeneity. For example, there may be omitted variables bias because certain universities shifted policy to increase their performance, which resulted in both Bitnet adoption and increased research output. Or maybe certain universities recruited young new faculty who had a taste for both electronic networking and collaboration. Or there may be endogeneity if universities adopted Bitnet *because* their collaborations were increasing.

We attend to these concerns in a number of ways. To address omitted variables bias, we first add the four covariates (X_{ijt}) described above that control for observable changes in department quality over time. The second column of Table 2 shows the results. The coefficient on *Both have Bitnet* is smaller in this regression, indicating that these controls explain some of the variation, but the relationship of interest is still statistically significant and economically important; the rate of collaboration increased by approximately 40 percent if both institutions were connected. We include these four controls in all subsequent specifications.

A second way we address omitted variables bias is by verifying that the measured impact of Bitnet did not begin prior to adoption. If the increase in collaboration is related to Bitnet adoption because middle-tier schools were improving in research and also investing in communications technology, then we would expect to observe an increase in collaboration in the years preceding adoption. To explore this possibility, we substitute the *Both have Bitnet* variable for a sequence of dummy variables for the years before and after adoption. Table 2, column 3, shows no preexisting trend toward increasing collaboration between schools that connect. Figure 1 gives a fuller specification with the predicted collaboration rates by year before and after adoption. Again, collaboration did not significantly increase in the years preceding Bitnet adoption. Collaboration rates began to rise in the year following adoption and then rose substantially two and three years after adoption. They then remained at a higher rate.

[21] We focus on the linear results for three reasons. First, OLS allows coefficients to be easily interpreted and compared across models. Second, linear regression allows for differencing out the mean fixed effects and using the full dataset. Third, while fixed effects logit and poisson regressions also allow differencing of mean effects, nonlinear methods are not necessarily consistent when there is a large number of zeros in the dependent variable (Gary King and Lanche Zeng 2001). The linear probability model is consistent and the estimated errors (with a heteroskedasticity correction) are correct. Jeffrey M. Wooldridge (2002) argues that the primary concerns about the linear probability model involve extreme values of the independent variables. He further argues that the case for using the linear probability model instead of a nonlinear model is strongest when the variables of interest are discrete, as is the case here. Our companion working paper, Agrawal and Goldfarb (2006), shows robustness to numerous other specifications in modeling (i.e., fixed effects (FE) probit, FE negative binomial, FE zero-inflated poisson, conditional FE logit, conditional FE poisson, and random effects poisson), independent variable choices, and samples.

TABLE 2—BITNET ADOPTION AND COLLABORATION USING INSTITUTION-PAIRS

	(1)	(2)	(3)	(4)	(5)
Dependent variable is *collaboration*	Main specification without time-varying institution characteristics	Main specification: Linear regression with a dummy for any collaboration as the dependent variable	Includes three years prior to adoption	Includes variable if just one institution has adopted	Includes variable if just lower-tier institution has adopted
Both have Bitnet	0.000852***	0.000667***	0.000701***	0.000673***	0.000890***
	(0.000198)	(0.000199)	(0.000209)	(0.000198)	(0.000297)
One or more has adopted Bitnet				−0.0000652	
				(0.000178)	
Only lower-tier institution has adopted Bitnet					−0.000396
					(0.000296)
One year before both have Bitnet			−0.00000781		
			(0.000252)		
Two years before both have Bitnet			0.0000992		
			(0.000258)		
Three years before both have Bitnet			0.000124		
			(0.000254)		
Sum of no. of single-institution papers		0.00000277	0.00000257	0.00000271	0.00000293
		(0.0000586)	(0.0000586)	(0.0000586)	(0.0000585)
Sum of R&D in electrical engineering (millions of $, lagged)		0.000146***	0.000147***	0.000146***	0.000146***
		(0.0000330)	(0.0000330)	(0.0000330)	(0.0000329)
Sum of no. of electrical engineering post-doctoral students present (lagged)		0.0000287	0.0000290	0.0000290	0.0000283
		(0.0000480)	(0.0000479)	(0.0000480)	(0.0000479)
Sum of no. of electrical engineering doctorates given (lagged)		0.0000435	0.0000437	0.0000434	0.0000427
		(0.0000349)	(0.0000349)	(0.0000349)	(0.0000348)
No. of observations	399,465	399,465	399,465	399,465	399,465
No. of groups	36,315	36,315	36,315	36,315	36,315
R^2 (within)	0.001	0.001	0.001	0.001	0.001
R^2	0.186	0.187	0.187	0.187	0.187

Notes: Regressions include year and institution-pair fixed effects. Robust standard errors (clustered by pair-Bitnet status) in parentheses.
*** Significant at, or below, 1 percent.
** Significant at, or below, 5 percent.
* Significant at, or below, 10 percent.

A final exploration of omitted variables bias is shown in Table 2, columns 4 and 5. These columns examine whether only one of the schools in the pair adopting Bitnet increased collaboration (column 4) and whether only the lower-tier school in the pair adopting Bitnet increased collaboration (column 5). If so, this would imply that Bitnet adoption was correlated with some other factor that influenced collaboration since both institutions needed to be connected to utilize the network as a collaboration tool. However, the coefficients on *One or more has adopted* and *Only lower-tiered institution has adopted* are neither significant nor large, and the coefficients on *Both have Bitnet* remain similar in significance and magnitude. This finding is consistent with the assertion that Bitnet facilitated collaboration by lowering communication costs between connected institutions. We believe that these combined results suggest that omitted variables bias is not a primary concern in our analysis.

To fully dispel endogeneity concerns, we would need a strong instrument that is correlated with adoption but not with the propensity to collaborate. Unfortunately, such an instrument is unavailable here. In its absence, we rely on the institutional history of the Bitnet connection process, which indicates that directors of university computing centers, rather than individual

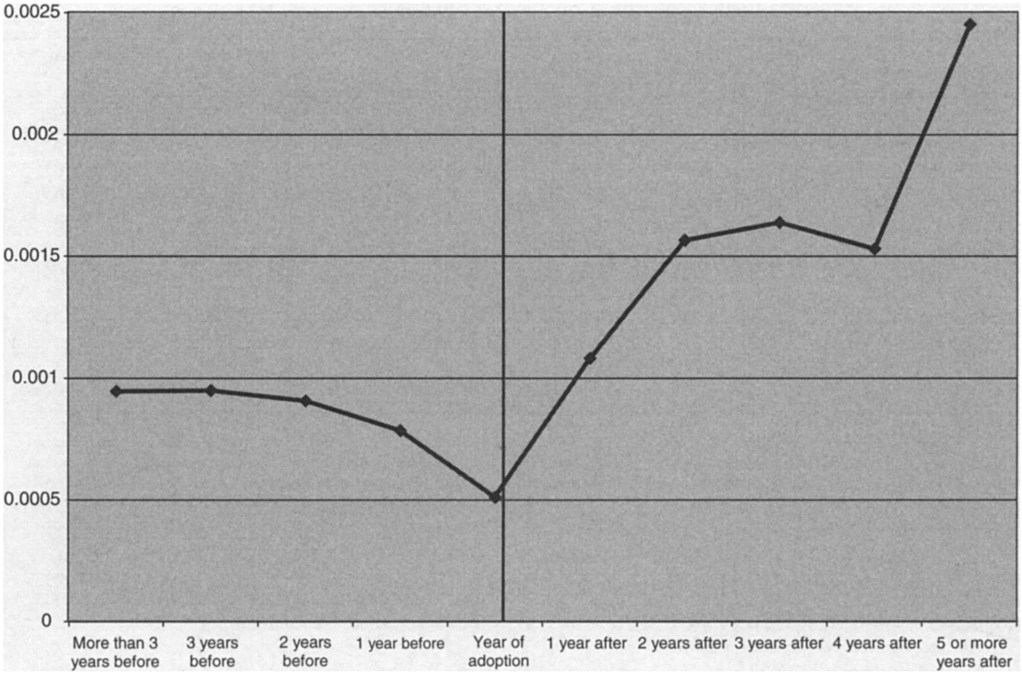

FIGURE 1. PREDICTED COLLABORATION RATES BY YEAR BEFORE AND AFTER ADOPTION*

* See Agrawal and Goldfarb (2006) for coefficient estimates.

professors from electrical engineering, drove institution-level adoption decisions.[22] Furthermore, the finding that collaboration increased only after Bitnet adoption (Figure 1) provides additional support. Moreover, the thrust of our argument is that Bitnet facilitated (rather than caused) an increase in cross-university collaboration. Researchers collaborate only if they want to. Even if the researchers studied here did influence their university's decision to adopt Bitnet so they could collaborate, the network succeeded in facilitating that collaboration.

B. *Does the Bitnet Effect Vary with Institution Quality?*

In addition to an overall rise in collaboration, a drop in communication costs might have led to a change in the relative roles of institutions of different qualities in research production. To explore this, we divide the university-pairs in our sample into six quality-type groups as catego-

[22] Ira Fuchs, the founder of Bitnet, described the many individual university adoption decisions he was familiar with as being made predominantly by computing center directors. At the conception of the network, for example, he personally sent letters to IT administrators (not researchers) at approximately 50 institutions and visited many more on a personal basis to convey the benefits of joining Bitnet. In addition, he lectured about the mechanics and attributes of Bitnet at public forums, such as EDUCOM, that were primarily attended by administrators. Dr. Fuchs described the "value proposition" that was used to persuade university administrators to connect as being largely predicated on the argument that "if nothing else, it will be very useful for aiding your IT staff to communicate with others" (personal interview, July 26, 2007). This description that emphasizes the role of administrators in making the adoption decision is corroborated by the dean of science at the National University of Singapore in an article that describes an administrator learning about Bitnet at EDUCOM and then championing his university's connection through the institution's bureaucracy (http://www.physics.nus.edu.sg/~phytanb/bitnet4.htm).

TABLE 3—BITNET ADOPTION, COLLABORATION, AND INSTITUTION-PAIR QUALITY

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable is *collaboration*	TIER 1 and TIER 1	TIER 1 and TIER 2	TIER 1 and TIER 3	TIER 2 and TIER 2	TIER 2 and TIER 3	TIER 3 and TIER 3
Both have Bitnet	−0.00164	0.00181***	0.0000951	0.000513*	0.0000359	−0.000404*
	(0.00156)	(0.000451)	(0.000300)	(0.000278)	(0.000235)	(0.000225)
Sum of no. of single-institution papers	0.0000117	−0.0000248	0.00000701	0.000390***	0.0000531	0.000391*
	(0.000165)	(0.0000845)	(0.0000362)	(0.000151)	(0.000130)	(0.000205)
Sum of R&D in electrical engineering (millions of $, lagged)	0.000312***	0.000124**	0.0000560***	−0.0000689	−0.0000225	−0.000162
	(0.000103)	(0.0000575)	(0.0000216)	(0.0000486)	(0.0000169)	(0.000226)
Sum of no. of electrical engineering post-doctoral students present (lagged)	−0.0000918	0.000103	0.0000674	−0.000106*	0.00000113	−0.0000474
	(0.000117)	(0.0000744)	(0.0000618)	(0.0000585)	(0.0000524)	(0.0000522)
Sum of no. of electrical engineering doctorates given (lagged)	0.000154	−0.00000932	−0.0000666**	−0.0000602	−0.000135**	0.0000691
	(0.000110)	(0.0000409)	(0.0000290)	(0.0000495)	(0.0000665)	(0.000133)
No. of observations	44,055	89,100	89,100	44,055	89,100	44,055
No. of groups	4,005	8,100	8,100	4,005	8,100	4,005
R^2 (within)	0.002	0.001	0.001	0.001	0.001	0.001
R^2	0.199	0.178	0.178	0.102	0.175	0.092

Notes: Regressions include year and institution-pair fixed effects. Robust standard errors (clustered by pair-Bitnet status) in parentheses. TIER 1, TIER 2, and TIER 3 based on NSF funding from 1977 to 1980.
*** Significant at, or below, 1 percent.
** Significant at, or below, 5 percent.
* Significant at, or below, 10 percent.

rized by ranking total NSF grants received by each university over the four years preceding our study: Tier 1–Tier 1, Tier 1–Tier 2, Tier 1–Tier 3, Tier 2–Tier 2, Tier 2–Tier 3, and Tier 3–Tier 3. Interestingly, only the coefficients on Tier 1–Tier 2 and Tier 2–Tier 2 pairs are significantly positive (Table 3). Tier 1–Tier 2 pairs in particular showed a substantial increase in collaboration rate after connection. For this subsample, both universities in the pair being connected increased the likelihood of collaboration by 133 percent over the average collaboration rate in the sample.

We next seek to better understand who benefits from collaboration between top-tier/middle-tier pairs. We analyze single-institution-level data to provide suggestive evidence that it is the middle-tier institutions that benefited most from top-tier/middle-tier collaboration. These are OLS regressions of total publications on *HasBitnet*, institution-specific covariates, year fixed effects, and institution fixed effects (differenced out).[23] Table 4 shows that Bitnet adoption is associated with an increase in total research output by middle-tier schools. This is not true of top-tier and lower-tier schools.

Overall, our results suggest the benefits of Bitnet adoption, measured by an increase in publications, likely accrued primarily to middle-tier schools (Table 4) due to collaboration with top-tier schools (Table 3). The reduction in communication costs associated with Bitnet seems to have

[23] The qualitative results of this table do not change if we use fixed effect poisson regressions instead. In fact, the significance of the Tier 2 results increases. We use a linear model to be consistent with the rest of the paper. For the regressions in Table 4 to identify the relationship between adoption and research production, we assume that we can decompose unobserved institution quality into an additively separable fixed component and a time-varying component that is constant across institutions. This assumption is questionable if Bitnet adoption is associated with an unobserved quality improvement. For this reason, we are especially cautious in our interpretation of the Table 4 results.

TABLE 4—BITNET ADOPTION AND TOTAL PUBLICATIONS, SINGLE-INSTITUTION DATA

	(1)	(2)	(3)
Dependent variable is *# of publications*	TIER 1	TIER 2	TIER 3
Has Bitnet	−0.0775	0.233*	0.0638
	(0.656)	(0.135)	(0.122)
R&D in electrical engineering (millions of $, lagged)	0.0856	−0.0142	0.611*
	(0.0805)	(0.0195)	(0.322)
No. of electrical engineering post-doctoral students present (lagged)	0.0344	0.218**	0.0897
	(0.0744)	(0.103)	(0.179)
No. of electrical engineering doctorates given (lagged)	−0.0284	0.0347	−0.116
	(0.0843)	(0.0421)	(0.0998)
No. of observations	990	990	990
No. of groups	90	90	90
R^2 (within)	0.06	0.04	0.04
R^2	0.808	0.447	0.679

Notes: Regressions include year and institution fixed effects. Robust standard errors in parentheses. TIER 1, TIER 2, and TIER 3 based on NSF funding from 1977 to 1980.
*** Significant at, or below, 1 percent.
** Significant at, or below, 5 percent.
* Significant at, or below, 10 percent.

led to a broadening of the institutions participating in the production of high-quality research, perhaps due to the benefits of gains from trade through cross-university collaboration.[24]

C. *Does the Bitnet Effect Vary with Distance?*

If the drop in communication costs did not have a uniform effect over distance on propensity to collaborate, then distance, like quality, may have mediated the effect of Bitnet, leading to a change in the spatial distribution of collaboration. To explore this, we employ a spline regression, grouping university pairs by the distance between them. Our results using all institution-pairs (Table 5 column 1) suggest that overall Bitnet adoption was associated with increases in both local and distant collaborations.[25]

Splitting the data by pair quality provides important detail on how distance was related to Bitnet adoption and collaboration. Columns 2 through 7 of Table 5 show that the greatest effect on multi-institutional paper production occurred for co-located, top-tier/middle-tier pairs. Middle-tier universities also increased their collaboration with non-co-located top-tier universities, but the effect of Bitnet was several times greater for those that were co-located.[26] These findings suggest that low-cost electronic communication, while perhaps a substitute for face-to-face interactions under certain conditions, is also an effective complement, reinforcing other factors that lead to agglomeration, including thicker labor markets and scale economies in capital-intensive equipment.

[24] The introduction of Bitnet was only one piece of an evolving US research infrastructure in the latter half of the twentieth century. For example, a report prepared for President Eisenhower (Glenn T. Seaborg 1960) explicitly called for an increase in research funding, especially for middle-tier research universities. Still, the specific timing of changes in collaboration patterns identified here is so tightly tied to the adoption of Bitnet that it seems probable the introduction of this communications technology was instrumental in unlocking the potential of middle-tier universities, though they may have been nurtured for some time prior through this and other policy initiatives.

[25] Defining "local" more broadly as 250 kilometers leads to a large and significant coefficient.

[26] Examples of collaborating, co-located, top-tier/middle-tier pairs include the Massachusetts Institute of Technology–Northeastern University in Boston/Cambridge and the University of Pennsylvania–Drexel University in Philadelphia.

TABLE 5—BITNET ADOPTION, COLLABORATION, INSTITUTION-PAIR QUALITY, AND DISTANCE

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent variable is *collaboration*	All Data	TIER 1 and TIER 1	TIER 1 and TIER 2	TIER 1 and TIER 3	TIER 2 and TIER 2	TIER 2 and TIER 3	TIER 3 and TIER 3
Distance is under 100 km and Both Adopted Bitnet	0.00347	−0.0179	0.01424***	0.00538	−0.00208	0.00657	−0.000317
	(0.00310)	(0.0143)	(0.00476)	(0.00511)	(0.00153)	(0.00511)	(0.000250)
Distance is between 100 km and 500 km and Both Adopted Bitnet	0.000380	−0.00347	0.00218***	−0.000578	0.00280**	−0.00138*	−0.000203
	(0.000559)	(0.00277)	(0.000830)	(0.000872)	(0.00122)	(0.000767)	(0.000175)
Distance is between 500 km and 1000 km and Both Adopted Bitnet	0.000133	−0.00411*	0.00143**	−0.000121	0.000570	0.000135	−0.000148
	(0.000355)	(0.00210)	(0.000670)	(0.000331)	(0.000464)	(0.000278)	(0.000166)
Distance is between 1000 km and 3000 km and Both Adopted Bitnet	0.000614***	−0.000319	0.00119***	0.000264	−0.0000372	0.000313	−0.000607*
	(0.000216)	(0.00166)	(0.000443)	(0.000378)	(0.000258)	(0.000271)	(0.000324)
Distance is over 3000 km and Both Adopted Bitnet	0.00131***	0.000101	0.00234***	−0.000146	0.000660	−0.000238*	−0.000320
	(0.000419)	(0.00178)	(0.000873)	(0.000526)	(0.000710)	(0.000144)	(0.000196)
Sum of # of single-institution papers	0.00000348	0.0000110	−0.0000230	0.00000720	0.000388**	0.0000534	0.000393*
	(0.0000585)	(0.000165)	(0.0000845)	(0.0000362)	(0.000151)	(0.000130)	(0.000205)
Sum of R&D in electrical engineering (millions of $, lagged)	0.000145***	0.000326***	0.000116**	0.0000569***	−0.0000685	−0.0000224	−0.000163
	(0.0000328)	(0.000103)	(0.0000556)	(0.0000218)	(0.0000485)	(0.0000167)	(0.000227)
Sum of # of electrical engineering post-doctoral students present (lagged)	0.0000270	−0.0000967	0.0000995	0.0000669	−0.000107*	0.00000358	−0.0000471
	(0.0000479)	(0.000117)	(0.0000743)	(0.0000618)	(0.0000590)	(0.0000531)	(0.0000521)
Sum of # of electrical engineering doctorates given (lagged)	0.0000429	0.000150	−0.00000595	−0.0000674**	−0.0000600	−0.000136**	0.0000710
	(0.0000350)	(0.000111)	(0.0000411)	(0.0000291)	(0.0000492)	(0.0000666)	(0.000134)
# of Observations	399,465	44,055	89,100	89,100	44,055	89,100	44,055
# of Groups	36,315	4005	8100	8100	4005	8100	4005
R^2 (within)	0.001	0.002	0.001	0.001	0.001	0.001	0.001
R^2	0.187	0.199	0.178	0.178	0.102	0.175	0.092

Notes: Regressions include year and institution-pair fixed effects. Robust standard errors (clustered by pair-Bitnet status) in parentheses. TIER 1, TIER 2, and TIER 3 based on NSF funding from 1977 to 1980.
*** Significant at, or below, 1 percent.
** Significant at, or below, 5 percent.
* Significant at, or below, 10 percent.

IV. Conclusions

Overall, these findings enhance our understanding of knowledge production. A sharp decrease in collaboration costs amplified the role of middle-tier universities in the production of high-quality research. In effect, Bitnet widened the circle of institutions participating in the national innovation system.[27] These findings offer meaningful insight since knowledge production ("innovation") is central to economic growth (Paul M. Romer 1990) and universities are an important component of the innovation system (Richard Nelson and Nathan Rosenberg 1993). Universities, of course, are not all the same; they are endowed with different levels of resources and different specialized expertise. Our results are indicative of a profound shift in the knowledge production system from previously unrealized gains from trade, possibly through the increased exploitation of underutilized research equipment and/or enhanced specialization.[28]

[27] Ira Fuchs, the founder of Bitnet, responded to these findings by stating that part of the *raison d'être* of Bitnet was to "democratize connectivity" beyond the defense research community (personal interview, May 25, 2006).

[28] Our findings are consistent with a *vertical* specialization of tasks, in contrast to the "O-Ring" theory of production in which workers match with other workers of equal quality (Michael Kremer 1993).

Due to the nature of our data, however, we are unable to comment on whether Bitnet delivered an overall productivity increase. We have no data on inputs, and our output measure—publications from a fixed set of journals—remains reasonably constant over time. To be clear, what we observe is that Bitnet facilitated a change in the relative roles of certain types of universities with respect to the production of high-quality research.

Moreover, institutions other than universities, such as those from the private sector, also became more involved in the collaborative production of knowledge. For example, in 1981, private firms did not contribute to the collaborative research output in our set of journals, whereas they contributed to 7 percent and 12 percent of the collaborative papers in 1986 and 1991, respectively. Thus, our findings provide only a partial picture of the evolution of knowledge production.

In terms of the generalizability of our results, E. Han Kim, Adair Morse, and Luigi Zingales (2006) report that in economics, the research productivity effect of being affiliated with an elite institution was significant in the 1970s, weakened in the 1980s, and disappeared in the 1990s; the timing of this relative rise of nonelite institutions in economics is consistent with our engineering results.[29] Furthermore, sociologists studying oceanography, mathematics, physics, chemistry, and experimental biology found a greater correlation between network use and productivity for "peripheral" scientists who had limited access to research resources (Hesse et al. 1993; Walsh and Bayma 1996). Collectively, these papers suggest that the findings we report here may apply in fields beyond electrical engineering.

REFERENCES

Agrawal, Ajay K., and Avi Goldfarb. 2006. "Restructuring Research: Communication Costs and the Democratization of University Innovation." National Bureau of Economic Research Working Paper 12812.

Aral, Sinan, Erik Brynjolfsson, and Marshall Van Alstyne. 2007. "Information, Technology and Information Worker Productivity: Task Level Evidence." National Bureau of Economic Research Working Paper 13172.

Athey, Susan, and Scott Stern. 2002. "The Impact of Information Technology on Emergency Health Care Outcomes." *RAND Journal of Economics*, 33(3): 399–432.

Audretsch, David B., and Maryann P. Feldman. 1996. "R&D Spillovers and the Geography of Innovation and Production." *American Economic Review*, 86(3): 630–40.

Baker, George P., and Thomas N. Hubbard. 2003. "Make Versus Buy in Trucking: Asset Ownership, Job Design, and Information." *American Economic Review*, 93(3): 551–72.

Dixit, Avinash K., and Victor D. Norman. 1980. *Theory of International Trade*. Cambridge: Cambridge University Press.

Gaspar, Jess, and Edward L. Glaeser. 1998. "Information Technology and the Future of Cities." *Journal of Urban Economics*, 43(1): 136–56.

Gurbaxani, Vijay. 1990. "Diffusion in Computing Networks: The Case of Bitnet." *Communications of the ACM*, 33(12): 65–75.

Hamermesh, Daniel S., and Sharon M. Oster. 1998. "Tools or Toys? The Impact of High Technology on Scholarly Productivity." National Bureau of Economic Research Working Paper 6761.

Helpman, Elhanan, and Paul Krugman. 1985. *Market Structure and Foreign Trade: Increasing Returns, Imperfect Competition, and the International Economy*. Cambridge, MA: MIT Press.

Hesse, Bradford W., Lee S. Sproull, Sara B. Kiesler, and John P. Walsh. 1993. "Returns to Science: Computer Networks in Oceanography." *Communications of the ACM*, 36(8): 90–101.

Kim, E. Han, Adair Morse, and Luigi Zingales. 2006. "Are Elite Universities Losing Their Competitive Edge?" National Bureau of Economic Research Working Paper 12245.

[29] Daniel S. Hamermesh and Sharon M. Oster (1998) and Gaspar and Glaeser (1998) also examine the field of economics and provide more general results that are consistent with our findings. Hamermesh and Oster show an increase in collaborative research in economics from the 1970s to the 1990s, while Gaspar and Glaeser find a rapid growth in local collaboration in economics since the 1960s.

King, Gary, and Lanche Zeng. 2001. "Logistic Regression in Rare Events Data." *Political Analysis*, 9(2): 137–63.

Kremer, Michael. 1993. "The O-Ring Theory of Economic Development." *Quarterly Journal of Economics*, 108(3): 551–75.

Mairesse, Jacques, and Laure Turner. 2005. "Measurement and Explanation of the Intensity of Co-publication in Scientific Research: An Analysis at the Laboratory Level." In *New Frontiers in the Economics of Innovation and New Technology: Essays in Honor of Paul David*, ed. Cristiano Antonelli, Dominique Foray, Bronwyn H. Hall, and W. Edward Steinmueller, 255–94. Cheltenham: Edward Elgar Publishing.

Merton, Robert K. 1973. *The Sociology of Science: Theoretical and Empirical Investigations.* Chicago: The University of Chicago Press.

Nelson, Richard R., and Nathan Rosenberg. 1993. "Technical Innovation and National Systems." In *National Innovation Systems*, ed. Richard R. Nelson, 3–28. Oxford: Oxford University Press.

Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy*, 98(5, 2): 71–102.

Rosenblat, Tanya S., and Markus M. Mobius. 2004. "Getting Closer or Drifting Apart?" *Quarterly Journal of Economics*, 119(3): 971–1009.

Seaborg, Glenn T. 1960. *Scientific Progress, the Universities, and the Federal Government.* Washington, DC: US Government Printing Office.

Walsh, John, and Todd Bayma. 1996. "The Virtual College: Computer-Mediated Communication and Scientific Work." *The Information Society*, 12(4): 343–63.

Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross Section and Panel Data.* Cambridge, MA: MIT Press.

Zucker, Lynne G., Michael R. Darby, and Marilynn B. Brewer. 1998. "Intellectual Human Capital and the Birth of U.S. Biotechnology Enterprises." *American Economic Review*, 88(1): 290–306.