ICIS 2003 Proceedings

International Conference on Information Systems (ICIS)

December 2003

So Close and Yet So Far: Information Technology and the Spatial Distribution of Customer Service

David Fitoussi
Massachusetts Institute of Technology

SO CLOSE AND YET SO FAR: INFORMATION TECHNOLOGY AND THE SPATIAL DISTRIBUTION OF CUSTOMER SERVICE

David Fitoussi
Sloan School of Management
Massachusetts Institute of Technology
Cambridge, MA USA
dbf@MIT.EDU

Abstract

Where should firms locate? As communication technologies spread across city, region,and country boundaries, and communication channels multiply, many firms can potentially relocate some of their activities to regions with lower costs. While manufacturing has long been globalizing, IT is enabling a new wave focused on services. Spatial relocation is attractive to companies faced with the incessant pressure to control costs and information technologies can help firms transcend location boundaries. But these same technologies may also confer renewed importance to local assets that are hard to replicate in remote locations.

This paper develops a framework to analyze these effects by estimating the regional demand for customer service representatives of a homogenous set of Fortune 1000 manufacturing firms. The model is estimated using firm-level data and the estimated demand structure is used to assess the effects of technology on customer volume, location choices and cost savings. A 10 percent increase in the use of Internet applications is found to lead to a 2.5 percent decrease in the firm's employment of agents nationally. Moreover, the same increase reduces the willingness of firms to pay for regional benefits (technology "levels" the field between regions). However, the cost savings from the associated relocation are surprisingly small, averaging 1.3 percent of unit costs. Finally, the research shows that regional preferences vary widely among firms, suggesting that sensitivity to cost is highly firm-specific and that the importance of local assets does not vanish. Overall, these results show a positive relationship between technology and firms' price sensitivity, but not on the scale of a massive spatial reorganization. Firm-specific regional preferences still matter.

Keywords: IT, decentralization, globalization, customer service, discrete choice models, random coefficients

Introduction

Technological advances in computing and communication technologies have offered companies more flexibility in organizing work. Starting in the late 1970s with manufacturing jobs in industries such as textiles, shoes, and electronics moving overseas, the trend has now reached an increasing number of services that are produced at a distance from their final marketplace. General Electric, for example, employs almost 6,000 scientists and engineers in 10 foreign countries so that it can tap the world's best talent. Most of the largest IT firms—Microsoft, HP, IBM, EDS, CSC, Accenture, Cisco—have now moved part of their software development efforts offshore to Bangalore, India. The phenomenon is highly advertised, and hyped, to the point where some claim that any organization that does not outsource will lose its competitiveness (reminiscent of the predictions regarding e-business in the late 1990s). Still, others emphasize the difficulties arising from cultural differences, time differences, language, reliability, and accountability, issues that are harder to address for services, especially the kind that require customer interaction and personalization (Macke 2003). In fact, during previous go-global drives, many companies ended up repatriating manufacturing and design work because they felt they were losing control of core businesses or found them too hard to coordinate

(Engardio et al. 2003). Recently, Allegheny Energy Supply, a utility engaged in the supply of electricity and energy-related commodities, has moved its trading operations (*a priori*, a location-free activity) back to Pennsylvania, in order to be closer to its physical generating plants in the Midwest and Mid-Atlantic markets. Clearly, while lower costs are important, companies may not be able to exploit these cost savings if this would entail decentralizing proprietary assets that are hard to manage remotely. The role of IT in enabling or constraining decentralization is, therefore, dependent on the importance of these proprietary assets for the firm's business. The next section presents the theoretical framework for this analysis.

Theoretical Framework

The theory of the multinational enterprise emphasizes the existence of proprietary assets for explaining the basis for horizontal multiplant enterprises. This approach, developed through the work of a number of authors including Caves (1971) and Hennart (1982), describes proprietary assets as the resources that the firm can use but not necessarily contract upon or sell. An asset might represent knowledge about how to produce a cheaper or better product at given input prices, or how to produce a given product at a lower cost than competing firms. Assets of this kind are closely related to the firm-specific resources in the resource-based view of the firm (Wernerfelt 1984). These resources form the basis for the firm's competitive advantage, since they hold a revenue productivity for the firm, closely akin to product differentiation. Proprietary assets might affect the ability of multi-national firms to locate production based on production costs. In fact, Maki and Meredith (1986) point out that multinationals might be able to transfer production from a low-cost to a high-cost location if their proprietary assets embrace the ability to transfer their source-country cost advantages. Similarly, the inability of a firm to transfer proprietary assets might hinder its ability to exploit cost differences and relocate to lower-cost regions. In fact, economists have long recognized that local conditions can generate benefits for firms that cannot be replicated elsewhere. Ellison and Glaeser (1999) and Kim (1999) find that natural advantages explain a significant fraction of industry localization and location patterns. More significantly, Rosenthal and Strange (2001) list localization externalities (proximity to other, similar firms) as a determinant of industry localization patterns. Clearly these resources cannot be transferred to another location. In that respect, the role of IT is ambiguous. If computer and communication equipment allows firms unprecedented flexibility in locating business units, these same technologies are also associated with very large intangible investments in proprietary assets that might not be easily dispersed. Indeed, recent studies showed that each dollar of installed computer capital in a firm is associated with up to $10 of market value, suggesting very large investments in other intangible assets (Brynjolfsson et al. 2002). For tasks to be easily isolated and run in a remote low-cost region without costly central control and supervision, they must be relatively well defined and structured. But these are also the tasks that are more apt to be automated (McGrath and Hollingshead 1994). In that case, IT could have little effect on the delocalization of tasks that are intangible-intensive and, in fact, increase their relative importance.

In order to crystallize this idea, let us look at the following simple general-equilibrium model based on the Hecksher-Ohlin model typically used in international economics. Suppose that there are two types of services. The first, basic services, are well-defined, repetitive tasks that can be easily monitored and adapted without physical interaction. The second type of service requires coordination among several divisions of the firms, producers, and suppliers, as well as the exchange of sometimes ambiguous and sensitive information. The first type of service is labor intensive whereas the second type of service is dependent on the assets of the firm, in particular its intangibles (for example, its reputation or business processes) but does not require a large amount of labor. Let us call the first type of service B (basic) and the second type C (complex). If a firm is limited to using local labor, demand conditions determine which point on the production function will be chosen. Consider now what happens when service B can be provided at a distance at a lower cost: as labor employed in service B is more abundant, output declines in locally provided service B and expands in service C. Local service B discharges a lot of labor, thereby raising the rental price of the firm's intangible assets that are location-specific. As the value of the local assets increases, the firm's valuation of location becomes more important and employment in service C goes up as well. This simple model shows that IT can have a positive effect on the relative value of local employment. In order to analyze these effects, an empirical evaluation of the role of IT on staffing decisions is necessary.

As shown in previous productivity studies, there are significant advantages to studying IT effects at the firm level whenever possible. Firm-level data analysis has unveiled the impact of IT on productivity where aggregate-level analysis had found a productivity paradox (Brynjolfsson and Hitt 2003). Even more insightful is studying a single, well-defined function or process across firms. Doing so increases the confidence one has in the accuracy of the econometric results (Ichniowski and Shaw 2003). For this reason, I focus on manufacturing firms and on one service, archetypical of information work: customer service representatives (CSRs) answering customer calls. This type of service consists of providing information to customers on the phone, including processing orders and providing solutions to common questions and inquiries. It is also a good example of a footloose process that can, in principle, be sited anywhere. This makes it an ideal candidate to investigate in order to gain an understanding of how information technologies may impact the location of information work. I analyze the choices of firms in locating customer service representatives across various U.S. geographical locations. I conduct an econometric study of 106 Fortune 1000 firms in six different manufacturing industries and combine data on firms' IT assets, firm characteristics, and the features of different geographical regions in order to estimate the demand for customer service representatives in different regions using a discrete-choice model with random coefficients. In this framework, firms choose where to locate their CSRs (akin to consumers choosing which products to buy in a differentiated product setting) and how many CSRs to hire at each location (how many units of each product to purchase).

The contributions of this paper are three-fold. First, by using a random utility/discrete-choice model, I am able to estimate labor demand at the level of the firm and of the region by drawing on detailed micro-level data. To the best of my knowledge, this is the first application of such models to hiring decisions by firms (most of the applications of discrete-choice models consider the purchase of differentiated products by consumers). This demonstrates the usefulness of this kind of methodological framework for a variety of settings. As I will explain below, the discrete nature of the decisions involved (where to locate and how many employees to hire) makes alternative modeling frameworks unsuitable. Second, I derive an expression for firm profits from customer service—the latent variable—from the premises of a queueing model. The resulting profit function provides an approximation to the revenue generated by customer service activities that could be used to estimate the value of these activities (estimating the value of customer service is notoriously difficult). Third, the estimated demand allows me to evaluate the impact of technology on staffing requirements and location decisions. I find that a 10 percent increase in Internet-based applications leads to a 2.5 percent decrease in the number of CSRs. Presumably, the substitution effect between channels (between self-serve customer support over the Internet and customer service over the telephone) outweigh any awareness effect. I also find that firms that use more Web-based technologies or delegate purchasing decision rights more often have statistically significantly different location patterns (they are less willing to pay for local quality, suggesting that technology may increase competition between regions and encourage factor-price equalization). I evaluate the average cost savings that derive from these relocation patterns: a 10 percent increase in Web-based applications or in decision rights leads, respectively, to a 1.3 percent and 0.5 percent decrease in total unit costs. Finally, I find significant variation in firms' valuations of regional characteristics. This indicates that idiosyncratic preferences of regional characteristics (such as proximity to complementary activities at the same firm) play a significant role in location decisions, casting doubt on the predictions of a massive exodus of service operations to low-cost countries.

The rest of this paper is structured as follows: in the next section, I present the methodology, a variant of discrete-choice models with random coefficients. I then develop the basic model and the estimation method. As the model is highly nonlinear, I resort to simulation techniques as described in Pakes and Pollard (1989). The data is then described and sample statistics provided. Finally, the results of the estimation are presented and discussed.

The Methodology

The analysis of localization and CSR staffing decisions by firms cannot be accomplished by simply specifying a down-sloping demand curve for CSRs and estimating the parameters of this demand curve. The localization pattern of firms is a complex one: firms choose a variety of locations with different numbers of CSRs at each location to provide customer assistance. Specifying a demand curve for each region is unpractical: it would need to incorporate in each equation both the regional unit costs and the unit costs of every other region as dependent variables. The number of parameters to be estimated would be a quadratic function of the number of regions and, in general, unmanageable. An additional problem is that dependent variables (in that case, the number of CSRs at a location) are discrete and truncated at zero, leading to truncation bias (Amemiya 1974). The alternative used in this paper is to put some structure on the demand problem by assuming that regional characteristics drive demand patterns. The approach of product characteristics (Lancaster 1979) applied to geographical regions assumes that regions can be characterized by a set of common attributes (CSR average wage, communication infrastructure level, population and industrial concentration). In this case, a firm's valuation of a regional worker is a function of these attributes, of firm-specific taste parameters, and of a small set of parameters to be estimated. The demand patterns that we observe implicitly reflect a firm's profit maximization over its various localization and staffing alternatives. This framework is an extension of the classical discrete-choice model (DCM) allowing for multiple units (in our case, workers) to be chosen in addition to choices between brands (in our case, regions). Unlike classical discrete-choice models, however, it uses micro-level data. In general, in the absence of consumer-level data, DCMs aggregate consumer choices in an aggregate market demand and the estimation process relies on product market shares. This requires making assumptions regarding the distribution of consumer characteristics, assumptions that are not required if micro-data are available. I follow the framework developed by Hendel (1999) in his estimation of multiple-discrete choice models. However, I do not assume an arbitrary profit function but develop it from the premises of a queueing model (see the section on estimation later in this paper). The basic building blocks of the model are regional characteristics, firms' characteristics, preferences, and this profit function. I describe each of these in detail below.

All regional characteristics are assumed to be fully observable by firms when making their decisions. However, for each region, there exists an attribute ζ_i that is not observable (see Berry 1994). Failure to control for this unobservable characteristic may leave out important characteristics that firms consider in their decision process but that are not available to the researcher, and may also cause an endogeneity bias, as unit costs and this unobservable attribute are likely to be correlated. Although every firm faces the same regional characteristics, valuations of these attributes will likely vary across firms. The standard method to capture this feature is to introduce random coefficients. For each firm f, there exists a set of random coefficients β_f (the taste parameters) that reflect this firm's valuation of the characteristics. The random coefficients are assumed to be drawn for each firm independently from a normal distribution, whose mean and standard deviation are to be estimated.

$$\beta_f = [\beta_{f,1}, \ldots, \beta_{f,I+N-1}]$$

where the first I entries represent the firm's subjective quality perceptions about the I different regions and the last $N\text{-}1$ entries are the firm's valuations for the regional attributes.

Assume that V_i is a vector that contains regional characteristics, including the regional dummies. Then $V_i\beta_f$ represents the firm's valuation of a regional CSR. Since the coefficients on attributes are random, I can only estimate the mean valuation of each regional characteristic, controlling for firms' characteristics. In other words, I can derive an estimate for the mean valuation of each regional characteristic reflecting the average willingness to pay for that characteristic by firms, conditional on their characteristics. Let D_f be the set of these characteristics. Then $<B_f, D_f>$ completely specifies the behavior of the firm f.

The Model

In this section, I specify the profit environment that firms face in providing customer service to callers. The profit function plays the role of the utility function in classic discrete choice models, but unlike many of these models in which the utility function is a simple arbitrary linear form of the product characteristics, the functional form of the profit function is derived directly from the primitive components of a stochastic queueing model.

Calls to service representatives are generated randomly among current customers. I also assume that the time intervals between calls are independent and identically distributed (arrivals are Poisson). Mandelbaum et al. (2000) present some empirical evidence in support of this distributional assumption in a call center environment. I also allow for differences in call arrivals based on firms' vertical and horizontal differentiation as captured by the number of different industries in which the firm operates, A, on the firm's usage of the Internet as a proxy of its communication capabilities, and on the magnitude of its sales. Specifically, call arrivals follow a Poisson process with arrival rate $\lambda_f = \phi_f * Sales$, where ϕ can be described as:

$$\varphi_f = f_0 + f_1 * In_f + f_2 * A_f \tag{1}$$

This expression specifies the proportion of revenue that is associated with calls arriving at the locations that we observe as a function of the firm's parameters. Since I cannot control for international locations, variations in λ represent variations in the number of calls received at U.S. locations, which could be the effect of a reduction in the total number of calls a firm receives or an increase in the use of international agents.

Calls are routed to one of the available agents (regardless of location) or are placed in the queue, waiting for the next available agent. Given a firm's staffing strategy and system load, callers may have to wait in line for the next available agent a long time, and some will drop. Negative impact on customer satisfaction that results from long waiting times and its counterpart, higher retention rates from quick and efficient service, define a measure of revenue from customer service, R_f. We consider a call that drops a loss in revenue. Firms then choose the number of agents and their location to minimize staffing costs and revenue loss from dropped calls.

$$\pi_f = R_f - \sum_i X_i P_i \tag{2}$$

where i is indexing the different regions where the firm locates its CSRs, X is the choice variable (number of CSRs in region i), and P is the unit cost of locating a CSR in region i.

The revenue function R_f can be given an explicit form by using the Pollaczek-Khinchin formula (see Gallagher 1996) that relates, in an M/G/1 queue, the expected queueing time for a calling customer to the expected service time. An M/G/1 queue is a lower-bound approximation of an M/G/n queue. (In other words, n servers in parallel can be approximated as one server with a service rate bounded from above by the sum of the service time of the individual servers. In the derivation of the profit function below, I will approximate this compounded service time by a concave polynomial whose order α will be estimated.) Specifically:

$$\overline{W} = \frac{\lambda E[Z^2]}{2(1-\lambda E[Z])} \tag{3}$$

in which $\overline{W}$ is the time average waiting time and Z is customer service time. The customer service time is a function of the number of service representatives and their characteristics. The valuation of these characteristics by a firm is denoted μ_{if}. In each region i, μ_{if} is a firm's f valuation of the region's CSRs. μ_{if} incorporates the interactions between firm and region characteristics (the $X_i\beta_f$ in Berry 1994) and is defined as:

$$\mu_{if} = \max(0, B_f \cdot V_i)^{m(D_f)} \tag{4}$$

For instance, firms could value regions differently based on whether or not they already have operations in the region, or whether there is a good fit between the region's characteristics and the production process of the firm. The term $m(D_f)$ captures a form of vertical differentiation between firms in the sense that firms with similar value for the quality of the region might still differ in their willingness to pay for this quality.

Consider now what happens when a customer does not receive adequate service and its revenue is lost. Suppose that *(1-N)* is the proportion of incoming calls that are not answered (or that are given inadequate service). Then, the average revenue for the firm is $R_f = N^*\lambda_f$. Notice that N, the proportion of incoming calls that do not drop, corresponds to the survival rate of the queueing system. Mandelbaum et al. study these survival rates and show that they are exponentially decreasing functions of the average waiting time. After some algebraic manipulations (see Appendix A), the profit function becomes:

$$\pi_f = \lambda_f e^{-\frac{\lambda_f}{2(\sum_i \mu_{if}(X_i)^\alpha)^2}} - \sum_i X_i P_i \tag{5}$$

Firms choose the number of CSRs at different locations in order to maximize the profit function in equation (5). This problem is a discrete (integer) problem and therefore cannot be solved by standard optimization techniques. It is instructive though, and ultimately useful for solving the maximization problem, to temporarily ignore the integer constraint and derive the optimal number of agents in the relaxed problem. Appendix B shows the derivation of the optimal number of agents at the different locations i from the first-order conditions of the profit function. This derivation yields two interesting insights. First, the relative number of agents between different locations i and j is given by

$$\frac{X_{jf}}{X_{if}} = \left[\left(\frac{\mu_{jf}}{P_j}\right)\left(\frac{P_i}{\mu_{if}}\right)\right]^{\frac{1}{1-\alpha}} \tag{6}$$

Equation (6) shows that a firm's relative valuation of different regions can outweigh cost differentials (i.e., firms would not locate in a region with lower unit costs if their valuation of the regional characteristics were much lower relative to other regions). But it also shows the factors that affect the relative importance of local characteristics and firm characteristics (as reflected in the μ's) versus local unit costs (the P's). The model captures a kind of vertical differentiation between firms: firms with similar valuations of regional characteristics but different own characteristics (e.g., differences in technological investments or organizational practices as reflected in the term $m(D_f)$) will exhibit different willingness to pay and, thus, different localization strategies. Notice that regional choice is affected by $m(D_f)$ but not by the scale factor λ_f: the latter determines the number of CSRs at a chosen location. This is why I can identify both functions of firms' characteristics, by using data on location choices on the one hand, and the number of agents in a region on the other hand.

Second, the derivation of the optimal number of agents in the relaxed problem suggests an approach for solving the integer problem. I can derive the optimal number of CSRs without the integer constraint. For each firm, I select one region (without loss of generality, I select the region *i* where the firm located the highest number of CSRs X_{if}) and use equation (6) to compute the ratios of X_{jf} / X_{if} for the remaining regions *j*. This provides an analytical expression for each *X* as a function of X_{if}. I then use the first-order condition of the profit function with respect to X_{if} to derive a closed-form solution for X_{if}, and therefore for all of the X_{if}. Using this non-integer solution, I search for the optimal vector of integers by means of a standard branch-and-bound algorithm. The procedure is fast enough for the limited number of regions in the sample (19 regions, see the sample description below). The outcome of this procedure is a vector X^e of predicted CSRs employment in every one of the different regions in the choice set for a given set of random coefficients and parameters. The next section describes the role of this predicted vector in estimating the parameters using the method of simulated moments.

The Estimation

The model predicts the number of agents $X_f(D_f, \beta_f, \theta)$ at each location for a firm *f* as a function of observed firm characteristics D_f, random coefficients β_f, and the vector of parameters to be estimated θ. Let $X_f(D_f, \beta_f, \theta) = (X_{f,1}, \ldots, X_{f,j})$. The expectation of X_f, X^e_f, is given by:

$$X_f^e(D_f, \theta) = \int_{-\infty}^{\infty} X_f^*(D_f, \beta_f, \theta)\, h(d\beta_f \mid D_f, \theta) \tag{7}$$

where *h* is the density of the random parameters β_f conditional on the information D_f.

Given these predicted staff assignments and the observed number of agents X_f at the different locations for firm *f*, let us define the prediction error $\varepsilon_f(D_f, \theta)$ as:

$$\varepsilon_f(D_f, \theta) = X_f - X_f^e(D_f, \theta) \tag{8}$$

At the true parameter θ_0, the moment of the prediction error is identically zero:

$$E(\varepsilon_f \mid D_f, \theta_0) = 0 \quad for\ f = 1, \ldots, F \tag{9}$$

Any function $g(D_f)$ of the conditioning variables must also be uncorrelated with this error. As a result, the value of θ, say $\hat{\theta}$, that sets the sample analog of this moment

$$G_F(\theta) = \frac{1}{F} \sum_1^F g(D_f) \otimes \varepsilon_f(D_f, \theta) \tag{10}$$

equal to zero or as close as possible to zero is a consistent estimator of θ_0. Under appropriate regularity conditions, asymptotic normality of $\hat{\theta}$ is ensured (see Hansen 1982). If the number of moment conditions is larger than the number of parameters to be estimated (the model is over-identified), an efficient estimator is found by combining the moment conditions through a weighing matrix *V*. The efficient weighing matrix as suggested by Hansen (1982) is

$$V = E((g(D)\varepsilon)(g(D)\varepsilon)') \tag{11}$$

$\hat{\theta}$ is then asymptotically normally distributed with mean θ_0 and asymptotic variance-covariance matrix

$$\Omega = \left(\left(\frac{\partial G(\theta_0)}{\partial \theta} \right)' V^{-1} \frac{\partial G(\theta_0)}{\partial \theta} \right) \tag{12}$$

Unfortunately, the function $X_f^e(D_f, \theta)$ is not known analytically. Unlike classic discrete-choice models in which latent variables are simple linear functions of characteristics and error terms are assumed to have a specific structure (extreme value distribution, for example, in the logit specification), the profit function above is highly nonlinear and the integrals are not easily computable. When analytic expressions are not available, it is possible to obtain simulation-based estimates of the distributions as suggested by McFadden (1989) and by Pakes and Pollard (1989). The straightforward way of simulating the expectation $X_f^e(D_f, \theta)$ is by averaging the underlying random function over a set of random draws. The resulting estimator of $X_f^e(D_f, \theta)$ is trivially an unbiased estimator of the true expectation $X_f^e(D_f, \theta)$. McFadden and Pakes and Pollard prove that the MSM estimator that sets the simulated moment as close as possible to zero is typically consistent for a finite number of simulation draws (the intuition is that the simulation error averages out over observations as $N \to \infty$). To conduct the simulation, it is therefore enough to draw $F \times S \times K$ normals where K is the number of random coefficients per firm. The resulting values represent the random components of *f*'s preferences. I will return to the actual estimation procedure after describing the data in the next section.

The Data

The data consist of a sample of 106 Fortune 1000 firms in the manufacturing sector (six different SIC two-digit codes corresponding to industries such as machinery, computer and electric/electronic equipment, food, chemicals, etc.). These industries (Table 2) were selected for their relatively homogeneous customer-support activities (all are manufacturing firms and offer sales and post-sales support of consumer or industrial goods).

Table 1. Industries

SIC 2-Digit Class	SIC Description	SIC industries	Corporations in Sample
20	MFG-FOOD PROD	BEVERAGE-TOBACCO-FOOD	10
28	MFG-CHEMICALS	CHEMICAL, MEDICAL, MISC	26
35	MFG-MACHINERY	COMPUTER-ELECTRONIC, MACHINERY	35
36	MFG-ELEC EQUIP	COMPUTER-ELECTRONIC, ELECTRICAL	20
38	MFG-INSTRUMENTS	COMPUTER-ELECTRONIC, MACHINERY, MEDICAL, MISC	13
39	MFG-MISC	MISC	2

The sample is constructed using data from Harte-Hanks, a company that collects detailed data about U.S. firms and their computing and communication equipment. For each firm, I have data on firm characteristics that include the total number of employees at the firm, its annual sales, and its main sector of activity (SIC two-digit industry). I also have data at the establishment level (a single firm has several establishments across the United States) that allow me to compute three additional firm-level metrics: the number of sectors (SIC three-digit codes) in which a firm operates, a measure of Internet usage at the firm, and a decision-rights measure. The number of sectors is found by aggregating SIC three-digit codes over the sites of the firm. I calculated the Internet metric by aggregating the number of different types of Internet applications used by the firm (across its sites). This index assigns one point for every type of Internet application used. Internet applications encompass what Harte-Hanks codes as Internet server applications, Internet applications, Internet/Web programming languages, Internet/Web servers, and Internet/Web software. Examples of these applications are e-commerce, technical support, Web development, Java, and Web server. I divided this index by sales to estimate the intensity of Internet usage per billion dollars of revenue. The decision-rights measure is the proportion of sites for which IT (PCs, non-PCs, and telecommunications) purchase decisions are made locally (computer purchase decisions are either made locally or at the parent/headquarters). The vector D_f has the following five components: D_f = {*Emp, Sales, Activities, Internet, Rights*}.

The establishment-level data are also the source for CSRs employment in different regions. Sample statistics for all of these variables are presented in Table 2.

Table 2. Sample Statistics for Firms

Variable	Obs	Mean	Min	Max
E: Emp (thousands)	106	37.96	1.95	316.3
S: Sales (billions)	106	9.70	1.22	88.4
A: Activities	106	39.2	1	266
Internet	106	76.82	3	354
In: Internet/Sales	106	15.77	1.10	93.69
R: Rights	106	0.47	0.13	1

The choice set consists of nineteen areas of the continental United States. The areas correspond to nine subregions of the United States that are further subdivided into two (and in one case, into three) areas based on population. The choice set reflects differences in regional characteristics (for example, New England versus Pacific) and in labor force between states within a subregion. Each area is characterized by four variables: the monthly unit costs of locating a CSR in the area, the labor force in the area, telecom infrastructure, and industrial presence. Monthly unit costs of locating a CSR include monthly wages, obtained from the Bureau of Labor Statistics data on occupations (occupation code 43-4051, customer service representatives), and monthly commercial real estate rents in the area, obtained from the Society of Industrial and Office Realtors. The BLS data is at the metropolitan statistical area level so the measure used is calculated as the weighted average of the data at the metropolitan statistical area (for the wage data) and residential area (for the rent data) with weights equal to the relative number of CSRs in the different metropolitan statistical areas of the region. Telecom infrastructure is obtained from data provided by the Federal Communications Commission regarding telecom penetration in different states (that I aggregated at the region level). The penetration rate compiled by the FCC is a percentage between 89.88 and 97.05 rescaled, for convenience, to a range from 0 to 7.17. Industrial presence is calculated from the data as the number of sites of Fortune 1000 manufacturing firms (the ones in the six SIC codes chosen above) in the area. This measure captures sources of positive externalities that derive from proximity to other firms or to other activities of the same firm. Descriptive statistics can be found in Table 3.

Results

The estimation of the model proceeds based on the method of simulated moments as explained in the estimation section above. I first compute the predicted staff assignments across regions, and then calculate the function *G*. The process is iterated for different values of the parameters, as described below. The full specification of the model requires functional forms assumptions and distributional assumptions for the random coefficients. The random coefficients are assumed to be normally distributed with mean and variance to be estimated. I use regional dummies to avoid potential endogeneity biases of the kind studied by Berry (1994) that arise when prices (unit costs) are correlated with unobserved characteristics (the unobserved quality). By using dummies, one does not need the inversion procedure proposed there. Fixed effects are also useful to capture the other regional factors that may affect profits. Notice that dummies affect the mean utility level of the region, but have no effect on the substitution patterns between regions (which are affected by observed characteristics of the region). Once I use regional dummies, taste coefficients on regional attributes cannot be estimated directly as they vary with the regional dummy. However, they can be retrieved from the estimated dummies using a minimum distance procedure (Chamberlain 1982). This is identical to performing a regression of the estimated dummy coefficients on observed regional attributes. Assuming that the unobserved regional attributes are uncorrelated with the observed attributes, the coefficients on observed characteristics are unbiased and consistent (this regression and the associated coefficients are presented at the end of this section).

I use random coefficients for each of the regional dummies and for the variable *PI* (industrial presence). The functional forms are as follows. For λ_f, I use equation (1). I assume that $m(D_f)$ is a linear function of the firm's characteristics as follows (the intercept is normalized to 1):

Table 3. Sample Statistics for Regions in Choice Set

Area	Range	Labor Force (average)	Monthly Unit Costs (average)	Telecom Infrastructure	Industrial Presence
E-N-Cen(IL,MI,OH)	3	4692	2271.9	3.89	850
E-N-Cent (IN,WI)	2	2093.5	2265.8	4.75	288
E-S-CEN (AL,TN)	2	1864	1980.2	3.80	254
E-S-CEN (KY,MS)	1	657	1993.5	1.38	144
MID-ATL (NY)	4	8485	2268.6	5.18	212
MID-ATL (NJ, PA)	3	4816	2252.5	5.72	362
MOUN (AZ,CO)	2	2148.5	2026.7	5.22	157
MOUN (ID,UT)	1	567	1876.4	3.80	89
NEW-ENG (MA)	3	3271	2482.3	4.72	138
NEW-ENG (CT,NH)	2	1010.5	2493	7.17	115
PACIFIC (CA)	4	16728	2519	5.95	603
PACIFIC (OR,WA)	2	1901.5	2242	4.97	116
SOU-ATL (FL)	4	7500	1983	2.22	203
SOU-ATL (GA,SC,MD,NC,VA)	2	2092.6	2030	3.83	679
SOU-ATL (DE)	1	392	2518	6.38	17
W-N-CEN(MN, MO)	2	2411	2039	6.72	256
W-N-CEN (IA,KS,ND,NE)	1	479.5	2071	6.15	244
W-S-CEN (TX)	4	9195	1845	3.65	489
W-S-CEN (AR,OK)	1	841	1968	0	148

$$m(D_f) = 1 + m_1 In + m_2 R \qquad (13)$$

The estimation was performed in Matlab. I started the estimation with a small number of draws (three) and increased the number of draws to 10 (S = 10) to increase efficiency. I computed the predicted vector of CSRs (the vector has 19 elements, one per region) as explained above in the estimation section. Following McFadden (1989) and Pakes and Pollard (1989), I held the draws constant over different function evaluations (to avoid infinite jumpiness) and used different simulation draws for different observations to make the simulation error average out faster. The instruments that I used were a constant, the number of sectors a firm operates in, and the decision rights at the firm. The profit function in equation (5) is non-differentiable for any finite number of simulation draws. Therefore, I used the Nelder-Meade non-derivate simplex search algorithm to minimize the function (Hendel 1999). To ease the search, I broke the problem into two sub-problems, estimating the dummy coefficients and then the other parameters, before estimating the whole coefficient vector. Estimates of the parameters can be found in Table 4 (parameters that are significant are in bold).

Table 4. Estimates of the Parameters

	Coefficient	Std. Dev.
f_0	**$0.68*10^{-3}$**	$0.21*10^{-3}$
f_1	**$-0.0028*10^{-3}$**	$0.0008*10^{-3}$
f_2	$-0.00031*10^{-3}$	$0.00096*10^{-3}$
m_1	**-0.0092**	0.0043
m_2	**-0.027**	0.012
$Var(B_i)$	**9.21**	1.73
$Var(B_p)$	**3.95**	1.48
A	**0.87**	0.05

The significant f_1 captures the relationship between Internet applications at the firm and the proportion of income associated with customer calls. The sign of the coefficient is negative, pointing to a negative relationship between the average number of calls at the locations we observe and the use of Internet applications by the firm. As mentioned above, this could be the result of fewer callers (customers of firms with extensive Internet presence substitute Web-based service to call agents), or to the outsourcing of customer service to locations that we do not observe, for instance, overseas. The significant f_1 shows that there is a small impact of Internet business applications on these practices, even though the data cannot distinguish between the two effects. The intercept, f_0, is also significantly different from zero. The third coefficient f_3, on activities, is not significant: differentiation cannot be shown to affect the average volume of calls (in income units) based on this dataset.

To evaluate the magnitude of the Internet effect, I use the first-order conditions in Appendix B to find the elasticity of X with respect to the use of Internet applications. To keep things simple, I assume that the value of $m(D_f)$ stays constant (i.e., there is a compensating change in *rights* with the change in *Internet* so that $m(D_f)$ remains constant). Given this assumption, the elasticity of X with respect to the use of Internet applications is $f_1 \; \varphi^{\frac{\alpha-1}{\alpha+1}} (1-\alpha)^{-1} In$. Evaluated at the sample means, this elasticity has a value of –0.247. That is, a 10 percent increase in the index of Internet application use is associated with a 2.5 percent decrease in the national employment of CSRs.

The coefficients m_1 and m_2 are also significant, supporting the hypothesis that Internet usage and decision rights affect location choice patterns. The coefficients are significant and negative, suggesting that firms with higher Internet usage, or more dispersed decision rights, are less sensitive to quality differences between regions and more price sensitive. In other words, Internet-based applications and distributed decision rights reduce the vertical differentiation with respect to quality between firms. In addition, the ratio of CSRs between two regions becomes closer to 1, suggesting that the average number of location at which the firm will locate its CSRs increase. To get a sense of the magnitude of these effects, I calculated the change in average costs from a change in $m(D_f)$. Given a valuation ratio k_i (between a region i and a reference region j), the elasticity of relative regional employment (between i and j) with respect to m_1 is $\varepsilon_i = m_1 * ln\ k_i * In$. The change in total costs TC'/TC is then:

$$\frac{TC'}{TC} = \frac{\sum_i \varepsilon_i k_i P_i}{\sum_i k_i P_i} \tag{14}$$

Evaluated at the regional dummy coefficients (the mean utility of the different regions) and sample mean of *Internet*, the change in total costs is equal to -0.13. Thus, a 10 percent increase in the number of Internet applications (per sales) is associated with savings of 1.3 percent from the unit costs of CSRs. The intuition behind this result is that firms take advantage of lower unit costs by locating their staff in regions that would not have been attractive without IT, presumably because of coordination and informational costs. The same technique gives an estimate of the impact of a change in decision rights allocation between subsidiaries and headquarters (replace *m1* by *m2* and *Internet* by *rights* in ε_i). The resulting value is –0.05, which implies that increasing the number of sites to which purchasing decision rights are delegated (or the number of decisions at a site) by 10 percent results in a 0.5 percent reduction in unit costs. An interesting result concerns the coefficients Var(B_i) and Var(B_p): both are significantly different from zero (9.21 and 3.95, respectively), indicating heterogeneity in tastes among firms between regions and in valuing urbanization externalities. This validates the use of random coefficients for regional dummies and for industrial presence, and shows that idiosyncratic differences among firms have a significant impact on valuation of regions and, ultimately, location decisions.

Finally, I estimate the individual effect of regional characteristics on valuation by regressing the regional dummy coefficients on the observable regional attributes. The regression is

$$\delta = -\underset{(0.025)}{0.09} + \underset{(0.029)}{0.078}\,LR + \underset{(0.007)}{0.006}\,TI + \underset{(0.00018)}{0.00094}\,PI$$

The coefficients on every observable variable are significant except for the coefficient on telecom infrastructure. This might be a result of poor data, the FCC data being an aggregated index that covers residential, rural, and business telecommunication lines and is perhaps not sufficiently correlated with the portfolio and price of telecom services offered to businesses. Also, since local residential markets have not become as competitive as business and long-distance markets, the index might not reflect true telecom costs for businesses. The variable that captures externalities due to proximity is, however, highly significant. This shows that firms value proximity to customers (as proxied by the labor range, LR) and proximity to other firms (PI). The R^2 of this regression is 0.12, suggesting that random valuation accounts for most of the variation in regional preferences.

Conclusion

The adoption of information technologies and communication technologies on a worldwide scale presents both challenges and opportunities for firms. This paper considers the ability of firms to exploit regional cost differentials and save costs by locating their customer-service function in low-cost regions. The demand estimation is based on a novel application of multiple-discrete choice models to firms' location and employment strategies, using micro-data. The results show a statistically significant effect of technology on both customer calls and location patterns but the impact is economically small. For managers, the estimation demonstrates the importance of balancing region-specific preferences in deciding where to locate business functions and suggests that cost is not always the main determinant of location. Furthermore, the results establish that better communications can change the dynamics of location. But the vision of technology enabling firms to relocate activities on the basis of cost alone has yet to materialize. This presents a challenge for researchers who may have been premature in declaring the "death of distance."

Acknowledgements

I am grateful to Erik Brynjolfsson, Igal Hendel, Daniel Klein, Jeffrey Hu, participants at the MIT Applied Econometrics Seminar, and two anonymous referees for helpful comments, and to the National Science Foundation, grant number IIS-0085725, and the CSK Corporation through the Center for E-Business at the Massachusetts Institute of Technology.

References

Amemiya, T. "Multivariate Regression and Simultaneous Equations Models When the Dependent Variables are Truncated Normal," *Econometrica* (42:6), 1974, pp. 999-1012.

Berry, S. "Estimating Discrete-Choice Models of Product Differentiation," *Rand Journal of Economics* (25:2), 1994, pp. 242-262.

Brynjolfsson, E., and Hitt, L. "Computing Productivity: Firm-level Evidence," unpublished manuscript, Massachusetts Institute of Technology, Cambridge, MA, 2003.

Brynjolfsson, E., Hitt, L. and Yang, S. "Intangible Assets: Computers and Organizational Capital," *Brookings Papers on Economic Activity* (I), August 2002, pp. 137-98.

Caves R. E. "International Corporations: The Industrial Economics of Foreign Investment," *Economica* (38:149), 1982, pp. 1-27.

Chamberlain, G. "Multivariate Regression Models for Panel Data," *Journal of Econometrics* (18:1), 1982, pp. 5-46.

Ellison, G., and Glaeser, E. "The Geographic Concentration of Industry: Does Natural Advantage Explain Agglomeration?," *American Economic Review, Papers and Proceedings* (89:2), 1999, pp. 311-316.

Engardio, P., Bernstein, A., and Kripalani, M. "Is Your Job Next?," *Business Week*. February 3, 2003, pp. 50-60.

Gallagher, R. *Discrete Stochastic Processes*, Kluwer Academic Publishers, Boston, 1996.

Hansen, L. "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica* (50:4), 1982, pp. 1029-1054.

Hendel, I. "Estimating Multiple-Discrete Choice Models: An Application to Computerization Returns," *Review of Economic Studies* (66:2), 1999, pp. 423-446.

Hennart, J.-F. *A Theory of Multinational Enterprise*, University of Michigan Press, Ann Arbor, MI, 1982.

Ichniowski, C., and Shaw, K. "Beyond Incentive Pay: Insiders' Estimates of the Value of Complementary Human Resource Management Practices," *Journal of Economic Perspectives* (17:1), 2003, pp. 155-180.

Kim, S. "Regions, Resources, and Economic Geography: Sources of U.S. Regional Comparative Advantage, 1880-1987," *Regional Science and Urban Economics* (29:1), 1999, pp. 1-32.

Lancaster, K. *Variety, Equity and Efficiency: Product Variety in an Industrial Society*, Columbia University Press, New York, 1979.

Macke, G. "Apres l'industrie, le secteur des services commence a delocaliser," *Le Monde* (Paris), April 24, 2003, p. 18.

Maki, D. R., and Meredith, L. N. "Production Cost Differentials and Foreign Direct Investment: A Test of Two Models," *Applied Economics* (18:13), 1986,pp. 1127-34.

Mandelbaum, A., Sakov, A., and Zeltyn, S. "Empirical Analysis of a Call Center," unpublished manuscript, Technion, Haifa, Israel, 2000.

McFadden, D. "A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration," *Econometrica* (57:5), 1989, pp. 995-1026.
McGrath, J., and Hollingshead, A. *Groups Interacting with Technology*, Sage Publications, Thousand Oaks, CA, 1994.
Pakes, A., and Pollard, D. "Simulation and the Asymptotics of Optimization Estimators," *Econometrica* (57:5), 1989, pp. 1027-1057.
Rosenthal, S., and Strange, W. "The Determinants of Agglomeration," *Journal of Urban Economics* (50:2), 2001, pp. 191-229.
Wernerfelt, B. "A Resource-Based View of the Firm," *Strategic Management Journal* (5:2), 1984, pp. 171-80

Appendix A. Derivation of the Profit Function

$\overline{W} = \frac{\lambda E[Z^2]}{2(1-\lambda E[Z])} \cong \frac{\lambda\left(\frac{1}{\sum X_i^\alpha \mu_i}\right)^2}{2\left(-\ln\left(\frac{\lambda}{\sum X_i^\alpha \mu_i}\right)\right)} = 2\left[\frac{\ln\lambda}{\lambda\left(\frac{1}{\sum X_i^\alpha \mu_i}\right)^2} + \frac{\left(\sum X_i^\alpha \mu_i\right)^2 \ln \sum X_i^\alpha \mu_i}{\lambda}\right]^{-1} \approx \left[\frac{2\left(\sum X_i^\alpha \mu_i\right)^2}{\lambda}\right]^{-1}$. Then the survival rate is: $e^{-\frac{\lambda}{2(\sum X_i^\alpha \mu_i)^2}}$ and=

the associated expected revenue is $R = \lambda e^{-\frac{\lambda}{2(\sum X_i^\alpha \mu_i)^2}}$.

Appendix B. Optimal Number of Agents (Without the Integer Constraints)

The FOC of the profit function with respect to X_i are $\lambda e^{-\frac{\lambda}{2(\sum X_i^\alpha \mu_i)^2}}\left(\frac{\lambda}{\left(\sum \mu_i X_i^\alpha\right)^2}\right)\alpha\mu_i X_i^{\alpha-1} - P_i = 0 \quad \forall i$, which

implies that $\frac{X_i}{X_j} = \left(\frac{P_i}{P_j}\frac{\mu_j}{\mu_i}\right)^{\frac{1}{\alpha-1}} = \left[\left(\frac{\mu_i}{P_i}\right)\Big/\left(\frac{\mu_j}{P_j}\right)\right]^{\frac{1}{1-\alpha}}$.