

Software Versioning and Quality Degradation? An Exploratory Study of the Evidence

Anindya Ghose Arun Sundararajan

Leonard N. Stern School of Business, New York University

aghose@stern.nyu.edu, asundara@stern.nyu.edu

Working Paper # CeDER-05-20, Center for Digital Economy Research

July 2005

We present a framework for measuring software quality using pricing and demand data, and empirical estimates that quantify the extent of quality degradation associated with software versioning. Using a 7-month, 108-product panel of software sales from Amazon.com, we document the extent to which quality varies across different software versions, estimating quality degradation that ranges from as little as 8% to as much as 56% below that of the corresponding flagship version. Consistent with prescriptions from the theory of vertical differentiation, we also find that an increase in the total number of versions is associated with an increase in the difference in quality between the highest and lowest quality versions, and a decrease in the quality difference between "neighboring" versions. We compare our estimates with those derived from two sets of subjective measures of quality, based on CNET editorial ratings and Amazon.com user reviews, and discuss competing interpretations of the significant differences that emerge from this comparison. As the first empirical study of software versioning that is based on both subjective and econometrically estimated measures of quality, this paper provides a framework for testing a wide variety of results in IS that are based on related models of vertical differentiation, and its findings have important implications for studies that treat web-based user ratings as cardinal data.

Keywords: software quality, vertical differentiation, price discrimination, quality distortion, information goods, Internet, electronic commerce, economics of IS, econometrics, analytical modeling, salesrank.

⁰We thank Michael Smith, seminar participants at New York University, and participants at the *First Annual Symposium on Statistical Challenges in E-Commerce* at the University of Maryland for their feedback. We also thank Rong Zheng for excellent research assistance. This research was partially supported by a grant from the NET Institute.

1 Introduction

Many manufacturers create product lines by first developing a flagship product with an optimal level of features and functionality, and then creating one or more inferior versions by deliberately reducing the quality of this flagship product. This practice is commonly referred to as *quality degradation*, and many instances of this have been documented in practice. For example, the Intel 486SX processor was initially produced by beginning with a fully functioning 486DX processor, and then disabling the math coprocessor. In 1990, IBM introduced the LaserPrinter E, which printed at 5 pages per minute, half the printing speed of the superior IBM LaserPrinter, but used the same printing engine, though with added firmware that inserted wait states to slow printing. A number of other examples of “damaging” a flagship product to create inferior versions exist across a variety of industries, ranging from disk drives to overnight delivery service (Deneckere and McAfee, 1996).

Quality degradation is ubiquitous in the software industry. There are multiple versions of large number of popular desktop software packages that differ only in their quality or number of features (rather than in their development or release date), and which are sold at different prices. At any point in time, one can find different versions of popular software titles like Adobe Acrobat, TurboTax, Microsoft Money and Norton AntiVirus available. These are examples of software titles for which a firm has developed a flagship version, disabled a subset of the features or modules of this version, and released both the higher quality version and one or more lower quality versions simultaneously. Correspondingly, a large variety of software manufacturers make a limited functionality version of their product available for free, and charge a positive price for a full-featured version (Eudora and Eudora Light being a popular example). This prevalence is not surprising, given that the cost of disabling features or removing software modules is relatively low, as are the variable costs of producing software.

The theory used to study software versioning and quality degradation typically draws from second-degree price discrimination (Mussa and Rosen, 1978, Maskin and Riley, 1984). The key idea drawn from this theory is that of market segmentation using multiple versions that are quality-

differentiated substitutes, and which are created by strategically distorting down the quality of a flagship version. Many recent papers in information systems have used this theory to study the optimal number of versions for a seller of information goods (Jones and Mendelson, 1998, Raghunathan, 2000, Varian, 2000, Bhargava and Choudhary, 2001, 2004, Weber 2002), often concluding that a single version is optimal. Moreover, this underlying model of market segmentation using quality distortion is used in an increasing number of related IS studies that investigate, among other things, licensing in software contracts (Zhang and Seidmann, 2002), optimal software upgrade paths (Bala and Carr, 2004, Sankaranarayanan, 2005), pricing of online services (Bhargava and Sundaresan 2003), efficiency and pricing of interorganizational (IOS) systems (Barua, Kriebel and Mukhopadhyay 1991, Nault 1994), free download policies (Boom, 2005, Tang and Cheng, 2003) and managing digital piracy (Chen, Wu and Anandalingam, 2003, Snir, 2003, Sundararajan, 2004, Chellappa and Shivendu 2004).

To summarize, some notion of quality assessment by consumers, and of strategic degradation of quality levels by sellers is embedded in a variety of IS studies that use the underlying model of vertical differentiation. Since each of these studies makes different managerial and policy prescriptions based on their models, it seems important to determine how to actually measure levels of quality and quality degradation predicted by such models, towards setting up a way of testing their theories empirically, and towards exploring whether quality degradation estimates based on the model of vertical differentiation seem reasonable. Moreover, since there are a growing number of Internet-based resources that report more subjective measures of software quality (editorial reviews and user reviews are the two most common), it is likely that there is some information about actual customer perceptions of quality differences between software versions contained in such ratings.

The central objective of our study is therefore to empirically estimate the measures of software quality and quality degradation predicted by the commonly-used economic theory of vertical differentiation, to assess how they vary across software titles, and to contrast these estimates with those based on subjective Internet-based ratings. We do so by making the following contributions:

(a) We develop a method for directly estimating the extent of quality degradation based on the framework of price discrimination using vertical differentiation, and using publicly available pricing and demand data.

(b) We provide the first systematic estimate of the extent of quality degradation in the software industry, using a 7-month, 108-version panel of demand and pricing data gathered from Amazon.com.

(c) We contrast these economic estimates of quality degradation with two independent subjective assessments of software quality: editorial ratings gathered from CNET, and average consumer ratings gathered from Amazon.com.

Our estimates of quality degradation across software versions indicate that, relative to the assessed quality of the flagship version, the quality levels of inferior versions are degraded from little as 8% to as much as 56%, and that the extent of degradation varies quite widely across software titles, and within sets of titles with two versions and three versions. Moreover, we find that an increase in the total number of versions is associated with an increase in the difference in quality between the highest and lowest quality versions, and a decrease in the quality difference between successive (or neighboring) versions. This is consistent with the predictions of the theory of vertical differentiation.

However, we also find that the economic estimates of quality degradation are significantly *different* (and significantly higher) than those assessed from subjective ratings. Put differently, when data about the actual purchasing behavior of customers is embedded into the economic model, it predicts very different levels of perceived quality differences than those suggested by the subjective ratings that these customers and other experts assign to the different versions of a software title. There are at least two possible interpretations of the differences we observe:

(1) The economic theory systematically models wider variations in software quality than are actually observed in practice. This may have important implications for the managerial and policy prescriptions derived from models that are based on this theory.

(2) The numbers/ratings that subjectively measure quality differences between software versions tend to systematically understate the actual differences, where by actual differences, we mean those based on economic measures of how much quality affects consumer willingness to pay. In other words, these ratings, while having an reasonable ordinal interpretation, are not robust cardinal measures of quality.

Our study makes other research contributions as well. By transforming the parameters of a commonly-used analytical model into those that can be estimated from demand data, in particular, to assess quality distortion and quality ratios across versions directly from demand and price data, we provide a new framework for future empirical studies of software versioning. In our concluding section, we discuss many directions for future research that might use this framework. We also report on a fairly comprehensive new method for converting Amazon sales rank data into demand data, that uses a combination of purchasing experiments and analysis of the ranking time series, and provides the first such calibration for the computer software industry. Moreover, our analytical model is developed in a manner that enables one to estimate customer distribution characteristics from widely available demand data in a straightforward way. This makes future empirical studies of pricing and quality differentiation in other IT industries more easily feasible. Thus, our paper also adds to the new emerging stream of literature which has used e-commerce based panel data to conduct industry specific studies (Ghose, Smith and Telang 2004) and new instances of existing phenomena such as auctioneer/bidder strategies and price formation in online auctions (Bapna, Goes and Gupta 2004, Bapna, Gank and Shmueli 2004).

The preceding discussion has highlighted a fraction of the IS literature that is related to our current paper in its approach to modeling quality degradation. To our knowledge, ours is the first paper that attempts to empirically validate this modeling approach using data in the software industry. Research has assessed quality degradation in other industries include studies of the automobile industry (Kwoka, 1992), the airline industry (Borenstein and Rose, 1994), and the cable television industry (Crawford and Shum, 2005), although none of these papers contrast econometri-

cally estimated quality levels with subjective measures. Additionally, there is an impressive body of literature on software quality (for instance, Kemerer and Porter 1992, Chidambaram and Kemerer 1994, Banker and Slaughter 1998, Slaughter, Harter and Krishnan, 1998, 2000, Krishnan et. al. 2000). Much of this literature is empirical, although our paper differs from theirs on two important dimensions. First, these tend to study quality issues for large-scale specialized software implementations in organizations, rather than measuring quality for mass market shrink-wrapped software like those in our data. Second, they tend to assess and study software quality using supply side measures – intrinsic measures such as reliability and integrity of the source code, and the number of defects per function point – while our approach is focused on different demand-side measures of software quality.

The rest of this paper is organized as follows. Section 2 presents our analytical model that relates pricing and quality degradation to customer characteristics, and describes how to connect the equations of this model to our demand data. Section 3 describes our data set, our method for converting sales rank data into demand data, and some details on how we estimate our model’s key parameters. Section 4 presents the results of our estimation of quality degradation in the software industry, and contrasts them with the subjective measures of software quality. Section 5 concludes and outlines directions for future research.

2 Model

A monopolist sells n versions of a software product. This seller first develops the highest quality (or flagship) version, of quality s_1 , and then degrades the quality of this to create a set of inferior substitutes, of quality s_2, \dots, s_n , where $s_1 > s_2 > \dots > s_n$. The price charged by the seller for version i is denoted p_i .

Customers are heterogenous in their preferences for quality. A customer of type $\theta \in \Theta$ is willing to pay upto $U(s, \theta)$ for a version of quality s , and $U(s, \theta)$ is non-decreasing in both its arguments. The set Θ is discrete, with elements θ_i . Customer types are distributed according to a probability




Figure 1: Illustrates how the utility function varies with increases in quality s , and changes in customer type θ .

measure F over Θ , and for notational convenience, we denote the measure of customers of type θ_i as $f_i \equiv F(\theta_i)$. We make the assumption of discrete types for subsequent ease of estimation (more on this later). Our analysis could equivalently assume that customer types are uniformly distributed over some continuous interval.

Since the versions are substitutes, each customer purchases upto one version. A customer of type θ therefore purchases a version of quality s_i if version i yields the highest positive level of consumer surplus, that is if:

$$U(s_i, \theta) - p_i > U(s_j, \theta) - p_j \quad (1)$$

for each $j \neq i$, and if

$$U(s_i, \theta) - p_i \geq 0. \quad (2)$$

The seller's problem is to choose the optimal number of versions, the quality level for each version, and their associated prices. We assume that the utility function takes the following simple quadratic

form¹:

$$U(s, \theta) = \begin{cases} \theta s - \frac{1}{2}s^2, & s < \theta \\ \frac{1}{2}\theta^2, & s \geq \theta \end{cases}. \quad (3)$$

Equation (3) indicates that customers value increasing quality at a diminishing rate, and the highest quality level that a customer of type θ is interested in is $s = \theta$. Therefore, if $\Theta = \{\theta_1, \theta_2, \dots, \theta_n\}$, the socially optimal outcome involves the seller offering n versions, with quality levels $s_1 = \theta_1, s_2 = \theta_2$, and so on. We analyze the cases of two types, that is, $\Theta = \{\theta_1, \theta_2\}$ and three types, that is, $\Theta = \{\theta_1, \theta_2, \theta_3\}$, since our data set contains software titles with either two or three versions.

2.1 Two versions

We start by assuming that $\Theta = \{\theta_1, \theta_2\}$. The seller offers the flagship version of quality s_1 , and may offer a second version of quality $s_2 < s_1$. It is well known that the seller only need consider direct mechanisms, will design one quality-price pair for each type, such that these pairs are incentive-compatible (IC) and individually rational (IR). These conditions yield the following price equations:

$$p_2 = U(s_2, \theta_2), \quad (4)$$

$$p_1 = U(s_1, \theta_1) - U(s_2, \theta_1) + p_2. \quad (5)$$

The firm therefore chooses s_1 and s_2 to maximize

$$f_1[U(s_1, \theta_1) - U(s_2, \theta_1)] + U(s_2, \theta_2), \quad (6)$$

and the solution to this maximization yields the optimal quality levels:

$$s_1 = \theta_1, \quad (7)$$

$$s_2 = \max \left\{ \theta_2 - \frac{f_1}{f_2}(\theta_1 - \theta_2), 0 \right\}. \quad (8)$$

¹A more commonly used functional form is even simpler: $U(s, \theta) = \theta s$, although that predicts that for a good with constant variable costs, a single version is optimal, which, based on our data, is clearly inconsistent with the software industry.

Equations (7-8) indicate that the seller will offer two versions so long as the difference between the quality preferences of the two types is not too large, and there is a sufficient fraction f_2 of lower type customers. Notice that the flagship version is assigned the socially optimal quality level, while the quality of the lower version is distorted downwards. In this case, the corresponding prices as functions of the model's primitives are:

$$p_1 = \frac{f_1}{f_2}(\theta_1 - \theta_2)^2 + \frac{1}{2} [\theta_2^2 + (\theta_1 - \theta_2)^2], \quad (9)$$

$$p_2 = \frac{1}{2} \left[\theta_2^2 - \frac{f_1^2}{f_2^2} (\theta_1 - \theta_2)^2 \right]. \quad (10)$$

We will refer back to these expressions shortly to discuss how we use them to estimate quality.

2.2 Three versions

Next, we assume that $\Theta = \{\theta_1, \theta_2, \theta_3\}$. The seller will therefore offer the flagship version of quality s_1 , along with upto two other versions with quality $s_2 < s_1$ and $s_3 < s_2$. The corresponding price equations from the IC and IR conditions are:

$$p_3 = U(s_3, \theta_3), \quad (11)$$

$$p_2 = U(s_2, \theta_2) - U(s_3, \theta_2) + p_3, \quad (12)$$

$$p_1 = U(s_1, \theta_1) - U(s_2, \theta_1) + p_2. \quad (13)$$

The firm chooses s_1, s_2 and s_3 to maximize

$$f_1[U(s_1, \theta_1) - U(s_2, \theta_1)] + (f_1 + f_2)[U(s_2, \theta_2) - U(s_3, \theta_2)] + U(s_3, \theta_3), \quad (14)$$

yielding the following optimal quality levels:

$$s_1 = \theta_1, \quad (15)$$

$$s_2 = \max \left\{ \theta_2 - \frac{f_1}{f_2} (\theta_1 - \theta_2), 0 \right\}. \quad (16)$$

$$s_3 = \max \left\{ \theta_3 - \frac{f_1 + f_2}{f_3} (\theta_2 - \theta_3), 0 \right\}. \quad (17)$$

Assuming that all these quality levels are in fact non-zero, the corresponding expressions for prices as a function of the model's basic parameters are:

$$p_1 = \frac{f_1}{f_2}(\theta_1 - \theta_2)^2 + \frac{f_1 + f_2}{f_3}(\theta_2 - \theta_3)^2 + \frac{1}{2}[\theta_3^2 + (\theta_2 - \theta_3)^2 + (\theta_1 - \theta_2)^2], \quad (18)$$

$$p_2 = \frac{f_1 + f_2}{f_3}(\theta_2 - \theta_3)^2 + \frac{1}{2}[\theta_3^2 + (\theta_2 - \theta_3)^2 - \frac{f_1^2}{f_2^2}(\theta_1 - \theta_2)^2], \quad (19)$$

$$p_3 = \frac{1}{2} \left[\theta_3^2 - \left(\frac{f_1 + f_2}{f_3} \right)^2 \frac{(\theta_1 - \theta_2)^2}{2} \right]. \quad (20)$$

2.3 Linking this theory to prices and demand data

Rather than the numerical values of quality implied by the model, we are interested in the extent of quality degradation for different products: that is, in the ratios s_1/s_j , for each $j > 1$. Given a set of price data for each of the versions of a software product, one needs estimates of each of the θ_i and f_i parameters in order to use equations (7-8) or (15-17) to compute these quality ratios. Further, in the underlying model, notice that the demand for version i is simply f_i .

These observations lead to a natural way of linking the theoretical model to a data set of prices and demand. First, for each product, we observe the fraction of demand realized by each of its versions, in each of our time periods, and use this as an observation about the corresponding f_i . We use these observations to assess a maximum likelihood estimate of each f_i for each product (further details follow a description of our data). Given these estimates, we can use observed average prices and the system of equations (18-20) to estimate θ_1, θ_2 and θ_3 (or correspondingly, equations (9-10) in the two-version model to estimate θ_1 and θ_2). Notice that, given prices and the f_i/f_j ratios, (18-20) is a system of 3 equations in three unknowns. These estimates can then be used in equations (15-17) or in equations (7-8) to estimate the quality degradation associated with versions of the software title.

3 Data and Estimation

3.1 Overview of our data

We estimate our models using a panel data set compiled from publicly available information about software prices and sales rankings, gathered using automated Java scripts to access and parse HTML and XML pages downloaded from Amazon.com. The panel includes over 280 products, with an equal number from each of five major categories – Business and Productivity, Graphics and Development, Security and Utilities, Children’s Software and Operating Systems. These are major categories listed by Amazon.com, and resemble a parallel categorization by CNET.com, although we do not use this categorization in any substantive way.

Of our 280 products, we identify 108 as belonging to a family of different versions of the same product. In this context, it is important to distinguish between *versions* and successive *generations*. For instance, Adobe Standard 7.0 and Adobe Professional 7.0 are 2 different versions of Adobe Acrobat 7.0. Similarly, TurboTax Premier 2004, Deluxe 2004 and Standard 2004 are 3 different versions of TurboTax 2004. On the other hand, Adobe Illustrator 10.0 and Adobe Illustrator CS are successive *generations* of Illustrator, and while substitutes, were developed at different points in time and released over two years apart. As a consequence, this pair of products is not consistent with our underlying economic model of versioning, in which a seller develops a flagship version and then strategically degrades its quality to create inferior versions.

We separate our products into two sets. The first consists of all software titles which have 3 different versions. The ordering of these versions is naturally inferred from their titles (a typical labeling would be Premier, Deluxe and Standard, for instance, in decreasing order of quality). The second category consists of all products which have 2 versions (often labelled Professional and Standard, in decreasing order of quality). We end up with 27 software titles with two versions, and 18 software titles with three versions (for a total of $27 \times 2 + 18 \times 3 = 108$ versions).

We collected data every 8 hours, over a 7-month period (from November 2004 to May 2005).

Variable	Mean	Standard Deviation	Minimum	Maximum
Sales Rank	1649.61	1971.26	1	11622
List Price	99.2	226.2	19.99	1799.99
Amazon Price	95.53	208.57	10.95	1699.99
Customer Rating	3.14	0.99	1	5
Number of Reviewers	56.2	81.6	7	606

Table 1: Summary statistics of our data

Each observation contains the product’s list price, its Amazon retail price, its Amazon sales rank (which serves as a proxy for units of demand, as described further later), the date the product was released into the market, the average customer rating for the product, the number of reviews based on which the average rating was computed, and some secondary market data. The summary statistics of our data are in Table 1.

For benchmarking purposes, we have also collected similar data from Buy.com: sales ranks, list prices, retail prices and so on. Similar to Amazon.com, Buy.com provides sale rankings of all of its products publicly and these sales ranks are also based on actual quantities sold at their site. The Buy.com data exhibits qualitatively similar characteristics as the Amazon.com data, and since we do not use this data further in our analysis, it is not described.

3.2 Inferring software demand from sales ranks

A few recent papers have used the following Pareto relationship to infer product quantities from Amazon.com sales ranks

$$q = \delta(\text{SalesRank})^\beta \tag{21}$$

Chevalier and Goolsbee (2003) estimate the parameters of this equation for books by associating demand data with sales rank on the Wall Street Journal best-seller list, and by independently conducting a purchasing experiment on one book whose actual weekly demand was known to them,

observing the extent to which its sales rank reacted to their purchases. They estimate the value of $(1/\beta)$ to be -1.2 . Brynjolfsson, Hu and Smith (2003) provide an alternative estimate of the parameters of equation (21) for books, using data from a book publisher that maps observed sales rank to the number of copies the publisher sold to Amazon, and estimate $\beta = -0.871$ (this is the parameter β_2 in their model), $\log[\delta] = 10.526$ (this is the parameter β_1 in their model).

To our knowledge, there are no corresponding estimates available for software, and industry specific demand patterns preclude using estimates from book demand for the software industry. Moreover, in summer 2004, Amazon altered its sales rank system in the following way: they eliminated their three-tier system, updating ranks each hour for most products (rather than merely for the top products), and they moved to a system that uses exponential decays to give more weight in the sales rank to newer purchases.

We therefore conducted an independent analysis to convert our measured sales ranks into demand data. We retain the assumption of a Pareto relationship (21) between demand and sales rank. Over a two-week period in mid-June 2005, we collected hourly sales rank data for each of the 280 products in our panel, yielding a time series of 336 observations for each product. For products not ranked too high, a general trend in these time series is an extended downward drift in the rank value over many hours (that is, the rank becomes a progressively larger number), followed by intermittent spikes which result in a large upward shift in rank (that is, the rank became a smaller number suddenly). We interpret these spikes as reflecting time periods in which one or more purchases have occurred.

A day before this data collection began, we had purchased 1, 2 and 3 copies of different software titles over the course of an hour. We chose titles that had different initial sales ranks, and also those for which Amazon reported on the number of items they had in stock, towards confirming that ours were the only purchases of these titles during that hour. We then tracked the magnitude of the spikes in sales rank that corresponded to the purchase of one unit, two units and three units, for different starting sales rank values. This was towards being able to assign hour-over-hour units

of demand for corresponding changes in our sales rank time series.

This experiment led us to conclude that for products that were highly ranked (that is, which were ranked 150 or better), there would be a lot of noise in associating changes in sales ranks with units of demand. There were two reasons for this conclusion. First, the magnitude of change in sales rank induced by a purchase was relatively small for these highly-ranked products (less than 50). Second, and related, it was likely that the extent to which a product's sales rank changed was affected not just by how many units of it were purchased in the last hour, but by how many units of products ranked close to it were purchased – these are sales *ranks*, after all, and even with the same unit sales in an hour, a product's rank can move up more if lots of nearby products don't sell a unit during that hour. This is true of all products, clearly, but for products ranked below 150, the magnitude of the change in sales rank for a unit purchase is high enough to discount these smaller movements.

We therefore chose to focus on the 186 products whose average sales rank over the two-week period was between 500 and 5000 (a small number of products with ranks over 5000 had virtually no demand over the two weeks, and were therefore dropped). We verified that among these products, and across 336 observations for each products, there were only 5 instances of the sales rank dropping below 150, and it never dropped below 100. We tracked the changes in sales rank from period to period for each of these 186 products, and associated units of demand² with upward shifts of more than 50 (that is, changes for which $rank_{t-1} - rank_t > 50$).

This procedure yielded a data set of 186 observations, which associated a weekly demand level

¹A low numerical value of sales rank indicates higher demand. An increase in the numerical value of sales rank indicates a decrease in the rate at which a product is selling. We use the term "highly ranked" to refer to products that have higher demand, that is, those products whose sales ranks are low integer values. Therefore, a product with sales rank 50 is more "highly ranked" than a product with sales rank 5000.

²We also checked changes in rank magnitude between 25 and 50 by hand, towards verifying whether these might correspond to single-unit sales. Of the hundreds of such changes, only four appeared to be possible sales (in a sample over which there were thousands of unit sales in the two week period). This seemed like an acceptable level of noise, and we therefore ignored all salesrank movements below 50.

Variable	Estimated Value
$\log[\delta]$	8.352*** (0.042)
β	-0.828*** (0.032)
R^2	0.779
*** significant with $p \leq 0.001$	

Table 2: Mapping average sales rank to unit demand for software

with each average sales rank, for two successive weeks. Weekly unit sales ranged from 0 to 16. Using the implied pairs of average weekly demand and average sales rank, we then estimated the following OLS equation:

$$\log[q + 1] = \log[\delta] + \beta \log[rank], \quad (22)$$

where q is average weekly demand, and $rank$ is the corresponding average salesrank³. The results of this estimation are summarized in Table 2. To provide a sense for what this estimate implies: weekly sales of two units correspond to an average salesrank of about 3100, weekly sales of 10 units correspond to an average salesrank of about 440, and weekly sales of 25 units correspond to an average salesrank of about 150.

3.3 Estimating the customer type distribution

The preceding experiment enables us to associate our sales ranks with corresponding periodic unit demand levels. Now, consider a software title with n versions, and demand data over T periods. In any period t , let the demand for version i be q_{it} , and define the total demand for this title during period t as $q_t = \sum_i q_{it}$. One can model this demand as the result of q_t draws from the distribution

³Similar to Brynjolfsson, Hu and Smith (2003), we used White's heteroskedasticity-consistent estimator (see Greene, 2000, page 463) to estimate both parameters.

F over Θ , with the outcome reflecting q_{it} draws of type θ_i , for each version i of each software title.

Define Q_i as the random variable that takes the value 1 if a draw from F yields θ_i , and takes the value 0 otherwise. It follows that Q_i is a Bernoulli random variable with parameter f_i . Therefore, each realization of a unit of demand for any version of the product is an observation about the true value of f_i , and it is well known that with n such observations, the maximum likelihood estimator of f_i is simply the fraction of true realizations. As a consequence, once we have computed the periodic demand levels for each version of each title, the maximum likelihood estimate of f_i for a specific software title is simply:

$$f_i = \frac{\sum_{t=1}^T q_{it}}{\sum_{t=1}^T \sum_{k=1}^n q_{kt}}, \quad (23)$$

or the estimated demand for the version as a fraction of the total demand for all versions of the title. As described in section 2.3, once we have estimates of f_i for each version i of each product, we are able to compute the implied corresponding values of θ_i , and the corresponding quality degradation levels.

It is worth noting that our equations always involve a ratio of two f_i values (rather than an f_i value in isolation). Therefore, if one chooses the appropriate periodic demand rate associated with an average sales rank, these ratios can be computed directly from average sales ranks, since, based on equation (23):

$$\frac{f_i}{f_j} = \frac{\sum_{t=1}^T q_{it}}{\sum_{t=1}^T q_{jt}},$$

which, based on equation (21), and an appropriate normalization for the length of the time interval which cancels out in the numerator and denominator, simplifies to

$$\frac{f_i}{f_j} = \left(\frac{rank_i}{rank_j} \right)^\beta,$$

where $rank_i$ is the average sales rank of product i .

4 Evidence

Before we present the results of our estimated quality degradation, and contrast them with the subjective measures we have collected, it seems important to distinguish between quality degradation and quality distortion, since the latter term is used quite extensively in the price discrimination literature. Our measure of *quality degradation* for any version is simply the ratio of the estimated quality of the highest version to the estimated quality of the version in question. For instance, the extent of quality degradation for the second-highest quality version of a product with three versions is (s_1/s_2) , which based on (15-16), is:

$$\frac{s_1}{s_2} = \frac{\theta_1}{\theta_2 - \frac{f_1}{f_2}(\theta_1 - \theta_2)}. \quad (24)$$

A higher value of this ratio implies a higher difference in quality, and therefore, more significant quality degradation.

On the other hand, *quality distortion* refers to the extent to which the quality of an inferior version i has been distorted below the socially optimal level θ_i . We also report on our estimated percentage quality distortion levels, which are simply $(\theta_i - s_i)/\theta_i$, though we do not discuss them much. For instance, the percentage of quality distortion for the second-highest quality version of a product with three versions, based on (15-16), is:

$$1 - \frac{s_2}{\theta_2} = \frac{f_1}{f_2} \left(\frac{\theta_1}{\theta_2} - 1 \right). \quad (25)$$

4.1 Estimated quality degradation

Tables 3 and 4 summarize our estimates of quality degradation and quality distortion, for software titles with two versions and three versions respectively. For titles with two versions, we find that the quality ratios vary from as low as 1.09 to as high as 1.75. This reflects a downward degradation in the quality of the flagship version from as little as 8% to as much as 43%, with a mean degradation of about 27%. For titles with three versions, we find that the quality ratios for the medium quality version (that is, the ratios s_1/s_2) range from 1.08 to 1.46, thereby reflecting degradation of the

Product	Quality degradation ratio (s_1/s_2)	Quality distortion % $1 - (s_2/\theta_2)$
Adobe Acrobat 6.0 (WIN)	1.3	12%
Adobe Acrobat 7.0 (WIN)	1.31	13%
Adobe Acrobat 7.0 (Mac)	1.3	12%
Adobe Creative Suite 2 (WIN)	1.21	10%
Adobe Creative Suite 2 (Mac)	1.21	11%
Adobe Video	1.27	11%
Adobe Photoshop	1.4	16%
Macromedia Flash 2004	1.3	13%
Quicken 2001	1.39	18%
Quicken 2002	1.46	19%
Quicken 2003	1.44	17%
Quicken 2004	1.31	12%
Norton SystemWorks 2002	1.38	16%
Norton SystemWorks 2003	1.17	8%
Norton SystemWorks 2004	1.27	11%
Norton SystemWorks 2005	1.32	14%
Norton AntiVirus 2003	1.71	20%
Norton AntiVirus 2004	1.49	16%
Norton Internet Security 2004	1.75	21%
Symantec AntiVirus SBS	1.28	11%
McAfee VirusScan 7.0	1.31	13%
McAfee VirusScan 8.0	1.27	12%
McAfee Internet Security 2005	1.69	21%
Microsoft Encarta 2002	1.34	15%
Microsoft Streets 2005	1.65	24%
Microsoft Windows SBS 2003	1.09	5%
Microsoft Office 2003	1.15	7%

Table 3: Estimated quality degradation and downward quality distortion for software titles with two versions

Product	Quality degradation ratio		Quality distortion %	
	s_1/s_2	s_1/s_3	$1 - (s_2/\theta_2)$	$1 - (s_3/\theta_3)$
Encyclopaedia Britannica 2005	1.19	1.69	9%	28%
MS Encarta 2003	1.21	1.83	9%	31%
MS Encarta 2004	1.11	2.04	6%	39%
MS Encarta 2005	1.21	1.76	8%	33%
Quickbooks 2002	1.37	1.69	14%	22%
Quickbooks 2003	1.30	1.58	12%	22%
Quickbooks 2004	1.21	1.83	12%	21%
Quickbooks 2005	1.33	1.63	11%	21%
Quicken 2005	1.32	1.88	11%	29%
Turbo Tax 2003	1.34	1.63	14%	23%
Turbo Tax 2004	1.36	1.70	14%	24%
Microsoft Money 2004	1.08	1.68	3%	27%
Microsoft Money 2005	1.16	1.73	8%	29%
Zone Alarm	1.35	1.96	10%	34%
PC-Cillin Internet Security 2005	1.23	2.31	5%	39%
Art Explosion	1.28	1.81	12%	32%
Corel WordPerfect 12	1.46	1.86	20%	28%
MS Windows XP	1.3	1.69	15%	26%

Table 4: Estimated quality degradation and downward quality distortion for software titles with three versions

quality of the flagship version from as low as 7% to as high as 31% (with a mean of about 21%). The corresponding quality ratios for the low quality version (the ratios s_1/s_3) are between 1.63 to 2.31, which correspond to quality degradation ranging from 39% to as much as 57% (and mean of about 44%).

These estimates indicate that when a product line has three versions, the extent of quality degradation between the best version and the second-best version is significantly lower (both on average, and in its variance) than the extent to which the quality of a second-best version is degraded when a third version does not exist. However, the extent to which the quality of the lowest version is degraded when the product line has three versions is significantly more than when the product has just two versions. We verify these statements by testing the difference in mean between both pairs of data (finding significant t-statistics in each case). These results are interesting because they are consistent with what the theory of vertical differentiation would predict. All else being equal, an increase in the number of versions offered will increase the extent of quality degradation of the lowest quality version, but will also reduce the differences in quality between neighboring versions.

Exactly the same statements can be made about the percentage of quality distortion (highest for lowest of three versions, lowest for second of three versions, significant differences in means). We do not explore any welfare issues using these estimates, though this represents an interesting direction for future work.

4.2 Contrasting economic and subjective measures of quality degradation

We next report on our estimates of quality degradation based on two subjective measures of assessed quality, from CNET and from Amazon.com. A set of editors at CNET evaluate most software products according to a standard set of review criteria, and rank these products on a scale of 1-10. According to CNET, they judge a product on the quality and appropriateness of its feature set along with service and support provided by the firm. They also evaluate the number and severity of

any bugs as well as the overall ease of setup, configuration, and use. We use these summary scores from CNET as our first subjective measure of quality, and assess quality degradation by computing the ratio of scores for different versions of a title. Prior studies have used such rankings as an objective measure of software quality (for instance, Liebowitz and Margolis 1999). We collected these scores from CNET’s web site on a periodic basis. Since CNET also archives ratings for older products, we have been able to gather these ratings for most products in our dataset.

Our second source of subjective quality assessments is from reviews for each product provided by Amazon.com’s customers. These are part of Amazon’s customer review feature, and each review contains a written report as well as a numerical score on a scale of 1 to 5. We use the average numerical score associated with a product as our second subjective measure of quality. We collected longitudinal data on these ratings, along with the total number of reviewers based on which the average rating is published. We chose to drop those product ratings which were based on reviews by 5 or fewer customers. The average number of reviews for the remaining products is 56, and the number of reviews ranges from 7 to 605 (though most products have 20-50 reviews).

We find that the extent of quality degradation assessed from our economic estimates is significantly higher than those assessed from the subjective measures of quality. The mean quality degradation is significantly *higher* for comparisons of s_1/s_2 for products with two versions, and with three versions, and for comparisons of s_1/s_3 , for the subjective measures based both on CNET editorial ratings and on Amazon customer reviews. As illustrated in Table 6, these differences are somewhat higher for CNET than for Amazon. Furthermore, the differences were most stark when comparing the extent of quality degradation of the lowest quality version for products with three versions.

There are many ways in which one might interpret these findings. One interpretation might be that in models of vertical differentiation, the extent to which quality varies across versions in the model are far wider than are actually observed in practice. In other words, the magnitude of the optimal quality difference prescribed by the model’s quality parameters s_i and s_j may be higher

Product	CNET Ratings	Amazon Ratings
	s_1/s_2	s_1/s_2
Adobe Acrobat 6.0 (WIN)	1.04	1.17
Adobe Acrobat 7.0 (WIN)	1.08	1.11
Adobe Acrobat 7.0 (Mac)	–	–
Adobe Photoshop	1.14	1.23
Adobe Creative Suite 2(WIN)	1	–
Adobe Creative Suite 2 (Mac)	1	–
Adobe Video	1.125	1
Macromedia Flash 2004	1.125	0.8
Quicken 2001	1	0.71
Quicken 2002	1.08	1
Quicken 2003	1.14	1
Quicken 2004	1.14	1
Norton SW 2002	1.04	1.17
Norton SW 2003	1.14	1.17
Norton SW 2004	1.26	1
Norton SW 2005	–	0.69
Norton AntiVirus 2003	1	1
Norton AntiVirus 2004	1.0	1
Norton IS 2004	1.05	1.0
Norton AntiVirus SBS	–	–
McAfee VirusScan 7.0	1.04	–
McAfee VirusScan 8.0	1.04	1
McAfee IS 2005	1.04	–
MS Encarta 2002	1.125	1.5
MS Streets 2005	1.15	1.17
MS Windows SBS 2003	–	–
MS Office 2003	1	–

Table 5: Estimated quality degradation based on two subjective quality measures for software titles with two versions. Blank cells indicate that either no review was available for one or more of the versions of that product, or, in the case of the Amazon.com numbers, there were too few reviews (5 or fewer) to include the rating in our data set.

Product	CNET Ratings		Amazon Ratings	
	s_1/s_2	s_1/s_3	s_1/s_2	s_1/s_3
Encyclopaedia Britannica 2005	1	1.14	–	–
MS Encarta 2003	1	1.33	1	1.33
MS Encarta 2004	1	1.04	1	1.14
MS Encarta 2005	1	1.04	–	–
Quickbooks 2002	1.125	1.125	1.2	1.2
Quickbooks 2003	1.04	1.04	1.25	1.25
Quickbooks 2004	1.04	1.04	1.2	1.2
Quickbooks 2005	1	1.14	–	–
Quicken 2005	1.1	1.14	1.07	1.5
Turbo Tax 2003	1	1.08	1.5	1.2
Turbo Tax 2004	1	1.14	1.1	1.65
MS Money 2004	1.1	1.1	1.2	1.2
MS Money 2005	1.1	1.1	1.02	1.33
Zone Alarm	1.04	1.06	–	–
PC-Cillin Internet Security 2005	1	1.04	–	–
Art Explosion	–	–	1.14	–
MS Windows XP	1.00	1.125	1	1.13
Corel WordPerfect 12	–	1.04	–	–

Table 6: Estimated quality degradation based on two subjective quality measures for software titles with three versions. Blank cells indicate that either no review was available for one or more of the versions of that product, or, in the case of the Amazon.com numbers, there were too few reviews (5 or fewer) to include the rating in our data set.

	s_1/s_2 for products with two versions		s_1/s_2 for products with three versions		s_1/s_3 for products with three versions	
	Model	CNET	Model	CNET	Model	CNET
<i>Mean</i>	1.38	1.07	1.25	1.03	1.79	1.10
<i>Variance</i>	0.03	0.01	0.01	0.002	0.03	0.01
<i>Observations</i>	23	23	16	16	17	17
<i>t-statistic</i>	7.93		8.59		14.54	
<i>p</i>	5E-09		9E-09		1E-12	

	s_1/s_2 for products with two versions		s_1/s_2 for products with three versions		s_1/s_3 for products with three versions	
	Model	Amazon	Model	Amazon	Model	Amazon
<i>Mean</i>	1.37	1.03	1.25	1.14	1.75	1.29
<i>Variance</i>	0.02	0.03	0.01	0.02	0.02	0.02
<i>Observations</i>	19	19	11	11	12	12
<i>t-statistic</i>	6.25		2.01		8.13	
<i>p</i>	2E-07		0.03		3E-08	

Table 7: Summary of difference in means tests between the quality degradation levels we estimated from our data, and those from CNET’s editorial ratings, and Amazon’s average customer ratings respectively. The numbers under the label "Model" refer to those from our estimates described in Section 4.1.

than the actual quality difference that is required to obtain the appropriate optimal magnitude in value difference; the latter difference is what influences the willingness to pay of customers and the firm’s eventual success with price discrimination based on versioning. This would suggest that prescriptions from models of versioning or price discrimination that are based on the magnitude of the quality difference across versions should be interpreted carefully.

Another interpretation might be that the numbers that subjectively measure quality differences between software versions tend to systematically understate the actual differences, where by actual differences, we mean those based on economic measures of how much quality affects consumer

willingness to pay. These subjective ratings might therefore be a good way of ranking different versions, but their numerical magnitudes may not be appropriate cardinal measures of quality. This interpretation has important implications for future research, because editorial ratings have been used as measures of software quality in prior studies, and aggregate customer feedback measures from eBay, Amazon.com, and various other review sites are frequently used in IS research as cardinal measures of some form of quality, in studies of seller reputation, movie quality, used-good quality and so on.

A third interpretation might simply be that editors and customers have a different benchmark when assessing the quality of different versions, and that these benchmarks (or reference points) are affected by what the customer/editor expects from a specific version. For example, a rating of 5 on a Professional version might require a higher level of overall quality than a rating of 5 for a Standard version. This would cause a systematic overstatement of the quality of lower versions as measured by these average ratings or reviews, which in turn would lead to lower assessed quality degradation levels.

A preliminary analysis towards a better understanding of the relationship between these objective and subjective measures did not yield results that were significant enough to report⁴. Determining which of these interpretations might be the most valid remains an open question, though we believe that more data is required to answer this well.

⁴We estimated equations of the form

$$\Delta s_{eco} = a + b(\Delta s_{subj}) + \sum_i c_i version_i + \sum_i d_i(\Delta s_{subj})version_i,$$

where Δs_{eco} and Δs_{subj} were the economic and subjective quality ratios respectively, and $version_i$ was a dummy variable that captured whether the ratio was between versions 1 and 2 of a 2-version product, versions 1 and 2 of a 3-version product, or versions 1 and 3 of a 3-version product. We estimated these with and without interacting the dummy variables with Δs_{subj} . None of these generated significant coefficients for the CNET data. For the Amazon data, a couple of equations had significant coefficients, but since there was sufficient variation in coefficients across equations, we do not have a good interpretation of their relative magnitudes.

5 Conclusions and directions for future research

This paper has presented the first empirical study of versioning in the software industry. The contributions of this study are summarized below:

(1) In order to assess the success of a chosen versioning strategy relative to others, it is useful for firms to derive an economic measure of the relative quality of each version which has been created based on quality degradation. This represents a considerable challenge in the software industry, because while subjective assessments of software value from independent experts and from its end users are available, there are no natural objective measures of product size or quality (counting the number of features is not really sensible, for instance). Therefore, objective assessments of software quality based on economic demand-side measures of a product's quality can be of managerial value. We develop a framework for directly estimating the extent of software quality degradation based on a widely used model of price discrimination using vertical differentiation, and that can be estimated using pricing and demand data that is publicly available.

(2) We provide the first systematic estimate of the extent of quality degradation associated with versioning in the software industry. We do so by compiling and using a 7-month, 108-product panel of demand and pricing data for software sold on Amazon.com. Our results indicate that there is significant quality degradation associated with software versioning, and significant variations in its extent across software titles. Our estimates are consistent with theoretical predictions that an increase in the number of versions is associated with an increase in the quality difference between the highest and lowest quality version, but a reduction in the quality differences between neighboring versions.

(3) We provide new estimates of quality degradation between versions using two independent sources of subjective quality assessments: editorial ratings gathered from CNET, and average user ratings gathered from Amazon.com. We contrast these estimates of quality degradation with those from our economic model. We show that the estimates of quality degradation from the latter are significantly and consistently higher than those assessed from subjective measures of software

quality, and discuss different interpretations of this measured difference.

(4) We extend existing methods for imputing demand from Amazon.com’s sales rank information, and provide the first calibration of this relationship for the software industry.

Apart from providing a first step towards testing other existing IS theories that are based on models of vertical differentiation, our work suggests a number of new directions for future research, and provides an infrastructure that can be used to explore these directions. A natural question that arises from our study is whether software versioning is in fact an optimal strategy for sellers, and if so, measuring the extent to which it increases profits. It is likely that the benefits from versioning are related to both the category of software, and the extent to which the flagship version has been degraded to create each inferior version. Examining this relationship could be of particular interest to IS practitioners making pricing and product management choices.

We have also provided the first estimates of the extent of quality distortion for software (relative to the socially optimal quality level of a version). This is a first step towards assessing the magnitude of the welfare losses that ensue on account of this distortion. However, there are likely to be welfare gains from the prevalence of versioning, due to the expansion of the set of customers who can afford a version of the product. Comparing the relative gains and losses from quality distortion, given that the absence of this kind of distortion would lead to higher prices represents another promising line of research. A related study might examine whether there is a relationship between measured quality distortion and subjective measures of quality, based on the hypothesis that subjective ratings assess product quality relative to a benchmark for that kind of version, rather than relative to the flagship version, and therefore might measure distortion rather than degradation. An analysis of the text associated with editorial reviews might be instructive in this regard.

A preliminary exploration of whether there are variations in the differences between subjective and economic measures of quality degradation across each of our product categories did not yield significant results, though this may be a consequence of the fact that there are insufficient titles in each category for any systematic differences to show up. For instance, it may be relatively

straightforward for consumers and experts to assess the quality of finance and accounting software, based on their features and ease of installation. However, the quality of security software is much harder to assess, since it is contingent on future performance at detecting and suppressing viruses, minimizing the probability of a breach, or detecting an intrusion. Studying this in more detail, using a larger data set, or perhaps longitudinal data, seems like another interesting direction for future work.

Finally, while our study has been of the software industry, many of our techniques can generalize to other IT industries. Future empirical researchers might use our method to map sales ranks to demand for other categories of products, which would facilitate new industry-specific quality degradation studies that answer related strategic and welfare questions in other IT product industries. We hope that our study is a first step in this direction.

6 References

1. Bala, R. and S. Carr, 2004. Pricing and Market Segmentation with Software Upgrades. Available at http://www.anderson.ucla.edu/documents/areas/fac/dotm/bio/pdf_SC09.pdf
2. Banker, R. D., G. B. Davis, and S. A. Slaughter, 1998. Software Development Practices, Software Complexity, and Software Maintenance Performance: A Field Study. *Management Science* 44, 433-451.
3. Bapna, R., P. Goes and A. Gupta. 2003. Replicating Online Yankee Auctions to Analyze Auctioneers' and Bidders' Strategies. *Information Systems Research*, 14:(3) 244-268.
4. Bapna, R. W. Gank and G. Shmueli. 2004. Price Formation and its Dynamics in Online Auctions. Working Paper, University of Maryland.
5. Barua, A., C. Kriebel and T. Mukhopadhyay. 1991. An Economic Analysis of Strategic Information Technology Investments, *MIS Quarterly* 15 313-331.
6. Bhargava, H. and V. Choudhary, 2001. Information Goods and Vertical Differentiation. *Journal of Management Information Systems* 18, 89-106.
7. Bhargava, H. and V. Choudhary, 2004. One Size Fits All? Optimality Conditions for Versioning Information Goods. Available at <http://www.smeal.psu.edu/faculty/hob2/Pub/Papers/EC/mkt-segment-2002.pdf>

8. Bhargava, H. and S. Sundaresan 2003. Contingency Pricing for Information Goods and Services under Industry-wide Performance Standards. *Journal of Management Information Systems* 20 (2), 113-136.
9. Boom, A., 2004. "Download for Free" - When Do Providers of Digital Goods Offer Free Samples? Discussion Paper 2004/28, Free University of Berlin.
10. Borenstein, S. and N. Rose, 1994. Competition and Price Dispersion in the US Airline Industry. *Journal of Political Economy* 102, 653-683.
11. Brynjolfsson, E., Y. Hu, and M. Smith, 2003. Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety. *Management Science* 49, 1580-1596.
12. Chevalier, J., and A. Goolsbee, 2003. Measuring Prices and Price Competition Online: Amazon and Barnes and Noble. *Quantitative Marketing and Economics* 1, 203-222.
13. Chellappa, R. and S. Shivendu. 2005. Managing Piracy: Pricing and Sampling Strategies for Digital Experience Goods in Vertically Segmented Markets. *Information Systems Research* (forthcoming).
14. Crawford, G. and M. Shum, 2005. Monopoly Quality Degradation and Regulation in the Cable Television. Working Paper, University of Arizona.
15. Chidamber, S. and C. Kemerer, 1994. A Metric Suite for Object-Oriented Design. *IEEE Transactions on Software Engineering* 20, 476-492.
16. Denekere, R. and P. McAfee, 1996. Damaged Goods. *Journal of Economics and Management Strategy* 5, 149-174.
17. Ghose, A., M. Smith and R. Telang. 2004. Internet Exchanges for Used Books: An Empirical Analysis of Welfare Implications. Working Paper, New York University.
18. Greene, W. H. 2000. *Econometric Analysis*. Prentice-Hall, Upper Saddle River, NJ.
19. Jones, R. and H. Mendelson, 1998. Product and Price Competition for Information Goods. Working Paper, University of Rochester.
20. Kemerer, C. and B. Porter, 1992. Improving the Reliability of Function Point Measurement: An Empirical Study. *IEEE Transactions on Software Engineering* 18, 1011-1024.
21. Krishnan, M. S., C. H. Kriebel, S. Kekre, and T. Mukhopadhyay, 2000. An Empirical Analysis of Productivity and Quality in Software Products. *Management Science* 46, 745-760.
22. Kwoka, J., 1992. Market Segmentation by Price-Quality Schedules: Some Evidence from Automobiles. *Journal of Business* 65, 615-628.
23. Liebowitz, S. and J. Margolis, 1999. *Winners, Losers and Microsoft: Competition and Antitrust in High Technology* (2nd ed.), Independent Institute, Oakland, CA.

24. Mussa, M. and S. Rosen, 1978. Monopoly and Product Quality, *Journal of Economic Theory* 18, 301–317.
25. Nault, B. 1997. Quality Differentiation and Adoption Costs: The Case for Interorganizational Information System Pricing. *Annals of Operations Research* 71, 115-142.
26. Oestreicher-Singer, G. and . Sundararajan, 2004. Are Digital Rights Valuable? Theory and Evidence from the eBook Industry. *Proceedings of the 25th International Conference on Information Systems*, R. Agrawal, L. Kirsch, J. I. DeGross (Eds.), Washington, DC, 617-629.
27. Raghunathan, S., 2000. Software Editions: An Application of Segmentation Theory to the Packaged Software Market. *Journal of Management Information Systems* 17 (1), 87-114.
28. Sankaranarayanan, R., 2005. Excessive Software Upgrades: Reasons and Remedies. Available at <http://ssrn.com/abstract=764667>
29. Slaughter, S., D. Harter and M.S. Krishnan, 1998. Evaluating the Cost of Software Quality. *Communications of the ACM* 41 (8), 67-73.
30. Slaughter, S., D. Harter and M.S. Krishnan, 2000. Effects of Process Maturity on Quality, Cost and Cycle Time in Software Product Development. *Management Science* 46, 451-466.
31. Snir, E., 2003. The Record Industry in an Era of File Sharing: Lessons from Vertical Differentiation. *Proceedings of the 24th International Conference on Information Systems*, S. T. March, A. Massey, and J. I. DeGross (Eds.), Seattle, WA, 72-84.
32. Sundararajan, A., 2004. Managing Digital Piracy: Pricing and Protection. *Information Systems Research* 15, 287-304.
33. Tang, C. and H. Cheng, 2003. Free Trial or No Free Trial: Optimal Software Product Design with Network Externalities. Working Paper, University of Florida.
34. Varian, H., 2000. Buying, Selling and Renting Information Goods. *Journal of Industrial Economics* XLVIII, 473-488.
35. Weber, T., 2002. Mixed Differentiation of Information Goods under Incomplete Information. *Proceedings of the 23rd International Conference on Information Systems*, L. Applegate, R. D. Galliers, and J. I. DeGross (Eds.), Barcelona, 81-94.
36. Wu, S., P. Chen and G. Anandalingam, 2003. Fighting Information Good Piracy with Versioning. *Proceedings of the 24th International Conference on Information Systems*, S. T. March, A. Massey, and J. I. DeGross (Eds.), Seattle, WA, 617-629.
37. Zhang, J. and A. Seidmann, 2002. The Optimal Software Licensing Policy Under Quality Uncertainty. *Proceedings of the 23rd International Conference on Information Systems*, L. Applegate, R. D. Galliers, and J. I. DeGross (Eds.), Barcelona, 225-236.