The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners
Author(s): Manuel Trajtenberg
Source: *The Journal of Political Economy*, Vol. 97, No. 2 (Apr., 1989), pp. 444-479
Published by: The University of Chicago Press
Stable URL: http://www.jstor.org/stable/1831321
Accessed: 30/01/2009 05:49

The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners

Manuel Trajtenberg

Tel-Aviv University

The main goal of this paper is to put forward a methodology for the measurement of product innovations using a value metric, that is, equating the "magnitude" of innovations with the welfare gains they generate. This research design is applied to the case of computed tomography scanners, a revolutionary innovation in medical technology. The econometric procedure centers on the estimation of a discrete choice model (the nested multinomial logit), which yields the parameters of a utility function defined over the changing quality dimensions of the innovative product. The estimated flow of social gains from innovation is used to compute a social rate of return to R & D, to explore the interrelation between innovation and diffusion, and to trace the time profile of benefits and costs, the latter suggesting the possible occurrence of "technological cycles."

I. Introduction

One of the most distinctive features of western economies is their relentless ability to expand the universe of available goods and to improve the qualities of existing ones, that is, to engage successfully in product innovation. There is no doubt that this phenomenon has contributed enormously to our welfare; moreover, we have grown so accustomed to it that our present well-being depends in no small way

This article is based on my Ph.D. dissertation (Trajtenberg 1983). I acknowledge with deep gratitude Zvi Griliches's continuous support and guidance and the advice of Richard Caves, Daniel McFadden, and Michael Spence. I am also grateful to the National Bureau of Economic Research for its good auspices and institutional support.

[*Journal of Political Economy*, 1989, vol. 97, no. 2]

on the expectation that technical change will keep delivering an ever-improving menu of goods. Likewise, competition both in domestic and in international markets seems to be increasingly a matter of quality rather than of price: the quest for "better" appears to prevail over the traditional "more."

Yet there seems to be a serious imbalance between the perceived importance of product innovation as an economic activity and as a dimension of economic performance, and its place in economics. In particular, it is clear that the measures of overall performance commonly used (e.g., real income or productivity growth) can hardly capture the contribution of product innovations and may therefore be quite misleading. Clearly, the difficulties in getting a handle on product innovations stem mainly from the fact that they come about in dimensions other than prices and quantities, and when it comes to empirical studies, they require unconventional data, often far removed from the economist's turf.

In response to this challenge, the main goal of this paper is to conceptualize the notion of product innovations and provide it with operational content so as to enable their empirical measurement. The emphasis is on assessing the *value* of those innovations, that is, the social benefits stemming from improvements in the qualities of available products and from the appearance of new ones. The proposed methodology is then applied to the detailed study of computed tomography (CT) scanners, one of the most remarkable innovations in medical technologies of recent times.

As suggested above, the literature on product innovations is very scanty, with only a handful of studies attempting to tackle them empirically, most notably Mansfield et al. (1977) and Bresnahan (1986). The former analyzes the welfare impact of 17 product innovations but does so by treating them in fact as process innovations; that is, it computes the benefits from the implied price reductions in performing given services. Bresnahan's study, assessing the gains from computerization in financial services, comes closer in aim to the work here. Still, there are important differences in method: whereas Bresnahan bases his analysis on the observed changes in hedonic price indices, here there is an attempt to assess directly the welfare impact of changes in the qualities of products (at the cost of much heavier data requirements).

The paper is organized as follows. Section II outlines the theoretical foundations and elaborates on the measurement of product innovations in the context of a discrete choice model (the multinomial logit). The quality dimensions of CT scanners and the data are discussed in Section III, and in Section IV, the actual econometric analysis is performed, centered around the estimation of a nested multinomial logit model. The benefits from innovation are computed in Section V, which tackles on the way the findings of positive price coefficients and changing preferences. Section VI integrates the processes of innovation and diffusion in order to compute the total gains and then uses those figures in order to compute a social rate of return. The most telling aspect of the empirical findings emerges in the analysis of the time profile of those gains, which appear to follow a wavelike pattern; that is, they are very large—and rising—at first but then decline dramatically, dwindling to a trickle soon after. This pattern is interpreted in terms of the workings of nonlinear returns operating simultaneously in three dimensions, namely, R & D technology, utility, and market size. Section VII suggests some further applications.

II. The Measurement of Product Innovations

In view of the fact that the "output" of innovative activity does not present itself in countable units of any sort,[1] it is clear that a quantifiable dimension for innovations can be defined directly only in *value* terms, that is, in terms of their impact on social welfare. In other words, the question "how much innovation took place" in a certain field over a certain period of time can be interpreted only as asking "how much additional consumer and producer surplus was generated by technical advance in that field and time." Other measures such as patent counts, number of "important innovations," or rates of change of attributes could play at best the role of proxies, and their accuracy as such can be judged only by relating them to the value measures themselves (as in Trajtenberg [1987]).

The methodology to be used here in measuring product innovations draws from the "characteristics approach" to demand theory (e.g., Lancaster 1979), hedonic price functions (Griliches 1971; Rosen 1974), discrete choice models (McFadden 1981), and the welfare analysis that goes with the latter (Small and Rosen 1981). The basic idea is fairly simple: consider a technologically dynamic product class as it evolves over time and assume that the different brands in it can be described well in terms of a small number of attributes (i.e., quality dimensions) and price. Product innovation can then be thought of in terms of changes over time in the set of available products, in the sense both that new brands appear and that there are improvements in the qualities of existing products.

[1] Strictly speaking, such output consists of bits of new and socially useful knowledge: clearly, this is an all but impossible notion to quantify. Patents will not do either, given the enormous variance in their "value."

Now if the number of brands in the sets is finite and at least some of their attributes are "noncombinable" (hence ruling out a continuous budget set in characteristics space), then the appropriate framework for modeling the demand side would be discrete choice. Applying such models to data on the distribution of sales per brand and on their attributes and prices, one can estimate the parameters of the demand functions and, under some restrictions, of the underlying utility function. The magnitude of innovation occurring between two periods can then be calculated as the benefits of having the latest choice set rather than the previous one, in terms of the ensuing increments in consumer surplus. Clearly, the procedure involves some sort of weighting of the changes in characteristics by their marginal utilities; that is, the "amount" of innovation thus measured is meant to capture both the relative valuations of attributes by consumers and the physical changes in the products themselves.

To fix ideas, define $\mathbf{s}_i \equiv (\mathbf{z}_i, p_i)$, where p_i denotes the price and $\mathbf{z}_i = (z_{i1}, z_{i2}, \ldots, z_{im})$, the vector of relevant characteristics of product i in a given product class.[2] Thus the choice set from which the consumer selects the most preferred brand in period t is $S_t = (s_{1t}, s_{2t}, \ldots, s_{nt})$. In this setting product innovation is taken to mean simply that changes occur over time in the vector $\mathbf{z}_i$ and in n_t and hence that the choice set changes from S_{t-1} to S_t. Given a "social surplus" function $W(S)$[3] and under the assumption that the changes in S are discrete, the magnitude of innovations occurring from $t - 1$ to t will be measured by

$$\Delta W = W(S_t) - W(S_{t-1}). \quad (1)$$

Thus, for example, in the simplest possible case whereby S consists of just one product having a single quality dimension z, the gains from an innovation taking the form $\Delta z = z_t - z_{t-1} > 0$ would be (see, e.g., Willig 1978)

$$\Delta W = \int_{p}^{\infty} [x(v, z_t) - x(v, z_{t-1})]dv, \quad (2)$$

where $x(p, z)$ is the demand function, and income effects are assumed away. Simply put, ΔW measures the consumer's willingness to pay for

[2] By "product class" I mean a well-defined set of close substitutes located in a common space of attributes and constituting a separable utility branch. The practical problems of establishing the bounds of related product classes will be tackled here in the context of the nested multinomial logit model (see Sec. IV*B*).

[3] The function $W(S)$ is meant to comprise both consumer and producer surplus. However, since profit is a well-defined magnitude whose measurement does not pose special conceptual problems in the present context, ΔW will be associated with gains from innovations in terms of consumer surplus only.

the innovation; that is, it is the additional area under the demand function, brought about by its upward shift in response to the quality upgrade.

In general, the function $W(S_t)$ is to be obtained by integrating the corresponding demand system, whose features will depend in turn on the nature of the choice set. As suggested above, in the present context it seems appropriate to characterize those sets as discrete: R & D constitutes a fixed cost, and hence innovative sectors typically exhibit in equilibrium a finite, and usually not too large, number of differentiated products. I shall assume discreteness also in the sense that consumers purchase a single unit of a single product out of the set, thus making the choice problem exclusively qualitative (the analysis can be easily extended to accommodate cases of discrete/continuous choice as well; see Hanemann [1984]).

Before I turn to the formal derivation of the $W(\cdot)$ function, two a priori limitations of the proposed ΔW_t measure should be mentioned. First, ΔW_t does not allow one to assess directly the initial impact of radical innovations, since in order to compute that there has to be a nonempty set S and a function $W(S)$ to begin with, and neither exists (by definition) in the case of radical innovations. However, it is possible to obtain *indirect* estimates of those initial gains, as will be shown in Section VIA. The second limitation is that ΔW does not include "spillover" effects; that is, it leaves out the value of improvements in other product classes that may have resulted from the original innovations considered. This is not a shortcoming inherent to the approach itself but rather a result of the empirical difficulties in actually tracking down those externalities. The measure ΔW should therefore be regarded as a lower bound since the "true" social gains would include, conceptually at least, the unobserved spillovers.

A. The ΔW Measure in the Context of Discrete Choice Models[4]

The basic hypothesis underlying discrete choice models is that consumers maximize a random utility function, $U_i = U(\mathbf{z}_i, m; \mathbf{h}) + \epsilon_i$, subject to $s_i \in S$, and $p_i + m = y$, where m denotes a composite "outside" good, $\mathbf{h}$ a vector of observable attributes of the individual, and ϵ_i an independently and identically distributed (i.i.d.) random disturbance, encompassing unobserved attributes of the individual and the product. If we assume that ϵ_i conforms to the type I extreme-value (or Weibull) distribution, the maximization of U_i leads to prob-

[4] I rely here primarily on McFadden (1981) and Small and Rosen (1981).

abilistic demand functions of the form

$$\pi_i = \frac{\exp(V_i)}{\sum_{j=1}^{n} \exp(V_j)}, \quad i = 1, \ldots, n, \tag{3}$$

where V_i is the deterministic component of the conditional indirect utility function, and π_i are fractional demands (thus $\Sigma\, \pi_i = 1$). This is the well-known conditional multinomial logit model, which will be the cornerstone of the estimation procedures here[5] (eventually I shall use the nested multinomial logit model in order to bypass the limitations stemming from the assumption of "independence of irrelevant alternatives").

It is easy to prove that the n equations in (3) constitute a well-behaved demand system, and hence the notion of consumer surplus applies to it as well and can be computed by integration. To make the problem more tractable, income effects are assumed away; that is, the utility function is specialized to be additive separable in the group products (those in S) and in the outside good m, rendering $V_i = \alpha(y - p_i) + \phi(\mathbf{z}_i, \mathbf{h})$, where α stands for the (constant) marginal utility of income. Substituting in (3), we get

$$\pi_i = \frac{e^{\alpha(y-p_i)+\phi(\mathbf{z}_i;\mathbf{h})}}{\sum_{j=1}^{n} e^{\alpha(y-p_j)+\phi(\mathbf{z}_j;\mathbf{h})}} = \frac{e^{-\alpha p_i+\phi(\mathbf{z}_i;\mathbf{h})}}{\sum_{j=1}^{n} e^{-\alpha p_j+\phi(\mathbf{z}_j;\mathbf{h})}}, \quad i = 1, \ldots, n. \tag{4}$$

The identity of Hicksian and Marshallian demand functions in (4) allows one to obtain the surplus function $W(S, \mathbf{h})$ simply by integrating under these demand functions, the integral being path independent. If we ignore the constant of integration (y in this case), the result is[6]

$$W(S, \mathbf{h}) = \frac{\ln\left\{\sum_{i=1}^{n} \exp[-\alpha p_i + \phi(\mathbf{z}_i; \mathbf{h})]\right\}}{\alpha}. \tag{5}$$

This surplus function will be the key element in the upcoming empirical analysis (see Trajtenberg [1983, chap. 2] for a detailed examination of its properties).

[5] The probit, based on the assumption that ϵ is normally distributed, had to be discarded from the outset since the models estimated here included up to 20 choices, whereas at the time the probit could handle in practice only up to four or five alternatives.

[6] Note that when we divide by α, the function $W(\cdot)$ is being normalized so that it is now expressed in money terms. Notice also that $-\partial W/\partial p_i = \pi_i$, and hence eq. (5) is indeed the correct solution.

B. *Incorporating the Hedonic Price Function into the Multinomial Logit Model*

The discussion above overlooked an important feature of markets for differentiated products, namely, the fact that prices and attributes usually exhibit a systematic relationship, embedded in the hedonic price function:

$$p_i = p(\mathbf{z}_i) + \tilde{p}_i, \tag{6}$$

where $p(\mathbf{z}_i)$ is the systematic component, and $\tilde{p}_i$ is an i.i.d. error term (the "residual price"). The existence of such a relationship poses a serious multicollinearity problem in the estimation of the choice probabilities defined in (4): since both price and the vector $\mathbf{z}_i$ appear there as explanatory variables, their individual coefficients cannot be estimated with any precision. The solution suggested here involves incorporating the hedonic function into the consumers' indirect utility function (as a sort of budget constraint) and providing the latter with a more specific structure.

If we substitute (6) for p_i in V_i and ignore y and $\mathbf{h}$,

$$V_i = -\alpha[p(\mathbf{z}_i) + \tilde{p}_i] + \phi(\mathbf{z}_i) = \phi(\mathbf{z}_i) - \alpha p(\mathbf{z}_i) - \alpha\tilde{p}_i$$

or, if we define $V^n(\mathbf{z}_i) \equiv \phi(\mathbf{z}_i) - \alpha p(\mathbf{z}_i)$,

$$V_i = V^n(\mathbf{z}_i) - \alpha\tilde{p}_i, \tag{7}$$

where the term $V^n(\mathbf{z}_i)$ can be interpreted as the "net utility" conferred by product i (i.e., net of the *expected* price of the product). Thus the behavior of consumers is now seen to depend on $\mathbf{z}_i$ and $\tilde{p}_i$ rather than on $\mathbf{z}_i$ and p_i. In other words, given the existence of a hedonic function, p_i largely replicates the information already conveyed by $\mathbf{z}_i$; therefore, only the component of price that is orthogonal to $\mathbf{z}_i$, $\tilde{p}_i$, can affect behavior, qualifying as a legitimate explanatory variable in the choice model. Notice that $\tilde{p}_i$ is very likely to comprise omitted (presumably unobservable) quality dimensions, in which case the interpretation of its coefficient α as the marginal utility of income is called into question (more on this in Sec. V*A*).

In order for (7) to offer an actual solution to the multicollinearity problem, a suitable specification for $V^n(\mathbf{z})$ needs to be found. The following straightforward proposition furnishes the required structure: $V^n(\mathbf{z})$ can be closely approximated by the sum of a linear and a quadratic form, provided only that it has an interior maximum. More formally, $V^n(\mathbf{z}) \cong \mathbf{z}'\beta + \mathbf{z}'\mathbf{G}\mathbf{z}$, where $\mathbf{G}$ is a symmetric matrix, if there is a $\mathbf{z}^* > 0$ such that $\mathbf{z}^* = \text{argmax } V^n(\mathbf{z})$.

When this is so, the approximation $(\mathbf{z}'\beta + \mathbf{z}'\mathbf{G}\mathbf{z})$ obtains readily from a second-order Taylor expansion about $\mathbf{z}^*$. Normally we would expect $\phi(\mathbf{z})$ to be concave (or quasi-concave) and the hedonic function to be convex (as has been the case in many empirical studies), in which case $V''(\mathbf{z})$ would necessarily meet the required condition.[7]

The suggested specification of the net utility leads to the following model:

$$\pi_i = \frac{\exp V_i}{\Sigma \exp V_j},$$
$$V_i = \mathbf{z}_i'\beta + \mathbf{z}_i'\mathbf{G}\mathbf{z}_i - \alpha\tilde{p}_i, \tag{8}$$
$$\tilde{p}_i = p_i - p(\mathbf{z}_i),$$

implying a two-stage procedure: first, estimate the hedonic price function and compute the residuals $\tilde{p}_i$; second, enter $\tilde{p}_i$ as an independent variable in π_i and estimate the multinomial logit model. Clearly, this is not the most general formulation possible, and in some circumstances, a simultaneous equation framework might be more appropriate. What allows me to ignore here simultaneity issues is the working assumption that each choice set S_t, and hence each hedonic price function, is determined prior to the beginning of period t and does not change in the course of the period. This corresponds quite closely to the actual functioning of the market for CT scanners.

III. The Case Study: Computed Tomography Scanners

As mentioned in the Introduction, the methodology outlined above will be applied to the case of CT scanners (also known as computed axial tomography [CAT] scanners). Following is a brief description of the innovation, its main characteristics, the data gathered, and selected quantitative indicators.

The last two decades have seen dramatic advances in diagnostic medicine and, in particular, in imaging technologies, that is, those designed to provide visual information of the interior of the human body. Computed tomography scanners, widely hailed as one of the most remarkable medical innovations of recent times, came to epitomize this revolution in diagnostic technologies[8] and set the stage for subsequent innovations, such as nuclear magnetic resonance and positron emission tomography.

[7] The function $p(\mathbf{z})$ may be concave and the stated conditions could still hold: all that is required is, loosely speaking, that $p(\mathbf{z})$ be "more concave" than $\phi(\mathbf{z})$. If that is not the case, then there will be a corner solution, and hence the proposed approximation will probably be less precise; likewise, the estimated coefficients may turn out in such a case to have the wrong signs.

[8] To wit, the 1979 Nobel Prize in medicine was awarded to the two scientists that pioneered the system.

Research on computed tomography began in 1967 at the British electronic company EMI. The first operational prototype was built in 1971, and the first commercial system was installed in the United States in June 1973. Two types of scanners were developed: head systems, capable of scanning only the brain (the original scanners were of this type), and the more sophisticated body scanners, capable of scanning most body organs. Aware of the revolutionary nature and the vast potential of the innovation, nearly 20 firms rushed to enter the new field during 1974–77, among them major electronic conglomerates (General Electric, Siemens, Hitachi, and Philips), pharmaceutical companies (Pfizer, Searle, and Syntex), and small, specialized firms. They engaged in fierce product competition, each trying to capture a share of the growing market by offering ever-improving performance, thus bringing about a staggering pace of technological advance.

In contrast to the traditionally cautious attitude of the medical profession to the adoption of innovations, the diffusion of CT scanners in the United States proceeded at a very fast pace. This happened at a time of mounting concerns regarding the costs of health care, prompting the government to enact a series of regulatory controls. The full impact of those measures was felt in 1978–79, bringing about a sharp downturn in the market for CT scanners and the exit of many firms. This was, however, a short-lived occurrence: after a major reshuffling of the industry and the easing of regulatory controls, the market rebounded and its structure stabilized, with General Electric emerging as the well-established leader (holding about 50 percent of the U.S. market) and some 10 other firms scrambling for the other half.

A. *Quality Dimensions of CT Scanners*

Even though CT scanners are highly complex systems, it is relatively easy to identify their most important performance characteristics (i.e., the vector $\mathbf{z}$): scan time, image quality, and, as a far third, reconstruction time. A brief description of each follows.

Scan time refers to the speed at which a CT scanner can "take a picture" of a thin cross-sectional slice of the organ examined.[9] Since

[9] A scan is done by rotating around the patient a mechanism that has on one side an X-ray source and an array of detectors opposite to it. The X-ray tube emits a narrow beam that goes through the organ examined, and the detectors read the outcoming (attenuated) energy and send the digitized information to a computer. The procedure is repeated many times as the mechanism rotates, thus taking a very large number of readings. Finally, the data are processed by the computer, and the reconstructed image is displayed on a screen.

internal organs are subject to involuntary motions, the faster the CT scanner is, the less will be the distortion in the picture caused by those motions. The first scanners were very slow and hence could be used only for studies of the brain; as scan time dropped to a few seconds, body scanners became capable of rendering good images of almost any section of the body. Scan time will be denoted by SPEED and taken to be the minimum scan time technically possible in a CT scanner, in seconds.

The output of a CT procedure consists of a series of computer-generated pictures displayed on a television-like screen. Thus the acid test for the performance of a CT scanner is the quality of the received images, that is, the accuracy and richness in detail of those pictures. An important dimension of image quality is "spatial resolution," that is, the ability of the system to record detail. More precisely, it refers to the size of the smallest object that can be just visualized in the "best of conditions." Although this is only a partial indicator, measures of spatial resolution will be used here as a proxy for image quality.[10] It will be denoted by RESOL and measured in millimeters.

Reconstruction time refers to the time interval between the end of the scan and the display of the image, that is, how long it takes the computer to process the data generated during the scan and reconstruct the final picture. Clearly, the shorter the interval, the more efficient the scanner will be. Reconstruction time is denoted RTIME and is defined to be the minimum of the available range, in seconds. For future reference, notice that the three characteristics considered are defined so that less is better.

B. The Data: A First Look

I have gathered for this study a comprehensive data bank on CT scanners (the sources are detailed in the Appendix), covering the 9-year period since the announcement of the innovation in 1972 to the end of 1981 (some of the data extend further, up to 1985). The main data sets thus compiled are (i) the characteristics and prices of all CT scanners developed and marketed up to 1982 (a total of 55 different systems), by year; (ii) detailed sales data, that is, information on virtually all purchases of CT scanners in the United States, including the identity of each buyer (hospital or private clinic), the date on which the scanner was ordered and installed, and the precise scanner model

[10] I designed an econometric procedure to obtain more comprehensive measures of image quality and applied it successfully to experimental data on six scanners (see Trajtenberg 1984). However, similar data for all the other scanners marketed in the United States could not be obtained, and hence I had to content myself with measures of spatial resolution.

bought (over 2,000 observations); (iii) the attributes of hospitals, for example, size, affiliation, and expenses; and (iv) information about the manufacturers of CT scanners, including annual R & D expenditures on CT and patents granted.

As a first look at these data, table 1 presents various indicators of the technological and market evolution of CT scanners over time (all figures refer to the U.S. market only, but the trends displayed are fairly representative of the world market as well). Note that, ever since their introduction in 1975, body scanners came to dominate the scene, both in sales and in terms of technological advance. Notice also that the trends followed by the two types of scanners have been diametrically opposed: whereas head scanners became simpler and cheaper over time, body scanners exhibited a tremendous pace of technological advance and a corresponding steep rise in prices. Since 1981 the market segment occupied by head scanners has shrunk to negligible proportions, and it is very likely to remain there.

IV. The Econometric Analysis

The cornerstone of the econometric analysis consists, as already noted, of the estimation of the multinomial logit model of choice of CT scanners. The model will be estimated separately for each year in the period 1975–81,[11] using the two-stage procedure outlined in Section II*B*. In order to reduce computational complexities to a manageable level, the vector **h** of individual attributes will be ignored (except in the estimation of the multinomial logit model for 1975).[12] Thus the estimates will refer to the "representative" user of CT scanners, including both hospitals and private clinics.

An important issue in formulating the multinomial logit model is how to structure the choice set, in view of the existence of two types of scanners, head and body. The question is whether the two types constitute a single choice set or two separate subsets and, if the latter proves to be the case, whether or not the two systems actually belong at all to a common decision tree. These questions hinge on the pattern of substitution between scanners and, hence, on the compatibility of the various choice structures with the assumption of independence of irrelevant alternatives that underlies the multinomial logit model.

[11] In 1973 and 1974 there was effectively just one scanner in the market, and hence the model cannot be estimated for those years. For 1981 (the last year), the data cover only the first six months; thus (and unless stated otherwise) all figures for 1981 refer to the first half of the year only.

[12] The extra computational burden is not so much in the estimation of the multinomial logit model but in the computation of ΔW: if **h** were included, $W(S)$ would have to be computed separately for each buyer each period and then integrated over **h**.

TABLE 1

SELECTED INDICATORS OF CT SCANNERS, 1973–82

YEAR	PRICE* ($ Thousands)		SPEED† (Seconds)		RESOLUTION‡ (Millimeters)		UNIT SALES	BODY (Percentage of Sales)	NUMBER OF FIRMS	R & D§
	Head	Body	Head	Body	Head	Body				
1973	310	...	300	...	3.1	...	16	0	1	1.0
1974	370	...	300	...	1.7	...	74	0	1	7.8
1975	379	365	285	195.0	1.8	1.6	221	.45	4	28.6
1976	374	471	105	63.0	1.7	1.5	374	.76	9	58.2
1977	354	573	95	19.0	1.7	1.3	385	.85	12	46.6
1978	167	620	96	7.1	1.6	1.2	248	.72	10	37.0
1979	154	667	150	6.6	1.5	1.1	273	.70	9	33.7
1980	154	739	115	5.5	1.5	1.0	270	.80	8	29.6
1981	150	827	115	4.9	1.5	.8	392	.92	8	22.6
1982	150	850	115	2.6	1.5	.7	428	.94	8	n.a.

* Weighted average of all scanners in the market (annual sales as weights).
† Minimum scan time, simple average.
‡ Spatial resolution, simple average.
§ R & D expenditures on CT by U.S. firms, constant 1982 dollars (millions).

The search for the appropriate specification will proceed as follows: First, I estimate the multinomial logit model for the entire set of head and body scanners and for a restricted set of body scanners only, and I conduct an appropriate specification test. As the null hypothesis that independence of irrelevant alternatives holds for the larger set is rejected, a nested structure with two branches (one for head and one for body scanners) is specified next and estimated sequentially. The results indicate that the elasticities of substitution between the two types of scanners are nil and therefore that head and body scanners constitute in fact unrelated sets from the viewpoint of the choice process; this, in turn, determines the way in which the gains from innovation are to be computed.

A. *Estimating the Hedonic Functions and Specifying* V_i

The first stage of the econometric procedure consists of the estimation of hedonic price functions and the subsequent computation of residual prices. As in most hedonic price studies (see Griliches 1971), there are not in the present case any strong priors regarding the functional form of the hedonic equation, except perhaps for some plausible arguments favoring convexity. Thus the matter is to be decided by comparing the fit of alternative specifications and subsidiary considerations of a pragmatic nature. Three functional forms were considered (the double log, the semilog, and the linear log) and estimated both for the joint set of head and body scanners and for body scanners only, for every year in the period 1976–81. In comparisons of their fit using the Box-Cox transformation, the linear log emerged as the clear winner: it ranked first (in terms of minimum corrected mean square error) in one-half of the cases considered and second in one-third of them. The linear log also happens to be a very convenient specification here since its residuals are defined in the same units as price and can therefore be incorporated directly in the multinomial logit model as an independent variable.

As to the estimates of the hedonic equations per se (they are omitted here for lack of space), it is worth noting that the R^2 in all of them hovers about .85 and that the estimated coefficients are, with few exceptions, statistically significant. Thus the inclusion of both price and attributes in the multinomial logit model would have resulted indeed in a serious multicollinearity problem, and therefore the use of residual price instead is amply justified.

Now to the specification of the indirect utility function, V_i. As argued in Section II*B*, the net utility $V^n(\mathbf{z}) = \phi(\mathbf{z}) - \alpha p(\mathbf{z})$ can be closely approximated by $(\mathbf{z}'\beta + \mathbf{z}'\mathbf{G}\mathbf{z})$. However, if we further assume that the utility branch is of the form $\phi(\mathbf{z}_i) = \Sigma_{j=1}^{m} f^j(\mathbf{z}_{ij})$, where $f^j(\cdot)$ is a quasi-concave function of the jth attribute, and substitute the linear log hedonic price equation for $p(\mathbf{z})$ in $V''(\mathbf{z})$, then the net utility function simplifies to the simple quadratic form

$$V''(\mathbf{z}_i) = \sum_{j=1}^{m} [f^j(\mathbf{z}_{ij}) - \beta_j \log(\mathbf{z}_{ij})] \cong \sum_{j=1}^{m} (a_j \mathbf{z}_{ij} - b_j \mathbf{z}_{ij}^2). \tag{9}$$

Simply put, if both the utility function and the hedonic price equations are additive separable in (some transformation of) the $\mathbf{z}_{ij}$'s , then the interaction terms drop out of the quadratic form, leaving only the linear and squared terms. Clearly, the main reason to prefer (9) over the original quadratic approximation is that one gains in that way $m(m - 1)/2$ degrees of freedom in the estimation of the multinomial logit model. This will turn out to be a significant advantage in the implementation of the model, particularly since the attributes of individuals are not included.[13]

B. *A Test for Independence of Irrelevant Alternatives and the Nested Multinomial Logit Model*

A distinctive feature of the multinomial logit model, accounting for many of its advantages and limitations, is the assumption of independence of irrelevant alternatives. For our purposes here, the important point is that this assumption precludes the possibility of having a flexible pattern of substitution between alternatives. For example, one would expect the cross-price elasticities *within* the subset of head (or body) scanners to be higher than those *between* scanners of different types, but the assumption of independence constrains them to be equal (i.e., $\partial \ln \pi_i / \partial \ln p_j$ is independent of i). The question is how to test for this property and how to introduce more flexibility to the model if the test is rejected.

Hausman and McFadden (1981) put forward a specification test based on the comparison between the estimates obtained when applying the multinomial logit model to the full choice set, β_f, and those of a restricted set, β_r: if independence of irrelevant alternatives holds, both should be statistically equivalent. Formally, the statistic

$$S = (\beta_r - \beta_f)'(\text{cov}_r - \text{cov}_f)^{-1}(\beta_r - \beta_f) \sim \chi^2_{(m)} \tag{10}$$

[13] If the independent variables included in the multinomial logit model are "generic," i.e., they vary across alternatives but not across individuals, then the maximum number allowed is $n - 1$, n being the number of choices in the set. Clearly, the constraint does not hold when the personal attributes $\mathbf{h}$ are interacted with the characteristics of the choices ($\mathbf{z}$ and p in this case).

can be used to test the null hypothesis $(\beta_r - \beta_f) = 0$ (m is the rank of the cov matrix). In order to implement this test, the multinomial logit model was estimated for the full set of head and body CT scanners and for the restricted set of body scanners, for two far apart years 1977 and 1981. In both cases the estimated coefficients of the restricted and unrestricted models differ substantially, and the null hypothesis is indeed rejected: the critical χ^2 value at the .99 significance level is 18.5, whereas the values of the statistic in (10) are $S(1977) = 21.83$ and $S(1981) = 27.77$. Thus the data do not support the assumption for the full choice set; that is, the choice behavior of buyers does not allow for the symmetric treatment of head and body scanners.

In the light of those results, I formulate next a nested multinomial logit model having two branches, one for head (H) and the other for body (B) scanners. The probability of choosing the ith scanner of type h, h = H, B, can then be written as $\pi(i|h) = \pi(i|h)\pi(h)$, where

$$\pi(i|h) = \frac{e^{V_i}}{\sum_{j=1}^{n_h} e^{V_j}}, \tag{11}$$

$$\pi(h) = \frac{e^{\lambda W_h}}{e^{\lambda W_{\mathrm{H}}} + e^{\lambda W_{\mathrm{B}}}}, \tag{12}$$

n_h is the number of brands in cluster h, and W_h is the inclusive value of that cluster, that is,

$$W_h = \ln\left[\sum_{j=1}^{n_h} \exp(V_j)\right], \quad h = \mathrm{H, B}. \tag{13}$$

The key parameter in this model is λ, which is to be interpreted as a measure of substitutability of alternatives across clusters.[14] Thus if we take the limit cases (λ has to lie in the unit interval), when $\lambda = 1$, the cross-elasticities do not depend on the location of alternatives, and therefore the simple multinomial logit model applies to the whole set. On the other hand, $\lambda = 0$ means that the cross-elasticities between alternatives belonging to different clusters are zero, and hence each cluster should be regarded as a separate analytical unit or market. Closely related, the value of λ determines also the form of the $W(\cdot)$

[14] There has been some confusion in the literature regarding the role of λ, which at times has been interpreted as a measure of independence *within* rather than *between* clusters. The exact role of λ in determining substitution patterns is formally proven in Trajtenberg (1983, chap. 3).

function: if $0 < \lambda \leq 1$, then

$$W = \ln\left[\sum_h \left(\sum_i \exp V_{i,h}\right)^{\lambda}\right], \tag{14}$$

whereas if $\lambda = 0$, then

$$W = \sum_h W_h \pi(h), \qquad W_h = \ln\left(\sum_i \exp V_{i,h}\right), \tag{15}$$

where the probabilities $\pi(h)$ no longer depend on the attributes and prices of the alternatives as captured by the inclusive values in (13) but are instead exogenous to the model and can therefore be taken to be simply the observed frequencies in the population.

C. *Estimation Results*

I proceed now to estimate the nested multinomial logit model sequentially; that is, first (11) is estimated separately for head and body scanners, the estimated coefficients are used to compute the inclusive values in (13), and then those values are incorporated in (12) in order to estimate λ.[15] This is done for every year from 1976 to 1981 (1975 received separate treatment).[16] The small number of different brands of head scanners offered every year imposed several restrictions on the first-stage estimates of the head branch (recall n. 13): first, only four years (1976–79) could be estimated since in 1980 and 1981 the choice set contained too few scanners (only three). Second, V_i had to be specified as linear in the **z**'s (with price) rather than quadratic with residual price.[17] Third, only three variables could be included in 1979, and hence RTIME (the least important characteristic) was omitted. Similar problems arose in trying to estimate 1975 (when only a small set of both head and body scanners existed in the market),

[15] This procedure is not fully efficient, and hence the second-stage standard errors are not entirely reliable. An alternative would have been to estimate the system using full-information maximum likelihood methods, but the software available at the time was not good enough for that purpose.

[16] New models of CT scanners are usually introduced at the annual meetings of the Radiological Society of North America, which take place during the month of November of each year. Thus a year is taken to be the period November 1–October 31. Owing to specific features of the software used (the MLOGIT), the choice sets are allowed to expand within years (i.e., to include new entries that occur during the year), but not to contract.

[17] However, the correlations between prices and characteristics of head scanners were found to be systematically lower than those for body scanners, suggesting that the multicollinearity problem (which motivated the use of the quadratic) might be less severe for head scanners and hence that the linear form might in fact be appropriate for them.

TABLE 2

Multinomial Logit Estimates for Head CT Scanners, by Year (Linear Form)

	1976	1977	1978	1979
PRICE	−.748	−.709	−.893	−.818
	(−1.5)	(−1.9)	(−5.5)	(−2.4)
SPEED	.018	−.238	.303	−.318
	(.2)	(−1.4)	(1.9)	(−.9)
RESOL	−4.706	−5.565	−7.756	−10.307
	(−3.6)	(−2.9)	(−4.8)	(−4.7)
RTIME	−.619	−.366	1.119	...
	(−2.2)	(−1.2)	(2.7)	
$\rho^2 = 1 - [L(\beta^*)/L(\beta^0)]$	.131	.116	.42	.455
Corr(π^*, π)	.99	.910	.993	.998
	(.0001)	(.012)	(.0001)	(.0016)
Number of scanners	6	6	8	4
Number of observations	89	56	69	80

Note.—Asymptotic *t*-values are in parentheses.

forcing the use of a hybrid specification, described in detail in Trajtenberg (1983, chap. 5).

Table 2 presents the first-stage multinomial logit estimates for head scanners, and table 3 those for body scanners. I shall not analyze here those results in any detail but just note the following: first, the model fits fairly well,[18] even though only **z** variables have been used, with individual and firm-specific attributes omitted. Second, the estimated coefficients vary substantially from year to year, suggesting that preferences for attributes have been changing over time[19] and giving rise to a sort of index numbers problem in the computation of ΔW (see Sec. V*B*). Finally, note that the coefficients of RPRICE for body scanners are positive (except for 1981), whereas those for head scanners are negative and fairly stable; those findings will have far-reaching implications for the welfare analysis.

Now to the second stage of the nested multinomial logit model. With the inclusive values computed according to (13) and with the results from tables 2 and 3, the estimation of (12) renders the following estimates of λ (obviously, this can be done only for 1976–79, the

[18] The still scanty experience with these models indicates that values of ρ^2 ("McFadden's R^2") of the order of .20 represent fairly good fits (see Hensher and Johnson 1981). The other goodness-of-fit measure used here, corr(π^*, π) (i.e., the correlation between predicted and actual shares), is coarser but fairly informative of the performance of the model in the aggregate.

[19] Since the utility function is not linear, "preferences" are not uniquely defined in terms of the estimated coefficients, but depend on the level of the attributes. I have chosen to measure preferences for an attribute as the derivative of the choice probability with respect to that attribute, averaged over the choice set.

TABLE 3

MULTINOMIAL LOGIT ESTIMATES FOR BODY CT SCANNERS (Quadratic Form with Residual Price)

	1976	1977	1978	1979	1980	1981
RPRICE	11.252	.993	1.020	.485	.695	−.277
	(6.4)	(4.8)	(4.8)	(1.8)	(2.4)	(−2.5)
SPEED	−2.292	2.138	4.624	−8.669	11.347	−7.504
	(−7.3)	(2.8)	(1.0)	(−1.5)	(2.0)	(−.5)
$SPEED^2$	.236	−1.264	−8.283	31.292	−34.838	74.161
	(4.0)	(−3.4)	(−.6)	(1.9)	(−1.6)	(1.4)
RESOL	69.107	9.113	−34.126	−15.283	−18.129	32.877
	(7.3)	(2.4)	(−6.3)	(−5.0)	(−3.6)	(−3.9)
$RESOL^2$	−23.360	−2.533	15.096	6.291	7.738	−24.028
	(−7.6)	(−1.5)	(5.8)	(3.8)	(2.7)	(−4.2)
RTIME	−3.931	5.082	2.385	3.288	3.161	−2.591
	(−5.3)	(7.0)	(2.0)	(3.3)	(2.8)	(−2.8)
$RTIME^2$	1.054	−2.370	−1.511	−1.401	−2.093	5.560
	(4.5)	(−6.7)	(−2.0)	(−2.1)	(−2.2)	(3.9)
$\rho^2 = 1 - [L(\beta^*)/L(\beta^0)]$	.29	.12	.16	.16	.20	.14
$Corr(\pi^*, \pi)$	.999	.877	.900	.870	.722	.547
	(.0001)	(.0001)	(.0001)	(.0001)	(.0024)	(.082)
Number of scanners	8	15	16	16	15	11
Number of observations	285	324	164	177	193	153

NOTE.—Asymptotic t-values are in parentheses.

years in which the first-stage multinomial logit model for head scanners could be estimated).

Year	$\hat{\lambda}$	Standard Error	ρ^2
1976	.0205	.0021	.21
1977	.0784	.0065	.40
1978	−.2120	.0351	.12
1979	.0706	.0121	.10

The key result is that, except for 1978, the estimates of λ are very small, implying that the cross-elasticities between head and body scanners are nil and hence that the two types of scanners do not belong to a common decision tree. Without the *corrected* standard errors,[20] the hypothesis that $\lambda = 0$ could not be formally tested (using either a Wald or a Lagrangian test). However, the stated conclusion is not contingent on the acceptance of this hypothesis: even if the λ's had proved to be statistically different from zero, their small magnitude makes the between cross-elasticities negligible, and that is all that is required.[21]

The inference that head and body scanners do not belong to a common decision tree is hardly surprising: as suggested in Section III*B*, the evolution over time of the CT technology, of relative prices and capabilities, and of patterns for acquisition and use clearly indicates that head and body scanners diverged rapidly from each other, forming two highly segmented submarkets. More generally, it seems that the procedure followed above, centered around the estimation of λ, could be widely applied in tackling the all-pervasive problem of drawing market boundaries in empirical micro studies.

V. Computing the Gains from Innovation

The multinomial logit estimates obtained above provide us with the parameters of the utility functions needed to compute the social gains from innovation in CT scanners. To recall, those gains are defined as $\Delta W_t = W(S_t) - W(S_{t-1})$, where

$$W(S_t) = \frac{\ln(\Sigma \exp V_{it})}{\alpha}. \tag{16}$$

[20] In principle, these could have been obtained by further iterations using the Berndt et al. (1974) procedure; however, the necessary software was not readily available to me at the time.

[21] The negative value of λ for 1978 is symptomatic of a local failure of the conditions underlying the model (see McFadden 1984). Although firm conclusions cannot be drawn from such a finding, I take it as further evidence that head and body scanners could not be regarded as substitute systems.

Thus all that one needs is to compute the functions V_{it} using the estimated multinomial logit coefficients and the observed characteristics and prices of scanners in adjacent years, aggregate them as in (16), and take differences. There are, however, two important issues to be considered beforehand: first, the fact that the coefficients of RPRICE for body scanners were found to be positive (except for 1981) and, second, the fact that virtually all the estimated coefficients change significantly from year to year and hence that ΔW_t is not uniquely defined.

A. *Upward-Sloping Demand Curves and Welfare Analysis*

Recall from equations (4) and (5) that the price coeﬃcient, α, is supposed to estimate the marginal utility of income and that, if correctly measured, the latter plays the role of conversion factor between utility and money. However, a *positive* α cannot be regarded as an estimate of the marginal utility of income, and, more generally, one can no longer obtain a measure of consumer surplus by integrating under the observed, upward-sloping demand curve. Likewise, it is no longer clear whether the term $\alpha\tilde{p}_i$ should be included as such in V_{it} when computing (16): that will depend on the kind of phenomena that give rise to those positive price coefficients.

In the context of a simple model with perfect information, upward-sloping demand curves are readily obtained by having demand depend not only on price but on quality as well, provided that quality and price are positively associated (see, e.g., Spence 1973). Clearly, if one were to estimate such a model omitting quality, the price coefficient would be biased upward and could even become positive.

In the present case, quality was presumably included in the estimated model, in the form of the vector **z**. As mentioned earlier, though, the **z**'s do not exhaust the relevant quality space, and some of them can be taken only as proxies to the true performance dimensions (e.g., spatial resolution vis-à-vis "image quality"). It is important to emphasize that this is not just a problem for outside observers (such as ourselves) but rather a prime concern for the buyers of the systems themselves: observable attributes are only partially informative of the expected performance of a CT scanner (which is the case with most durable goods), and the remaining uncertainty leaves room for prices to play a role as signals of quality (see, e.g., Shapiro 1983). Or, what amounts to the same thing, the residual prices from the hedonic regressions may in fact incorporate unobserved quality dimensions, and hence the estimated α may also pick up their effect on utility.

Denoting the vector of unobserved attributes by $\mathbf{z}^u$ and assuming for simplicity that $V_i = \phi_1(\mathbf{z}_i) + \phi_2(\mathbf{z}_i^u) - \alpha p_i$ and that $\phi_2(\mathbf{z}_i^u) = \beta \tilde{p}_i$,[22] we get

$$V_i = [\phi(\mathbf{z}_i) - \alpha p(\mathbf{z}_i)] + (\beta - \alpha)\tilde{p}_i. \tag{17}$$

Recalling that the bracketed term in (17), defined as "net utility," was approximated by a simple quadratic form on **z** and denoting $\delta = (\beta - \alpha)$, we get

$$V_i = \sum_j (a_j \mathbf{z}_{ij} + b_j \mathbf{z}_{ij}^2) + \delta \tilde{p}_i, \tag{18}$$

which is the model actually estimated, except that now the price coefficient no longer stands for the marginal utility of income but incorporates also the marginal utility of the unobserved attributes, β (thus $\delta \lesseqgtr 0 \Leftrightarrow \beta \gtreqless \alpha$).

The problem is that one still needs to know the true marginal utility of income, α, in order to carry out the welfare analysis: using instead the observed coefficient δ would obviously be meaningless (i.e., it would be like integrating under the upward-sloping demand curve). Fortunately, the striking differences in the evolution over time of head and body scanners will allow us to decompose δ and associate α with the price coefficient of head scanners.

Consider the following straightforward proposition: The relative importance of unobserved versus observed quality dimensions (and hence of β) will be greater, the more technologically complex a product is, the less experience users have with it, and the faster is the pace of technological advance. Conversely, as the basic configuration and range of applications of a product stabilize and as experience with it accumulates, β will tend to vanish.

It is clear how this proposition applies to the two types of scanners: head scanners were introduced first, their applications remained unchanged (i.e., for brain studies), and, after an initial stage of improvements, the dominant tre·ıd was toward less expensive and simpler systems. On the other hand, the trend in body scanners was all along toward increased sophistication, with quantum technological jumps in the first couple of years and a slowdown in the pace of change afterward. Thus we would expect, first, that β will vanish early on for head scanners and hence that their estimated price coefficients will be negative and fairly stable after the first few years and, second, that the

[22] Strictly speaking, the inclusion of βp in the utility function implies that there is a signaling equilibrium in the market, i.e., that actual (or ex post) quality coincides on average with expected quality (the expectations having been formed on the basis of observed prices) and hence that residual prices do convey the right information. This is a plausible assumption for the case of CT scanners.

price coefficients of body scanners will be systematically higher than those of head scanners but that they will tend to converge toward the latter as the pace of innovation subsides. And indeed, a look at the following estimated price coefficients of the two types of scanners (taken from tables 2 and 3) strongly supports these conjectures.

	1975	1976	1977	1978	1979	1980	1981
Head	7.31	−.75	−.71	−.89	−.82	n.a.	n.a.
Body	10.18	11.25	1.00	1.02	.48	.69	−.28

Thus one can take the price coefficients of head scanners for 1976–79 as reliable estimates of α and hence as the appropriate parameters to be used in computing $W(S_t)$.[23] Likewise, the difference between the price coefficients of body and head scanners provides some idea of the magnitude of the residual uncertainty with respect to quality faced by buyers of body scanners. Although the foregoing discussion refers to a particular market, there is room to believe that the phenomenon addressed here is fairly widespread and that the approach used could be widely applied in dealing with quality uncertainty (and with upward-sloping demand functions) in segmented markets.

B. Alternative Measures of Welfare Gains

It was noted earlier that the coefficients estimated in the multinomial logit model change substantially from year to year, and therefore the welfare measures are not uniquely defined but depend on the choice of a reference year. Thus there are two alternative ways of assessing the value of a change in the choice set from S_t to S_{t+1}:[24] (i) ex ante: $\Delta W^a = W_t(S_{t+1}) - W_t(S_t)$, and (ii) ex post: $\Delta W^p = W_{t+1}(S_{t+1}) - W_{t+1}(S_t)$, where W_t stands for the surplus function using the coefficients estimated for year t. That is, the ex ante or forward-looking measure answers the following question: How much would the consumer be willing to pay for the option of facing next year's choice set rather than the present one, given his or her preferences today? On

[23] Since the price coefficients for head scanners could be estimated only for 1976–79 and in view of the fact that they are fairly stable, the 1976 coefficient will be used as an estimate of α for earlier years, and that for 1979 as the estimate for subsequent years.

[24] A further alternative would be to use a Divisia-like measure, i.e., to estimate each pair of adjacent years (or more) jointly and use the resulting "average" preferences to evaluate the changes from one year to the next. However, if tastes do change substantially from year to year, then suppressing those differences can hardly render a better measure; if they do not, then the ex ante and the ex post measures will be very similar, and the problem will not occur to begin with.

the other hand, the question posed by the ex post criterion is, How much income could be taken away from the consumer so as to leave him or her indifferent between facing today's and yesterday's choice sets, in the light of his or her present tastes?[25]

In general, the ex ante and ex post measures will provide different quantitative answers, and a priori it is not clear whether they would differ in a systematic way or which should be deemed to be more relevant. In the present case, though, the changes in preferences have followed a well-defined pattern, making the ex ante measure systematically higher than the ex post. This pattern emerged as the result of a "dual-inducement" process: strong preferences for a given attribute in one period seem to induce a relatively large improvement in that attribute in the following period, which in turn brings about a reduction in its marginal desirability afterward (for a detailed description of this mechanism, see Trajtenberg [1983, chap. 7]). Thus the value of the improvement will necessarily be larger if judged according to the original preferences (i.e., ex ante). Moreover, the same inducement mechanism seems to indicate that the ex ante measure is somewhat more appropriate from a normative viewpoint since consistency would require that any kind of change be judged according to the preferences that give rise to it rather than by hindsight. Nevertheless, given that the issue is inherently inconclusive,[26] both measures are computed whenever possible; as will be seen, though, the qualitative results hold equally well for both measures.

C. *Computations and Results*

First, note that the residual prices in (18) have to be obtained from the hedonic price function of the reference year since the nonstochastic component of the hedonic regression, $p(\mathbf{z})$, has to be added back to the quadratic approximation in order to retrieve the original utility function. That is, in computations of ΔW^a, the hedonic price function of year t is used to generate the $\tilde{p}$'s of both years (t and $t + 1$), and, likewise, in computations of ΔW^p, the residuals are obtained from the hedonic price regression estimated for year $t + 1$. Second, since the hypothesis that head and body scanners belong to a common prefer-

[25] This is not to be confused with the distinction between compensating and equivalent variations, or the Laspeyres-Paasche dichotomy: there tastes are held fixed, and the dilemma resides in choosing the reference utility level or consumption bundle (the ambiguity arising because of income effects), whereas here it is the taste parameters themselves that change (see Fisher and Shell 1971).

[26] A similar problem arises in the context of endogenous tastes, and there too the choice of a criterion for welfare analysis is not a close issue (see, e.g., von Weizsäcker 1971).

TABLE 4

WELFARE GAINS BY SCANNER TYPE AND OVERALL GAINS (in Constant 1982 Dollars [Millions])

	Head Scanners			Body Scanners			Overall Incremental Gains	
Year	ΔW_H^a	ΔW_H^p	π(H)	ΔW_B^a	ΔW_B^p	π(B)	Ex Ante	Ex Post
1974	...	1.204	1.00	...	1.204	0	8.713*	1.204
1975	...	.208	.55	...	.208	.45	1.509*	.208
1976	4.776	.319	.24	4.776	.632	.76	4.776	.557
1977	.056	.017	.15	1.097	.163	.85	.940	.141
1978	.361	.225	.28	.022	−.066	.72	.115	.014
1979	−.013	−.047	.30	.204	.055	.70	.140	.025
1980	−.019	...	.20	.097	.054	.80	.074	.010*
1981	.007	...	.08	.201	−.009	.92	.184	.026*
1982	−.003	...	.06	.209	...	.94	.195	.027*

NOTE.—ΔW_h^a = ex ante gains; ΔW_h^p = ex post gains. Overall incremental gains: $\Delta W = \Delta W_H \pi(H) + \Delta W_B \pi(B)$.
* Computed using the mean ratio $\Delta W^a/\Delta W^p$ for 1976–79.

ence tree was rejected (from 1976 on), the ΔW's are computed separately for each type of scanner in the period 1976–82. As for the initial years, the gains correspond to the joint set of head and body scanners since the parameters used are those of 1975.[27] For completeness, the 1975–76 ex ante joint gains are imputed to each type of scanner according to their respective ex post shares: $s_H = \Delta W_H^p/(\Delta W_H^p + \Delta W_H^p)$, $s_B = 1 - s_H$.

The computations are shown in table 4. Note that an overlapping series of ex ante and ex post pairs could be computed only for the four years 1976–79. As expected, ΔW^a was found in those years to be systematically higher than ΔW^p by a relatively constant factor: the mean of $(\Delta W_t^a/\Delta W_t^p)$ is 7.24, and the standard deviation is 1.34. This proportionality factor is used to obtain imputed values for the 1974 and 1975 ex ante gains (i.e., $\Delta W_t^a = 7.24 \cdot \Delta W_t^p$, $t = 74, 75$), and likewise for the 1980–82 ex post measures. The measure ΔW_t will be referred to as *incremental* gains, to be distinguished from the *cumulative* gains $CW_t = \Sigma_{\tau=1}^{t} \Delta W_\tau$.

VI. Diffusion, Social Returns, and the Pattern of Innovation

The ΔW's obtained above stand for the yearly incremental gains accruing to the representative buyer of CT scanners. The goal now is to

[27] That is, the 1975 $W(\cdot)$ function was used to compute the 1973–74 ex post gains and the 1975–76 ex ante gains; note, however, that the 1973 and the 1974 sets include only head scanners, one each year.

compute the yearly flow of *total* gains, relate them to R & D expenditures, and examine in detail the time profile of those benefits and costs. As will become clear shortly, the computation of those total gains has to do with an important phenomenon in the realm of technical change, namely, the dynamic interaction between innovation and diffusion.

A. Total Gains and Diffusion

As a first step, consider the problem of delimiting, in time and in technology space, the benefits from innovation accruing to a consumer who buys, say, a personal computer today. Are those benefits to be identified with the cumulative gains stemming from the long sequence of innovations in computers from the ENIAC on? Or perhaps from the first Apple onward? Or should we regard them to be just the last incremental gains, gauged by contrasting the 1988 versus the 1987 sets of personal computers in the market? Looking at it from a different angle, we can also phrase the same conundrum as follows: Are we to compute the total gains from the innovations embedded in the 1988 set of personal computers simply by multiplying the incremental gains, ΔW_{88}, times the number of buyers in 1988? Or rather times the projected number of buyers from 1988 onward? And what about those buying replacement units versus first-time buyers? Aside from the issue of choosing an appropriate baseline, it is clear that the key to the problem lies in the dynamics of demand and the impact of successive innovations on it. In order to address it, I formulate next a simple diffusion model incorporating the gains from innovation that amounts essentially to the reduced form of a full-fledged dynamic model.[28]

In order to grasp the nature of the model, consider the following polar versions of the diffusion process: in the first, diffusion is due entirely to the workings of the traditional demonstration effects, usually associated with learning, emulation, and rivalry. In other words, the process reflects only dynamic phenomena occurring within the population of potential adopters and is not affected by external forces. Thus, for example, such a model predicts that if technological change would have ceased after the introduction of the first CT scanner, the pattern of diffusion would have been exactly the same as it was in actuality. According to the opposite version, the diffusion path

[28] A thorough treatment of the problem would require estimating a dynamic discrete choice model, relying, e.g., on Heckman (1981). The total gains would then be obtained by integrating the ensuing intertemporal demand function over individuals and over time. Such a task is at present beyond the scope of this work.

is nothing but a temporal demand curve having only extensive margins; that is, it traces the distribution of some sort of reservation price in the population of potential adopters. Thus successive innovations that result in what can be thought of as reductions in "real" (i.e., quality-adjusted) prices trigger immediate adoption by inframarginal consumers. Consequently, if the process of technological advance were to come to a halt, diffusion would stop as well, all future purchases would be for replacement or capacity additions only, and hence the innovations that took place up to that point would cease to generate further benefits.

The implications of these alternative scenarios for the computation of total gains are immediate: if diffusion corresponds to the case of the distribution of reservation prices, then the total gains generated by innovations occurring at time t would just be $TW_t = \Delta W_t n_t$, n_t being the number of buyers in period t and TW_t total gains. On the other hand, if diffusion is due entirely to demonstration effects, then

$$TW_t = \Delta W_t N \int_t^{\infty} f(\tau) e^{-r(\tau - t)} d\tau, \tag{19}$$

where r is an appropriate discount rate, N is the size of the population of potential adopters, and $f(\cdot)$ is the marginal distribution of adoption times, corresponding to the cumulative distribution $F(\cdot)$. To make it clearer, let me ignore discounting for a moment and rewrite (19) as

$$TW_t = \Delta W_t N[1 - F(t)] = \Delta W_t\{n_t + N[1 - F(t + 1)]\} = \Delta W(n_t + n_t^f), \tag{20}$$

where $n_t^f = N[1 - F(t + 1)]$ stands for the number of future adopters. It is easy to see that the total gains would be larger in the demonstration effects model since they include also the benefits bestowed by current innovations to future buyers (with discounting n_t^f will be somewhat smaller).

Actual diffusion processes may correspond to either model or, most likely, to a combination of both, and it is of course an empirical matter to uncover the appropriate characterization. Thus I estimate next the aggregate diffusion process as a function of time and of the cumulative gains from innovation, $CW_t = \Sigma_{\tau=1}^{t} \Delta W_\tau$. As in traditional diffusion models, time is meant to capture the forces associated with demonstration effects, whereas CW_t can be thought of as tracing the cumulative changes in a quality-adjusted index, thus bringing in the scenario associated with the distribution of reservation prices. In particular, it is assumed that the diffusion path corresponds to a logistic distribution and that innovation affects the process by shifting the ceiling K, that is, $K = K(CW_t)$. As to the functional form of $K(\cdot)$, both a linear and a concave specification were considered, corresponding to an underlying uniform and exponential distribution of reservation prices, respectively. Since the two yielded very similar results, only those obtained with the linear form, $K(CW_t) = K_0 + kCW_t$, are shown here. The estimated equation is thus

$$F(t) = \frac{K_0 + kCW_t}{1 + \exp(\alpha - \beta t)}, \tag{21}$$

and the data consist of the monthly number of adopters of CT scanners in the United States from November 1972 to July 1981.[29] As a benchmark I estimate also a logistic equation without the term $k \cdot CW_t$ but with a free ceiling. The results are shown in the following table (the numbers in parentheses are asymptotic standard errors).

	$\hat{K}_0$	$\hat{k}$	$\hat{\alpha}$	$\hat{\beta}$	Residual Sum of Squares
(i)	.459	...	4.053	.071	.013
	(.004)		(.076)	(.002)	
(ii)	.074	.025	3.443	.06	.006
	(.003)	(.002)	(.07)	(.001)	

First, note that diffusion was strongly influenced both by demonstration effects (embedded in t) and by technological advance, as manifested in the fact that the estimates of both β and k are highly significant. Moreover, the fit improves greatly when one goes from (i) to (ii), implying that the traditional diffusion model (which ignores innovation) would have been widely off the mark. To put it in quantitative terms, $\hat{k} = .025$ means that for every million dollars' worth of improvements in CT, the number of adopters increased by 2.5 percent. Thus if innovation would have ceased just after the introduction of the first CT scanner, only 7.4 percent of the total population would have adopted throughout the period (since $\hat{K}_0 = .074$). In reality, the ceiling had climbed to 49 percent by 1982 as a result of the flow of innovations from 1974 on.

The estimated equation can also be used in order to obtain, albeit in an indirect way, a measure of the *initial* gains from innovation, that is, of the gains associated with the introduction of the first CT scanner (recall the limitations of ΔW discussed in Sec. II). The question can be formulated as follows: What would those first gains (ΔW_{73})

[29] The dependent variable in (21) is $F(t) = n_t/N$, where n_t is the number of new adoptions in month t, and the total population of potential adopters is $N = 3{,}457$ (see Trajtenberg and Yitzhaki, in press). The figures for CW_t are computed from table 4.

have to be in order to give rise to the initial ceiling K_0, given the estimated function $K(CW_t) = K_0 + kCW_t$? The answer is simply $\Delta W_{73} = \hat{K}_0/\hat{k} = .074/.025 = 2.99$; that is, and to put it carefully, equation (ii) above is equivalent to an equation in which K_0 is deleted, and ΔW_{73} is set equal to 2.99 rather than to zero. In other words, if the behavior underlying equation (ii) is stable, then the introduction of the first CT scanner had to be worth $3 million so as to induce 7.4 percent of hospitals and clinics to adopt it. Although ΔW_{73} will be used below along with the other measures, it should be borne in mind that this figure was arrived at in a very different fashion, probably making it a less reliable estimate.

A word about assigning benefits to purchases for replacement and additions to capacity: again, these can be properly dealt with only in the context of a full-fledged dynamic model, incorporating a capital accumulation process. Short of that and preferring to understate rather than overstate the total gains, I proceed on the assumption that the benefits accruing to repeat purchases at time t are just the incremental gains ΔW_t.

The yearly total gains can now be computed as (note that n_t includes also second scanners and replacements)

$$TW_t = \Delta W\left[n_t + N(K_0 + kCW_t)\int_{t+1}^{\infty} f(\tau)e^{-r(\tau-t-1)}\,d\tau\right], \quad (22)$$

where the value of the parameters K_0 and k and those of $f(\cdot)$ are taken from the diffusion equation estimated above, and the yearly discount rate is assumed to be .05 (since $f(\tau)$ is defined in months, the rate is actually .0041). The integral in (22) does not have a closed-form solution (because of discounting) and hence had to be solved numerically. The computations are presented in table 5, which is largely self-explanatory.[30] As suggested above, these figures are likely to be biased downward and should therefore be regarded as a lower bound: aside from the fact that only the incremental gains were assigned to additional scanners and replacements, I have ignored the gains stemming from upgrades, that is, from the retrofitting of older units by the manufacturers, a widespread practice in this field (by mid-1981, nearly 25 percent of all scanners installed had been upgraded). The only possible source of upward bias lies in the implicit assumption that the underlying distribution of reservation prices is a step function, that is, that inframarginal users are just at the margin

[30] The computations are carried out using the ex ante gains only: since those were found to be systematically higher than the ex post measures, the total gains corresponding to the latter would just be a relatively constant fraction of the former. The qualitative results are thus unaffected by the type of measure used.

TABLE 5

COMPUTATION OF TOTAL, EX ANTE GAINS (in Constant 1982 Dollars [Millions])

Year	ΔW_t (1)	$K(CW_t)$ (2)	$1 - F(t)$ (3)	$1 - F(t)$ Discounted (4)	n_t^f (5)	$\tilde{n}_t^f$ (6)	n_t (7)	$n_t + \tilde{n}_t^f$ (8)	TW_t (9)
1973	2.990	.074	.938	.772	240	197	16	213	638
1974	8.713	.290	.878	.751	879	752	43	795	6,926
1975	1.509	.327	.777	.685	879	775	221	996	1,503
1976	4.776	.445	.628	.568	967	874	374	1,428	5,959
1977	.940	.469	.451	.414	731	671	390	1,061	997
1978	.115	.471	.285	.265	465	431	250	681	79
1979	.140	.475	.162	.151	267	249	275	524	73
1980	.074	.477	.086	.081	142	133	271	404	30
1981	.184	.481	.024	.023	41	38	392	430	79
1982	.195	.486	.012	.011	20	19	428	447	87

NOTE.—In col. 1, ΔW_t = incremental gains. In col. 2, $K(CW_t) = K_0 + kCW_t = .074 + .025CW_t$. In col. 3, $1 - F(t) = \int_{t+1}^{\infty} f(\tau)d\tau$. In col. 4, $1 - F(t)$ discounted $= \int_{t+1}^{\infty} f(\tau)e^{-r(\tau-t-1)}d\tau$. In col. 5, $n_t^f = (2) \times (3) \times N, N = 3{,}457$. In col. 6, $\tilde{n}_t^f = (2) \times (4) \times N$. In col. 7, n_t = number of scanners sold in year t. In col. 8, $n_t + \tilde{n}_t^f$ = number of current and future discounted beneficiaries from the incremental gains at t. In col. 9, TW_t = total gains.

prior to purchase so that, as the latest innovations trigger adoption, they receive ΔW_t in full. Otherwise, only a fraction of the incremental gains would actually be realized. On the other hand, the restraining influence of technological expectations may significantly reduce and even reverse this potential bias.[31]

B. Computing a Social Rate of Return to Investment in Innovation

An immediate application of the total benefits obtained above is to compute some version of a social rate of return by relating those benefits to R & D expenditures in CT.[32] Using the R & D figures of table 1 and the *TW*'s from table 5, I obtain a capitalized benefit/cost ratio of 270 percent (this is the average between two alternative specifications, one using R & D by U.S. firms only, the other including R & D by foreign firms as well).

[31] That is, users may delay adoption on the expectation of further improvements, even though they have already passed their threshold. In that case the gains realized will exceed the latest incremental gains.

[32] It should be noted here that the notion of social benefits (and hence of a social rate of return) is not unequivocal when applied to markets for medical technologies, in view of the distortions prevalent in the health care sector (see Trajtenberg [1983, chap. 6] for a detailed discussion of this issue). Moreover, the computation of the *level* of benefits and costs involves as seen some inherent ambiguities (such as the ex ante, ex post dilemma), and hence the resulting rate of return should be taken with great caution.

Given the small number of studies in which rates of return to R & D have been computed for individual innovations, it is difficult to assess the significance of this particular figure. A natural benchmark is Griliches's (1958) seminal study on the social returns to R & D in hybrid corn, which he estimated to be about 700 percent (also using the benefit/cost ratio).[33] There is, however, an important difference between Griliches's and the present study in the treatment of future benefits: in the former the flow of benefits continues ad infinitum at the level determined by the ceiling of the diffusion curve; that is, whenever hybrid corn is planted, now and in the future, society gets the benefits of its superior yield vis-à-vis conventional corn. By contrast, in the model here, the benefits cease as the diffusion process dies off since the purchase of additional or replacement scanners in the future does not confer further gains. If the rate of return on hybrid corn is recomputed omitting future benefits (i.e., those accruing after 1955), the result is 350 percent, much closer to the figure for CT scanners. Still, in order to draw solid inferences, one would need further empirical studies of a similar nature and a much better understanding of the determinants of returns to R & D.

C. The Time Profile of Benefits and Costs

Moving beyond the summary view provided by rates of return, I proceed now to examine in detail the evolution over time of both social gains and R & D expenditures: it is hoped that this will shed some light on the dynamics of the innovative process itself. To that end, figure 1 presents the time path of incremental gains, both the actual figures and a 3-year moving average,[34] and figure 2 the profiles of total gains and total R & D expenditures (in logs). To recall, the incremental gains reflect advances in the technology itself (i.e., the social valuation of those advances), whereas total gains incorporate also the effect of market size.

The first feature to note is that the gains generated in the first half of the period are far larger than those in the second half. In fact, the smoothed-out profile resembles a sort of lognormal distribution; that is, it starts high, rises still further during the initial period, and then

[33] Social rates of return to R & D were computed also in Mansfield et al. (1977): they found that the median rate of return in a set of 17 innovations was 56 percent, the highest being 307 percent. However, their results are not comparable with those here since they used internal rates of return rather than benefit/cost ratios.

[34] There is always an element of randomness in the precise timing of specific innovations. Thus, given the arbitrariness of any discrete partition of the time dimension (e.g., in calendar years), a moving average may better capture the essence of the underlying process.

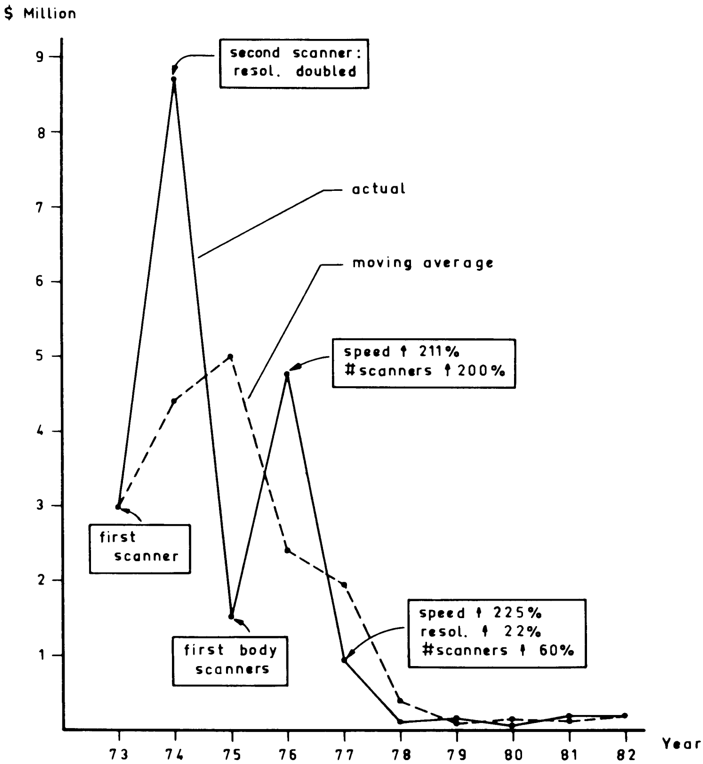

FIG. 1.—Time profile of incremental gains (yearly figures and 3-year moving averages)

declines rapidly, carrying a low-level tail into the future. Such a pattern is in fact highly plausible and may be accounted for by a generalization of a common feature of economic processes, namely, an initial phase of scale economies, promptly followed by the setting-in of sharply diminishing returns. What is peculiar in the context of innovation is that these nonlinearities appear to "take place" in three different dimensions: in the production of innovations (i.e., the R & D technology), in the utility or value generated by them, and in the determination of the size of the market for them.

With respect to R & D, usually it is relatively easy to improve the performance of a technology during its initial stages, from the viewpoint both of the resources needed and of the enabling scientific and technological principles. Later on, though, as the obvious advances are no longer there to be made and the technology is pushed to its limits, the marginal cost of further improvements rises, and hence the rate of innovation tends to abate. As to the second dimension, there is often a "threshold effect" in the utility derived from the characteris-

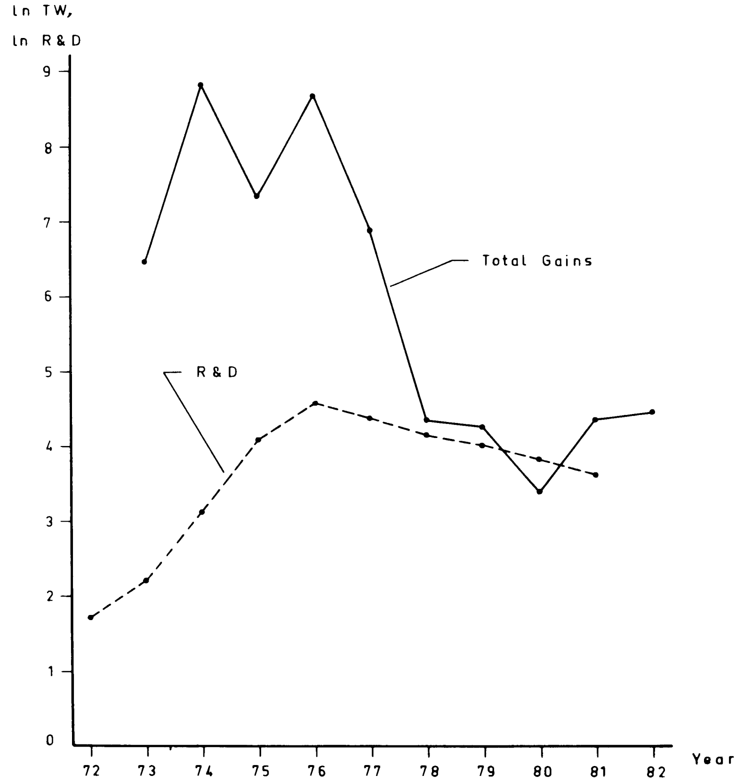

FIG. 2.—Total gains and total R & D expenditures (logarithmic scale). Original figures for total gains and R&D are in constant 1982 dollars (millions).

tics of new products, in the sense that below a certain level of performance the product is pretty much worthless (e.g., minimal resolution in CT scanners). That would account for an initial phase of "increasing returns" in the valuation of technical improvements, but soon after, diminishing marginal utility prevails. Thus, for example, increasing the speed of CT scanners from, say, 12 to 2 seconds was not nearly as valuable as going from 60 to 10 seconds. As to the third dimension, market size, consider how n_t^f evolves over time (recall that n_t^f stands for the number of users that benefit from the incremental gain ΔW_t as the diffusion process unfolds): from (20)–(22), assuming for simplicity that $K_0 = 0$, we get

$$\frac{d \ln n^f}{dt} = \frac{d \ln CW_t}{dt} - h(t), \quad h(t) = \frac{f(t)}{1 - F(t)}; \tag{23}$$

that is, the behavior of n_t^f over time depends on the rate of technological improvement vis-à-vis the hazard rate of the diffusion process,

$h(t)$. As argued above, one would expect the former to be nonincreasing, whereas if diffusion follows, for example, a logistic pattern, $h(t) = \beta F(t)$, and hence it increases monotonically. Thus n_t^f will in general be a concave function and have a maximum; that is, even if innovation proceeds at a constant pace, total gains will eventually decrease because the diffusion process exhausts itself faster than the rate at which the market expands as a result of technical advance.

As a convenient illustration of those effects, consider the scanner that generated the largest benefits in the history of CT, namely, the second CT scanner (the EMI 1000). This system was very similar to the original EMI Mark I except that it had much better resolution, a change that greatly improved the ability to visualize brain pathologies. Thus if the first scanner proved the *feasibility* of CT, the second transformed it into a useful diagnostic tool that could be widely applied. This seems to be a fairly general phenomenon in product innovations: the first models of a new product are rarely more than the embodiment of a potentially useful idea. The real leap forward (and the concomitant benefits) comes with the advent of a model in which some key attributes are greatly improved, turning it into a *practical* rather than just an ingenious device that can command wide appeal. Examples abound: the Ford Model T in cars, the DC-3 in aircraft, and the UNIVAC I in computers.

This early emergence of a big winner that brings about the most gains can thus be seen as the lumpy realization of increasing returns that operate simultaneously along the three dimensions mentioned above. Likewise, the rather dramatic drop in the flow of gains from innovation occurring later on can be attributed to diminishing returns setting in *at the same time* in those same dimensions. As a result, ΔW_t displays a definite wavelike pattern over time. It is tempting to take a step further and relate this finding to Schumpeterian theories that assert that the process of technological change at large follows a cyclical pattern (e.g., Schumpeter 1939). However, this being just a case study, all that can be said is that the evidence presented, and its underlying rationale, may be seen perhaps as enhancing the plausibility of technology-driven cycles.

A few remarks about the time profile of R & D: As figure 2 reveals, the flow of R & D outlays is *not* correlated with either incremental or total gains, reflecting the increasing/decreasing returns sequence in the technology dimension. Moreover, the ratio of gains to R & D for the period 1968–77 was a staggering 80 to 1, whereas for 1978–82 it was a bare 1.4 (in 1980 R & D actually exceeded the benefits). A straightforward implication is that *average* social rates of return to R & D are not very informative, certainly not as guides for policy. The question is not so much whether public support to R & D is warranted (it is not too difficult to make a case for it), but rather until when along the innovation cycle such support should be provided. In the case of CT, for example, it is clear that it was socially desirable at first to promote research, but that was certainly not the case after 1976.[35] The problem is that, if one were to wait for hard evidence on returns to R & D in order to decide what fields to support, the result might be a standoff; that is, once the evidence is available, the support may no longer be necessary. Thus learning more about time profiles of innovations may substantially contribute to the design of public policy in this area.

These last remarks highlight one of the key features of the present study, namely, that it looks into the very emergence of a new, high-tech market. The results, if typical, suggest that the bulk of the innovative "action" in those sectors takes place very early on. Thus studies in this area, focusing mostly on well-established or "mature" sectors, are likely to be missing a great deal of the phenomena sought. Moreover, since innovation plays a crucial role in the dynamics of technologically advanced sectors, it is crucial to have a better grasp (at least factual) of their initial stages in order to understand how they evolve into their long-term equilibria, and the normative aspects of those equilibria.

VII. Further Applications

There are, to be sure, many interesting ways in which the measures of welfare gains from innovation can be applied, aside from the straightforward applications reported in the previous section. To mention just a few: First, the ΔW's can easily be related to the dynamics of market structure (e.g., to changes in concentration or in monopoly power), thus providing a way to address one of the key issues in the economics of innovation. Preliminary results show a strong positive correlation between innovation and competition and suggest simultaneity rather than simple-minded, one-way causal links. Second, these measures can shed light on the usefulness of patent counts for the study of innovation: as shown in Trajtenberg (1987), patents weighted by citations closely follow ΔW_t over time, whereas simple patent counts trace well the level of innovative activity (e.g., of R & D

[35] It is worth noting that the British Department of Health supported the development of the CT scanner at EMI in 1970–72 and that apparently such support was vital for carrying out the project. On the other hand, the American neurologist W. H. Oldendorf had independently developed a similar system earlier on but was unable to pursue the project for lack of support in the United States, from either the government or private firms.

expenditures). Third, the ΔW's can be used to construct quality-adjusted price indices (or, more precisely, cost-of-living indices) far superior to those that can be computed from just hedonic price regressions. This would seem enough for now as stimuli to the research agenda on technological change.

Appendix

The main sources used in gathering data on CT scanners were (*a*) a questionnaire sent to all CT manufacturers and follow-up personal contacts with officers from those companies; (*b*) articles in the scientific literature (i.e., in radiology and medical physics), journals related to the medical sector such as *Modern Health Care* and *Diagnostic Imaging*, reports from consulting companies (e.g., A. D. Little and Eberstadt), and a variety of other publications; (*c*) public and government agencies, primarily the Bureau of Radiological Health at the Food and Drug Administration, the Office of Technology Assessment of the U.S. Congress, the U.S. Patent Office, and the American Hospital Association; and (*d*) a telephone survey of a few hundred hospitals and personal contacts with faculty and researchers at various medical schools.

All but two of the companies that were still active in the CT market by 1981 answered the questionnaire, at least partially: Elscint, General Electric, Omnimedical, Picker, Siemens, Technicare (of Johnson and Johnson), and Toshiba. The two that did not were CGR (from France) and Philips; the latter had to be excluded from the analysis for lack of reliable information, but not much is lost since the company had sold fewer than 20 scanners in the United States by 1981. Of those that had exited the market, only one (Varian) responded to the questionnaire, whereas the other six (Artronix, AS & A, EMI, Pfizer, Searle, and Syntex) could not be reached. I am quite certain that the installation data set, comprising more than 2,000 observations (including upgrades), covers about 98 percent of all CT installations in the United States up to July 1981 and that it is very accurate.

References

Berndt, E. K.; Hall, B. H.; Hall, Robert E.; and Hausman, Jerry A. "Estimation and Inference in Non-linear Structural Models." *Ann. Econ. and Soc. Measurement* 3 (October 1974): 653–65.

Bresnahan, Timothy F. "Measuring the Spillovers from Technical Advance: Mainframe Computers in Financial Services." *A.E.R.* 76 (September 1986): 742–55.

Fisher, Franklin M., and Shell, Karl. "Taste and Quality Change in the Pure Theory of the True-Cost-of-Living Index." In *Price Indexes and Quality Change: Studies in New Methods of Measurement,* edited by Zvi Griliches. Cambridge, Mass.: Harvard Univ. Press, 1971.

Griliches, Zvi. "Research Costs and Social Returns: Hybrid Corn and Related Innovations." *J.P.E.* 66 (October 1958): 419–31.

———, ed. *Price Indexes and Quality Change: Studies in New Methods of Measurement.* Cambridge, Mass.: Harvard Univ. Press, 1971.

Hanemann, W. Michael. "Discrete/Continuous Models of Consumer Demand." *Econometrica* 52 (May 1984): 541–61.

Hausman, Jerry A., and McFadden, Daniel. "Specification Tests for the Multinomial Logit Model." Working Paper no. 292. Cambridge: Massachusetts Inst. Tech., 1981.

Heckman, James J. "Statistical Models for Discrete Panel Data." In *Structural Analysis of Discrete Data, with Econometric Applications,* edited by Charles Manski and Daniel McFadden. Cambridge, Mass.: MIT Press, 1981.

Hensher, David A., and Johnson, Lester W. *Applied Discrete-Choice Modelling.* London: Croom Helm, 1981.

Lancaster, Kelvin J. *Variety, Equity, and Efficiency: Product Variety in an Industrial Society.* New York: Columbia Univ. Press, 1979.

McFadden, Daniel. "Econometric Models of Probabilistic Choice." In *Structural Analysis of Discrete Data, with Econometric Applications,* edited by Charles Manski and Daniel McFadden. Cambridge, Mass.: MIT Press, 1981.

———. "Econometric Analysis of Qualitative Response Models." In *Handbook of Econometrics,* vol. 2, edited by Zvi Griliches and Michael D. Intriligator. New York: North-Holland, 1984.

Mansfield, Edwin; Rapoport, John; Romeo, Anthony; Wagner, Samuel; and Beardsley, George. "Social and Private Rates of Return from Industrial Innovations." *Q.J.E.* 91 (May 1977): 221–40.

Rosen, Sherwin. "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition." *J.P.E.* 82 (January/February 1974): 34–55.

Schumpeter, Joseph A. *Business Cycles: A Theoretical, Historical, and Statistical Analysis of the Capitalist Process.* New York: McGraw-Hill, 1939.

Shapiro, Carl. "Premiums for High Quality Products as Returns to Reputations." *Q.J.E.* 98 (November 1983): 659–79.

Small, Kenneth A., and Rosen, Harvey S. "Applied Welfare Economics with Discrete Choice Models." *Econometrica* 49 (January 1981): 105–30.

Spence, A. Michael. "Price, Quality, and Quantity Interdependencies." Memorandum no. 161. Stanford, Calif.: Stanford Univ., Center Res. Econ. Growth, 1973.

Trajtenberg, Manuel. "Dynamics and Welfare Analysis of Product Innovations." Ph.D. dissertation, Harvard Univ., 1983.

———. "The Use of Multivariate Regression Analysis in Contrast-Detail Studies of CT Scanners." *Medical Physics* 11 (July/August 1984): 456–64.

———. "Patents, Citations, and Innovations: Tracing the Links." Working Paper no. 2457. Cambridge, Mass.: NBER, December 1987.

Trajtenberg, Manuel, and Yitzhaki, Shlomo. "The Diffusion of Innovations: A Methodological Reappraisal." *J. Bus. and Econ. Statis.* (in press).

von Weizsäcker, Carl C. "Notes on Endogenous Change of Tastes." *J. Econ. Theory* 3 (December 1971): 345–72.

Willig, Robert D. "Incremental Consumer's Surplus and Hedonic Price Adjustment." *J. Econ. Theory* 17 (April 1978): 227–53.