QUALITY-ADJUSTED PRICE INDICES AND THE MEASUREMENT OF ECONOMIC GROWTH

M. Trajtenberg
Tel-Aviv University

ABSTRACT

The goal of this paper is to address the problem of "product innovations" (i.e., new goods, increased variety, and quality change) in the construction of price indices and, by extension, in the measurement of economic growth. The premise is that a great deal of technical progress takes the form of product innovations, but that conventional economic statistics (e.g. "real product", productivity growth, and the like) fail by and large to reflect them. The approach suggested here consists of two stages: first, the benefits from innovation are estimated with the aid of discrete choice models and, second, these benefits are used to construct "quality-adjusted" price indices. Following a discussion of the merits of such approach vis-a-vis hedonic price indices, I apply it to the case of CT Scanners. The findings suggest that conventional economic indicators may, indeed, be missing a great deal of the welfare consequences of technical advance, particularly during the initial stages of the product cycle of new products.

I. INTRODUCTION

The key to the problem is that conventional index-number methods cannot possibly capture quality change and that hedonic price indices may offer some paliative, but by no means a full cure. As it stands now, then, there is no proven way of incorporating product innovations into measures of economic performance and, therefore no way of assessing the possible discrepancies that may therefore exist between "real" and conventionally measured aggregate product and growth.

In order to gain an intuitive idea of the problem, consider for example the case of personal computers (PCs). Year after year, the range of products available in this market keeps improving, displaying enhanced performance along various quality dimensions (speed, memory, portability, etc.), increased variety, new and better peripherals, and so forth. Presumably, consumers are better off as a consequence of these advances, probably to a significant extent. One way to capture their increased satisfaction is to ask how much they would be willing to pay for the privilege of being able to choose from, say, the 1990 range of PCs offered in the market rather than from the 1989 range. If such a measure were available, we would undoubtly want to incorporate it in any aggregate measure of economic performance, certainly in measures of aggregate growth: after all, the bettering of the universe of goods at our disposal is as important for our welfare (and for judging the performance of the economy) as is the expansion in the quantities of given products. Moreover, the innovations embedded in the 1990 range of PCs are the result of costly investments in R&D, the social returns to which have to be sought in the welfare gains stemming from the enhanced qualities; seeking those returns in the quantity dimension instead (and, of course, not finding them there), would lead to the totally unwarranted conclusion that the resources committed to those innovations have gone to waste and, hence, that the productivity of the economy has deteriorated.

Unfortunately, the fact is that there are no overt manifestations of those gains (not even remote proxies for them) that the statistical agencies in charge of producing economic indicators could go out and gather. Still, personal computers do leave an imprint in the statistics through the observed changes in their prices, provided, of course, that PCs have already been incorporated in the sampled basket of goods (something that typically happens a few years after the appearance of a new product). Suppose, now, that the average price of PCs went up from 1989 to 1990; that will have two closely related consequences: the inflation rate as conventionally measured will be higher (how much higher depends on the weight of PCs in the CPI or in the PPI) and the computed total factor productivity (TFP) growth rate will be lower than what we would have obtained had we ignored PCs. On both accounts, then, we would seem to be absolutely *worse off* whereas, as we said before, in reality consumers are likely to be much better off in view of the changes embedded in the 1990 range of PCs.

Furthermore, suppose that some macropolicies respond to those mismeasured indicators as, for example, if social security benefits (or taxes) are indexed to the inflation rate, or if there is a change in investment incentives (e.g. an increase in investment tax credits) in order to counteract the apparent (extra) slowdown in TFP. Clearly, those would be awfully wrong policies that will lead to a misallocation of resources. Mismeasurement, then, is not only a matter of perception or of historical record, but can have (adverse) real consequences as well.

The plan for this paper is as follows. First, acting on the belief that the "goodness" of a deflator is to be judged according to its ability to capture changes in consumers' welfare, I sketch an econometric approach for measuring *directly* the benefits from product innovations, which I have laid out in detail in previous work (Trajtenberg, 1989, 1990b). The proposed method draws primarily from discrete choice models and from the "characteristics approach" to demand theory, leading to estimates of the preferences for the attributes of products, and thence to value measures of quality changes. The novelty here resides in using those measures in order to construct "real" (or "quality-adjusted") price indices. That is, having obtained money measures of the gains (in terms of consumer surplus) from product innovations, I show a way to express those gains as changes in "real" prices. Following a discussion of the merits of such an approach vis-a-vis the use of hedonic price indices, I apply it to a specific innovation, namely CT (Computed Tomography) Scanners. The main finding is that the rate of *decline* in the quality-adjusted price of CT scanners was staggering (averaging 55 per cent over 9 years), particularly during the first few years following the introduction of the innovation. By contrast, a hedonic-based "real" index captures just a small fraction of the decline, not to speak of the unadjusted price index, which shows a substantial price *increase* over the same period.

Thus, conventional indices may, indeed, be missing many of the welfare consequences of technical advance, particularly during the initial stages of the product cycle of new product[1]. Rather than merely stating once again the suspicion that we may, therefore, be mismeasuring growth, the approach taken here is a constructive one and offers a pragmatic way of dealing with the problem. True, its application requires both the gathering of more extensive and detailed data (primarily on the quality dimensions of products and on market shares) and the use of more advanced econometric techniques (e.g. discrete choice models). However, it is my belief that both tasks are well within the realm of the feasible, and that it is increasingly important to do this if the presumed link between aggregate economic measures (such as GNP) and "economic well-being" is to be preserved.

II. THE ASSESSMENT OF PRODUCT INNOVATIONS

In view of the fact that the "output" of innovative activities does not present itself in countable units of any sort, innovations can only be quantified directly in value terms, i.e. in terms of their impact upon social welfare. Thus, the question "how much innovation has taken place" in a certain field over a certain period of time, can only be interpreted as asking "how much additional consumer and producer surplus was generated by technical advance in that field and time"[2]. If the innovation takes the form of cost reductions in the production of given products, then the assessment of its value is conceptually straightforward, involving the displacement of cost functions along a fixed demand schedule (see, e.g. Griliches, 1958). On the other hand, if the innovation consists of the introduction of new products or changes in the quality of existing ones, then its value to consumers cannot be represented simply as a cost saving but requires, instead, a more elaborate framework.

The methodology for the assessment of *product* innovations put forward in Trajtenberg (1989, 1990b) draws primarily from the "characteristics approach" to demand theory and from the econometrics of discrete choice models[3]. The basic idea is as follows: 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 and price. Product innovation can then be thought of in terms of change over time in the range of available products, both in the sense that new brands appear and that there are improvements in the quality of existing products. Applying discrete choice 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 social value of the innovations occurring between two periods can then be calculated as the benefits of having the latest choice range rather than the previous one, in terms of the ensuing increments in consumer and producer surplus.

To fix our ideas, define $s_i \equiv (z_i, p_i)$, where p_i stands for the price and z_i for the vector of characteristics of product (brand) i in a given product class. The choice set from which the consumer selects the most preferred brand in period t is thus $S_t = (s_{1t}, s_{2t}, \ldots, s_{nt})$. In this setting, product innovation is taken to mean simply that changes occur over time in the vectors 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)[4] and assuming that the changes in S are discrete, the magnitude of innovations occuring from t-1 to t will be measured by,

$$\Delta W = W(S_t) - W(S_{t-1}) \qquad [1]$$

The main problem, then, is to find a suitable specification for the function $W(S_t)$, and be able to retrieve its parameters from observable data. In principle, this is to be done by integrating over the underlying demand function, the features which would depend, inter alia, on whether the choice set is continuous or discrete. In the context of technologically-progressive products, it seems appropriate to characterise those sets as discrete: R&D constitutes a fixed cost, and hence innovative sectors typically exhibit in equilibrium a finite and not-too-large number of differentiated products[5]. Discreteness is assumed also in the sense that consumers purchase a single unit of a single product, thus making the choice problem exclusively qualitative (the analysis can be easily extended to accommodate cases of discrete/continuous choice as well). Those assumptions allow one to resort to discrete choice models, and to make use of the associated welfare analysis (see McFadden, 1981).

The basic hypothesis underlying discrete choice is that consumers maximise a random utility function, $Ui = U(z_i, m; h) + \epsilon_i$, subject to $s_i \in S$, and $p_i + m = y$, where m denotes composite "outside" good, h a vector of observable attributes of the individual, and ϵ_i an i.i.d. random disturbance. Assuming that ϵ_i conforms to the type I extreme-value (or Weibull) distribution, the maximisation of U_i leads to probabilistic demand functions of the form.

$$\pi_i = \exp V_i \;/\; \sum_j^n \exp V_j \;,\; i = 1, \ldots, n \qquad [2]$$

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 (MNL). It is easy to prove that the n equations in [2] 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, i.e. the utility function is specialised to be additive separable in the group products (those in S) and in the outside good m, rendering $V_i = \alpha(y - p_i) + \phi(z_i,h)$, where α stands for the (constant) marginal utility of income. Substituting in [2]

$$\pi_i = \exp[-\alpha p_i + \phi(z_i; h)] \,/\, \sum_{j=1}^n \exp[-\alpha p_j + \phi(z_j; h)] \qquad [3]$$

The identify of hicksian and marshallian demand functions in [3] allows one to obtain the surplus function W(S,h) simply by integrating under these demand functions, the integral being path independent. Ignoring the constant of integration, the result is [6],

$$W(S, h) = \ln\left[\sum_{i=1}^n \exp(-\alpha p_i + \phi(z_i; h))\right] / \alpha \qquad [4]$$

This surplus function is, then, the key element in assessing the value of product innovations: after estimating the choice probabilities in [3], one can retrieve the parameters of [4], and compute the benefits from innovations occurring between any two adjacent years, as in [1]. One of the problems that may arise in estimating [3] is that prices and characteristics are typically highly correlated, and the ensuing multicollinearity makes it very difficult to obtain reliable estimates of the parameters. The solution put forward in Trajtenberg (1990b), involves the use of residuals from estimated hedonic price functions in a multi-equation context (see the Appendix). Thus, hedonic price functions may still have an important (albeit more indirect) role to play in assessing product innovations, even though they may not be sufficient by themselves as indicators of quality changes.

III. THE CONSTRUCTION OF QUALITY-ADJUSTED PRICE INDICES ON THE BASIS OF ΔW

Suppose that we have estimated the multinomial logit model as in [3] and computed the yearly gains ΔW_t from [4] and [1]; the question now is how to construct on the basis of those ΔW values a "real" price index that would faithfully reflect the innovations that had occurred from one period to the next. The procedure suggested here involves relying on the expenditure function dual to [4], and using it to compute the hypothetical price change that would have resulted in the same welfare effect (measured by ΔW) as the innovations that actually took place[7]. In this sense, the proposed index belongs to the class of "cost-of-living" — or Konüs — indices (see Diewert, 1987).

As shown in Trajtenberg (1990a), one can write [1] as

$$\Delta W_t = e(p_{t-1}, Z_{t-1}) - e(p_t, Z_t) \qquad [5]$$

where p stands for the *vector* of prices of all brands in S, Z for the matrix of their attributes, and $e(\cdot)$ is the expenditure function dual to [4] (we omit the reference utility level, V^o, since $e(\cdot)$ is in this case linear additive in V^o). Thus, after having obtained estimates of ΔW_t using the method outlined in Section II, one can construct two different price indices that would reflect the quality changes embedded in S_t *vis-à-vis* S_{t-1}. The first requires that we solve for δ_t out of,

$$\Delta W_t = e[p_{t-1}, Z_{t-1}] - e[(1-\delta_t)\cdot p_{t-1}, Z_{t-1}] \qquad [6]$$

That is, δ_t is the hypothetical average price reduction that would make consumers indifferent between $[(1-\delta_t)\cdot p_{t-1}, Z_{t-1}]$ and $[p_t, Z_t]$. From a computational viewpoint, the values of δ_t can be obtained from [6] with methods of iterative search. However, using a more restrictive notion of "average price change" allows one to compute δ_t in a much simpler way: the price of each brand at time t can always be written as $p_{it} = \bar{p}_t + \Delta p_{it}$ where $\bar{p}_t$ is the average across brands; now, suppose that the changes in prices take the from,

$p_{it} = (1 - \delta_t)\bar{p}_{t-1} + \Delta p_{it-1}$, that is, the distribution of prices moves leftwards by a factor of $(1 - \delta_t)$, but the variance remains the same. It is easy to show that in such a case [6] simplifies to,

$$\Delta W_t = \delta_t \bar{p}_{t-1} \qquad [7]$$

and hence δ_t can be computed immediately as the ratio

$\Delta W_t / \bar{p}_{t-1}$.

Having arrived at the series $\{\delta_t\}$, a quality-adjusted price index can then be computed simply as $I^1_t/I^1_{t-1} = (1 - \delta_t)$, with $I^1_0 = 100$ (the superscript is used to distinguish between the two alternative indices). The second price index obtains by solving for φ_t from,

$$\Delta W_t = e[1+\varphi_t)\cdot p_t, Z_t] - e[p_t, Z_t] \qquad [8]$$

where $(1 + \varphi_t) . \bar{p}_t$ can be interpreted as the reservation price for the innovations embedded in S_t. Thus, if prices of the improved products had been $(1 + \varphi_t)$ times higher than actual prices, the implied percentage price reduction of $\delta'_t = \varphi_t/(1 + \varphi_t)$ would be welfare-equivalent to the quality improvements that took place. Assuming again that the price change consists just of a shift of the mean price, φ_t would be obtained simply from

$$(1 + \varphi_t) = (\Delta W_t + \bar{p}_t) / \bar{p}_t \Rightarrow \varphi_t = \Delta W_t / \bar{p}_t,$$

implying a percentage price reduction of,

$$\delta'_t \equiv \varphi_t / (1 + \varphi_t) = \Delta W_t / (\Delta W_t + \bar{p}_t) \qquad [9]$$

and the associated price index would be

$$I^2_t / I^2_{t-1} = 1/(1+\varphi_t) = (1-\delta'_t).$$

Clearly, the two indices are equally legitimate and have equally well defined interpretations. However, a technical difference between them precludes the use of the first index when innovations are "drastic" i.e. when $\Delta W_t > \bar{p}_{t-1}$ (and hence $\delta_t > 1$). The meaning of such a case is that consumer would prefer to have the new range of products and pay their full price, even if the products that existed in period t-1 were to be sold at *zero* price. The index I^1_t is then undefined, since $\delta_t > 1$ would imply a negative value for it. On the other hand, if ΔW_t is larger than $\bar{p}_t$ and hence $\varphi_t > 1$, the second index is still well defined: the hypothetical reservation prices that would make the consumer indifferent between the improved (but more expensive) products and the older range can be as high as necessary. Thus, if innovations in a given field are very substantial there is no choice but to use the second index only. On the other hand, if a field consistently displays just incremental innovations, it may be worth considering some sort of average between the two indices[8] and/or using the average of the mean price in the two periods to compute either index.

IV. ΔW-BASED INDICES VERSUS HEDONIC PRICES

Having thus put forward price indices based on the measures ΔW, it is important to step back and ask whether one really needs the rather complicated method outlined above in order to obtain reasonably good deflators for rapidly changing goods; could it not be that indices based on hedonic price regressions would do the job just as well?[9]. It is important to note that this question is in fact equivalent to asking whether or not there is a meaningful distinction between process and product innovations: as I shall argue below, the use of hedonic price indices (in lieu of ΔW-based indices) is justified only when "quality" is merely a redefinition of quantity, and hence "product innovation" is just process innovation in disguise.

Quality-adjusted price indices in the "repackaging" case

The answer to the question just posed can essentially be found in the classic work of Fisher and Shell (1972), on the theory of price indices (even though the question was not quite put in those terms there): hedonic-based price indices (or a price/performance ratio if "quality" is one-dimensional) would suffice to account for quality change only in the "repackaging" case. If the choice set consists of one good only (say, good 1), and "quality" can be fully accounted for with one parameter θ, "repackaging" implies that the corresponding argument in the utility function is just θx_1. That is, θ is sort of the amount of services provided by the good, and hence "quality change" (meaning $\theta_t > \theta_{t-1}$) amounts essentially to a redefinition of units. In such a case one can define a "price-performance" ratio p_1/θ such that, for any θ,

$$e(V^o, p_1, p_2, ..., p_n; \theta) = e(V^o, p_1/\theta, p_2, ..., p_n) \qquad [10]$$

and the implied "quality adjusted" price index would simply be $(p_{1t}/\theta_t) / (p_{1t-1}/\theta_{t-1})$. Thus, if θ were easily observable (as when it is indeed just a matter of redefining units), accounting for "quality change" would be a very simple matter. It is important to note that in such a case the distinction between process and product innovations all but vanishes (as does the quality-quantity dichotomy): defining the relevant price as p_1/θ, rather than just p_1, it is clear that technical change that brings about a reduction in costs leading in turn to a decrease in the unadjusted price p_1 (i.e. a *process* innovation) is exactly equivalent to a "product" innovation that results in the enhancement of θ.

When the choice set consists of $n > 1$ brands, "repackaging" implies that the corresponding branch of the utility function takes the form $U(\Sigma^n_i \theta_i x_i)$. Clearly, if $U(\cdot)$ is common to all consumers, then in order for more than one brand to be purchased in a cross-section it must be that $p_i/p_j = \theta_i/\theta_j$. Denoting by $\tilde{p}_o$ the quality-adjusted price of the reference variety, one can always write $p_i = \tilde{p}_0\ \theta_i$. Furthermore, if θ_i is not one-dimensional, but depends upon a vector of attributes $\bar{z}_i$, then (see for example Deaton and Muellbauer, 1980),

$$\log p_i = \log \tilde{p}_0 + \log \theta(\bar{z}_i) \qquad [11]$$

which is one of the forms that estimated hedonic price functions commonly take. In a two-year panel, for example, the term $\log \tilde{p}_0$ would be obtained as the coefficient of a time dummy variable, and can be taken as a sufficient price index in the same of [10] above (i.e. $\tilde{p}_t$ would be the equivalent in this context of the price-performance ratio p_t/θ_t)[10]. To insist, the point is that the hedonic price function by itself is able to account for more than one attribute in computing price indices, but such indices can serve as sufficient indicators of "quality" change only in the highly restrictive context of the repackaging case.

Assessing the performance of hedonic-based price indices

One of the intended uses of the price index based on the mesures ΔW_t's, is for it to serve as a test criterion for other indices and, in particular, for hedonic-based indices. This is essentially an empirical task; however, in order to have a better feel for what the empirical

Figure 1
ALTERNATIVE EFFECTS OF INNOVATION ON HEDONIC PRICE FUNCTIONS

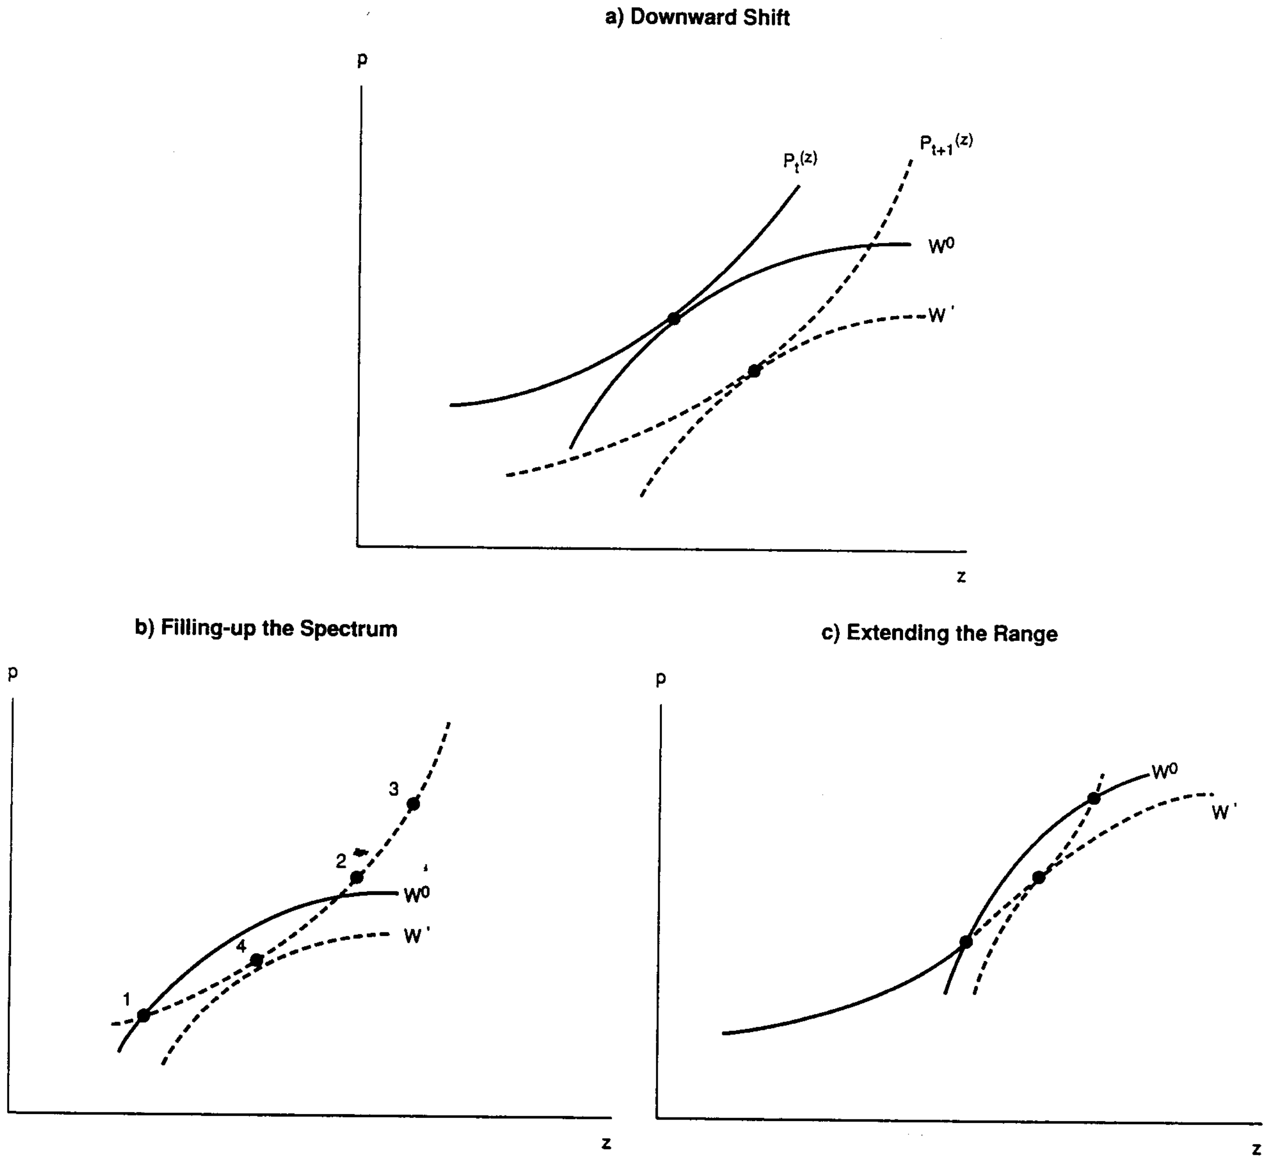

comparisons may entail, it is worth examining first in a heuristic manner how hedonic price indices are likely to perform in various stylized situations.

Quite clearly, if a price index is to account faithfully for quality change, it should measure the "distance" (in money metric) between the attainable utility level before and after the innovation. Consider the case where innovations occur so that there is a downward shift in the hedonic function, as shown in Figure 1.a. In the simplist possible situation (abstracting from discreteness, aggregation problems and income effects), the distance between the indifference curves labelled W^o and W' would be a good approximation to the monetized welfare gains associated with the innovations that induced the shift in the hedonic function. Thus, the coefficient of a time dummy in a hedonic regression pooling adjacent years will accurately measure those gains, and the resulting quality-adjusted price index could thus be taken as a faithful indicator of the changes which have occurred.

By contrast, now consider Figure 1.b: innovation in this case consists of the filling-up of the spectrum of products, e.g. in the base period only brands 1, 2 and 3 exist, but in the second period products such as 4 and 5 are added to the choice set. As the Figure suggests, in this case there will be no change whatsoever in the hedonic price function, and hence a price index based on it will altogether fail to register the occurrence of the innovations. On the other hand, a measure such as ΔW will certainly be positive, and could in fact be quite large. Figure 1.c illustrates a similar situation, except that innovation takes there the form of extending the range of available products, i.e. higher quality brands are introduced, priced (approximately) in accordance to the base hedonic function. Again, this type of innovations will leave no trace in the price index, whereas the actual gains may be substantial. Moreover, in the last two cases ΔW may be positive, and at the same time the hedonic-adjusted price index may actually *increase*, suggesting the occurrence of *negative* innovations (for an empirical finding of that nature, see Alexander and Mitchel, 1985).

It should be clear that the three stylized types of changes described are equally legitimate as instances of product innovations, and a priori it would apepar that they are equally likely. However, there is some evidence to the effect that the last two are more prevalent during the initial stages of the "product cycle", whereas the first tends to occur later on, in the wake of widespread imitation and price competition. If so, adjusting for quality changes with the aid of hedonic price functions may be a reasonable first approximation for well-established sectors, but not for tracing the emergence of new ones. As shown in Trajtenberg (1989), the bulk of the gains from innovation in the case of CT scanners occurred very early-on in the development of the field. If these results are typical then the picture painted by hedonic-based price indices may systematically understate a great deal of the "action" occurring in the technologically progressive sectors of the economy.

ΔW-based indices versus hedonic price indices in the case of CT scanners

Having measured the welfare gains from innovation using the approach of Section II in one particular case, namely CT Scanners, it is now possible to assess how far off-the-mark a hedonic-based price index would have been in this case and, thus, get a sense for the extent to which prevalent indices might be presenting a distorted image of the dynamic performance of high-tech sectors (a full discussion of these comparisons is in Trajtenberg, 1990a).

First, a few words about the innovation studied: Computed Tomography (CT) is a highly sophisticated diagnostic technology that produces cross-sectional pictures of internal organs of the body, using a special configuration of x-rays, detectors and computers. It has been hailed as one of the most remarkable medical innovations of recent times, comparable to the invention of radiography (the 1979 Nobel Prize in Medicine was awarded to the two scientists who pioneered the system). Originally developed at the British firm EMI in the early seventies, CT soon attracted some twenty other firms worldwide, and the fierce competition that ensued brought about a breathtaking pace of technical advance. The diffusion of the new systems proceeded very fast as well: first introduced in the US in 1973, by 1985 almost 60 per cent of hospitals (with more than 100 beds) had at least one system installed. The pace of innovation in CT subsided in the mid-eighties, as the technology matured and ceded its dominant place to new technological developments, particularly to Magnetic Resonance Imaging.

Table 1 shows the unadjusted index (i.e. one that just tracks the changes in the nominal average price of CT), a hedonic-based index computed on the basis of a weighted pooled hedonic regression (double log), and the ΔW-based index. The differences between them are striking: if no correction is made at all, one would conclude that CT Scanners were about 2.5 times more expensive in 1982 than a decade earlier and, hence, that we are significantly worse off on that account. Using the

Table 1. **Comparing Various Indices for CT Scanners**

	Unadjusted[a] Index		Hedonic[b]		ΔW-based	
1973	10 000		10 000		10 000	
1974	11 940		10 770		800	
1975	12 000		6 130		240	
1976	14 450		4 600		31	
1977	17 450	*100*	3 850	*100*	15	*100*
1978	15 940	*91*	3 050	*79*	13	*87*
1979	16 610	*95*	2 780	*72*	11	*73*
1980	20 190	*116*	2 840	*74*	10	*67*
1981	24 840	*142*	3 020	*78*	8	*53*
1982	25 940	*149*	2 730	*71*	7	*47*

a) $\bar{p}_t/\bar{p}_{73}$, where $\bar{p}_t$ is the weighted mean price in year t.

b) The hedonic index is based on a weighted pooled hedonic regression, using annual sales as weights.

hedonic technique significantly alters that initial assessment: the quality-adjusted hedonic index goes down from 100 to 27, implying an average price *decrease* of 13 per cent. Still, that is a far cry from the actual pace of technical advance that took place in CT: the ΔW-based index goes down from *10 000 to 7*, implying a staggering real price reduction of 55 per cent per year on average! It is important to note that, if one were to start the measurements, say in 1977, the extent of the discrepancies would be greatly attenuated, as can be inferred from the figures in italics in Table 5. However, rather than finding comfort in those figures, they should serve as a warning, i.e. the hedonic method may not do so badly when it comes to technologically mature industries, but it seems to be completely off mark early on, when it is needed the most.

V. CONCLUDING REMARKS

This paper offers a way of tackling the long-standing problem of quality change, that is both do-able and well grounded in welfare economics. One of the features of the method that I would like to stress is the fact that, contrary to common practice in economics, here we go from welfare measures to price indices rather than the other way round. This reflects, not a peculiarity of the method, but a basic characteristic of what we are trying to apprehend: product innovations "occur" in dimensions other than the traditional economic variables of prices and quantities and, hence, the assessment of their value (the only way of quantifying innovations) has to take place in the space spanned by those dimensions. One of the implications of this shift in focus is that *real* prices lose their classic role as a "primitive" notion, becoming instead one of the many possible derivedconstructs. Any method that purports to capture the impact of product innovations on real economic variables has to take this into consideration, rather than trying to force observed prices to "confess" something that they could not possibly know.

A comment about accounting for the infamous "productivity slowdown", or addressing the apparent paradox of explosive technical change, on the one hand, and "low" conventionally measured growth rates, on the other. As has been repeatedly pointed out (see, e.g. Baily and Gordon, 1988), it is not enough to uncover (yet another) source of mismeasurement: one has to show, first, that we are worse at measuring things now than before and, second, that the problem is widespread and substantial enough to make a real dent in the growth statistics. On both accounts, all I can offer at this stage is an intuitive feel for these issues, rather than hard evidence. First, it is often claimed that technical change has been increasingly taking the form of product rather than process innovation, particularly so since the advent of electronics. If this is the case (and I tend to believe so), then conventional price indices will, indeed, be less and less capable of capturing technical advance, and the gap between "real" and perceived growth will increase over time. Clearly, though, one would have to offer convincing quantitative evidence for this alleged change in the "mix" of technical advance in order to make the argument stick[11]. Second, the evidence presented above regarding the extent of mismeasurement refers, as said, to just one case study. Again, I believe that the qualitative phenomena uncovered in that study (particularly the fact that the largest gains from innovation occur at the very beginning, when the mismeasurement problem is most acute), may hold for a great many high-tech products. However, many more studies of this sort will be needed if that belief is to be substantiated.

Finally, it should be clear that whatever the quantitative impact of product innovations their mismeasurement could not possibly account for short-term fluctuations in growth rates. What we are referring to here is to long-term consequences, to be thought of in terms of decades rather than years; in that respect, at least, mismeasurement cannot serve as an excuse for macro-economic mismanagement.

Appendix

INCORPORATING THE HEDONIC PRICE FUNCTION INTO THE MNL MODEL

The discussion in Section II 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:

$$p_i = p(z_i) + \tilde{p}_i \qquad \text{[A.1]}$$

where $p(z_i)$ is a systematic component, and $\tilde{p}_i$ an i.i.d. error term (the "residual price"). The existence of such a relationship poses a serious multi-collinearity problem in the estimation of the choice probabilities of equation [3]: since both price and the vector 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.

Substituting [A.1] for p_i in V_i, and ignoring y and h,

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

or, defining $V^n(z_i) \equiv \phi(z_i) - \alpha p(z_i)$,

$$V_i = V^n(z_i) - \alpha\,\tilde{p}_i. \qquad \text{[A.2]}$$

where the term $V^n(z_i)$ can be interpreted as the "net utility" conferred by product i (that is, net of the *expected* price of the product). Thus, the behaviour of consumers is now seen to depend upon z_i and $\tilde{p}_i$, rather than upon z_i and p_i. In other words, given the existence of a hedonic function, p_i largely replicates the information already conveyed by z_i. Therefore, only the component of price that is orthogonal to z_i, $\tilde{p}_i$, can affect behaviour, qualifying as a legitimate explanatory variable in the choice model.

In order for [A.2] to offer an actual solution to the multi-collinearity problem, $V^n(z)$ needs to be given more structure. This is easily done with the aid of the following straightforward proposition: $V^n(z)$ can be closely approximated by the sum of a linear and a quadratic form, provided only that it has an interior maximum. More formally, $V^n(z) \equiv z'\beta + z'Gz$, where G is a symmetric matrix, if there is a $z^* > 0$, such that: $z^* = \arg\max V^n(z)$. When this is so, the approximation $(z'\beta + z'Gz)$ obtains readily from a second-order Taylor expansion about z^*. Normally we would expect $\phi(z)$ to be concave (or quasi-concave), and the hedonic function to be convex (as has been found in many empirical studies), in which case $V^n(z)$ would necessarily meet the required condition. The suggested specification of the "net utility" leads to the following model,

$$\begin{aligned} \pi_i &= \exp V_i \,/\, \Sigma \exp V_j, \\ V_i &= z'_i\beta + z'_i G z_i - \alpha\tilde{p}_i, \qquad \text{[A.3]} \\ \tilde{p}_i &= p_i - p(z_i) \end{aligned}$$

which can be estimated either simultaneously or using a two-stage procedure (i.e. first estimate the hedonic price function and compute the residuals p_i; second, enter $\tilde{p}_i$ as an independent variable in π_i and estimate the MNL model). If each choice set (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, then the latter method is appropriate; otherwise a framework of simultaneous equations is required.

NOTES

1. This is above and beyond the problem of the long delays in incorporating new goods in the computations of, say, the CPI. That is, even if new goods were incorporated right away in existing price indices, the problem of mismeasurement will remain.
2. Alternative measures such as patent counts, counts of "important innovations", rates of change of attributes, etc. could at best play the role of proxies, and their accuracy as such can be judged only by relating them to the values measured themselves (see e.g. Trajtenberg, 1990c).
3. Here I just sketch the essence of the approach; for a full discussion see Trajtenberg (1990b), Ch. 1.
4. 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.
5. Rosen (1974) analysed the continuous case and laid out the basis for the econometric estimation of such a system. However, the implementation of Rosen's approach poses serious difficulties, as discussed in detail by Epple (1987).
6. Note that, in dividing by α, the function $W(\cdot)$ is being normalised so as to express it in money terms. Notice also that $-\partial W/\partial p_i = \pi_i$ and, hence, [4] is, indeed, the correct solution.
7. The full derivation of the index is given in Trajtenberg (1990a); here I just provide an overview.
8. It can be shown that $\delta'_t \leq \delta_t$, i.e. the first index, will always show a larger "quality-adjusted" price reduction.
9. The hedonic method is certainly much simpler, its data requirements are more modest, it is well-known and commands wide acceptance. Moreover, since the early sixties, various government agencies have been considering using it in the construction of price indices and, in fact, the BEA recently started to compute an hedonic index for computer, in collaboration with IBM. Thus, if both methods were roughly equivalent, surely one would not hesitate to opt for the hedonic approach.
10. Even this simple case is subject to several qualifications. In particular, if the budget constraint in attributes space is non-linear (as it is most likely to be), then the estimation of [11] involves what can be construed as errors of aggregation.
11. See Scherer (1984) for some tentative evidence pointing in that direction.

REFERENCES

Alexander, A.J., and Mitchell, B.M., (1985), "Measuring Technological Change of Heterogeneous Products", *Technological Forecasting and Social Change*, 27(2/3):161-195.

Baily, M.N., and Gordon R.J., (1988), "The Productivity Slowdown, Measurement Issues, and the Explosion of Computer Power", Brookings Papers on Economic Activity, 2, 347-431.

Deaton, A., and Muellbauer, J., (1980), *Economics and Consumer Behavior*, Cambridge: Cambridge University Press.

Diewert, E.W., (1987), "Index Numbers", *The New Palgrave: A Dictionary of Economics*, Vol. 2, J. Eatwell, M. Milgate and P. Newman (eds.), London, Macmillan Press, 767-780.

Epple, D., (1987), "Hedonic Prices and Implicit Markets: Estimating Demand and Supply Functions for Differentiated Products", *Journal of Political Economy*, 95(1):59-80.

Fisher, F.M., and Shell K., (1972), *The Economic Theory of Price Indices: Two Essays on the Effects of Taste, Quality, and Technological Change*, New York: Academic Press.

Griliches, Z., (1958), "Research Costs and Social Returns: Hybrid Corn and Reelated Innovations", *Journal of Political Economy*, 66(5):419-431.

Griliches, Z., (ed.), (1971), *Price Indexes and Quality Change*, Cambridge, Massachusetts: Harvard University Press.

McFadden, D., (1981), "Econometric Models of Probabilistic Choice", in C. Manski and D. McFadden, (eds.), *Structural Analysis of Discrete Data with Econometric Applications*, Cambridge, Massachusetts: The M.I.T. Press.

Rosen, S., (1974), "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition", *Journal of Political Economy*, 82(1):34-55.

Scherer, F.M., (1984), "Using Linked Patent and R&D Data to Measure Interindustry Technology Flows", in Z. Griliches, (ed.), *R&D, Patents and Productivity*, Chicago: Chicago University Press.

Trajtenberg, M., (1989), "The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners", *Journal of Political Economy*, April, 97(2), 444-479.

Trajtenberg, M., (1990a), "Product Innovations, Price Indices, and the (Mis)Measurement of Economic Performance", Harvard Business School, Working Paper 90-040.

Trajtenberg, M., (1990b), *Economic Analysis of Product Innovation — The Case of CT Scanners*, Cambridge, Mass: Harvard University Press.

Trajtenberg, M., (1990c), "A Penny for your Quotes: Patent Citations and the Value of Innovations", *The Rand Journal of Economics*.

Triplett, J.E., (1975), "The Measurement of Inflation: A Survey of Research on the Accuracy of Price Indexes", in P.H. Earl, (ed.), *Analysis of Inflation*, Lexington, Massachusettes: Lexington Books.