Robert J. Gordon is Stanley G. Harris Professor of the Social Sciences at Northwestern University and is a research associate of the National Bureau of Economic Research.

Commentary

Robert J. Gordon

Charles Hulten's original and challenging article reaches two original conclusions and stresses a third conclusion based on the work of others. First, he introduces a new and unconventional approach to the theory of quality change, based on his new parameter μ, defined as the elasticity of the marginal cost of producing a good with respect to an improvement in its quality. He concludes that this "neglected cost dimension" may contribute a significant upward bias to the CPI in general, and his interpretation implies that the hedonic regression technique is fundamentally flawed as a tool of quality adjustment. Second, and unrelated to the first point, he argues that the link method presently used to adjust the CPI for quality change may create a downward bias in the CPI large enough to "dominate the quality-cost bias." A third point, but one that is not original to this paper, is that the CPI currently embodies a large adjustment for quality.

Hulten's first and second conclusions—that there are major sources of both upward and downward bias in the CPI and that these biases were ignored by the Boskin Commission—deserve careful consideration because they could be important. In these comments, however, I shall conclude that, to the contrary, the sources of bias he discusses are not new; they are already incorporated into the estimates presented in the Commission's report.

Let me defer the quantification of CPI bias in order to focus first on Hulten's theoretical contribution. When closely examined, Hulten's "μ factor" does not turn out to represent a novel contribution; rather, it is a rediscovery that the elasticity of marginal costs of quality characteristics, as measured directly by cost comparisons or indirectly in hedonic regressions, does not have to equal unity. An elasticity of unity has never received any special attention in the literature on the hedonic regression model, and for a very good reason: It has no theoretical importance, nor is there any empirical reason for an elasticity of unity to be a standard empirical result. Thus, Hulten's theoretical framework does not contribute a new source of bias in the CPI for quality adjustments based on the hedonic or direct-cost-comparison methods.

CAN THE COST ELASTICITY DIFFER FROM UNITY?

Hulten follows Triplett in treating quality differences between different models of the same good as embodying different quantities of what he calls "efficiency units" and what Triplett calls "characteristics." In this view, a good, Y, containing $(1+\theta)$ as many characteristics—e.g., computer MIPS—as another good, X, can be added together with units of X. As Hulten states, "this approach amounts to converting differences in the quality of goods into an equivalent difference in quantity." I prefer to refer to X and Y not as "goods" but rather as "models" of the same good—e.g., automobiles or computers. The "repackaging" assumption required to make direct comparisons between X and Y is more plausible for models of the same good than for different goods (e.g., houses and medical insurance) that embody totally different dimensions of characteristics.[1]

There are two problems with Hulten's analysis. First, he does not distinguish between the slope of the supply curve of characteristics and shifts in that curve made possible by technological change. Second, he does not interpret correctly the implication of a cost elasticity of less than unity. Let us consider first a world of con-

[1] Hulten's own discussion of "New Goods" recognizes that the repackaging assumption is of no use when one attempts to measure the value of totally new goods, such as the invention of television or of the VCR.

stant technology, in which a variety of products and a variety of models of each product are produced. Quality differences among these models are costly to produce. A large refrigerator costs more than a small one, and a 200 MHZ computer chip costs more than a 100 MHZ chip. In every type of commodity, from string beans to housing to automobiles to health insurance, we find many different varieties of the same good coexisting in the marketplace.

The appropriate framework for analysis is a simple diagram of supply and demand, in which the nominal price of a characteristic (or efficiency unit) is plotted on the vertical axis and the number of characteristics is plotted on the horizontal axis. The supply curve of characteristics is upward sloping, with a slope that expresses the marginal cost of producing models containing additional characteristics. The hedonic regression approach is one method to estimate the slope of the supply curve of characteristics; the BLS method of estimating production costs directly is a second method of estimating the slope. There is no expectation that the elasticity of this slope would be unity—i.e., that model *Y*, containing double the number of characteristics as model *X*, would cost twice as much. For instance, a computer operating at a speed of 200 MHZ might cost 50 percent rather than 100 percent more than another model operating at a speed of 100 MHZ.[2]

INTERPRETATION OF THE HEDONIC REGRESSION MODEL

Thus far I agree with Hulten that his μ coefficient, the elasticity of cost to changes in the quantity of characteristics, can differ from unity. At that point, however, Hulten diverges from the standard interpretation of the hedonic regression model. To be concrete, let us assume that there are two models on the market and that the second model, *Y*, incorporates $(1+\theta)$ times as many characteristics as the first model, *X*. We adopt Hulten's own numerical example, $\theta = 0.1$, and we assume that $\mu = 0.5$, so that model *Y* costs $(1+\mu\theta)$ or 1.05 as much as model *X*. How can the two models exist on the market at the same time? The larger model, *Y*, offers the characteristic in question at a price of $(1+\mu\theta) / (1+\theta)$ or 1.05/1.10 = 0.95 times the price per characteristic of model *X*.

Here is where Hulten goes astray. He states, "The result is that the old good, *X*, is cost-ineffective and therefore disappears from the market." When a 100 MHZ computer costs $1,000 and coexists on the market with a 200 MHZ computer costing $1,500, then the smaller model offers a price per MHZ of $10, while the larger model offers MHZ at a lower price of $7.50. If each unit of MHZ yielded the same marginal utility, then the smaller, more expensive, model would be driven from the market by the larger model, according to the "law of one price."

Yet Hulten does not recognize a basic aspect of markets in the real world: that different models of the same good, with different sizes and different prices per unit (e.g., price per MHZ of computers or price per ounce of dishwasher detergent or corn flakes), can and do exist on the market at the same time. Therefore, it must be true that the elasticity of utility to the quantity of characteristics is less than unity when the cost elasticity, μ, is less than unity.

In fact, this was the point of Rosen's famous 1974 article (which Hulten neither cites nor discusses) on the theory of the hedonic regression model—that the price-characteristic surface traced out by the hedonic function represents a series of tangency points between cost curves and indifference curves. Some purchasers forego the chance to buy the larger computer, despite the lower $7.50 price per MHZ, because for them the marginal utility schedule for units of MHZ declines even faster than the price. Another reason that packaged goods such as detergent can coexist on the market in different size packages at different prices per ounce is that the choice of a detergent is based on a number of factors: the detergent itself, the energy involved in carrying it home, the storage space needed, and the strength

[2] Here we ignore the role of multiple additional characteristics of the typical computer—e.g., memory size, type of CD-ROM drive, etc.

required to pick up the bottle and pour the detergent into the washing machine. So some people may face these tradeoffs and choose the smaller bottle even though it is more expensive per ounce.

This brings us to Hulten's Equation 1, intended to assess the accuracy of different methods of quality adjustment. In Hulten's formulation, the observed difference in price between one period and the next—when one model, $Y(t)$, replaces another model, $X(t-1)$, consists of the product of three terms: the ratio of differences in the quantity of characteristics, $(1+\theta)$, times the ratio of the difference in cost to the difference in characteristics $(1+\mu\theta) / (1+\theta)$, times the "pure price change" between the two periods $(1+\rho)$. Taking logs and designating the observed log price change between $X(t\text{-}1)$ and $Y(t)$ as α we have an alternative and simplified version of the relationship:

$$(1) \qquad \alpha = \mu\theta + \rho.$$

Contrary to Hulten's assertion, the overlap, hedonic, and direct cost methods all provide the correct answer: that the difference in utility yielded by the two models is $\mu\theta$ and is measured by their price differential on the market, which in turn is equal to their cost differential. To the contrary, Hulten claims that θ rather than $\mu\theta$, must be the correct answer when two models with different implicit prices of characteristics coexist on the market (that more than a single model exists at a given time is, of course, a necessary condition for a hedonic regression to be estimated).

We know in the real world that many models of a given product coexist, but also that models disappear. This occurs not because μ is less than unity, but because the supply curve shifts when new technology is introduced. The whole menu of choices between small and large models of, say, computers shifts in a favorable direction, with a decline in the price per characteristic of small and large models alike. With the new technology, μ may still be unchanged at a number like 0.5, the old menu of models is obsolete and disappears from the market, and small models incorporating the new technology coexist with large models once again.

MEASURING THE PURE PRICE CHANGE

Returning to our supply-demand diagram, until now we have suggested no reason for the supply curve to shift, but of course it can shift up or down, as it must if we are to register any pure price change (ρ) other than zero. Pure macroeconomic inflation would shift the supply curve up, and the technological change that has been typical in most electronics products, including computers and cellular phones, would shift the curve down. The hedonic regression technique is designed, of course, to disentangle *shifts* in the supply curve (pure price change) from *movements along* a single supply curve (changes in the quantity of characteristics in a particular good that alter both its utility and its cost).

Consider a shift in the supply curve in a diagram which plots the nominal price of a characteristic against the quantity of characteristics. There is a single unshifting downward-sloping demand curve. The supply curve shifts down because of technological change, and a new market equilibrium occurs at a lower price and larger quantity of characteristics, but the observed decline in price is not the pure change in price. By allowing us to estimate the slope of the supply curve, the hedonic technique enables us to add to the observed price decline the extra component of price decline attributable to the new technology. Consider an example of a typical model, in which we observe a new equilibrium involving a 10 percent price reduction, a 20 percent increase in the quantity of characteristics, and a supply curve slope (μ) once again at 0.5. We can use Equation 1 above to extract the pure price decline:

$$(2) \quad \rho = \alpha - \mu\theta = -0.10 - (0.5)0.2 = -0.20.$$

Hulten fails to recognize that the coexistence of models under these circumstances can occur only if there is a diminishing marginal utility of characteris-

tics for different potential purchasers (or if there is a joint product aspect, as in the above example of differently sized packages of detergent) and that models disappear not because the cost elasticity μ is less than unity, but because technological change can shift the supply curve of characteristics.

Although Hulten cites my treatment of this topic (Gordon, 1990, Chapter 2) as the source of the distinction between "proportional" and "nonproportional" quality change, I used those terms differently to distinguish, respectively, changes in quality caused by movements along the supply curve and shifts in the supply curve. In contrast, Hulten uses the word "proportional" to designate a μ elasticity of unity and "nonproportional" to designate μ less than unity, while he never considers the effects of a shift in the supply curve and cannot fit it into his framework as it stands. His example of "costless" quality change is a μ value of zero, whereas for me it is a downward shift in the supply curve that allows more characteristics to be produced at an unchanged price per characteristic.

HULTEN'S INTERPRETATION OF THE CPI BIAS

Hulten is wrong that "all methods used to construct CPI implicitly assume that μ is equal to one." By "all methods" he refers to the overlap, hedonic, and cost-comparison methods. But there is no special role of $\mu = 1$ in the hedonic regression model. Hulten states in his review of BLS methods that price hedonics provide an estimate of θ, whereas in fact the hedonic coefficients provide estimates of $\mu\theta$. The elasticity of cost and price to changes in the quantity of the characteristics entering into the hedonic regression equation is not unity; it is whatever it is, and there are plenty of examples in which this elasticity is below unity. Because existing methods are based on estimating the slope of a supply curve for a good when different-size models coexist, there is no bias in these methods. Thus Hulten's conjectural estimates of a major upward bias in the CPI due to a basic flaw in these methods are groundless, and we are back to the product-by-product examination recently debated by the Boskin Commission and its critics.

Hulten's second main point is that the BLS linking method creates a major downward bias in the CPI. This assertion is not based on any solid evidence but rather on a circumstantial "suspicion" growing out of the fact that the link method attributed virtually all of raw price changes to quality improvements in 1995. He computes a "link" bias by reducing the fraction of raw price change attributed to quality change in the link method (near unity) to the much lower fraction of 37 percent attributed to quality change by the direct quality adjustment method.

The much lower fraction of raw price change attributed to quality change in the direct quality adjustment method may reflect, as Hulten admits, "inherent quality differentials in the items handled with the various methods." Yet he proceeds to ignore this possibility by attributing the lower quality attribution to the link method, emerging with huge estimates of the downward bias.

There are two important reasons to disregard Hulten's "link bias" estimates. First, the evidence provided in the final version of the Moulton-Moses (1997) paper is that truncation of items with quality increases larger than 100 percent or decreases larger than 50 percent reduces the overall CPI correction for quality change from 1.65 to 0.30 percent. Obviously, most of the quality correction in the CPI is for items that are very different sizes, not what we mean by "quality" in the Boskin Commission or in other discussions of quality change.

Just as important, we have plenty of outside, independent evidence on the importance of any "link" bias. Because it collects prices monthly, the CPI is much more likely to allow items that may be temporarily out of stock to disappear, with the resulting need to link, than are alternative research results based on yearly or half-decadal comparisons of price and quality change. In my study of durable

goods prices based on the Sears catalog and *Consumer Reports*, only a small fraction of models of the typical product needed to be replaced from year to year, and there was no tendency for price changes to be concentrated only on new models rather than continuing models. Alternative studies of price and quality change, based entirely on information independent of CPI data sources, are consistent in finding an upward bias in the CPI, especially for appliances, radio, TV, computers, apparel (since 1985), drugs, and medical care. This result suggests that the linking bias cannot be major.

Hulten's third conclusion (listed first in his concluding section) is that the CPI "currently embodies a large adjustment for quality." This simply repeats the results of Moulton and Moses, but as they show, "large" becomes "small" (0.3 percent per year) when price increases of more than double or price decreases of more than half are deleted from the sample. These large changes, presumably due mainly to changes in package size, do not represent adjustments for quality change in the sense that most economists use that term.

CONCLUSION

To conclude, none of Hulten's three conclusions is valid, and none requires a reconsideration of the conclusions of the Boskin commission. The "neglected cost dimension" is not a neglected dimension at all, for it is Hulten who fails to consider why models with different prices per unit of characteristic coexist on the market, and it is he who fails to consider the distinction between movements along the supply curve and shifts in the supply curve. The "link bias" is not a bias at all, but rather a corollary of the fact that most of the quality adjustments in the CPI are made for differences in package size rather than actual differences in quality. And the conclusion that the CPI already makes large quality adjustments turns out to be untrue when large price differentials reflecting changing package size are discarded.

REFERENCE

Rosen, Sherwin, "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," Journal of Political Economy 82 (January-February 1974), pp. 34-55.

Hulten's Reply to Gordon

Robert Gordon has made many contributions to the field of price measurement. In the process, he has developed a particular way of looking at the problem. Unfortunately, he has tried to impose his framework on mine. The resulting criticisms of my article are essentially an attempt to pound my square pegs into his round holes.

Take, for example, Gordon's comparison of a 100 megahertz (MHZ) computer that costs $1,000 with a 200 MHZ model costing $1,500. He correctly concludes that both models may coexist in the market even though the 100 MHZ model delivers less computing power per dollar than its 200 MHZ competitor ($10 per MHZ versus $7.50 per MHZ). Gordon incorrectly believes that this example shows that my formulation, based on the cost-elasticity μ, is wrong: He believes that I do not recognize that different models of the same good coexist, without disappearing, even though they are simultaneously on sale at different prices per characteristic.

The problem with this criticism lies in Gordon's focus on price per characteristic. My model does not make use of this concept. Instead, I use the marginal rate of substitution (MRS) between goods of different qualities and, while the MRS is related to the characteristics of the different goods, it is far from the same thing. Gordon's problem is resolved in my framework by recognizing that the MRS between his two models is not necessarily equal to 200/100, even though this is the

ratio of the MHZ characteristic. Indeed, the MRS between the 100 MHZ and 200 MHZ models may well be 1.5 despite the greater difference in MHZ. In this case, the MRS is just equal to the ratio of marginal costs (or marginal rate of transformation) and both models survive in the market "at different prices per characteristic" as Gordon's example requires. On the other hand, if the MRS between the two models happens to be greater than 1.5, the 100 MHZ model would disappear because it is not cost effective. The importance of such corner solutions for the new-goods problem was recognized as early as Hicks (1940).

I also plead "not guilty" to the charge that I confuse shifts in the supply and movements along the supply curve. I assume that each model of a good is produced at a constant marginal cost and that quality improvements shift the cost function by the factor $\mu\theta$. Gordon's confusion arises, again, from his orientation to the characteristics of goods rather than to the goods themselves. There may, of course, be instances where the marginal cost of goods is rising, not constant as I assume. The parameter is not necessarily constant in this case, and competition among producers may cause it to rise to the point where it equals one (though this is not necessarily the only outcome). I allude to this possibility on pages 85-86 of my article. But surely no one would argue that the parameter μ is always equal to one, since this would mean that old varieties of goods never become obsolete.

Once it is admitted that μ can be less than one, it must also be admitted that the CPI may be biased on this account. This possibility is most apparent in the "cost of production" method used by BLS to handle quality change. In this procedure, producers are asked to estimate the direct production cost involved in bringing a new model to market. The change in cost is entirely attributed to quality change. This involves no bias if costs rise in exact proportion to (my definition of) quality—i.e., when μ is equal to one. However, when μ is less than one, quality rises more rapidly than costs and the BLS procedure understates this true increase. I believe that the other BLS procedures for handling quality change are also subject to a μ-bias, for the reasons given in the paper.

Gordon also challenges my contention that there might be a "link bias" operating in the opposite direction of the μ-bias. He states at the end of his comments that "The 'link bias' is not a bias at all . . ." This is surely too strong. Manufacturers sometimes wait until the introduction of a new model to introduce a planned price increase, and this delay in repricing gets attributed to "quality" by the link method. Gordon disputes this possibility, but the quote by Reinsdorf, Liegey, and Stewart (1996) cited in the text of my article suggests that the issue is a not a settled matter. This point is seconded by Triplett (1997), who has long been convinced that the link method over-adjusts for quality change. Gordon himself seems to suggest that there might be as much as a 0.30 percent bias, the residual left over after the Moulton-Moses truncation procedure. While I have reservations about the arbitrary truncation of the price data by Moulton and Moses (truncation is neither necessary nor sufficient to identify "true" quality change), the 0.30 percent estimate is not so very different from my "guesstimate" of 0.485 percent presented on page 94.

REFERENCES

Hicks, John R. "The Valuation of Social Income," *Economica* (Vol. 7, 1940), pp. 105-24.

Reinsdorf, Marshall, Paul Liegey, and Kenneth Stewart. "Avoiding Downward Bias from Linking in the U.S. Consumer Price Index," Bureau of Labor Statistics Working Paper 276, April 1996.

Tripplet, Jack E. "Measuring Consumption: The Post-1973 Slowdown and the Research Issues," in this issue of the *Review,* pp. 9-42