



WORKING PAPERS

RESEARCH DEPARTMENT

WORKING PAPER NO. 17-24
ACCOUNTING FOR GROWTH IN THE AGE OF THE
INTERNET: THE IMPORTANCE OF OUTPUT-SAVING
TECHNICAL CHANGE

Charles Hulten
University of Maryland and NBER

Leonard I. Nakamura
Research Department
Federal Reserve Bank of Philadelphia

July 2017

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/research-and-data/

Accounting for Growth in the Age of the Internet

The Importance of Output-Saving Technical Change

Charles Hulten, University of Maryland and NBER

and

Leonard Nakamura, Federal Reserve Bank of Philadelphia¹

July 2017

ABSTRACT

We extend the conventional Solow growth accounting model to allow innovation to affect consumer welfare directly. Our model is based on Lancaster's New Approach to Consumer Theory, in which there is a separate "consumption technology" that transforms the produced goods, measured at production cost, into utility. This technology can shift over time, allowing consumers to make more efficient use of each dollar of income. This is "output-saving" technical change, in contrast to the Solow TFP "resource-saving" technical change. One implication of our model is that living standards can rise at a greater rate than real GDP growth.

JEL No. O4, O3, E01

¹ The views expressed in this paper are solely those of the authors and should not be attributed to any organization with which they are affiliated, including the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Earlier versions of this paper were presented at the meeting of the Society for Economic Measurement, July 2016, in Thessalonica, Greece; the Federal Reserve *Day-Ahead Meeting on Productivity*, September 2016, in Cincinnati; and the 2017 ASSA meetings in Chicago. We thank John Fernald, Dan Sichel, Chad Syverson, and Roger Betancourt for their comments on earlier drafts. Any remaining errors are our own.

Accounting for Growth in the Age of the Internet

The Importance of Output-Saving Technical Change

I. Introduction

The digital revolution presents an interesting paradox. On the one hand, the revolution has transformed the economic landscape and has had a powerful impact on daily lives. On the other hand, real GDP growth has slowed in recent years despite the evident boom in information technology. Per capita GDP declined from its 1995-2006 rate of 2.2% to 1.3% from 2010 to 2015. Various explanations of this seeming paradox have been offered. This sharp and prolonged decline is seen by some as pointing to a more serious problem than a prolonged recession. Robert Gordon (2016), for example, has argued that the decline reflects the relatively anemic character of the digital revolution compared with earlier technological revolutions.

The disconnect between macroeconomic estimates of GDP and microeconomic analyses of innovation is reminiscent of the famous Solow (1987) paradox: “you can see the computer age everywhere but in the productivity statistics.” Solow’s remark was interpreted by many as a mild rebuke to those enthusiasts who over-hyped the impact of computers on productivity growth. It could also be interpreted as an observation about the failure of national statistics to capture the true impact of the computer revolution, a position championed by Alan Greenspan around the same time. We are now in a similar debate about the later stages of the digital revolution, again raising the question of whether there is less than meets the eye because there really is less of an impact on true GDP than enthusiasts imagine, or whether the impacts are concealed by the mismeasurement of real GDP.

We suggest that both may be true to some extent, and that the impact of the digital revolution cannot be properly assessed by focusing exclusively on how innovation affects the supply side of the economy. There is a growing conviction in the recent literature on growth accounting that at least part of the current round of innovation is more apt to affect the

consumer directly and thus does not appear even in a “correct” measure of GDP.² If this is true, a declining rate of real GDP growth may be consistent with the perception of a vibrant technological environment and the microeconomic analysis that supports it. And, if this is true, then a theoretical framework is needed that at least allows for an alternative non-GDP channel through which innovation operates. In this paper, we propose an extension of the conventional Solow production-function approach to growth analysis that permits consumers to make more efficient use of each dollar of income and allows for the possibility that living standards can be rising at a greater rate than is signaled by the growth rate of real GDP.

Our model is based on Lancaster’s *New Approach to Consumer Theory* (1966b), which we adapt to the growth accounting problem in a way consistent with the assumptions of the Solow model. In the Lancaster framework, there is a separate “consumption technology” that transforms the goods acquired from their producers, measured at production cost, into consumption “activities” or “commodities” that give utility based on their characteristics. We draw from the Lancaster model the idea that the utility function can shift over time as the consumption technology becomes more efficient. Efficiency can increase through costless improvements in product quality that allow better products to be purchased for the same amount of money or through an increase in effective information that allows the consumer to get more utility from a given amount of expenditure.³ These effects are separate from the resource-saving technical change of the Solow TFP model, and they can improve the standard of living even if those effects were static. They are, in effect, “output saving” technical change. They are particularly relevant for understanding the growth dynamics of the consumer-oriented digital age.

² See Ahmad and Schreyer (2016), Brynjolfsson and Oh (2013), Nakamura, Samuels, and Soloveichik (2016), and Varian (2009)

³ Search engines provide a concrete example of how the internet makes consumer choice more efficient. A consumer faced with a choice between different products can often find information about product specification and capabilities, the experience of other consumers, and explicit comparisons from rating organizations. Someone looking to buy a particular product can go on Amazon, for example, and see not only the price and availability of that item, but also a range of similar items that may turn out to be preferable. And this can be done while shopping in a store to see if a better price is available online, using a smartphone or other mobile device. GPS and traffic maps are often of great utility when traveling, as is immediate access to health information in times of need.

There is, however, an important empirical asymmetry between the two sides of the growth account: unlike GDP, utility is not directly observable. This leads us to reformulate our expanded growth model in terms of the associated expenditure and indirect utility functions, since they underpin much of the recent empirical literature on valuing the internet and other seemingly free goods. This literature approaches the problem from the price side rather than the quantity side, using compensating and equivalent variation concepts.

The paper then moves beyond costless technical change to allow for resource-costly innovation. Where costless innovation envisions technical progress on the supply side as a process based on inspiration, learning, and knowledge spillovers, the alternative view sees innovation as a matter of systematic investments in technology, including, for example, expenditures for R&D. These intangible inputs essentially “produce” innovation using resources that must be paid for one way or another. From a welfare standpoint, the gains from fully costed innovation are of a different nature; innovation of the costly sort does not convey the same benefits as the costless “Manna from Heaven” sort, be they output-saving or resource-saving technical change.

This paper does not attempt to resolve the debate over whether the benefits of economic growth are actually understated by the way GDP is measured. Rather, it attempts to extend the conventional growth accounting framework in such a way that the debate might, in time, be resolved. However, while the paper is essentially about theory, we do offer some brief comments on the growing body of empirical work on the boundaries of the digital economy to indicate both the current state of play and some of the orders of magnitude involved in including improvements in consumption technology in assessing the gains from innovation. This empirical research provides valuable information about the benefits of various aspects of the digital economy, and the goal of this paper is to provide an expanded conceptual framework into which the various contributions can be integrated.

To this end, the paper is largely illustrative and based on rudimentary mathematical modeling aimed at providing an intuitive foundation for thinking about the way innovation affects the economy and the welfare of the population. We have therefore included an

appendix setting out the geometry of our model, illustrating the implications of adding the consumption technology to the usual model of aggregate growth accounting.

II. Information, Utility, and Innovation

In their book *How Google Works*, Schmidt and Rosenberg (2014) argue that the world has entered an era in which “the internet has made information free, copious, and ubiquitous” to the consumer. This is one of the defining characteristics of what they call the “Internet Century.” At the same time, there are many other sources of economic growth that affect consumer well-being, and this raises the question of how to measure the contribution of “free, copious, and ubiquitous” information to GDP, and its relative importance compared to other factors. The question currently on many minds is whether the contribution is large enough to offset what appears to be a relative decline in real GDP, but there is the larger theoretical question of how, and whether, consumer information should be included in measured GDP.⁴

Where in the models of standard growth theory does an increase in information enter the analysis? This question has a long history, and the answer given by Hayek in 1945 is that it is largely absent. He argues that the standard model of economic theory is so closely wedded to the formal mathematics of optimization that it takes as given the information needed for the optimization process. Hayek frames his dissent from the prevailing theoretical orthodoxy in the following way:

“... the economic calculus which we [economists] have developed to solve this logical [optimization] problem, though an important step toward the solution of the economic problem of society, does not yet provide an answer to it. The reason for this is that the ‘data’ from which the economic calculus starts are never for the whole society ‘given’ to a single mind which could work out the implications, and can never be so given” [page 519].

⁴ The focus of this paper is on consumption technology and the increase in consumer welfare arising from the Information Revolution. This revolution also affects the production side of the economy, making processes more efficient and increasing the quality of many capital goods and intermediate inputs. However, when the supply-side effects of information technology are ignored or mismeasured, the result is usually suppressed into the productivity residual and measured indirectly (see Hulten (1992) for the case of capital goods). This is an important topic in its own right but is outside the scope of this paper.

No individual consumer can hope to possess all the information relevant to fully rational choice, or even to form preferences for items or circumstances never before encountered and not likely to be encountered in the future. In either case, the provision of “free, copious, and ubiquitous” information has ample opportunity to increase consumer utility.

Stigler (1961) proceeds along much the same conceptual path in his analysis of price dispersion and the prevalence of advertising expenditures. He takes academic economists to task for failing to recognize the importance of information:

“One should hardly have to tell academicians that information is a valuable resource: knowledge *is* power. And yet it occupies a slum dwelling in the town of economics. Mostly it is ignored: the best technology is assumed to be known; the relationship of commodities to consumer preferences is a datum. And one of the information-producing industries, advertising, is treated with a hostility that economists normally reserve for tariffs or monopolists” [page 213].

Both Hayek and Stigler emphasize that the link between consumer goods and consumer preferences cannot be treated as “a datum.” Five years later, Lancaster (1966b) went further in his *New Approach to Consumer Theory*, in which utility depends on the characteristics of goods consumed and not the goods themselves, and which introduced the concept of a “consumption technology.” He also proposed, in a companion paper (1966a), that this technology could change over time.

The goal of this paper is to incorporate these ideas into conventional growth accounting analysis in order to expand the discourse on how innovation can affect consumer welfare. We stress the term “conventional” since we do not take on the thorny problem of modeling decision-making under uncertainty and partial information. These subjects have received a lot of attention since the 1960s but have largely not found their way into conventional growth accounting, which has followed the neoclassical model developed by Robert Solow (1957), with a path-breaking extension by Jorgenson and Griliches (1967).⁵ This model intentionally abstracts from many hard real-world problems like imperfect information or uncertainty that

⁵ An account of the development of the Solow growth accounting model and the extensions that followed is given in Hulten (2001). The model is largely non-stochastic, but some randomness does creep into the model through fluctuations in demand, adjustment costs, and the discount and revaluation rates in the cost of capital variable. Information, in the form of R&D inputs, found its way into growth analysis in the 1960s.

make empirical work difficult – indeed, in the first sentence in the 1957 article on the residual, Solow acknowledges that “... it takes something more than a ‘willing suspension of disbelief’ to talk seriously of the aggregate production function.” Thirty years later, in his Nobel Lecture, he added “ ... I would be happy if you were to accept that [growth accounting results] point to a qualitative truth and give perhaps some guide to orders of magnitude” (Solow (1988)).

The usefulness of this model in providing insights into the process of economic growth has been widely accepted. It has become an official program at the Bureau of Labor Statistics and the mainstay of the current debate over the causes of slower growth. The question raised in this paper is whether the conventional framework, by itself, continues to provide a useful guide for understanding the digital economy. We suggest that this may no longer be the case in its current form, and that it should be extended to allow for “free, copious, and ubiquitous” information, whose benefits go directly to the consumer.

III. The Lancaster Model and Its Application

The essential feature of the Lancaster model is the specification of a utility function whose arguments are the “characteristics” of items that provide utility rather than the goods and services that enter the conventional utility function. Lancaster uses the example of a meal, which is more than just the items of food consumed, but a complex interaction of various factors.⁶ In its fullest form, the conceptual model is quite complex. The model he actually works with is a simplified form in which he assumes that characteristics, c , are functionally connected to outputs, q . In this case, $c = Bq$, where B is a set of parameters that define the consumer’s “technology” for transforming a collection of goods into the bundle of characteristics that provide utility. The associated utility function is then $u(c) = u(Bq)$. In the conventional formulation of utility theory, goods and commodities are identical and $B = 1$. In a

⁶ “A meal (treated as a single good) possesses nutritional characteristics but it also possesses aesthetic characteristics, and different meals will possess these characteristics in different relative proportions. Furthermore, a dinner party, a combination of two goods, a meal and a social setting, may possess nutritional, aesthetic, and perhaps intellectual characteristics different from the combination obtainable from a meal and a social gathering consumed separately” (Lancaster (1966 b), page 133). Subjective factors like ambiance, mood, and novelty matter.

more general form, one that will be used in this paper, the consumption technology is $c = g(q)$. It indicates that different levels of utility can be obtained from a given amount of q , depending on the efficiency with which the transformation occurs.

The consumption technology is central to the concerns of this paper. The availability of reliable information is clearly an important determinant of effective decision-making, and once this is accepted, it is but a straightforward extension to accept the possibility that increases in information could lead to increases in utility $u(g(q))$ for a given amount of goods, q . If technical innovation can shift the structure of production toward greater productivity, why cannot it also shift the productivity of consumers in converting expenditure to well-being using the information disseminated via the internet? As Stigler points out, the utility function is a process in which choices are made, and not a given “datum.”

The Lancaster framework is also valuable in sorting out the issue of product quality (Triplett (1983)). To say that one model of a particular q -good is better in the eyes of the consumer than a similar model is to say that it has more of a desirable characteristic and thus conveys more utility. Or, equally, that the consumer is willing to pay a price premium for the superior good based on the difference in marginal utilities. It is thus natural to regard product quality differentials as one factor determining b , and costless quality change one reason for b to change over time. The Bq formulation is thus a way to introduce product quality into growth accounting models, since it interprets “better” as “more.”⁷

IV. Generalized Growth Accounting

Innovation operates through many “micro” channels and affects the consumption technology in many complex ways, but the same can be said of the conventional Solow-Jorgenson-Griliches-BLS growth accounting model on the production side. Indeed, technical change in the aggregate production function is necessarily macroeconomic in its nature and is thus something of a black box that sweeps together microeconomic changes in technology along with much

⁷ Since the objective of this paper is to introduce consumer utility considerations into the conventional growth framework and examine its implications, we do not go into the many important issues raised by the characteristics approach for price indexes or for consumer demand and expenditure (e.g., Deaton and Muellbauer (1980)).

else.⁸ Since this paper extends this model to allow for the consumption technology in a way consistent with its assumptions, we treat the consumption technology as a black box as well.

The conventional growth accounting model is based on a constant-returns-to-scale production function in which output is produced with labor and capital, and in which the level of technical efficiency has the Hicks-neutral multiplicative form:

$$(1) \quad q_t = e^{\lambda t} f(\ell_t, k_t).$$

The growth rate of output (denoted with an over-dot) is then equal to the growth rate of labor, η , weighted by its output elasticity α , the growth rate of capital, κ , weighted by its output elasticity, $1-\alpha$, and the rate of growth of technical efficiency λ :

$$(2) \quad \dot{q} = \alpha \eta + (1 - \alpha) \kappa + \lambda.$$

The model is made empirically operational by assuming that labor and capital are paid the values of their marginal products in competitive markets (no price distortions) and that firms maximize profit. In this case, the output elasticities are equal to the factor shares in income, s_L and s_K (which sum to one under constant returns to scale). All the elements of the preceding equation are thus obtainable from data, leaving productivity change, λ , to be inferred as a residual:

$$(3) \quad R = \dot{q} - s_L \eta - s_K \kappa.$$

The R is the Solow residual measure of λ , obtained by substituting the factor shares for the output elasticities. The growth rates of the inputs are treated as exogenously given, and productivity change is assumed to occur without a resource cost.

⁸ The “technology” parameter in the aggregate production is a productive efficiency effect that refers to the output obtainable from a unit of input. It is, in effect, resource-saving “technical change.” However, productive efficiency is more than pure technology. It also includes organizational efficiency and management quality, worker effort, the effects of economic fluctuations and shocks, and background inputs like physical and regulatory infrastructure. When measured as a residual, as is conventionally the case, measurement error also comes into play, and the result is what Abramovitz (1956) terms a “measure of our ignorance.”

Our extension of this model to include the consumer side has two components. First, we assume that growth accounting should be extended to include the utility function, and second, that the utility function includes a consumption technology. The first part moves growth accounting from an exercise based on a metric that is objective and in principle measurable — output — to one that is subjective and for which no directly measurable yardstick is available — utility. However, the fact that utility is subjective and impalpable does not mean that it can be ignored in an analysis of how innovation affects well-being, particularly when there is reason to believe that this is how many of the benefits of the digital revolution are realized.

The incorporation of a utility function into growth accounting (with or without the consumption technology) is perhaps the largest deviation from economic orthodoxy, but it is not as heterodox as it might appear, since the Solow model implicitly exists in the context of a utility function. In welfare economics, the objective of economic activity is to maximize utility, whose determinants are the quantity or quality of the goods consumed. The level of real output q as determined by the production function (1) feeds into the consumption side of the economy, giving $U = U(q)$.⁹ However, an expanded growth accounting based on $U(q)$ rather than q is a large step beyond the conventional approach and is sometimes challenged on the grounds that GDP is a measure of resource use, not a measure of welfare. This is certainly true, q is not $U(q)$. Indeed, that is precisely the point of this study: there may be welfare effects of innovation that are not reflected in GDP. Both need to be included in a full assessment of innovation, and the welfare effects should be treated separately and not be shoe-horned into an expanded measure of GDP.

⁹ The link between the supply side of economic growth, as represented by the production function, and consumer side, represented by the utility function, has not received adequate attention in the literature on growth accounting. The Solow residual defined in equation (3) is a differential equation whose solution is a matter of line integration (Hulten (1973)). This can be accomplished by using the production function as the requisite “potential” function, as with the Solow residual, but it could also be accomplished using the utility function as the potential function (Hulten (2001)). It was the contribution of Jorgenson and Griliches (1967) that rooted the residual deeply in production theory. The 2002 paper by Basu and Fernald provides a valuable elaboration of the difference between technology and welfare growth.

Once growth accounting is expanded to include both production and consumption, there is then the question of how an increase in q affects $U(q)$. This is a variant on the question of the marginal utility of income in social welfare and income redistribution theory. An increase in q is generally assumed to have a positive marginal utility, but the rate of change in marginal utility is ambiguous, though generally assumed to be negative, supporting the case for progressive income taxation. In the growth context, a negative marginal utility of aggregate income implies that a steady rate of real GDP growth brings progressively less well-being.¹⁰ This may partly be what Robert Gordon has in mind when he says that the digital revolution is less significant in terms of well-being than inventions like the flush toilet, etc. (although he seems to suggest that this is also due to a declining rate of growth of q).

In keeping with the illustrative nature of this paper, we allow for the marginal utility problem in a simple one-parameter framework in which output growth is weighted by the parameter μ in the simple model,

$$(4) \quad U(q_t) = m q_t^\mu ,$$

The μ plays essentially the same role as the marginal utility of income in static theory. If $\mu < 1$, sustained output growth brings progressively less utility benefit. On the other hand, a value $\mu = 1$ implies that the growth rate of utility is identical to the growth rate in real GDP, in turn implying that the latter is a valid proxy for the former. This is, in a sense, the underlying assumption of standard growth accounting, since it renders unnecessary the utility side of the growth problem. It is also possible that μ could plausibly be greater than one over certain periods of rapid technological innovation. A rush of new goods might conceivably interact in ways that make the whole greater than the sum of its parts.

Introducing a consumption technology into the model adds a degree of complexity to the link between GDP and utility. Even if the parameter μ equals one in this case, a positive

¹⁰ This conclusion assumes that the relative distribution of real income remains unchanged as income increases. A more thorough analysis would allow for a shifting distribution, with the possibility that high income earners might well benefit more from growth than the poor, at least in the short and intermediate time horizons. However, since heterogeneity is not germane to the main point of this study, we will assume that there are many consumers with the same utility functions and incomes.

shift in the consumption technology can cause the standard of living to increase more rapidly than the growth rate of q . The shift in that technology can be modeled in different ways, but for the purposes of this paper, we will again adopt a minimalist specification that preserves symmetry with the growth accounting model of equations (1) and (2) — basically a multiplicative form in which shifts are the equivalent of Hicks-neutral productivity change in which the general information effect, Ω_t , is multiplicative, as is the embodied product quality effect, b_t :

$$(5) \quad U(c_t) = \Omega_t (b_t q_t)^\mu .$$

In this one consumption good model, the parameter b_t is interpreted as a product quality index. In the multiproduct case, with many individual b_t , the parameters can no longer be regarded exclusively in terms of product quality.¹¹ Equation (5) can be expressed in its growth rate form as

$$(6) \quad \dot{u} = \omega + \mu \beta + \mu \dot{q} ,$$

where β is the rate of change of embodied product quality and ω is the rate of change of the information parameter. The economic forces underlying both β and ω are assumed, here, to occur autonomously and without resource cost.¹²

¹¹ The intuition behind the multiplicative information effect, Ω , is illustrated in a two-good version of the model in Figure A3 of the appendix. It must also be acknowledged that the one-parameter multiplicative approach is little more than an acknowledgment of the general importance of information for consumer choice. Again, it is essentially a black box, but, then, so is the treatment of technical efficiency, λ . A more realistic treatment would allow for search costs, information technology, and uncertainty, but would add little to the general point that the growth in output is not the only source of increased consumer welfare. Much the same can be said of the β parameter, whose effect can also be shown in Figure A3.

¹² The increase in information associated with the growth of the digital economy, as captured by the parameter ω , is in reality not an entirely autonomous or costless event. It is linked to the development of complex software and mobile communication devices, as well as internet service and other communication service providers. These can be regarded as produced goods that are acquired directly, or indirectly, and combined with other goods (e.g., automobiles) to produce a consumption outcome (e.g., trips). In other words, the consumption technology treats the q -goods as “inputs” and transforms them into a c -good. The transformation process is treated in our simple model as costless — the use of the free information provided on demand. A more realistic approach would allow for costs associated with the shift in the technology and perhaps even its use, including the time cost involved. This would introduce a cost element in (5) and (6), but it would still leave a costless component because of the way access to information and some product quality change is priced, as well as the fact that information is essentially

This last equation makes the central point of this paper: real GDP growth alone is not a sufficient statistic for assessing the impact of technological revolutions on the standard of living, nor does a slowdown in the growth of real GDP necessarily imply that the standard's growth has slowed.

The full sources of the growth model can be derived by combining the Solow supply-side equation and the Lancaster utility-side growth equation. However, this is not a straightforward matter of inserting the production-side equation (2) into (6). The presence of capital in the production function (1), and its growth rate in (2), poses a conceptual problem because neither capital stock nor the flow of investment appears in the utility function, which is based on the amount of goods consumed, not on the amount produced. Capital is a produced means of production, and investment is thus an element of aggregate output, q , but from a consumption standpoint, it represents deferred consumption in a multiperiod "intertemporal" utility function.

We will return to the problem of capital in a subsequent section, but since it is not the crux of the Lancaster problem, we will ignore it for now and work with a simple one consumption output and one labor input framework. The production function in this case has the Cobb-Douglas form $q_t = e^{\lambda t} \ell_t^\alpha$, where α equals one in the constant returns case and q is the output of the consumption good. The associated production growth account is then

$$(7) \quad \dot{q} = \lambda + \eta.$$

Combining (7) and (6) gives the full growth account:

$$(8) \quad \dot{u} = \omega + \mu \beta + \mu \lambda + \mu \eta.$$

In this form, the growth rate of utility is decomposed into three sources of innovation plus the growth in factor input. The first two sources of innovation, $\omega + \mu \beta$, are output saving, while $\mu \lambda$ is resource saving.

a non-rival good. On the other hand, symmetry with the production function suggests that the time involved in the consumption of the c -good should not be treated as a cost, for the same reason that the time involved in the consumption of conventional q -goods is not currently considered part of GDP. The issue of time as an input is revisited in a subsequent section of the paper.

V. The Treatment of Product Quality and the Solow Residual

The taxonomy of innovation implied by (8) associates product quality change with the consumer technology. This is not the way it is treated in the conventional growth accounting model, where changes in product quality are regarded as a shift in the production function, reflecting an increase in *effective* output per unit of produced output — in effect, treating better output as more output and thereby portraying productivity quality change as resource-saving innovation rather than output-saving innovation as in the preceding section.¹³ In the absence of an explicit consumption technology, the better-is-more approach permits a potentially important source of innovation to enter the analysis. However, when consumption technology is made explicit, as above, this approach is less appealing, since a costless increase in product quality does not involve the use of resources, by definition, but it does increase consumer utility, as in equation (5). Costless quality change therefore belongs in the consumption technology as *output-saving* technical change — as an increase in the effective output consumed. This distinction defines the boundary between the production and consumption technologies.

Another important boundary question involves the household production that takes place outside the market sector. There is a risk of confusing the production of goods within the household with their consumption, since both occur within the economic veil of the home and often involve the same people. However, while a meal cooked at home is subject to the “technology” of recipes, ingredients, kitchen appliances, and the skill of the cook, it is also subject to Lancaster’s point about the “aesthetic characteristics” of the consumption of the meal as a separate event, which is part of his consumption technology. Many of these “aesthetic characteristics” occur whether the meal is prepared at home or a similar one purchased in a restaurant. Moreover, from the standpoint of pure theory, there is no essential

¹³ The intuitive difference between the two ways of looking at product quality change can be seen by comparing Figures A1 and A3 of the appendix. Analytically, when product quality change is treated as resource-saving, the modified model (7) is expanded to include β and the right-hand side becomes $\beta + \lambda + \eta$. The shift in the production function is then $\beta + \lambda$. This combined effect is what gets measured by the TFP residual.

difference between market and non-market production *per se*, since both are resource-using and involve the technological transformation of these resources into a product. The technologies of market and non-market meal production are far more similar to each other than to the technology for building, say, jet aircraft.¹⁴ It is the market versus non-market *distribution channels* of the meals that are different, and even here, theory would simply assign a shadow price to the homemade meal reflecting many of the factors that determine the market price of meals.

VI. The Price Dual

The Solow productivity residual works empirically because both the left-hand-side variable of the production function, output, and the right-hand-side factor input variables are observable. The technology parameters can then be estimated non-parametrically via the Solow residual, or parametrically using econometric methods. This is not the case with the expanded growth accounting model, since the left-hand-side variable is now consumer utility. Because the utility variable is subjective and not directly observable, it is useful to recast the analysis in one that is: the expenditure function. Under certain restrictive conditions, the utility formulation can be represented by its price dual, the expenditure function, and the production function by the dual factor price frontier.

The expenditure function associated with (8) has the form $e(p_c, u^*) = \Omega_t^{-1} \xi(p_c, u^*)$. This is the minimum expenditure needed to maintain utility at the level u^* when the price of the commodity, p_c , changes. The expenditure function shifts downward when costless information increases, since the minimum expenditure needed to finance a given level of utility decreases.¹⁵ The expenditure function can also be expressed in terms of the observable transaction prices of the good, q . The expenditure of the quality-adjusted good, c , is the same as that for the

¹⁴ Owner-occupied housing is another example that could be cited. Buying a house that you had previously rented moves you across the market/non-market conceptual boundary, since you now pay an implicit rent to yourself as owner-occupier in place of the rent paid to the previous landlord. Note, however, that the implicit rent is treated as part of GDP in the national accounts.

¹⁵ This is shown in Figure A3 of the appendix as a downward shift in both the utility function and the budget line.

unadjusted q , implying $p_c c = p_q q$, where p_q and p_c denote corresponding prices, and $p_c = B^{-1} p_q$. The expenditure function then becomes

$$(9) \quad e(p_c, u^*) = \Omega^{-1} \xi(B^{-1} p_q, u^*).$$

The growth rate of expenditures over time depends on the expenditure-share weighted growth rates of the prices, the negative of the growth rate of b_i , the negative growth rate of the information parameter Ω , and the rate of change of u^* . This analysis can also be framed using the indirect utility function associated with the expenditure function, in which case the growth rate of the utility index is decomposed into the negative of the expenditure-share weighted growth rates of the b_i , the negative growth rate of the information parameter Ω , the rate of change of real GDP. This is the price analogue of (8).

The factor price frontier associated with a neoclassical constant-return Hicks-neutral production function $q = e^{\lambda t} f(x)$ has the form $p_q = e^{-\lambda t} \phi(p_x)$, with p_x and x denoting the price and quantity of inputs, and p_q is the *resource-cost* price of output.¹⁶ Substitution into (8) gives

$$(10) \quad e(p_c, u^*) = \Omega^{-1} \xi(B^{-1} e^{-\lambda t} \phi(p_x), u^*).$$

The minimum expenditure needed to support u^* falls with an increase in information, when goods get better, when they are produced more efficiently and their price falls, or when the consumer moves to a higher or lower indifference curve (as shown in the appendix). In its growth rate form, equation (9) is the dual counterpart of the primal form (8).

The expenditure function offers a natural way to think about the consequences of innovation, since it defines the compensating variation (CV) and the equivalent variation (EV) of consumer surplus theory. When the price of a good changes from one period to the next, leading to a change in the level of utility from u_0 to u_1 , the CV is the amount of expenditure needed to regain the old utility level at the new prices p_1 :

¹⁶ In the terminology of Triplett (1983), p_c is *user-value* of the good and p_q is the *resource-cost* price of output. It should be noted, here, that the *resource-cost* price p_q is the market price at which the producer sells a unit of the good they produce. It is this price that determines the value of the marginal product of labor. The *user-value* p_c is the shadow price that the consumer uses in maximizing utility and the price that enters the expenditure function.

$$(11) \quad CV = e(p_1, u_1) - e(p_1, u_0).$$

The EV is defined with respect to the original prices. Together, they provide a willingness-to-pay metric on the change in utility resulting from changes in the various sources of growth. The willingness-to-pay metric will tend to exceed the GDP metric during periods of positive growth because it encompasses both the consumption and production technologies. The utility function may shift even if the production possibilities are unchanged, yielding a positive CV (EV). The issue at hand is how much of the total CV (EV) associated with the digital revolution is due to the production versus the consumption technology.

The expenditure function formulation of our model, with its CV and EV, is important for the empirical implementation of our utility-based framework. Utility cannot be observed directly, but there is a long tradition in environmental, health, and regulatory economics, as well as benefit-cost analysis, of measuring non-market mediated benefits using some form of the willingness-to-pay approach. These approaches have been applied to a variety of goods in a variety of ways.¹⁷ Some forms of the expenditure function model can also be estimated directly using econometric techniques applied to functional forms, as in Redding and Weinstein (2016). Looking ahead, the expansion of the accounts implied by (8) would almost certainly require new (or expanded) surveys in areas like time use, technology utilization, and perhaps willingness to pay. The advent of cloud computing may, for example, provide an opportunity to observe the “quantity” and “rental” price of the various programs and applications downloaded from the cloud. It is worth noting, here, that the GDP accounts developed in the 1930s and 1940s required the development of new data sources or the adaptation of existing ones.

VII. Consumer Surplus and New Technology

¹⁷ The literature ranges over a number of technology goods and approaches. A partial summary includes the papers that can be classified as using consumer surplus by Noll et al. (1973), Hausman (1996, 1999), Greenstein and McDevitt (2011), Brynjolfsson and Gannamaneni (2016); papers based on the value of time by Brynjolfsson and Oh (2012), Chen et al. (2014), Goolsbee and Klenow (2006); and papers using indicators or other approaches by Nevo et al. (2016), Chernew et al. (2016), and Highfill and Bernstein (2015).

Special mention should be made of consumer surplus in light of the preceding discussion (and footnote 18). It is a partial equilibrium approach that provides a monetary metric of the utility arising from the consumption of a good, one that is closely related to the CV and EV. The consumer surplus approach is particularly important for estimating the benefits associated with the arrival of new goods in the marketplace. A new good is one with characteristics that have no near precedent in the choice space of the consumer, as opposed to a good whose quality has improved.¹⁸ Given its prior absence, how should the introduction of this contribution be valued? How much does GDP change as a result of its arrival in the market place?

Valuing a new good at its observed price when it appears may understate the true benefits it brings, since this entry price will reflect (in part) a cost of production that may be low compared to the value of the innovation. A new vaccine may cost little in the way of resources but bring enormous benefits. The theoretical solution advanced by Hausman (1996, 1999) is to estimate the Hicksian “reservation” price of the new good, the price at which the quantity demanded of the good is zero (i.e., the price at which the demand curve intersects the price axis). The result is essentially a consumer surplus solution (Hausman (2003)) and can also be thought of in terms of the compensating variation (Romer (1994)). In terms of the aggregate price index needed to convert nominal GDP/GDI into real GDP, the reservation price serves as a quality correction, one that indicates a higher value to the consumer per units of resource cost. However, while consumer surplus is a valuable technique for getting at a difficult measurement problem, it is a partial equilibrium technique and it is not clear that it can capture all of the benefits of a general increase in information (see appendix Figure A5).

There is, however, a difficulty in applying techniques like consumer surplus, expenditure functions, price hedonics, or matched-price models to measure the benefits of the digital economy: they presume the existence of a market price. This is not always the case with some of the most important digital economy goods. The internet and its apps are not priced as individual goods with a market price per unit consumed. Instead, a general internet access fee

¹⁸ At a conceptual level, a newly available item is a “new” good at a low level of aggregation (Windows 7 versus 10), but a higher-quality good at a more aggregated level (productivity software).

may be charged by a service provider (though access may also be freely available in some cases). Internet applications are widely available without a direct use charge (again with exceptions). Many applications are supported by the marketing revenues they can generate, or are provided *pro bono publico*. The absence of observable unit prices, or an artificial zero price, has led researchers to use alternative measurement strategies, like the valuation of time and the use of indirect payments.

A deeper conceptual issue underlies some of these problems: many digital economy products exhibit the public good characteristics of non-rivalness and non-appropriability of property-rights. The “consumption” of information by one person does not diminish the amount available to anyone else and, in its purest form, each consumer gets the same quantity.¹⁹ In this case, each consumer would theoretically pay a price that reflects marginal willingness to pay, and that total willingness to pay (the sum of the individual prices) would equal marginal cost at the optimum. The fact that, in theory, there is no single price per unit makes modeling exercises more complicated and the use of consumer surplus techniques more difficult. There is also a free-rider problem, arising from the fact that it is often difficult to establish and enforce property rights over intellectual property (information, technology, product design, and artistic originals). This leads to efforts to protect intellectual property through patents, copyrights, and secrecy, and to the extent this is successful, a degree of monopoly power rises that must be brought into the model, driving a wedge between price and marginal cost.

Moreover, many new economy goods are also subject to non-rivalness in their production. Product research and development, as well as design and marketing, are largely overhead costs, and once incurred, can be spread over the units of output actually produced.

¹⁹ A further complication arises with “network” goods. These goods have the property that the individual demand for the good depends on its demand by other people. Social networks and other communication media, collaborative information gathering, and common standards for operating systems and productivity software are cases in which individual demand increases with the number of other users. However, the network effect may also be negative. Some types of information are most valuable when possessed by only a few people (first mover decisions in finance, access to scarce resources). In either case, the network effect adds another layer of complexity to the analysis.

Then there is the fact that information and other digital goods like software, music, and video products can be reproduced at little or no marginal cost. The result is that digital goods tend to be characterized by strongly diminishing average cost, and often zero marginal cost. This is another reason why many are distributed without a direct charge, with access fees or advertising revenues covering the cost plus a markup, and why it is difficult to use conventional demand and supply curves to measure consumer surplus.²⁰

VIII. Contingent Goods

The wedge between the output of a good and its outcome to the consumer is another source of difficulty in assessing the value of output. This wedge arises from the contingent nature of many goods and is particularly prevalent in those involving the “expert” industries of the service sector — medical, financial, and legal services, education, management consulting — that collectively account for about half of private GDP (Hulten (2015)). A visit to the doctor is usually in response to some perceived health problem, but you do not buy an improvement in health *per se*, you buy advice and perhaps an intervention that may or may not cause an improvement. That outcome will depend on the initial state of health, the doctor’s input, and the actions you take in response. Other expert services in education, legal matters, and finance can be modeled using a similar framework, since they represent an attempted transition from one state of being to another, appropriately defined.

Contingent goods have a natural interpretation in the Lancaster framework. The output from the standpoint of the consumer is the improvement (objective and subjective) in health status and not the medical expenditure, which is the basis for measuring the output of the doctor. The former could be thought of as *c*-output and the latter as *q*-output. The main addition is the dependence of *c*-output on an initial state, which means that a given expenditure on purchased health services may produce different levels of *c*. The total

²⁰ New digital-economy goods are not unique in this regard. Public goods constitute much of the output of the government sector, with the quantity determined in the political process and finance based on taxes not linked to benefits received, though some local taxes are loosely connected to benefits. Parts of the financial intermediation sector are also subject to the problem, with firms providing services at or near zero cost in exchange for the spread they earn on deposits. Both cases are notorious for the problems they pose to the measurement of GDP.

expenditure on q , p_qq , is generally used as a proxy for $p_c c$, although this omits any ancillary expenditure (joining a health club, dietary changes, etc.). However, the real problem lies in measuring real GDP, that is, in separating prices in p_qq and $p_c c$. This is even more difficult given that there are usually no observable (or even natural) units of measurement for either.²¹

The problems are even more difficult when much of the q -output involves information. Doctors bill by the visit, hour, or procedure, and not by the bits or bytes of expert information, nor by the outcome, which is uncertain and contingent on patient inputs.²² Moreover, the expert nature of information is an important part of the overall value. Information that is not organized or focused on informing a specific question or issue is of limited (sometimes negative) value. The professional organization of information, be it in healthcare, law, education, or finance, is the greater part of its value *vis a vis* information acquired by individuals without expertise. The latter is, however, not without value, since professional opinion is not infallible, and an informed consumer is usually better off than one that is uninformed. This is one important source of value of the internet, which hosts applications that organize information on various topics, and of mobile devices and search engines that put information into the hands of users more or less immediately.

IX. The Role of Capital Formation

We have thus far put off the question of capital formation. The expanded growth account implied by equation (8) is derived for a world of consumption output and labor input. Both are contemporaneous flows with no direct link between one period and the next, other than the

²¹ It is possible for c -output to appear to decline for the population as a whole, even if an innovation increases the quality of a medical procedure. This paradox can arise in the contingent state model if a medical advance leads to an increase in the number of cases that can be treated, but such cases have a lower probability of success than the average of the pre-innovation population. Total welfare has increased, so care must be taken in how c -output is measured.

²² Cutler and Berndt (2001), in the introduction to their conference volume, point to what they call the “outcomes movement” in health economics, which is the attempt to measure the health impact of medical care rather than the amount expended. As Hulten (2015) notes, anyone who remembers a visit to the dentist in the 1950s can testify to the enormous gains in efficacy and patient comfort that have occurred. Laparoscopic surgery has shortened recovery periods and brought major advances in patient comfort. These improvements in subjective outcomes are hard to measure, but progress is being made.

exogenous growth rates of labor and technical innovation. This changes when capital goods are added to the model. Capital is a produced means of production and can be regarded as an intertemporal intermediate input whose services are delivered to production in future years. On the consumption side, capital formation represents a deferral of consumption to future years, enabled by the increase in productive capacity.

The capital stock typically appears in the production function as a proxy for the quantity of capital services, as in equation (1). However, the output of this production function includes investment as well as consumption. This is not a problem when growth accounting is restricted to the production side of the economy, but it does pose a problem for the specification of utility when the model is expanded to include the consumption side. There is no provision in the utility function (4) and (5) for investment, only for consumption. Since investment is deferred consumption from the standpoint of the consumer, it is natural to replace these static utility functions with one based on consumption at different points in time. The “intertemporal” utility function is typically formulated as

$$(12) \quad U(C_1, C_2, \dots, C_N) = \sum_{\tau=1}^N \frac{u(C_\tau)}{(1+r)^\tau} .$$

This objective function is maximized subject to the constraints imposed by the annual accumulation equation, $K_t = I_t + (1-*)K_{t-1}$, the initial and terminal levels of the capital stock (assumed here to be zero), and $*$, the rate of depreciation. The production function in each year is:

$$(13) \quad Q_t = C_t + I_t = e^{\lambda t} F(L_t, K_t).$$

Under certain assumptions about “well-behaved” functions, a solution exists that defines the optimal paths of consumption, investment and capital, for given exogenous paths of labor and technical efficiency.²³

²³ This is the discrete-time form of the “textbook” intertemporal optimization problem, with initial and terminal capital stocks set to zero. It may be worth noting that stochastic variants of this model have also been around for a long time (e.g., Brock and Mirman (1972)).

The implications of this solution for growth accounting are discussed in Hulten (1979). Using a slight variant of the constraint function, this paper shows that the Divisia index (the share-weighted growth rates) of optimal consumption (C^*_1, \dots, C^*_N) equals the Divisia growth rates of labor input (L_1, \dots, L_N) and productive efficiency (A_1, \dots, A_N). This simplifies greatly when labor grows at a constant rate η and productive efficiency at a constant rate λ , and consumption grows at a constant steady-state rate. In this case,

$$(14) \quad \dot{c} = (\lambda / \alpha) + \eta.$$

The term λ/α is the Harrodian rate of productivity change, equal to the Hicksian rate λ in (1) and (13), divided by labor's share of income. The difference between the two rates is no accident. Capital is an intertemporal intermediate input and it largely disappears in the solution to the intertemporal optimization problem. Largely, but not entirely, because, as an intertemporal intermediate good, the productivity term must reflect the Domar expansion effect of the theory of contemporaneous intermediate goods (Hulten (1978)). In the intertemporal context, the expansion effect magnifies the annual productivity effects.²⁴

This basic conclusion is that the addition of capital to the original consumption-only model does not change the basic nature of the original analysis, at least when the consumption technology is ignored. Adding capital to the analysis changes the specification of the production function, whereas passing from (4) to (5) by adding the consumption technology does not involve capital directly in technology, but instead is based on the consumption emerging from the production side, whether or not capital is used to produce that consumption.

The problem with capital formation lies elsewhere, in the very nature of the intertemporal optimization problem. The idea that optimization deep into the future is possible requires an even greater suspension of disbelief than the period-by-period account of growth given by the aggregate Solow model and our proposed extension. It does not strain

²⁴ Equation (14) can also be derived using Golden Rule steady-state growth. Either way, technical change affects long-run equilibrium growth via λ/α , while it affects year-to-year changes in growth via λ . In the consumption-only case, this is apparent when comparing equations (7) and (14).

credulity to argue that the internet and its applications can improve consumer utility by informing near-term decisions. However, the strain on credulity begins and grows as the decision horizon is pushed further and further into the future (even more so when Hayek's point is recalled). Still, it is comforting to know that the logical properties of the dynamic optimization system are such that, if allowed to fully play out, it could reach a certain final state in which the addition of capital to the model does not change the essential results of the simpler consumption-labor version of the model.

X. Technical Innovation with a Resource Cost

Our model has thus far treated innovation as a costless process. This is appropriate for many situations, since information has many public good characteristics, and technology often diffuses at a low cost to users. Costless innovation also occurs through pure inspiration and creativity. It also occurs through what Eric von Hippel (2016) calls "free innovation," "innovations developed and given away by consumers" free of charge (p. 1). This said, innovation in the digital economy increasingly results from systematic efforts to bring about change. Firm-specific investment in intangible capital has become the dominant form of business investment in the U.S. and is at the forefront of this kind of change.²⁵ This kind of investment is responsible for the buildup of a firm's intellectual property, its product base, and it contributes to the know-how that defines the firm's overall capabilities. It is the source of much of the systematic effort at product and process innovation. Expenditures on intangible capital include payments to R&D workers, software engineers, product designers, marketers, and IT implementers who work with the firm's existing intangible capital base, as well as the necessary tangible capital (ITC equipment, instruments, research laboratories, etc.).

Costly and costless innovation can occur simultaneously, with implications for the model of the preceding sections. On the output side, costly innovation in products and processes are "purchased" at the cost of the resources used. The opportunity to buy "better" goods comes at a higher price that reflects that cost. Moreover, the transaction units of the improved good

²⁵ Corrado, Hulten, and Sichel (2005, 2009). See also Nakamura (2001).

embody the effects of the innovation. In terms of the preceding model in which $c = bq$, it is the c that is produced and purchased. Product quality change, β , is implicit in the transaction units of c and not produced directly or recorded separately. It is, nevertheless, a magnitude that should, in principle, be included in a theoretical taxonomy designed to get at the extent of innovation in any year, even though this version of the β is implicit and is resource costly. However, because of this, the β arising in this way should be assigned to the resource-using side of the taxonomy.

On the input side, the resources used in the production must be expanded to include intangible capital, both as an output flow, N_t , and as a stock, R_t . These are connected by a new accumulation equation in which $R_t = N_t + (1 - \delta_R)R_{t-1}$. The production function (13) then becomes

$$(15) \quad Q'_t = C_t + I_t + N_t = A_t F(L_t, K_t, R_t),$$

where Q'_t is greater than Q_t by the amount of investment in intangibles. In the past, N_t has been regarded as an intermediate input to production and omitted from GDP (though this is beginning to change).²⁶ The optimization analysis underlying the production function in (15) is further complicated by the fact that, as noted above, N_t and R_{t-1} tend to have the public good characteristics associated with intellectual property. Thus, the inclusion of intellectual property introduces an important pricing issue, since the decision to invest in this kind of capital presupposes a sufficient degree of monopoly power over property rights to make the investment profitable. This, in turn, introduces non-marginal cost pricing into the growth accounting model, in contravention to the perfect competition assumption of the basic Solow model.

Finally, there is another dimension of the innovation-cost problem that should be mentioned. As previously noted, it might plausibly be argued that the consumption technology requires a time input, just as time-use enters the production technology through labor input.

²⁶ The Bureau of Economic Analysis, the agency responsible for the U.S. national accounts, has recently moved to treat expenditures for R&D and artistic originals as capital investment. Computer software has been treated this way since the 1980s.

The advent of new or improved goods may well involve start-up costs and a learning curve, and the consumption of goods takes time (as in watching television). The latter is generally considered to be leisure time and to have an opportunity cost. Indeed, the consumption of each good in the utility function could be regarded as having an associated time requirement. We considered formulating our problem in this way but opted for a more simplified treatment that highlights our key point about the need for a consumption technology (however formulated) in growth accounting. Moreover, some output-saving innovations could involve less consumer or household time (search engines, e-commerce), others more time (social media, surfing the net), and some little time at all (a quick look at emails, learning how to use a new model of a familiar good). Time use could also involve more timely information (stock prices, email) rather than more or less time used. In other words, both time-complementarity and time-substitutability could be in play, depending on the application.²⁷

From a welfare standpoint, the most important difference between costless versus costly innovation lies in the way they increase utility. Costless technical change, in its original neoclassical conception of “Manna from Heaven,” brings about a shift in the production (and in our model) utility functions that is purely welfare enhancing. Because it is paid for with resources, costly innovation involves no Manna free-lunch.

XI. Summary and Final Thoughts

We have proposed an extension of the conventional growth accounting model that incorporates an explicit utility function. It owes a very large debt to Lancaster’s idea of a consumption technology, but much of our modeling holds without the characteristics part of his analysis, though it is useful for interpretation.²⁸ The key feature of our model is the

²⁷ Viewed from the standpoint of the conventional work-leisure decision, hours worked in the production of q goods may be affected by a shift in the consumption technology. A positive shift in that technology increases the utility of each unit of q output, implying an increase in the c -denominated wage for each level of the q -wage. This alters the standard work-leisure tradeoff, and depending on the relative strengths of the income and substitution effects, hours worked in the production of q -output may increase or decrease.

²⁸ The essential steps in developing equations (4) through (8) in Section IV involve expanding the production-side growth equations to include a utility function, and then adding a time dimension with parameters. These steps involve Lancaster’s idea of a consumption technology and the wedge between the resource cost of a good and its

possibility that innovation can affect the standard of living directly, above and beyond its effect on the production function — in other words, that innovation can be both resource-saving, as in conventional TFP, and output-saving. It thus implies that GDP may not be an adequate statistic for measuring the extent of innovation, and that some of the welfare lost by the recent slowdown in the growth rate of real GDP may be offset by growth in the alternative output-saving consumption channel.

However, we also want to emphasize that, while GDP may not be sufficient for fully characterizing economic growth in the age of the internet, it remains an essential tool for understanding the evolution of the market economy and for important policy issues involving the employment of resources, trade policy, and much more. We regard our proposed extension not as a substitute, but as a complement to the existing GDP-based accounts, one which would allow users to combine the various elements of the expanded accounts in ways best suited to their needs.

We have not attempted to test the sufficiency hypothesis, but our extended model does provide a conceptual underpinning for the recent empirical literature that does. The evidence in this literature on the relative size of these non-GDP welfare effects is mixed, with most estimates in the range of \$100 billion to \$1 trillion. Viewed against the overall size of GDP, currently around \$18 trillion, the effects seem relatively small. In summarizing his findings, Syverson (2016) concludes that:

“... estimates from the existing research literature of the surplus created by internet-linked digital technologies fall far short of the \$2.7 trillion or more of ‘missing output’ resulting from the productivity growth slowdown. The largest — by some distance — is less than one-third of the purportedly mismeasured GDP.”

Byrne, Fernald, and Reinsdorf (2016) reach a similar conclusion about mismeasurement as an explanation of the slowdown in GDP growth:

“While we find considerable evidence of mismeasurement, we find no evidence that the biases have gotten worse since the early 2000s.”

utility to the consumer. Our point, here, is not to minimize our debt, but to emphasize that the validity of our approach is not necessarily pinned to the acceptance of the Lancaster characteristics view of products, though it is a useful way of thinking about the issues we raise.

These are very reasonable assessments given the size of current GDP, even diminished as it is by slower growth. However, it is inherently difficult to talk about mismeasurement when even the approximate size of the “correct” measure is unknown. The effects of the digital revolution are pervasive and touch many aspects of economic life in ways that are hard to spot, much less measure. The unmeasured output-saving value of information alone is potentially very large, and it has grown rapidly in recent years. How much is yet to be determined.

A larger point is that any attempt to assess the effects of a broad revolution in technology by looking at each innovation separately, on a case-by-case basis, risks understating the totality of the transformation they brought about. Can the impact of the internal combustion engine be captured by quality adjustments to autos? Or compared to horse-drawn carriages? Can the impact of electricity be measured in terms of candles? Is not the overall impact of a wave of innovations like those of the Industrial Revolutions greater than the sum of the parts?

The lessons of economic history suggest that it is. Large-scale “tectonic” changes in technology have broad effects on the worldview and social structure of a society as they unfold.²⁹ Individual utility functions adapt to the logic and possibilities of the new technological era and, in turn, condition these possibilities. Luxuries appear and become necessities. The true “size” of a major technological revolution virtually defies measurement, given that the metrics for measuring size are themselves changing. This observation applies to the tectonic changes that are currently under way, with the immense possibilities of the information revolution — a world without work, artificial super intelligence, genomics, a fully networked society and economy.

²⁹ In his study of the history of lighting, Nordhaus (1997) distinguishes among three types of technical change: run-of-the-mill changes, seismically active sectors, and tectonic changes. He defines the latter as the situation in which “changes in production and consumption are so vast that the price indexes do not attempt to capture qualitative changes.” He argues that official price statistics missed much of the gain from the tectonic revolution in lighting.

REFERENCES

- Abramovitz, Moses, "Resource and output trends in the United States since 1870," *American Economic Review*, 46(2), 1956, 5–23.
- Ahmad, Nadim, and Paul Schreyer, *Measuring GDP in a Digitalized Economy*, OECD Working Paper, April 2016.
- Autor, D. H., F. Levy and R. J. Murnane, "The Skill Content of Recent Technological Change: An Empirical Investigation," *Quarterly Journal of Economics*, November, 118(4), 2003, 1279–1333.
- Basu, Susanto, and John G. Fernald, "Aggregate productivity and aggregate technology," *European Economic Review*, 46, 6, June 2002, 963–991.
- Brock, William A., and Leonard J. Mirman, "Optimal Economic Growth and Uncertainty: The Discounted Case," *Journal of Economic Theory*, 4, 1972, 479–533.
- Brynjolfsson, Eric and Avinash Gannamaneni, "Measuring Changes in Consumer Surplus in the Digital Economy," Working Paper, MIT, July 2016.
- Brynjolfsson, Erik, and JooHee Oh, *The Attention Economy: Measuring the Value of Free Digital Services on the Internet*, 2012, www.informatik.uni-trier.de/~ley/pers/hd/b/Brynjolfsson:Erik.
- Brynjolfsson, Erik, and Andrew Saunders, "What the GDP Gets Wrong (Why Managers Should Care)," *MIT Sloan Management Review*, 51(1), 2009, 95–98.
- Byrne, David M., John G. Fernald, and Marshall B. Reinsdorf, "Does the United States Have a Productivity Slowdown or a Measurement Problem," *Brookings Papers on Economic Activity*, March 2016.
- Chen, Yan, Grace Jeo, and Yong-Mi Kim, "A Day Without a Search Engine: An Experimental Study of Online and Offline Searches," *Experimental Economics*, 2014, 17, 4, 512–536.
- Chernew, Michael, David M. Cutler, Kaushik Ghosh, Mary Landrum, "Understanding the Improvement in Disability Free Life Expectancy in the U.S. Elderly Population," NBER Working Paper 22306, June 2016.
- Corrado, Carol, Charles Hulten, and Daniel Sichel, "Measuring Capital and Technology: An Expanded Framework." In *Measuring Capital in the New Economy*, C. Corrado, J. Haltiwanger, and D. Sichel, eds., Studies in Income and Wealth, Vol. 65. Chicago: The University of Chicago Press, 2005, 11–41.
- Corrado, Carol A., Charles R. Hulten, and Daniel Sichel, "Intangible Capital and U.S. Economic Growth," *Review of Income and Wealth*, 55, 3, September 2009, 661–685.
- Cutler, David M., and Ernst. R. Berndt, eds., *Medical Care Output and Productivity*, Studies in Income and Wealth, Vol. 62, Chicago, Illinois: University of Chicago Press, 2001.

Deaton, Angus, and John Muellbauer, *Economics and Consumer Behavior*, Cambridge University Press, 1980.

Goolsbee, Austin, and Peter J. Klenow, "Valuing Consumer Products by the Time Spent Using Them: An Application to the Internet," *American Economic Review Papers and Proceedings*, 96(2), 2006, 108–113.

Gordon, Robert J., *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War*, The Princeton Economic History of the Western World, Princeton, New Jersey: Princeton University Press, January 2016.

Greenstein, Shane, and Ryan C. McDevitt, "The broadband bonus: Estimating broadband Internet's economic value," *Telecommunications Policy*, 2011, 35, 7, 617–632.

Hausman, Jerry, "Valuation of New Goods Under Perfect and Imperfect Competition," T. Bresnahan and R. J. Gordon, eds., *The Economics of New Goods*, Chicago, Illinois: University of Chicago Press, 1996, 209–37.

Hausman, Jerry, "Cellular Telephone, New Products, and the CPI," *Journal of Business & Economic Statistics*, 17(2), 1999, 188–94.

Hausman, Jerry, "Sources of Bias and Solutions to Bias in the Consumer Price Index," *The Journal of Economic Perspectives*, 17, 1, Winter, 2003, 23–44.

Hayek, F.A., "The Use of Knowledge in Society," *The American Economic Review*, 35, 4, September 1945, 519–530.

Highfill, Tina and Elizabeth Bernstein, "Using Disability Adjusted Life Years to Value the Treatment of Thirty Chronic Conditions in the U.S. from 1987-2010," BEA Working Paper 2014-12 (updated April 2015).

Hulten, Charles R., "Divisia Index Numbers," *Econometrica*, 41, 6, November 1973, 1017–25.

Hulten, Charles R., "Growth Accounting with Intermediate Inputs," *Review of Economic Studies*, 45, 3, October 1978, 511–518.

Hulten, Charles R., "On the 'Importance' of Productivity Change," *American Economic Review*, Vol. 69, No. 1, March 1979, 126–135.

Hulten, Charles R., "Growth Accounting When Technical Change Is Embodied in Capital," *American Economic Review*, Vol. 82, No. 4, September 1992, 964–980.

Hulten, Charles R., "Total Factor Productivity: A Short Biography," in *New Developments in Productivity Analysis*, Charles R. Hulten, Edwin R. Dean, and Michael J. Harper, eds., Studies in Income and Wealth, Vol. 63, The University of Chicago Press for the NBER, Chicago, 2001, 1–47.

Hulten, Charles, "Measuring the Economy of the 21st Century," NBER Reporter, 4, December 2015, 1–7.

Jorgenson, Dale W., and Zvi Griliches, "The explanation of productivity change," *Review of Economic Studies* 34, July 1967, 349–83.

Lancaster, Kelvin J., "Change and Innovation in the Technology of Consumption," *American Economic Review*, 56, 1/2 March 1966a, 14–23.

Lancaster, Kelvin J., "A New Approach to Consumer Theory," *Journal of Political Economy*, 74, 2, April 1966b, 132–157.

Nakamura, Leonard, "What Is the U.S. Gross Investment in Intangibles? (At Least) One Trillion Dollars a Year!," Federal Reserve Bank of Philadelphia Working Paper No. 01-15, 2001.

Nakamura, Leonard, "Hidden Value: How Consumer Learning Boosts Output," *Business Review*, Federal Reserve Bank of Philadelphia, Q3 2014, 9–14.

Nakamura, Leonard, and Rachel Soloveichik, "Valuing 'Free' Media Across Countries in GDP," Working Paper 15-25, Federal Reserve Bank of Philadelphia, July 2015.

Nevo, Aviv, Jon L. Turner, and Jonathan W. Williams, "Usage Based-Pricing and Demand for Residential Broadband," *Econometrica*, 84 (2), March 2016, 411–443.

Noll, Roger G., Morton J. Peck, and John J. McGowan, *Economic Aspects of Television Regulation*, The Brookings Institution: Washington DC, 1973.

Nordhaus, William D., "Do Real Output and Real Wage Measures Capture Reality? The History of Lighting Suggests Not," T. Bresnahan and R. J. Gordon, eds., *The Economics of New Goods*, Studies in Income and Wealth, Vol. 58, University of Chicago Press for the National Bureau of Economic Research, Chicago, 1997, 29–66.

Redding, Stephen J., and David E. Weinstein, "A Unified Approach to Estimating Demand and Welfare," Working Paper, June 2016.

Romer, Paul M., "Increasing returns and long-run growth," *Journal of Political Economy*, 94 (5), 1986, 1002–37.

Romer, Paul, "New goods, old theory, and the welfare costs of trade restrictions," *Journal of Development Economics*, 43, 1994, 5–38.

Schmidt, Eric, and Jonathan Rosenberg (with Alan Eagle), *How Google Works*, Grand Central Publishing, September 2014.

Solow, Robert M., "A contribution to the theory of economic growth," *Quarterly Journal of Economics*, 70, February 1956, 65–94.

Solow, Robert M., "Technical change and the aggregate production function," *Review of Economics and Statistics*, 39, August 1957, 312–20.

Solow, Robert M., *Growth Theory: An Exposition*, Oxford University Press, New York and Oxford, 1988.

Solow, Robert M., "We'd better watch out," *New York Times Book Review*, July 12, 1987.

Solow, Robert M., "Reflections on Growth Theory," *Handbook of Economic Growth*, Volume 1, Part A, Philippe Aghion and Steven N. Durlauf, eds., North-Holland Elsevier, 2005, 3–10.

Stigler, George J., "The economics of information," *The Journal of Political Economy*, 1961, 213–225.

Stigler, George J., and Gary S. Becker, "De Gustibus Non Est Disputandum," *The American Economic Review*, Vol. 67, No. 2, March, 1977, 76–90.

Syverson, Chad, "Challenges to Mismeasurement Explanations for the US Productivity Slowdown," NBER Working Paper 21974, 2016.

Triplet, Jack E., "Concepts of Quality in Input and Output Price Measures: A Resolution of the User-Value Resource-Cost Debate," in *The U.S. National Income and Product Accounts: Selected Topics*, Murray F. Foss, ed., Studies in Income and Wealth, Vol. 47, University of Chicago Press, 1983, 269 – 312.

Varian, Hal, "Economic Value of Google," 2009: <https://cdn.oreillystatic.com/en/assets/1/event/57/The%20Economic%20Impact%20of%20Google%20Presentation.pdf>.

von Hippel, Eric, *Free Innovation*, MIT Press, 2016.

APPENDIX

The Geometry of Costless Innovation

1. The Conventional Geometry

The standard textbook representation of a general equilibrium in a two-good economy is shown in Figure A1, for goods X and Y , along with the production possibility frontier (PPF) and the utility function of the representative consumer (U). Equilibrium is at the tangency point A , where the ratio of marginal costs equals the ratio of marginal utilities. The equilibrium quantities are X_0 and Y_0 , and the relative prices are defined by the slope of the tangent line at A , the line that defines GDP.




Figure A1




Figure A2

An improvement in the productivity process with which X and Y are produced shifts the PPF outward (analogous to the λ in the one sector model of equation (1)). This is shown in Figure A1 as a shift from PPF_0 to PPF_1 , for the case of price-neutral shift in the technologies of the two sectors, holding capital and labor constant. The output bundle shifts from (X_0, Y_0) at the point A to (X_1, Y_1) at the point B , and real GDP increases from GDP_0 to GDP_1 . In view of the discussion of the link between growth accounting and the change in utility, it is worth noting that the level of utility increases from U_0 to U_1 . When the marginal utility of output is unitary ($\mu=1$), the change in real GDP is a sufficient statistic for measuring the change in welfare, and there is no need to separately account for utility.

The framework of Figure A1 can also do duty in describing changes in product quality. Suppose that rather than a costless change in productivity, a costless change in product quality occurs at a rate b in both goods. This is shown in Figure A2, which now portrays the commodity space in

both the units of the goods produced and the efficiency units that enter the utility function. The economy is initially at the point A , at which both efficiency and production units are the same, (X_0, Y_0) . After the costless change in product quality, the production of (X_0, Y_0) units of the goods is now the equivalent of (bX_0, bY_0) units from the standpoint of the utility they provide. In other words, at point B , (X^e_1, Y^e_1) is equivalent to (bX_0, bY_0) . If there is no change in the productivity with which the goods are produced, actual output remains at the point A , with (X_0, Y_0) still produced. PPF_0 is still the production possibility frontier of the economy and relative prices and GDP are unchanged. However, PPF^e is now the locus of attainable production combinations from the utility standpoint, and B on the new effective-output bundle that provides the new (and higher) level of utility, U_1 . Nominal GDP is unchanged, but real GDP has risen by the factor b , while the corresponding prices have fallen by this factor.

2. The Geometry of Innovation in the Consumer Technology

One question posed in this paper is whether an increase in information should be scored as a supply-side innovation or as a consumer-side innovation related to a shift in the consumption technology. The perspective from the standpoint of the consumption technology is shown in Figure A3. Instead of an outward shift in the PPF due to costless technical change, as in Figure A1, a neutral change in information increases the amount of utility attainable from a given bundle of X and Y . It now appears as a downward shift in the utility function in Figure A3. The old U_0 shifts downward, from U_0^{old} to U_0^{new} , and the output required for the latter is now (X_1, Y_1) . This is the *output-saving technical change* described in the paper.

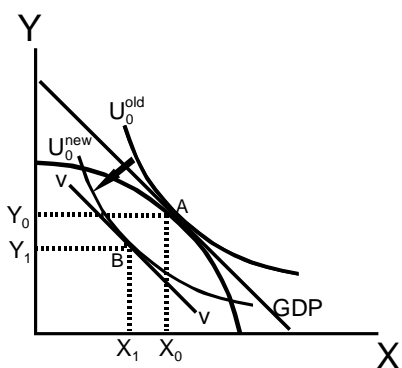


Figure A3




Figure A4

The *PPF* is unchanged as is GDP (real and nominal), and (X_0, Y_0) continues to be produced at the point A. What is different in Figure A4 is that the indifference curve tangent to the *PPF* at A is now U_1 , not U_0^{old} . The consumer still buys X_0 and Y_0 units of the goods, but now gets the higher utility. Again, the output required to support the old level of utility is (X_1, Y_1) , located (with unchanged prices) at B on the line *vv*. The distance between *vv* and *GDP* is the compensating variation (and the equivalent variation in this case).

As with Figure A2, this diagram can do double duty in representing product quality change. In this case, a quality change e in both X and Y change means that $(X_1, Y_1) = (eX_0, eY_0)$, with $e < 1$, and that (X_1, Y_1) now yields the same level of utility as the (X_0, Y_0) prior to the change. The bundle (X_0, Y_0) in Figure 4 continues to be produced, but now gives the higher utility U_1 .

3. Implications for Consumer Surplus

Figures A3 and A4 have an interesting implication for consumer surplus. Since production continues to take place at the point A, observed prices and income are unchanged. Thus, the implicit supply and demand curves for the two goods do not shift, implying that the area below the demand curve and the market price, the standard conception of consumer surplus, is similarly unchanged. What has changed is that utility obtainable per dollar of surplus is greater, reflecting the output-saving nature of costless innovation. In other words, more utility is “packed” into the consumer surplus area.

This is modified somewhat when the impact of innovation is not neutral. This is the situation shown in Figure A5, which portrays the case in which an innovation in information affects one good (X) more than the other. In this case, the indifference curve,




Figure A5

U_0 , twists in the direction of X as well as shifting downward, as shown in Figure A5. The non-neutrality of the innovation causes a twist in the indifference map, resulting in a new equilibrium point C on the PPF. The implicit supply and demand curves now shift, with the demand curve for X generally increasing and Y decreasing. A new consumer surplus occurs, but, while it captures the “twist” effect on welfare, it does not pick up the shift effect between C and D . The demand for Y may decline, but it still delivers more utility per unit than before.

4. Price Distortion Also Affects the Wedge Between Welfare and GDP

The analysis of the divergence between GDP and welfare has a long history in the literature on price distortions. The Harberger Triangle is but one part of this history, albeit a famous one. The paper by Basu and Fernald (2002) frames the problem in the general equilibrium context of Figure A1. Their paper deals with the wedge between the level of utility in a distorted economy (U' in Figure A6) and the maximal utility that could be obtained in an undistorted economy (U). The distorted equilibrium is at point B , supported by a wedge between price and cost. Move to the undistorted point A increases utility without a change in technology. Once at A , and along the dynamic path on which A lies, there are no further welfare gains.




Figure A6

Figure A6 also illustrates other forms of inefficiency: departures from potential output because of distortions in factor prices and fluctuations over the business cycle. Both can cause the economy to locate at a point like D inside the production possibility frontier.

In sum, the change in utility and the change in productivity can diverge along the growth path of the economy for two reasons: the *consumption technology* effect of this paper and the *distortions* that locate the economy away from its optimal equilibrium. Add to these the *productivity* effects of an outward shift in the PPF due to process-oriented technical change, as shown in Figure A1, and the outward shift in the PPF due to *growth in the factor inputs* (not shown), and much of the story of the growth accounting in this paper is told.