# ACCOUNTING FOR GROWTH IN THE AGE OF THE INTERNET: THE IMPORTANCE OF OUTPUT-SAVING TECHNICAL CHANGE[1]

Charles Hulten and Leonard Nakamura[2]

## 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. The output-saving formulation is a natural way to think about the free information goods available over the internet which bypass GDP and go directly to the consumer. One implication of our model is that living standards can rise at a greater rate than real GDP growth, which may shed light on the question of how the latter can decline in an era of rapid innovation.

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 Society for Economic Measurement, July 2016, in Thessalonica Greece; the Federal Reserve *Day-Ahead Meeting on Productivity,* September 2016, Cincinnati; the 2017 ASSA meetings in Chicago; and at the CRIW Workshop at the NBER Summer Institute, July 18, 2017. We thank Roger Betancourt, John Fernald, Brent Moulton, Dan Sichel, Chad Syverson, and the participants at the various presentations, for their comments and suggestions on earlier drafts of an NBER working paper. Remaining errors and interpretations are our responsibility.

[2] Corresponding author: Leonard Nakamura, Federal Reserve Bank of Philadelphia, 10 Independence Mall, Philadelphia PA 19106, fax: (215) 574-4303, phone: (215) 574-3804, email: leonard.nakamura@phil.frb.org.

## 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 growth declined from its 1995-2006 rate of 2.3% to 1.5% 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) has also 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.[3] 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

---

[3] Greenspan's concerns were first expressed in remarks at an FOMC meeting in late 1996 in regard to a staff analysis of sectoral productivity trends (Corrado and Slifman, 1999).

supply-side of the economy. There is a growing conviction in the recent literature on growth accounting that the current round of innovation is not adequately captured by conventional real GDP, particularly that which is available without a direct cost, and there is also an emerging view that it may bypass GDP entirely.[4]

How might this happen? The Internet accelerates the flow of information, and the increased flow can increase the utility that a consumer derives from a given amount of income. The mechanisms at work, here, include an improved consumer awareness of alternative options, more timely access to information, and superior matching of goods to wants. An important implication is that a general increase in the availability of information can increase consumer utility *without an increase in GDP*. The economy may remain at the same equilibrium point on its production possibility frontier but the utility associated with being at this equilibrium has increased. Moreover, the growth in consumer welfare over time may reflect both improvements in the efficiency of production and improvements in the efficiency of consumption.

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

---

[4] In her book on the history of GDP, Coyle (2014 ) concludes that "Gross domestic product is a measure of the economy best suited to an earlier era (p. 125)." Feldstein (2017) reaches a similar conclusion: "A great deal of effort and talent has been applied over past decades to the measurement of real income and inflation. These problems are extremely difficult. In my judgement, they are far from being resolved, and as a result, substantial errors of unknown size remain in our ability to measure both real output and inflation (p. 161)". The point was put even more forcefully by Nordhaus (1997) in his analysis of the history of lighting, where he suggests that official price indexes may well "miss the most important revolutions in economic history" because of the way they are constructed (pp. 54-55). Others point to the need to look beyond GDP (as, for examples, Ahmad and Schreyer (2016), Brynjolfsson et al (2017), Nakamura (2014), Nakamura, Samuels and Soloveichik (2016), Hulten (2015), and Varian (2009, 2016).

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* (1966a), 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.[5] 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.

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,

---

[5] 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 on line, using a smartphone or other mobile device. GPS and traffic maps are often of great utility when travelling, as is immediate access to health information in times of need. Timely access to general medical information can also be of great value.

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 slowdown in real GDP growth, but there is the larger theoretical question of how, and whether, consumer information should be included in measured GDP.

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 was that it was largely absent. He argued that the standard model of economic theory was so closely wedded to the formal mathematics of optimization that it took as given the information needed for the optimization process. Hayek framed 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].

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) proceeded along much the same conceptual path in his analysis of price dispersion and the prevalence of advertising expenditures. He took 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 this 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 and monopolists” [page 213].

Both Hayek and Stigler emphasized that the link between consumer goods and consumer preferences cannot be treated as “a datum”. Five years later, Lancaster (1966a) 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 (1966b), 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).[6] This model intentionally abstracts from many hard real world problems like imperfect information or uncertainty that 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.

---

[6] 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.

## 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.[7] 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, $C_t$, are functionally connected to outputs, $Q_t$. In this case, $C_t = BQ_t$, 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 $U(C_t) = U(BQ_t)$. In the conventional formulation of utility theory, goods and commodities are identical and $B = 1$. In a more general form, one that will be used in this paper, the consumption technology is $C_t = g(Q_t)$. It indicates that different levels of utility can be obtained from a given $Q_t$, 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 straight-forward extension to accept the possibility that increases in information could lead to increases in utility $U(g(Q_t))$ holding $Q_t$ constant. 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 wellbeing using the information

---

[7] "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 (1966a), page 133). Subjective factors like ambiance, mood, and novelty matter.

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".[8]

## 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 else. 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.

---

[8] Since the objective 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).

The standard version of the growth accounting starts with the aggregate production function. In this paper, we assume this function has the Cobb-Douglas form with constant-returns-to-scale and Hicks'-neutral technical change:

$$(1) \qquad Q_t = e^{\lambda t} (R_t)^{\alpha} (E_t)^{\delta} (S_t)^{\pi} (L_t)^{1-\alpha-\delta-\pi}.$$

This function relates the units of output produced ($Q$) to the inputs of intangible capital (the stock, $R$), tangible Information and Communication Technology capital (ICT) equipment (the stock, $E$), and other non-ICT capital *(*the stock, $S$*)*, as well as labor input ($L$). Output grows over time as the inputs increase or as technical change improves the productivity of outputs (here at the rate $\lambda$).[9] The parameter $\alpha$ is the intangible capital's output elasticity, $\delta$ is ICT's elasticity, and $\pi$ is the non-ITC elasticity. Under the assumption of constant returns to scale, elasticities sum to one and ($1-\alpha-\delta-\pi$) is the residual labor elasticity. In this case, the production function (1) can be expresses in "intensive' form as

$$(1') \qquad Q_t/L_t = e^{\lambda t} (R_t/L_t)^{\alpha} (E_t/L_t)^{\delta} (S_t/L_t)^{\pi}.$$

The growth equation associated with *(1')* can then be expressed in terms of output per worker as:

$$(2) \qquad q - \ell = \lambda + \alpha (r-\ell) + \delta (e-\ell) + \pi (s-\ell) .$$

Lower case letters denote, here, rates of growth. This formulation is based on the output elasticities ($\alpha$, $\delta$, and $\pi$), but could equally be formulated in terms of the corresponding shares in factor income ($v_R$, $v_E$, and $v_S$) under the assumption of competitive factor pricing. This is the way

---

[9] As a conceptual level, the technical change parameter $\lambda$ allows for costless improvements in productive efficiency. It includes resource-saving technical change arising from the diffusion of production techniques and organizational practice. This diffusion drives a wedge between the private return to the innovator and the overall social rate of return, and the literature survey by Hall et al. (2010) found the social return to R&D was "almost always estimated to be substantially greater than the private returns (page 1073)". However, because it is measured as a residual, the estimated $\lambda$ also includes the effects of economic fluctuations and shocks, and background inputs like physical and regulatory infrastructure, as well as pure measurement error. The result is what Abramovitz (1956) terms a "measure of our ignorance".

Solow (1957) proceeds in his derivation of the TFP residual. The shares can be computed from accounting data, and are the inputs and output, leaving the shift factor $\lambda$ to be estimated as a residual.

The growth in output per worker is often used as an indicator of the growth in wellbeing enabled by the process of economic growth. Equation *(2)* indicates the growth in output per worker will increase when there is an increase in the productivity with which resources are used, $\lambda$, and when there is more capital per worker, in its various forms, weighted by their respective output elasticities (or income shares).

In this framework, technological innovation, in its broadest sense, involves the first three terms on the right-hand side of *(2)*: costless increases in productivity, $\lambda$; and the deepening of intangible capital stocks like R&D and coinvestments in ICT, $\alpha(r-\ell)$ and $\delta(e-\ell)$. The first is "resource-saving" innovation associated with the shift in the production function ($\lambda$); the second and third are "resource-using" innovation associated with the growth in intangible capital and ICT equipment. However, innovation also occurs in non-ITC capital via embodied technical change.[10]

The taxonomy of innovation based on the decomposition shown in *(2)* follows the conventional practice of focusing on the sources of output growth originating on the supply side of the economy. This focus implicitly ignores the possibility that innovation can also occur in the consumption of goods, and specifically, the consumption technology might also shift over time. It ignores, in effect, the "free, copious, and ubiquitous" information of Schmidt and Rosenberg's Internet Age. Our proposed remedy has two components. First, we assume that conventional

---

[10] It might also be noted that costless increases in the quality of capital and intermediate goods that go unmeasured will appears as increases in $\lambda$.

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 measureable —units of output largely transacted in markets — to one that is subjective and for which no directly measureable 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 wellbeing, particularly when there is reason to believe that this is how many of the benefits of the digital revolution are realized.[11]

The incorporation of a utility function into growth accounting (with or without the consumption technology) is perhaps the largest deviation from orthodox growth accounting, 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 $Q$ determined by the production function *(1)* feeds into the consumption side of the economy, giving $U = U(Q)$.[12] 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,

---

[11] 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 (1957) residual is derived from the shift in the production function, and Jorgenson and Griliches (1967) rooted the residual even more deeply in production theory and are largely responsible for its current form. The residual, as defined in equation *(3)*, is a differential equation whose solution is a matter of line integration. 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). The 1992 paper by Basu and Fernald provides a valuable elaboration of the difference between technology and welfare growth.

[12] Since our interest in this paper is on the aggregate output of growth, not how the benefits of growth are distributed across a population of heterogeneous agents, our formulation of utility therefore adopts the single representative agent approach.

*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 shoehorned into an expanded measure of GDP.

Once growth accounting is expanded to include utility and the consumption on which it is based, 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. In that theory, an increase in *Q* is generally assumed to have a positive marginal utility, but the rate of change in marginal utility is ambiguous. It is generally assumed to be negative, supporting the case for progressive income taxation, but in the growth context, a declining rate implies that a steady rate of real GDP growth brings progressively less additional wellbeing to the representative agent. Conventional growth theory, on the other hand, implicitly assumes that this marginal utility is unitary, implying that the growth rate of utility is identical to the growth rate in real GDP. This, in turn implies that the latter is a valid proxy for the former.

We introduce the marginal utility of income, whatever its magnitude, into our framework using a simple one-parameter approach in which utility of output is an exponential function of the marginal utility parameter $\mu$. A simple representative agent model of average utility per capita in the aggregate economy might then have the form

$$(3) \quad U(C_t/N_t) \; = \; m\,(C_t/N_t)^{\mu} = m[\rho_t\,(1-\sigma_t)(Q_t/L_t)]^{\mu}$$

This equation links consumption per capita, $C_t/N_t$, to output per worker, $Q_t/L_t$. The rate of saving, $\sigma$, is needed, here, because utility depends on consumption and not total output. A fraction $\sigma$ of the output $q_t$ is diverted to capital formation in the form of investment in order to build up the capital stocks in *(1)* and thus allow for an increase in future consumption. Thus, contemporaneous $C_t = (1\text{-}\sigma_t)Q_t$. Moreover, the population variable $N_t$ differs from the labor variable $L_t$, but they are connected by the labor-force participation rate, $\rho_t$, so that $L_t=\rho_t N_t$. We assume, for simplicity, that $\sigma$ is constant and the same for all agents, allowing us to avoid the problem of modeling the utility of future consumption by making saving a fixed proportion of output. We also assume that ρ is constant, implying that a consumer's endowment of time is allocated in a fixed proportion between work and leisure and that the labor-force participation rate is also constant.

Introducing a consumption technology adds a further degree of complexity to the link between GDP and utility. 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, $e^{\omega t}$, is multiplicative, as is the embodied costless product quality effect, $e^{\beta t}$:

$$(4) \qquad U(C_t/N_t) = m\, e^{\omega t} e^{\beta t} [\rho_t (1\text{-}\sigma_t)(Q_t/L_t)]^{\mu}\,.$$

The utility function in *(4)* is a straightforward extension of *(3)* that allows for a shift in the Lancaster consumption technology due to an increase in the amount of consumer information, $e^{\omega t}$, and an increase in costless product quality, $e^{\beta t}$. A more realistic treatment would allow for multiple goods, as illustrated in a two-good version of model in Figure A3 of the appendix, and would also allow for search costs, an explicit information technology, and uncertainty.

When the utility function *(4)* is expressed in growth rate form, the result is an expanded growth accounting equation that combines the Solow growth-accounting equation *(2)* with the costless improvements in product quality and consumer information, $\beta$ and $\omega$. Under the assumption that the marginal utility associated with growth, $\mu$, is equal to one, the expanded sources-of-growth account has the form:

$$(5) \qquad u = \omega + \beta + (q-\ell) = \omega + \beta + \lambda + \alpha\,(r-\ell) + \delta\,(e-\ell) + \pi\,(s-\ell)\ .$$

The first equality in *(5)* indicates that the growth rate of utility is driven by the shift in the consumption technology (the first two terms) and the growth rate of output per worker. This expands the discourse on the benefits of technology and real GDP sources of growth beyond the conventional output effect to include some of the main non-GDP benefits of the information revolution. Equation *(5)* 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 decomposition implied by *(5)* draws a boundary between resource cost and non-cost improvements in welfare, that is, between production and consumption technologies. In *(5)*, costless quality change is assigned to the consumption technology side of this boundary, whereas GDP *measured at resource cost* is located on the other side of the boundary. This is not the way it is treated in the conventional GDP accounting, where GDP is measured (in principle) in units of *effective* output rather than units that reflect the cost of production. The former tend to be the units actually sold in the market place (transaction units), whereas effective output is defined with respect to the benefits received by the user and therefore includes costless changes in product quality. The two output concepts are linked by the rate of costless quality change $\beta$. To make this explicit, we now denote units of the output produced in *(1)* by $Q^{r}_{t}$ rather the $Q_t$ of the preceding sections; we then use the notation $Q^e_t$ to represent the effective units of output as experienced by the consumer. Formally,

$$(6) \qquad Q^e_t = e^{\beta t}\, Q^r_t = e^{\beta t}\, e^{\lambda t} f(L_t, R_t, E_t, S_t).$$

(We have used a generic technology *f* because the issue addressed here is more general than the Cobb-Douglas-form of this paper). Costless changes in product quality are regarded in *(6)* 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]

The shift in accounting convention from $Q^r_t$ to $Q^e_t$ changes the measured quantity of the goods produced. However, while the quantity real GDP is changed, nominal GDP is not. Even though $Q^e_t$ delivers more effective output than $Q^r_t$ when $\beta$. is positive, and a transaction unit can still be purchased for its cost, $P^r_t$, so the effective price must fall in order to maintain the equality $P^e_t Q^e_t = P^r_t Q^r_t$. This implies that nominal price GDP is invariant to the adjustment for product effectiveness, even though real GDP has risen.

The growth equations associated with *(6)* change in a way that mixes costless product and process innovation:

$$(7) \qquad q^e_t = [\beta + \lambda] + \alpha\,(r-\ell) + \delta\,(e-\ell) + \pi\,(s-\ell)\ .$$

Again, this formulation of GDP treats product quality change as a shift on the production function, in keeping with the implicit "better is more" view of product quality. A Solow residual

---

[13] The intuitive difference between the two ways of looking at product quality change can be seen by comparing appendix Figure A2, where the production function (here for two goods) shifts upward, with Figure A3, where the utility function shifts inward. Analytically, when product quality change is treated as resource-saving, and in Figure A2, the shift in the production function is then $\beta + \lambda$. This combined effect is what gets measured by the TFP residual.

based on *(7)* therefore conflates costless changes in process productivity with costless changes in benefits of product innovation as perceived by the consumer.

Much of the existing analysis of recent growth trends uses $Q_t$ as the baseline rather than $Q^r$. This line of argument presumes that $Q^e{}_t$ is, in fact, the appropriate concept of output, and that a better product can be expressed in equivalent units of the older, less preferred, output. This is, at best, a necessary fiction. A new medicine that cures 80% of the population of a particular disease with one pill a day is generally not equivalent to two pills that can reach only 40% of the population. A modern metal tennis racket isn't equivalent to N wooden rackets, etc. The "better-is more" assumption has some resonance in digital goods subject to the improvements enabled by Moore's law, but even there, a dramatic increase in computing power typically enables new applications, not just faster execution of existing ones.

How well the goal of measuring $\beta$ and $q^e$ is actually achieved is also problematic. The estimation of $\beta$ is largely approached from the price side of the accounts by adjusting "official" price indexes for quality change. A significant body of research has examined the question of systematic bias in these indexes, and three major commissions have been empanelled to study the problem and make recommendations.[14] In addition, Moulton and Moses (1997) and Groshen et al. (2017) have offered appraisals from a statistical agency perspective. The general thrust of this literature is that official estimates of $\beta$ are biased downward and thus understate the quality change that has occurred, but there is less consensus about the magnitude of the bias. At one end

---

[14] The commissions include the 1961 *Price Statistics Review Committee*, or Stigler Commission; the 1996 *Advisory Commission to Study the Consumer Price Index*, or Boskin Commission; and the 2002 National Research Council *Panel on Conceptual, Measurement and Other Statistical Issues in Developing Cost-of-Living Indexes*, or Schultze commission. Other general appraisals have been made by Shapiro and Wilcox (1996), Nordhaus (1997), Bils and Klenow (2001), Lebow and Rudd (2003), Bils (2009), and Feldstein (2017).

of the range, Nordhaus suggests that it could be extremely large, whereas Groshen et al. suggest that the main source of bias may be in medical goods, while the bias in digital goods may be relatively small (they cite the studies by Lebow and Rudd, 2003, and Greenstein and McDevitt, 2011). Groshen et al. point to the small GDP share of consumer electronics and the internet (0.6%), so that even though the Lebow-Rudd estimate of the bias itself grows at a very high rate of 6.5% per year, the product of two is too small to affect the overall growth of real GDP in economy.

However, one of the advantages of the framework set out in this paper is that the $\beta$'s are located in the welfare side of our accounts, where the GDP share *per se* is not relevant and a 6.5% annual bias can be of much greater significance when combined with a larger welfare share. This also applies to health care. An effective vaccine for a serious affliction (e.g., polio) may have a minimal GDP share but convey very large welfare benefits. A similar argument applies to the information available over the internet at a marginal price of zero.

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 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.[15] 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 home-made meal reflecting many of the factors that determine the market price of meals. The larger point is that there are two technologies associated with the household, one for production and the other for consumption, and that both can shift.

## V. 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 *(5)*, since the left-hand side variable is consumer utility. Because the utility variable is subjective and not directly observable, it is useful to recast the analysis in one that is: consumer expenditure. 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 generic expenditure function associated with the utility function *(3)* has the form $e(P_c, U^*) = e^{-\omega t}\,\xi(P_c, U^*)$. 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

---

[15] Owner-occupied housing is another example that could be cited. Buying a house that you had previously rented moves you across the market/nonmarket 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.

when costless information increases, since the minimum expenditure needed to purchase a given level of utility decreases.[16] 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 unadjusted $q$, implying $P_cC = P_qQ$, where $P_q$ and $P_c$ now denote corresponding prices, and the prices are related by $P_c = e^{-\beta t} P_q$ . The expenditure function then becomes

$$(8) \qquad e(P_c, U^*) = e^{-\omega t}\ \xi(e^{-\beta t} P_q, U^*).$$

The growth rate of expenditures over time depends on the expenditure-share weighted growth rates of the prices, the negative growth rate of the parameters $\omega$ and $\beta$, and the rate of change of $U^*$. This analysis can also be framed using the indirect utility function associated with the expenditure function.

The factor price frontier associated with a generic constant-return Hicks'-neutral production function $Q = e^{\lambda t} f(L,K)$ has the form $P_q = e^{-\lambda t}\ \phi(P_L, P_K)$.[17] Substitution into *(8)* gives

$$(9) \qquad e(P_c, U^*) = e^{-\omega t}\ \xi(e^{-\beta}\ e^{-\lambda t}\ \phi(P_L, P_K), U^*).$$

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 *(4)*.

---

[16] This is shown in Figure A3 of the appendix as a downward shift in both the utility function, with the result that the original level of utility can be achieved with less income and expenditure. Alternatively, the original level of income and expenditure can finance a higher level of utility (Figure A4), giving "more bang for the buck".

[17] 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.

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 $P^*_c$:

$$(10) \qquad CV = e(P^*_c, U_1) - e(P^*_c, U_0)$$

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.

## VI. Consumer Surplus and New Technology

Special mention should be made of consumer surplus in light of the preceding discussion. 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 the arrival of new goods in the market place. 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

---

[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).

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. Again, 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.

Consumer surplus is a valuable technique for getting at a difficult measurement problems, and it has been applied to the problem of estimating the consumer benefit.[19] However, it is also 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, moreover, 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

---

[19] A number of recent studies have applied consumer surplus techniques in various ways to the digital economy, some involving the valuation of time and others willingness to pay. A partial list includes Goolsbee and Klenow (2006), Varian (2009, 2016), Greenstein and McDevitt (2011), Brynjolfsson and Oh (2012), Chen et al. (2014), Nevo (2016), Brynjolfsson et al. (2017), and Cavallo (2017). For an expenditure function approach, see Redding and Weinstein (2017). This list does not include references to the large literature on health care.

per unit consumed (or observable units). Instead, a general internet access fee 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 are able to generate, or are provided *pro bono publico* through such activities as crowd-sourcing. 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.[20] 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

---

[20] 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.

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. 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.[21]

## VII. Contingent Goods

The private service-producing sectors of the U.S. economy have grown significantly as a share of GDP over the last half century, from around 50% in 1960 to 68% in 2015 according to estimates by the Bureau of Economic Analysis. These data also show that the GDP share of services that involve "expert" advice, information, or interventions — finance and insurance, professional and business, education, health, and information services — rose from around 14% to 33%. These are part of what Griliches (1994) called the hard-to-measure sectors of the economy, and he also observed in his 1992 paper that a measurement "problem arises because in many services sectors it is not exactly clear what is being transacted, what is the output, and what

---

[21] 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.

services correspond to the payments made to their providers" (page 7). There is a wedge between the output of a good and its outcome, arising from the contingent nature of many goods, and is particularly prevalent in those involving these "expert" industries (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 is contingent on the initial state of health, the doctor's input, and the actions taken by the patient 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. In the case of education, schools may provide educational services, but learning also depends on student motivation, ability, and family inputs.

Contingent goods have a natural interpretation in the Lancaster framework, separate from, but related to, the quality change parameter $\beta$. The output from the standpoint of the consumer is the improvement (objective and subjective) in the initial state of being (health, knowledge, financial or legal status). This initial state, call it $\Theta_{t-1}$, can be regarded as a determinant of utility in the Lancaster framework, and the utility function expanded accordingly. One way to approach this in our simple model is to make utility a function of the final state $\Theta_t$, and link $\Theta_t$ to the direct resources used to produce the service, $Q^r{}_t$, the rate of costless quality change, $\beta$, and a variable $X_t$ representing ancillary actions and expenditures (health club membership, new mattress, stopping smoking, taking medications on time). The expanded function might then have the form $U(\Theta_t) = U(e^{\beta t} Q^r{}_t, X_t; \Theta_{t-1})$. This expansion introduces a degree of recursion into the problem, since the initial state variable $\Theta_{t-1}$ was itself determined by the utility process of the preceding year. Its purpose is to show that the framework of the consumption technology can be applied to an important class of services, and is not restricted to information goods or digital goods undergoing quality change. It also suggests that part of the welfare gain associated with goods like health care belong to the welfare side of our expanded accounting framework rather that the production side (the wedge noted by Griliches,1992). Indeed, it may be helpful in advancing what Cutler and Berndt (2001) call the "outcomes movement" in health economics, which is the attempt to measure the health impact of medical care rather than the amount expended.

This perspective is all the more important because there are cases in which the size of the wedge on the welfare side of the account can be large despite the GDP share being small.. The simple and essentially costless practice of having doctors wash their hands prior to performing operations, which began to gain traction only in the latter part of the 19th Century, produced large gains in surgical outcomes and patient wellbeing. Moreover, vaccines that prevent terrible diseases like small pox and polio are almost trivial in their resource cost compared to the huge welfare gains produce in avoiding death, crippling incapacitation, and disfigurement.

One final point should also be made about the nature and value of expert advice. Information that is not organized or focused on informing a specific question or issue is often 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, and this is an important source of value of the internet.

## VIII. Technical Innovation with a Resource Cost

Our formulation of the expanded growth account *(5)* includes both costless and resource costly innovation. Because of their different implications for welfare, as well as growing importance of the latter, a deeper look is in order. Costless innovation arises from several sources. First, there are spillovers arising from the failure to protect property rights to costly innovations where rights are hard to enforce because of the non-rival nature of the good and the free-rider problem. This often occurs with information goods like the internet, as well as with the technology involved in product and process innovation. The non-rival nature of goods and the free-rider problem may even lead producers to distribute their products without a direct charge, recouping the cost of production with other revenue sources like advertising. Second, there is what Eric von Hippel (2016) calls "free innovation". This includes contributions to the common good through crowd-sourcing and *pro bono* innovations like open-source software. Finally, there is just plain inspiration and creativity. Costless innovation appears in the expanded growth account *(5)* through the term $\lambda$ in the production technology and the terms $\omega$ and $\beta$ of the consumption technology.

Costly innovation, on the other hand, results from systematic investments in innovation. Firm-specific own-account investment in intangible capital like R&D and its co-investments has become the dominant form of business investment in the U.S. and is at the forefront of this kind of change.[22] The outcome associated with resource-costly investments in R&D and organizational development is often an improvement in the processes of production, $\lambda$, and product quality and new goods, $\beta$. However, unlike before, costly innovations in products and processes are "purchased" at the cost of the resources. This presents the consumer with the opportunity to buy "better" goods at a higher price that reflects that cost (they are not a gift,

---

[22] Corrado, Hulten, and Sichel (2005, 2009). See also Nakamura (2001).

"Manna from Heaven," as before). This implies that the transaction units of the improved good already embody the effects of the innovation, which, in turn, implies that the *produced* $\beta^*$ belongs on the resource side of the expanded growth account in *(5)*, whereas the costless $\beta$ belongs on the welfare side. Unfortunately, costly product quality change, $\beta^*$, is typically buried in prices of the transaction units and not recorded separately. It should, nevertheless, treated as a separate effect and recorded on the resource side of the taxonomy.

Figure 1 compares the average annual growth rates of the different types of technical change for the U.S. private business sector. This figure goes beyond the "usual" sources-of-growth representation in that it includes the full range intangible capital studied by the Corrado, Hulten, and Sichel (2009), although the estimates shown in the figure are based on Corrado and Hulten (2010, 2014). The $\lambda$ is proxied by the Solow residual estimate of TFP (adjusted for the presence of intangible capital) and is the costly part of innovation by the variables $\alpha(r- \ell)$ and $\delta(e-\ell)$. The TFP series is very volatile, generally moving up and down around a trend following the business cycle, even when shown as a long moving average. The trend declined sharply from the mid-1950s through the early 1980s, recovered somewhat during the 1980s, then followed a volatile boom and bust path, ending with a sharp drop after 2005.

The Solow residual sweeps together many factors other than $\lambda$, like fluctuation in economic activity but others as well. However, it is evident that $\lambda$ has trended downward since the 1950s. In contrast, the growth rate of the share-weighted stock of intangible capital, $\alpha(r-\ell)$ has trended upward in recent years and has grown to the point it rivals TFP and offsets part of the decline.

The upper line in the figure combines these two measures and adds investment in ICT, $\delta(e-\ell)$, to obtain an even broader measure of innovative activity in the digital economy. The time path of the broader measure is still dominated by the volatility of TFP, and the downward trend through the 1970s is apparent. However, it is also apparent that innovative activity from 1985 through 2005 experienced two peaks that are comparable to those in the 1950s and 1960s. This suggests that innovation, as it affects real GDP, is not necessarily in decline, but instead that its character has changed from the free "Manna" of TFP growth toward the costly systematic investments needed to obtain it.

What is missing from Figure 1 is a systematic estimate of output-saving innovation. Estimates for selected components of this kind of innovation have appeared in recent years, but comprehensive time series estimates do not exist. The general theme of this paper is that the picture of innovation is incomplete without these non-GDP contributions to welfare, and that an expanded version of Figure 1 derived from a more sophisticated variant of the expanded accounting model set out in equation *(5)* of this paper is needed. What the historical statistics of an alternative Figure 1 would show is thus unknown. In light of Gordon (2016), it may well the the case the output-saving technical change was larger in the past than in the present, despite the recent growth of the internet and it applications.

We note, finally, that the time-cost associated with consumption is also missing from the figure. We have finessed the work-leisure decision by assuming that time is allocated in fixed proportions between work and leisure, but recognize that a more sophisticated version of our analysis would recognize that the consumption technology requires a time input, just as time-use enters the production technology through labor input. The information revolution has reduced the time required for many activities and thus a saving in time cost. However, the advent of new or improved goods may also involve start-up costs and a learning curve, and the consumption of goods takes time (as in watching television or communicating via social media). An extension of our model, perhaps along the lines of Stigler and Becker (1977), would enrich our analysis, but would not change our basic conclusion about the importance of the consumption technology for understanding the effects of the information revolution.

## IX. 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.[23] The key feature of our model is the 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 alone may not be a sufficient 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

---

[23] The essential steps in developing equations *(3)*, *(4)*, and *(5)* 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 involve the wedge between the resource cost of a good and its 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.

users to combine the various elements of the expanded accounts in ways best suited to their needs.

We have not attempted to estimate the magnitude of output-saving technical change, 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."

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 mismeasurment when even the approximate size of the "correct" measure is unknown. The unmeasured output-saving value of information alone is potentially very large, and it has grown rapidly in recent years. According to Census data, the fraction of U.S. households with a computer at home rose from 23 percent in 1993 to almost 80 percent in 2012, and the fraction with internet use at home went from 18 percent in 1997 to nearly 75 percent in 2012. The PEW Research Center also found that the percentage of adults who use at least one social media site increased from seven percent in 2005 to 65 percent in 2015, and other PEW surveys found that the market penetration of smart phones more than doubled from 2011 to 2016, from 35 percent to 77 percent. The rapid uptake of these goods may not be enough to offset the declining growth rate of measured real GDP given their relatively small share of GDP. But, as noted earlier, the GDP share (small or large) may not be the correct indicator of their importance to the consumer given the way these goods are priced and, specifically, the fact that much of the information now available to consumers is distributed without a direct charge.

Moreover, a piecemeal analysis of what Nordhaus (1997) termed a "tectonic" revolution may miss the bigger picture. A tectonic revolution is a situation in which "changes in production and consumption are so vast that the price indexes do not attempt to capture qualitative changes." We now appear to be in the midst of another tectonic revolution based on the rapid advances in the generation, transmission, use, and storage of information. 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. As before, parts of the statistical system have struggled with new modes of production and distribution, as well as with "free" information goods. Part of the struggle arises from the attempt to shoehorn the "new" digital economy into a GDP framework designed to measure the "old". This effort reflects, in part, the enormous challenges involved in developing the necessary metrics of a new conceptual framework. One step forward is to recognize that GDP by itself may not be a sufficient yardstick and that other metrics of consumer benefit are needed.

Charles Hulten, University of Maryland and NBER

Leonard Nakamura, Federal Reserve Bank of Philadelphia

# 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.

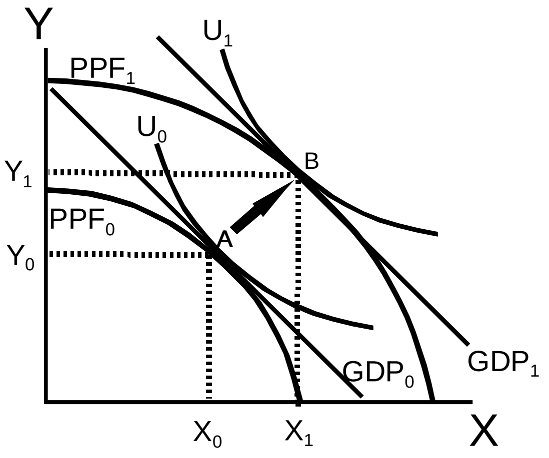

**Figure A1**

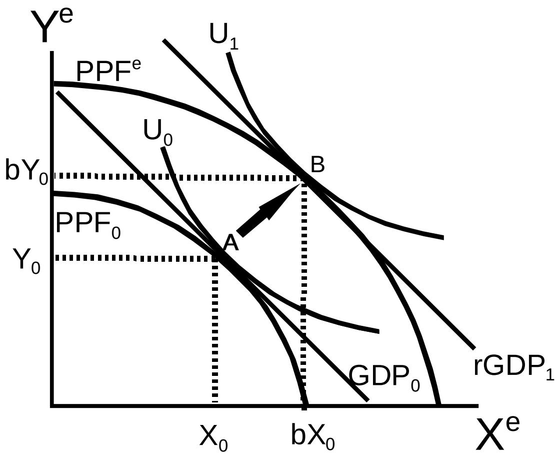

**Figure A2**

An improvement in the productivity process with which *X* and *Y* are produced shifts the PPF outward (analogous to the $\lambda$ 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 the 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 $\beta$ 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 $(\mathrm{X}_0,\mathrm{Y}_0)$ units of the goods is now the equivalent of $(b\mathrm{X}_0,b\mathrm{Y}_0)$ units from the standpoint of the utility they provide. In other words, at point *B*, $(X^e{}_1,Y^e{}_1)$ is equivalent to $(b\mathrm{X}_0,b\mathrm{Y}_0)$. 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 a 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 latter is now $(X_1,Y_1)$. This is the *output-saving technical change* described in the paper.

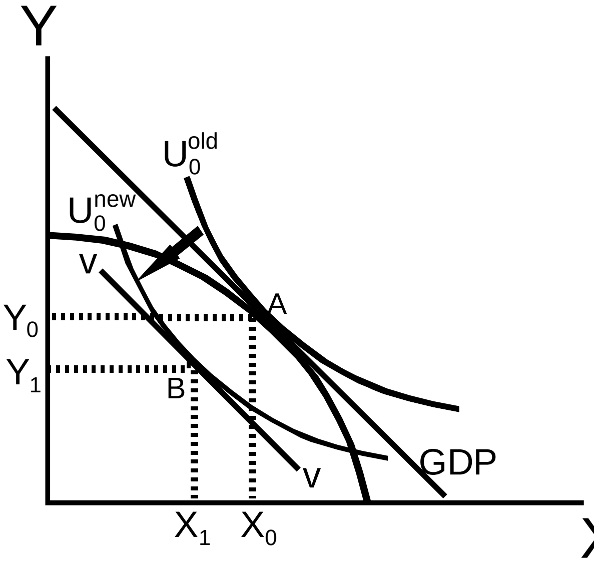

**Figure A3**

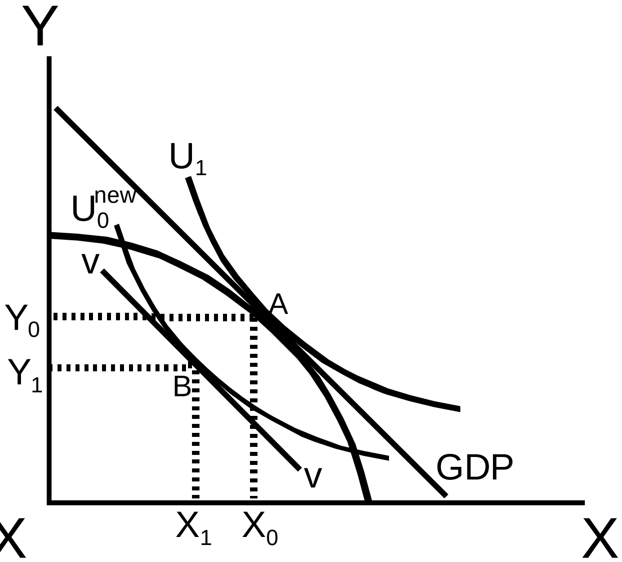

**Figure A4**

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 4 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,

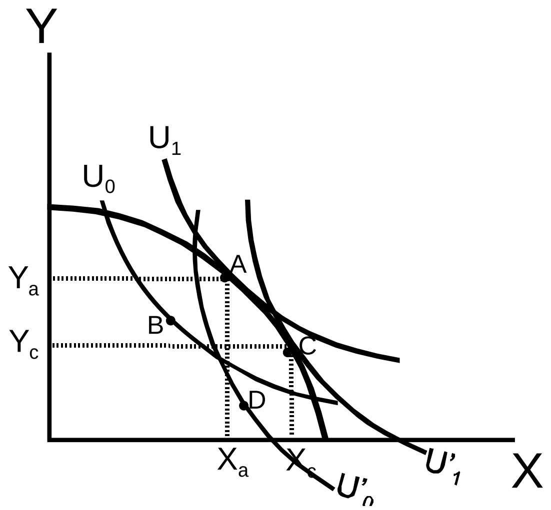

**Figure A5**

$U_0$, 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 Affect 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.

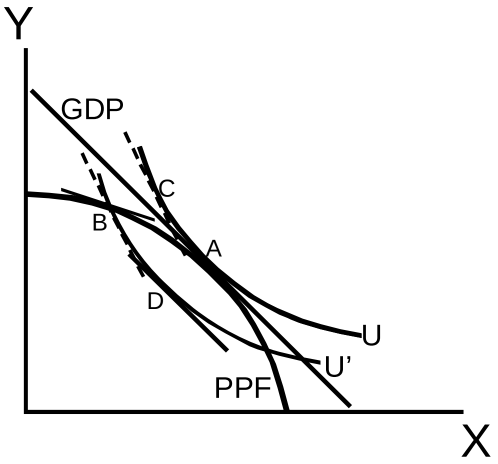

**Figure A6**

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.

# REFERENCES

Abramovitz Moses, “Resource and output trends in the United States since 1870”, *American Economic Review*, 46(2), 1956, 5–23.

Advisory Commission to Study the Consumer Price Index, T*oward A More Accurate Measure Of The Cost Of Living*, *Final Report to the* Committee on Finance, United States Senate, U.S. Government Printing Office, 36-352-CC, Washington D.C., December 1996.

Ahmad, Nadim, and Paul Schreyer. “Measuring GDP in a Digitalized Economy,” OECD Working Paper, April 2016.

Anderson, Monica,. “Technology Device Ownership: 2015,” Pew Research Center, October 29, 2015.

Basu, Susanto, and John G Fernald, “Aggregate productivity and aggregate technology”, *European Economic Review*, 46, 6, June 2002, 963-991.

Bils, Mark, “Do Higher Prices for New Goods Reflect Quality Growth or Inflation?, *The Quarterly Journal of Economics*, 124, 2, May 2009, 637-675.

Bils, Mark, and Peter J. Klenow. "Quantifying Quality Growth." *American Economic Review,* 91(4): May 2001, 1006-1030.

Brynjolfsson, Erik, Felix Eggers, and Avinash Gannamaneni, “Using Massive Online Choice Experiments to Measure Changes in Well-being,” MIT Working Paper, September 2017.

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

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.

Cavallo, Alberto, “Are Online and Offline Prices Smilar? Evidence from Large Multi-Channel Retailers,” *American Economic Review* 107, January 2017, 283-303.

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 .

Corrado, Carol A. and Charles R. Hulten, “How Do You Measure a Technological Revolution?,” *American Economic Review*, May 2010, 99-104.

Corrado, Carol, and Charles Hulten, “Innovation Accounting”, in *“Measuring Economic Progress and Economic Sustainability,”* in D.W. Jorgenson, J.S. Landefeld, and P. Schreyer, eds., Studies in Income and Wealth, vol. 72, The University of Chicago Press for the National Bureau of Economic Research, Chicago, 2014, 595-628.

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.

Corrado, Carol and Larry Slifman “A Decomposition of Productivity and Costs”, *American Economic Review* 89 (2, May), 1999, 328-332.

Coyle, Diane,, *GDP: A Brief But Affectionate History*, Princeton: Princeton University Press, 2014.

Cutler, David M., and Ernst R. Berndt, eds., “Introduction”, *Medical Care Output and Productivity*, Studies in Income and Wealth, Vol. 62 , Chicago, Illinois: University of Chicago Press, 2001, 1-11.

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

Feldstein, Martin, “Underestimating the Real Growth of GDP, Personal Income, and Productivity”, *Journal of Economic Perspectives*, Volume 31, Number 2, Spring 2017, 145–164.

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.

Griliches, Zvi, “Introduction”, *Output Measurement in the Service Sectors*, Zvi Griliches, editor, Studies in Income and Wealth, No. 56, University of Chicago Press, 1992, 1-22.

Griliches, Zvi, "Productivity, R&D, and the Data Constraint," *American Economic Review*, 84, 1, March 1994, 1-23.

Groshen, Erica L., Brian C. Moyer, Ana M. Aizcorbe, Ralph Bradley, and David M. Friedman, “How Government Statistics Adjust for Potential Biases from Quality Change and New Goods in an Age of Digital Technologies: A View from the Trenches", *Journal of Economic Perspectives*, 31, 2, Spring 2017, 187-210.

Hall, Bronwyn, Jacques Mairesse, and Pierre Mohnen, "Measuring the Return to R&D," *Handbook of the Economics of Innovation*, Chapter 24, B. H. Hall and N. Rosenberg (eds.), Elsevier-North Holland, 2010, 1033-1082.

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 and 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.

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”, *The Journal of Political Economy*, 74, 2, April 1966b, 132-157.

Lebow, David, and Jeremy B. Rudd., “Measurement Error in the Consumer Price Index: Where Do We Stand?” *Journal of Economic Literature*, 41(1), 2003, 159–201.

Moulton, Brent R. and Karin E. Moses, “Addressing the Quality Change Issue in the Consumer Price Index,” *Brookings Papers on Economic Activity 1*, 1997, 305-366.

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, Jon Samuels and Rachel Soloveichik, “Valuing “Free” Media: An Experimental Approach,” Working Paper 16-24, Federal Reserve Bank of Philadelphia, August 2016.

National Research Council,, *At What Price? Conceptualizing and Measuring the Cost-of-Living and Price Indexes,* Panel on Conceptual, Measurement and Other Statistical Issues in Developing Cost-of-Living Indexes, Charles L. Schultze and Christopher Mackie, eds. Committee on National Statistics, Division of Behavioral and Social Sciences and Education. Washington, DC: National Academy Press, 2002.

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

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.

Perrin, Andrew, "Social Networking Usage: 2005-2015," Pew Research Center. October 2015.

Pew Research Center, "Mobile Factsheet," January 12, 2017.

Redding, Stephen J., and David E. Weinstein, , "A Unified Approach to Estimating Price and Welfare Measurement", Working Paper, March 2017..

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., "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.

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

Stigler, George, ed., "The Price Statistics of the Federal Government: Review, Appraisal, and Recommendations." *Report to the Office of Statistical Standards*, Bureau of the Budget. New York: National Bureau of Economic Research, 196`.

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.

Triplett, 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.

U.S. Department of Commerce, Bureau of Economic Analysis, GDP-by-Industry, Industry Data, *Value Added by Industry as a Percentage of Gross Domestic Product*, November 2016.

U.S. Department of Commerce, Bureau of Economic Analysis, GDP-by-Industry, Industry Data, *Chain-Type Price Indexes for Value Added by Industry*, November 2016.

U.S. Department of Commerce, Bureau of the Census, “Computer and Internet Access in the United States: 2012”, February 2014.

U.S. Department of Labor, Bureau of Labor Statistics, *Trends in Multifactor Productivity, 1948-81*, Bulletin 2178, USGPO, Washington D.C., September 1983, and updates.

Varian, Hal, “Economic Value of Google”, 2009: http://cdn.oreillystatic.com/en/assets/1/event/57/The%20Economic%20Impact%20of%20Google%20Presentation.pdf

Varian, Hal. 2016, “A Microeconomist Looks at Productivity,” BEA Advisory Council Meeting, December https://www.brookings.edu/wp-content/uploads/2016/08/varian.pdf

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

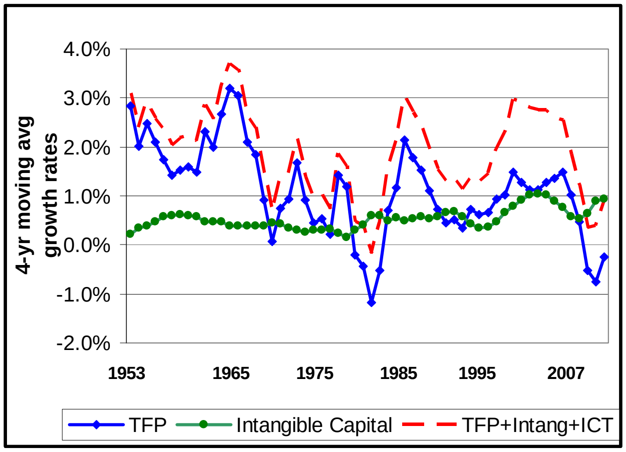

**FIGURE I**

**Growth Rates of TFP and Intangible and ICT Capital per Worker Hour**

**U.S. Non-Farm Business, 1953-2010**

**Source: Data underlying Corrado and Hulten (2010, 2014)**