

New Products and Long-term Welfare

Philip Trammell*

August 7, 2021

1 Introduction

A middle-class member of the developed world today has access to foods, medicines, electronics, and more to which not even the world's richest had access, say, five hundred years ago. These new goods and services plausibly leave her better off than the kings of the past, even though the kings of the past had access to dramatically more of the products available at the time. This is so even though these figures often held assets and enjoyed consumption baskets whose values at current prices far exceed modern middle class net worths and consumption expenditures, and even though these figures are, accordingly, typically considered among the richest people ever to have lived.

A common interpretation of this observation is that inflation is underestimated. I will argue that this interpretation is mistaken. In some cases, and likely in the empirical case, no adjustments to price indices can allow for welfare-relevant unidimensional consumption comparisons across periods following the introduction of new products. Failing to recognize this impossibility confuses our attempts to understand the extent to which people today are better off than in the past—and, more decision-relevantly, the extent to which changes in the rate of economic growth will affect how much better off people are in the future.

Instead, therefore, I propose a basic framework for understanding growth which can better capture the full range of possible relationships between growth and long-term welfare in light of new product introduction. I then find the conditions under which the framework is compatible with the relevant stylized facts of growth, and I give the resulting model a simple microfoundation. Finally, I note an immediate implication of a framework of this form: it successfully predicts a difference, in the observed direction, between relative risk aversion and intertemporal substitution elasticity—as, for instance, appears to underlie the equity premium puzzle—without

*Global Priorities Institute and Department of Economics, University of Oxford.
Contact: philip.trammell@economics.ox.ac.uk.

having to invoke nonstandard preference structures, such as Epstein-Zin preferences, that separate these parameters explicitly.

2 Framework

Given a continuum of consumption goods $i = 0$ to N , consider the following preferences:

$$u(\vec{c}) = \int_0^N v(c_i) di; \quad v(c_i) \triangleq \max\left(0, \frac{c_i^{1-\eta} - 1}{1-\eta}\right); \quad \eta > 1, \quad (1)$$

where $v(0)$ is defined to be 0. With these preferences, it is clear that the utility level cannot exceed $\frac{N}{\eta-1}$. Fixing N , therefore, utility is bounded. Introducing new products, however, raises N and raises this upper bound without limit.

We will assume that the prices of the goods are equal (and, by convention, equal to 1). It is therefore optimal to spread one's consumption evenly across the goods of which one purchases a positive quantity. Let us denote the common consumption level resulting from spreading consumption evenly across all goods \bar{c} . Spending on good i generates no utility, however, until $c_i > 1$. Consumption will therefore be spread across all available goods only if the utility $v(\bar{c})$ achieved by purchasing the common consumption level on a given good i is at least as great as that generated by spending nothing on i and increasing consumption on another good instead. That is, we must have

$$v(\bar{c}) \geq \bar{c}v'(\bar{c}) \iff \bar{c} \geq \eta^{\frac{1}{\eta-1}}, \quad (2)$$

where

$$c \triangleq \int_0^N c_i di, \quad \bar{c} \triangleq \frac{c}{N}. \quad (3)$$

Unless otherwise stated, we will assume that this condition holds.

Given $\eta > 1$, v must incorporate the floor of 0 in order to avoid the conclusion that, upon introducing a new product (i.e. upon raising N), some of the new product must be consumed in order to avoid a decrease—indeed, an infinite decrease—in utility. This complication stems ultimately from the fact that, given $\eta > 1$, goods are, at least in an interior solution, gross complements; the elasticity of substitution between them is $1/\eta$, which is bounded below 1.

More standard frameworks for new product introduction in growth models avoid this issue by assuming a continuum of goods with (typically common and constant) elasticity of substitution greater than 1. Under this assumption of gross substitutability, however, for any given utility level achievable by the consumption of a range of goods, the same utility level can also be achieved by consuming a sufficiently large

quantity of any narrower range of goods. That is precisely the condition we wish to avoid here. Utility function (1) allows for the intuition that no quantity of the consumption goods available in the past, however large, can reach a utility level as high as that reached by a sufficiently large basket of modern conveniences.

The utility kink at $c_i = 1$ implied by (1) is not necessary for this result, of course. This precise functional form was chosen for simplicity, but v could be replaced by a smooth, concave function with increasing relative risk aversion, so that $v(0) = 0$ even though $\lim_{c_i \rightarrow \infty} v(c_i) < \infty$. Still, it may be worth noting that the assumption that there are fixed costs to adopting a new class of consumption goods is often realistic. One cannot buy a nickel's worth of central air conditioning, for instance, but given that one is spending enough to cover the fixed costs of a basic installation, there is essentially a continuum of quality available. And empirically, of course, people do not spread their consumption over all good types available; as the poor get wealthier, they consume more (and more varieties of) food and clothing, but until they get sufficiently wealthy they consume no air conditioning at all.

3 Inflation measurement

Suppose N_t can grow across periods. We will define $u(c, N)$ to be the maximum utility achievable given total consumption c across the range of N products. That is, assuming as usual that condition (2) holds,

$$u(c, N) \triangleq N \frac{\left(\frac{c}{N}\right)^{1-\eta} - 1}{1-\eta} = N \frac{\bar{c}^{1-\eta} - 1}{1-\eta}. \quad (4)$$

As above, let us stipulate that the price level remains constant at 1, for each good type, at all times. It is then straightforward to see that inflation as measured by a standard price index will always equal 0, and measured consumption will equal c as defined in (3) (though this quantity may change over time). New product introduction will play no direct role in either inflation or consumption figures.

Nevertheless, there is an important welfare difference between an increase in c consisting in the consumption of more products and an increase consisting in more consumption of the same products. Consider consumption growth from period 0 to period 1, where $c_1 = 2c_0$. If $\bar{c}_1 = \bar{c}_0$ but $N_1 = 2N_0$, we have $u_1 = 2u_0$. On the other hand, if no new products are introduced—i.e. if $N_1 = N_0$ and $\bar{c}_1 = 2\bar{c}_0$ —then we have $u_1 < 2u_0$, at least for sufficiently large values of \bar{c}_0 . Given new product introduction, therefore, true utility gains over time may be higher than those estimated under the faulty assumption of isoelastic utility in consumption with fixed N .

In order to construct a measure of consumption that allows for welfare-relevant consumption comparisons across periods, therefore, some may be tempted to measure inflation by a cost-of-living index π_t such that $u(c_1, N_1) = u(c_1 \frac{\pi_0}{\pi_1}, N_0)$. That is, we

might posit deflation across periods, even though the prices of all goods available in both periods are equal, and define real consumption at period t as

$$\tilde{c}_t \triangleq \frac{c_t}{\pi_t}. \quad (5)$$

Real consumption in period 1 will then be measured to be more than double real consumption in period 0.

When the range of utility values attainable at $t = 1$ is not greater than that attainable at $t = 0$, this can be done. In a framework like the one presented here, however, in which consumption at $t = 1$ allows for utility levels higher than any attainable at $t = 0$, it is not generally possible. If $u(c_1, N_1) > N_0/(\eta - 1)$, any attempt to define consumption across the two periods in common units must absurdly conclude either that $c_0 = 0$ or that c_1 is (more than) infinite.

In other words, in this setting there is no welfare-relevant consumption aggregator. In an important sense, the concept of “consumption growth” must be left undefined.

Note that this problem arises only when we try to quantify consumption in a later period in the units of consumption in an earlier period, such as with a Laspeyres cost-of-living index or with any index that takes the Laspeyres as an input. That is, the problem only arises “looking forward”. If we restrict ourselves at t to “looking backward”, we can define π such that real consumption in a previous period $s < t$ is measured to be that which would enable utility level u_s given product range N_t . That is, we can, in period t , define π such that $u(c_s, N_s) = u(\frac{c_s}{\pi_s}, N_t) = u(\tilde{c}_s, N_t)$.

A fully Paasche-style, backward-looking index still faces two problems in this setting, however.

First, it cannot consistently be used to compare consumption levels both across individuals within a period and across periods. However we define the consumption level $\tilde{c}_s > 0$ of a typical consumption basket at time s , arbitrarily many consumption baskets at s will correspond to a lower utility level than the finite consumption level defined to be \tilde{c}_t , as long as $\tilde{c}_t > N_s/(\eta - 1)$. That is, if we say that a middle-class American enjoys consumption n times higher than a middle-class ancient Roman, for any n , we can imagine a Roman consuming $2n$ as much as her poorer compatriot in the consumption units of her day and yet on a lower indifference curve than the middle-class American.

Second, and more importantly for our purposes, this index will produce misleading conclusions about future welfare. Defining \tilde{c}_s such that $u(c_s, N_s) = u(\tilde{c}_s, N_t)$, we might construct a series of \tilde{c}_s across times $s \leq t$. But however we project this series into the future, because we must maintain predictions of $\tilde{c}_s < \infty$ across $s > t$, we will predict future utility levels to remain below $\lim_{\tilde{c}_s \rightarrow \infty} u(\tilde{c}_s, N_t) = N_t/(\eta - 1)$. In reality, of course, future utility may be arbitrarily large, if N_s can grow without bound.

Relatedly, note also that the undefinability of forward-looking welfare-relevant consumption measures necessarily arises only in the long run. In the short run, cost-of-living indices can generally allow for welfare-relevant consumption comparisons across pairs of adjacent periods. For illustration, suppose $u(c_2, N_2) < N_1/(\eta - 1)$, and $u(c_1, N_1) < N_0/(\eta - 1)$, but $u(c_2, N_2) \geq N_0/(\eta - 1)$. Consumption growth from $t - 1$ to t may be defined as that which would be necessary to reach utility level u_t given the period $t - 1$ product range—i.e. $\tilde{g}_t = \frac{\tilde{c}_t - \tilde{c}_{t-1}}{\tilde{c}_{t-1}}$ —for both $t = 1, 2$. Given continuous changes to either N_t or c_t , a sufficiently fine time grid will always allow for welfare-relevant consumption comparisons across adjacent periods in this way. As the above example illustrates, however, the resulting consumption growth rates cannot legitimately be chained across periods. Multiplying consumption at $t = 0$ by $(1 + \tilde{g}_1)(1 + \tilde{g}_2)$ does not yield a consumption level high enough to achieve u_2 , given N_0 ; as stipulated, no such consumption level is high enough.

4 Stylized facts and the equity premium

To satisfy a household's intertemporal Euler equation, we must have

$$r = \delta + \frac{d}{dt} \frac{\partial u(c_t, N_t)}{\partial c_t} = \delta + \eta g_{\bar{c}}, \quad (6)$$

where r is the interest rate, δ is the rate of time preference, and $g_{\bar{c}} = g_c - g_N$ is the rate of growth in consumption per product.

Two of Kaldor's (1957) stylized facts of growth are of course that (standardly measured) consumption growth g_c is roughly constant in the long run and that r is roughly constant as well. Given constant δ , these stylized facts imply that the preferences given by (4) require constant g_N . Note that any constant $g_N \leq g_c$ is in principle compatible with the framework, but $g_N > g_c$ would eventually produce a violation of (2), after which the introduction of new products would not increase the range of products actually consumed.

This framework also, however, produces a difference between the curvature of the utility function within a period, as measured by η , and the curvature that would be implied by fitting intertemporal consumption data to a standard model of isoelastic utility without new product introduction. Let us denote the latter quantity by $\tilde{\eta}$. It can be inferred from the standard Ramsey formula:

$$r = \delta + \tilde{\eta} g_c \quad \implies \quad \tilde{\eta} = \frac{r - \delta}{g_c}. \quad (7)$$

In the model of new product introduction presented here, we have

$$\tilde{\eta} < \eta = \frac{r - \delta}{g_c - g_N}. \quad (8)$$

The rates of return available from risky investments, such as equities, are far higher than the rates available from safer investments, such as bonds. Inferring the curvature of typical households' utility functions from the risk preferences implied here yields remarkably high values of η —sometimes even above 50. Estimating the curvature of household utility functions from their intertemporal consumption decisions, however (i.e., in the risk- and fluctuation-free framework here, from the Ramsey formula), yields much lower estimates—sometimes even below 1. The puzzling presence of such high equity returns, assuming that lower estimates of η are more accurate, has been called the equity premium puzzle (since, seminally, Mehra and Prescott (1985)).

By now it is clear how new product introduction can account for this discrepancy. With g_N arbitrarily close to g_c , η may indeed be arbitrarily high, while remaining compatible with reasonable values of r and δ . At the same time, we maintain $\tilde{\eta} = (r - \delta)/g_c$. Individuals may value marginal consumption in the future only slightly less than consumption in the present, because the future comes with the opportunity to spend on more products, and at the same time they may be highly risk-averse with respect to consumption in any period.

The observed discrepancy between η and $\tilde{\eta}$ may alternatively be reconciled with complex preference structures, such as Epstein-Zin preferences (Epstein and Zin, 1989), designed to separate coefficients of relative risk aversion (here, η) from inverse elasticities of intertemporal substitution (here, $\tilde{\eta}$). These preference structures can seem *ad hoc*, however, and they require us to abandon the simple and (at least to some) intuitively appealing assumption that individuals aim to maximize their expected discounted flow utility, somehow defined. By contrast, simply allowing for new product introduction, in a way that allows future consumption to achieve higher utility levels than past consumption, immediately produces a mechanism whereby the inverse EIS will tend to be higher than the (within-period) CRRA. Indeed, it would be surprising if we did *not* find this divergence reflected in investment behavior.

Finally, it may be worth noting that, given estimates of η inferred from individuals' within-period risk preferences, the rate of new product introduction can, in this context, be inferred to equal

$$g_N = g_c - \frac{r - \delta}{\eta}. \tag{9}$$

5 Endogenizing g_N

If we take this framework somewhat literally, we are left to explain why we should expect the rate of new product introduction to be a constant fraction of the rate of consumption growth. Indeed, assuming that new products are costly to develop, it is plainly not socially optimal to maintain a constant rate of new product introduction, at least while consumption is small. Given $g_c > g_N$, we must posit an early time t

at which $\bar{c}_t = \eta^{\frac{1}{\eta-1}}$ —the earliest time at which condition (2) is satisfied—and at t the marginal utility to introducing a new product is 0. Endogenizing a constant g_N on the assumption that new product introduction proceeds efficiently is possible for times $s > t$, but only given a complex and *ad hoc* assumption about the product development cost function.

If new products are introduced by profit-maximizing actors, however, there is a simple framework under which constant g_c gives rise to constant g_N as long as (2) is satisfied. Suppose that, once a new product i is introduced at some time t , its creator enjoys permanent monopoly rights over it. The overall consumption of i then equals \bar{c}_s at all periods $s \geq t$ (multiplied by the population size, which we will hold fixed). The present value of the profit earned by introducing i thus equals

$$\int_t^\infty e^{-r(s-t)} \bar{c}_s ds = \frac{\bar{c}_t}{r + g_N - g_c}, \quad (10)$$

which grows at rate $g_\pi = g_c - g_N$. (Note that, given $\delta > 0$ and $\eta > 1$, it follows from (6) that we will have $r + g_N - g_c > 0$.)

Without loss of generality, normalize N_0 to 1 and c_0 to $r + g_N - g_c$. Also, suppose that the cost of developing product i equals i^α for some $\alpha \geq 0$. We then want to find the product introduction rate g_N such that it profitable to develop i at time t such that $i = e^{g_N t}$: that is, at $t = \ln(i)/g_N$. This will hold precisely when

$$\frac{c_0}{r + g_N - g_c} e^{(g_c - g_N) \ln(i)/g_N} = i^\alpha \quad (11)$$

$$\implies i^{\frac{g_c - g_N}{g_N}} = i^\alpha \quad (12)$$

$$\implies g_N = \frac{g_c}{1 + \alpha}. \quad (13)$$

This will thus be the new product introduction rate observed in equilibrium. It follows from (9) that the interest rate, in turn, will equal

$$r = \delta + \eta(g_c - g_N) = \delta + \eta g_c \frac{\alpha}{1 + \alpha}. \quad (14)$$

The amount spent developing new products will grow at rate $g_c - g_N$ per product, so at rate g_c overall. The shares of income each period spent on new product development and on consumption will thus be constant.

To microfound g_c as simply as possible, we might simply wish to posit that production exhibits constant returns in a single factor, such as capital or effective labor, which exogenously grows at a constant rate. Because the elasticity of substitution between products at a given time is $1/\eta < 1$, however, monopolistic sellers will always be able to increase their profits by decreasing production, at least down to the threshold below which the product is no longer purchased. To describe this model more completely, therefore, we would have to introduce some further complication, such as a production subsidy. But we will not explore this further here.

6 Long-term welfare

This framework produces dramatically different implications from a standard growth framework both regarding the absolute welfare levels we should expect to attain in the long run and, more decision-relevantly, regarding the welfare implications of accelerating growth further.

The standard framework posits that the goods and services we enjoy can be aggregated into a unidimensional quantity, “consumption”, which can be compared across periods. It further supposes that preferences in consumption c are roughly described by isoelastic utility functions with upper bounds: i.e. that

$$u(c) = \frac{c^{1-\eta} - 1}{1-\eta}, \quad \eta > 1. \quad (15)$$

Under these assumptions, as noted above, utility per person can never exceed (or even equal) $\frac{1}{\eta-1}$, even as $c \rightarrow \infty$.

Furthermore, if we assume a constant baseline growth rate of $g > 0$, then even given an infinite time horizon and a zero rate of time preference, the cumulative welfare gains achievable by accelerating growth are finite. (Let us assume a fixed population size P , for simplicity.) To see this, consider the per-capita welfare gain to multiplying the growth rate by $k > 1$:

$$\int_0^\infty \left(\frac{e^{kgt(1-\eta)}}{1-\eta} - \frac{e^{gt(1-\eta)}}{1-\eta} \right) dt = \frac{k-1}{kg(\eta-1)^2} < \infty \quad (16)$$

As $k \rightarrow \infty$ —that is, as we approach the case in which we can give everyone infinite consumption forever—the total welfare gain approaches a mere

$$\frac{P}{g(\eta-1)^2}. \quad (17)$$

It may be worth taking a moment to put these magnitudes in perspective. In particular, let us estimate the possible welfare gains from accelerating growth, under these assumptions, accruing from increases to the future consumption of those already living comfortably (as distinct from the welfare gains that might result from accelerating the elimination of poverty).

Suppose η is as low as $5/4$. Also, let \$30,000 denote one unit of consumption, and suppose that the consumption level producing “zero welfare” is \$200/year, or $1/150$ of a unit. That is, assume that

$$u(c) = \frac{1}{4} \left(\left(\frac{1}{150} \right)^{-\frac{1}{4}} - c^{-\frac{1}{4}} \right). \quad (18)$$

Note that the added constant of

$$\frac{1}{4} \cdot \left(\frac{1}{150}\right)^{-\frac{1}{4}} \approx 0.87 \quad (19)$$

—which equals the welfare level $u(c)$ as $c \rightarrow \infty$ —does not affect the welfare difference calculations above, since it would be added and then subtracted.

Then the welfare level currently enjoyed at \$30,000/year, as a fraction of the welfare upper bound, is

$$\frac{\frac{1}{4} \left(\left(\frac{1}{150} \right)^{-\frac{1}{4}} - 1 \right)}{\frac{1}{4} \left(\frac{1}{150} \right)^{-\frac{1}{4}}} \approx 0.71. \quad (20)$$

That is, someone consuming \$30,000/year is already typically about 71% of the way from nonexistence to the welfare level she would approach if she had all conceivable wealth. On reflection, the conclusion that \$30,000/year typically offers welfare over 71% of the way to satiety, at least given the products currently available, strikes me as realistic. Is there any amount of wealth which it would currently be prudentially rational for a typical middle-class individual in the developed world to risk a 29% chance of death in order to attain?

Now suppose also that the baseline growth rate g is relatively low, at 0.016. Then the total welfare produced by bringing the entire population to satiety forever, as given by (17), is $1000P$. Since the “satiety level” here offers each individual approximately 0.87 units of welfare, by (19), the welfare benefit of infinite growth—of bringing everyone to satiety forever—equals approximately that of $(1000/0.87)P \approx 1143P$ years of life at satiety. That is, a classical utilitarian should be indifferent between (a) satiating the existing population forever and (b) leaving the existing population’s consumption growth path unchanged but creating an equal-sized population, a “parallel earth”, which enjoys full satiety and lasts slightly over a millennium.

Consider a (still highly optimistic, but at least somewhat more realistic) intervention: one that multiplies the growth rate by $k = 1.01$ forever. By (16), this produces approximately $9.9P$ units of welfare, which is in turn approximately as valuable as $(9.9/0.87)P \approx 11P$ years of life at satiety. That is, a classical utilitarian should be indifferent between (a) multiplying the growth rate by 1.01 forever and (b) leaving the existing population’s consumption growth path unchanged but creating an equal-sized population, a “parallel earth”, which enjoys full satiety and lasts slightly over a decade.

By contrast, as we have seen, a model in which new product introduction raises the utility upper bound allows utility per person to grow without bound. In fact, with g_N constant, the upper bound rises exponentially; and with $g_c \geq g_N$, utility stays near enough to its upper bound that it rises exponentially in the long run as well.

Clearly, such a model also allows growth accelerations to produce infinite increases to future welfare. In these respects, the model resembles a unidimensional model in which utility is isoelastic in consumption, but with $\eta < 1$.

Of course, even if we weaken the assumptions we have made to allow for infinite values (in particular an infinite horizon and no discounting), the point remains that there may be arbitrarily large differences between the welfare implications of a growth intervention under the standard framework and the welfare implications of a growth intervention under a framework in which new products allow for increases to the utility upper bound.

This observation is in some ways reminiscent of Lucas's (1987) comparison between the welfare implications of eliminating business cycles and the welfare implications of accelerating growth, within the standard framework. By his calculations, the latter swamps the former. My argument here is that he may not have gone nearly far enough. Even the latter—the welfare implications of accelerating growth, within the standard framework—may be swamped by the welfare implications of accelerating growth in reality.

7 Justifying the framework

So far we have considered only one simple approach to modeling the way in which new product introduction may allow for increases to the utility upper bound. Before closing, let us briefly consider how this approach might be generalized.

Consider an arbitrary utility function $u(c, N)$ that is (i) strictly increasing in the first argument (consumption) and (ii) weakly increasing in the second (product range). As long as

$$(iii) \quad u(N) \triangleq \lim_{c \rightarrow \infty} u(c, N) < \infty \text{ and}$$

$$(iv) \quad u(N) \text{ is strictly increasing in } N$$

for all N , u is a utility function for which new product introduction is ultimately necessary for welfare growth, and for which projecting growth in c alone is bound ultimately to produce arbitrarily large (in proportional terms) underestimates of future welfare growth.

Of course, many two-argument utility functions satisfy (i)–(iv). Suppose, therefore, that we also wish to impose the assumption (v) that $u(c, N)$ is isoelastic in c (holding N fixed), at least above some consumption threshold $c^*(N)$, with CRRA $\eta > 1$. Then, assuming this threshold is satisfied, we must have

$$u(c, N) = a(N) \frac{c^{1-\eta}}{1-\eta} + b(N) \tag{21}$$

for some functions $a > 0$ and b .

We would presumably like to capture the intuition that increases in N raise the marginal utility of consumption for any given consumption level $c > c^*(N)$. More product types, that is, allow for consumption in more enjoyable ways than fewer product types. Let us therefore impose condition (vi) that $a(N)$ is continuous and increasing in N , with $\lim_{N \rightarrow \infty} a(N) = \infty$ and $\lim_{N \rightarrow 0} a(N) = 0$. The limit conditions here come with no loss of generality, since we may impose upper or lower bounds on N . Requiring $a(N)$ to be continuous does constitute a loss of generality, but only in a way that is pessimistic about the value of new product introduction: it rules out the possibility that individual, measure-zero “products” may raise the marginal utility of consumption discontinuously.

Now, because $a(N)$ is an increasing and continuous function mapping $\mathbb{R}_{>0}$ to itself, we may, again without loss of generality, specify it. Until we have imposed further restrictions on u , this simply amounts to choosing a measure for the continuum of potential products. Let us choose $a(N) = N^\eta$, so that we have

$$u(c, N) = N \frac{\left(\frac{c}{N}\right)^{1-\eta}}{1-\eta} + b(N). \quad (22)$$

Next, suppose we wish to assume (vii) that the marginal utility of consumption falls at a constant rate over time, given constant g_c . (Recall that this is necessary to maintain a constant inverse EIS $\tilde{\eta}$, and that this in turn is necessary to fit the Kaldor facts of a constant interest rate and consumption growth rate alongside a constant rate of time preference.) Then, because $a(N)$ is specified to be a power function of N , we must have g_N constant as well, as shown in §4.

Finally, to maintain condition (ii) that $u(c, N)$ is weakly increasing in N , given that u takes the structure of (22) for $c > c^*(N)$, we must have

$$\frac{\partial u(c, N)}{\partial N} = \eta \frac{\left(\frac{c}{N}\right)^{1-\eta}}{1-\eta} + b'(N) \geq 0, \quad (23)$$

for all N at which b' is defined. (Strictly speaking, though b must be increasing in N , b may not be differentiable at a sparse set of values of N ; but at these values, $u(c, N)$ is necessarily increasing in N . For simplicity, let us proceed on the assumption that b is differentiable.) Taking

$$c^*(N) = N\eta^{\frac{1}{\eta-1}}, \quad (24)$$

as in (2), and observing that the derivative of $N\left(\frac{c}{N}\right)^{1-\eta}/(1-\eta)$ with respect to N is more strongly negative when c is lower, condition (23) on b reduces to

$$b'(N) \geq \frac{1}{\eta-1}. \quad (25)$$

Of course, this implies that among permissible forms for $b(N)$, given threshold level (24), the form *most pessimistic* about the relationship between new product introduction and welfare is $b(N) = N/(\eta - 1)$ (uniquely, up to an unimportant additive constant). As we can see, this is precisely the functional form presented in (4) and microfounded in (1).

Thus the only substantive decisions we made in the framework of this paper, beyond satisfying conditions (i)–(vii), are to set the minimum consumption threshold as in (24)—a weaker assumption would be to set it higher, or have it increase more quickly in N —and to have $b(N)$ increase in N as slowly as possible.

8 Conclusion

Many share the intuition that modern conveniences allow individuals to attain higher utility levels through their consumption than were available even to very wealthy individuals in the past. This intuition cannot be accommodated by a standard growth framework. It can, however, be accommodated by alternative frameworks. Here we have explored one such alternative, chosen for its simplicity and for its near-unique satisfaction of a list of relatively conventional desiderata.

As we have seen, it straightforwardly accounts for the intuition above. It also offers a natural resolution to a long-standing puzzle in financial economics, namely the fact that coefficients of relative risk aversion appear to be higher than inverse elasticities of intertemporal substitution. Finally, it suggests that long-term value of economic growth is conventionally understated, and perhaps severely so.

9 References

- Epstein, Larry G. and Stanley E. Zin**, “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, 1989, 57 (4), 937–969.
- Kaldor, Nicholas**, “A Model of Economic Growth,” *The Economic Journal*, 1957, 67 (268), 591–624.
- Lucas, Robert E., Jr.**, *Models of Business Cycles*, Oxford: Basil Blackwell, 1987.
- Mehra, Rajnish and Edward C. Prescott**, “The Equity Premium: A Puzzle,” *Journal of Monetary Economics*, 1985, 15 (2), 145–161.