

New Products and Long-term Welfare

Philip Trammell*

April 20, 2024

Abstract

Growth models typically assume an inaccurate equivalence between the consumption of greater quantities of existing products (as an individual achieves by growing richer, all else equal) and the consumption of new products. As a result, they typically arbitrarily understate the welfare benefits of growth. They also arbitrarily overstate the extent which future growth will motivate a substitution from consumption to other goods. Finally, a more realistic model of new product introduction can be shown to alleviate the equity premium puzzle: steeply diminishing marginal utility in within-period consumption is compatible with a high saving rate because the marginal utility of consumption will be higher when new products are available.

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.

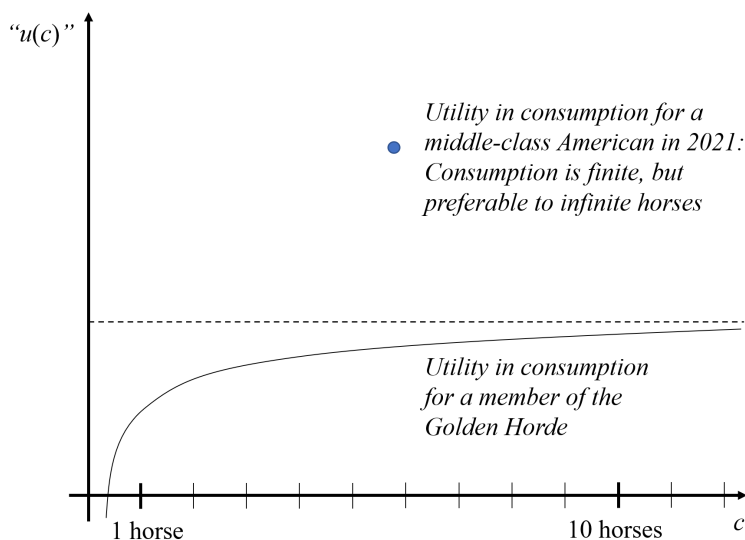
*Global Priorities Institute and Department of Economics, University of Oxford.
Contact: philip.trammell@economics.ox.ac.uk.

Failing to recognize this impossibility can carry two large costs. First, it can lead us to underestimate the extent to which people today are better off than people in the past—and, more decision-relevantly, the extent to which changes in the rate of economic growth will affect how well-off people are in the future. Second, it can lead us to overestimate the extent to which continued economic growth will decrease marginal utility in consumption and thereby motivate future individuals and societies to substitute from consumption to leisure, health (Jones (2016)), or existential security (Aschenbrenner (2020), Jones (2023)).

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

2 Framework

This paper is intended to explore a simple intuition, visually represented as follows. Consider an early society in which horse is the only consumption good:



If the economists of the Golden Horde had estimated the contemporaneous relationship between utility and consumption in their society, and used their estimates to

make claims about the welfare implications of consumption growth, they would have gone badly wrong.

In particular, a one-dimensional stylized utility function “ $u(c)$ ” forces one either to overstate the utility benefits of increasing consumption contemporaneously or to understate the utility benefits of increasing the consumption in the future, once new products are available. This issue is especially stark when the one-dimensional utility function in question exhibits constant relative risk aversion (CRRA), with coefficient of relative risk aversion (RRA) $\eta > 1$, as is commonly employed, in which case utility in consumption is assumed to have a hard upper bound (as in the figure above).

Nevertheless, policy recommendations predicated on analogous estimates of the welfare benefits of future consumption growth are pervasive today. [*Cite lots of examples.*]

2.1 Alternative stylized preferences

An alternative framework promises to model the long-term welfare implications of consumption growth only somewhat less simply and substantially more realistically.

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

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

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 marginal rate of transformation between any pair of goods equals 1, and thus that the prices of the goods are equal. 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.

Let us define $u(c, N)$ to be the maximum utility $U(\vec{c})$ achievable given total consumption c across the range of N products. That is,

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

Note that, given $\eta > 1$, v must incorporate the floor of 0 in order for us 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, in an interior solution, gross complements. The elasticity of substitution between them is $1/\eta < 1$.

More standard frameworks for new product introduction in growth models avoid this issue by assuming a range of goods with a (typically common and constant) elasticity of substitution greater than 1, as popularized by e.g. Romer (1990). Under this assumption of gross substitutability, however, for any given utility level achievable by the consumption of one 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.

2.2 Axiomatization of alternative stylized preferences

To motivate exploring the particular utility function given by (4), consider the following axiomatization.

Let $u(c, N)$ represent the utility level achievable by consumption level c , as conventionally measured, given a continuum of available products ranging from 0 to N . Also, let

$$\underline{c}_u(N) \triangleq \inf \{c : u^2(c, N) > 0\} \quad (5)$$

if this set is nonempty. (A superscript of $k > 0$ will be used to denote a function's partial derivative with respect to its k^{th} argument.)

Proposition 1. *Utility functions and new products*

Consider a continuously differentiable cardinal utility function $u(c, N)$ defined for $c > 0$ and $N > 0$. Suppose $u(\cdot)$ satisfies the following conditions on its domain:

- Unimportance of new products when impoverished: $\underline{c}_u(N)$ is strictly increasing in N from 0 to ∞ .
- Symmetry across products: $u^1(\underline{c}_u(N), N)$ is independent of N .

Given c, N such that $c > \underline{c}_u(N)$,

- Monotonicity: $u(\cdot)$ is strictly increasing in N .
- Isoelasticity: $u(\cdot)$ is CRRA in c , independent of N .

Then, over c, N such that $c > \underline{c}_u(N)$, $u(c, N)$ is equal (up to affine transformation) to $\tilde{u}(c, \underline{c}_u(Nz))$, where

$$\tilde{u}(c, M) \triangleq \begin{cases} M \frac{\left(\frac{c}{M}\right)^{1-\eta} - 1}{1-\eta}, & \eta \neq 1, \\ M \ln\left(\frac{c}{M}\right), & \eta = 1; \end{cases} \quad (6)$$

$$z \triangleq \begin{cases} \eta^{\frac{1}{\eta-1}}, & \eta \neq 1, \\ e, & \eta = 1; \end{cases} \quad (7)$$

and $\eta > 0$ is $u(\cdot)$'s RRA.

Proof: See Appendix A.1.

Two of these conditions may deserve explanation.

The “unimportance” condition states that, when even existing products are consumed in sufficiently small quantities, there is no benefit to introducing a new product. The aggregate consumption level required before existing products *are* consumed in large enough quantities that new product introduction would be beneficial is increasing in the number of products already being consumed. When there are almost no products available, new products are valuable even for people with very low consumption levels; and as the number of products already consumed rises without bound, the aggregate consumption of them required before adding a new product would increase utility also rises without bound.

The “symmetry” condition states that, for a consumer to be just indifferent to consuming a new product, the marginal utility of consuming existing products must

have been driven down to a constant level—i.e., a level that does not depend on the new product in question. This is arguably the least compelling of the conditions above, and relaxing it may produce fruitful generalizations of the framework presented here. Nevertheless, it formalizes a sense in which the utility function $\tilde{u}(\cdot)$ given by (6) is the simplest utility function defined over consumption and product range satisfying the other three conditions, and so perhaps the most natural place to begin an exploration in this direction.

Note that $\tilde{u}(\cdot)$ does indeed satisfy all four desiderata, with $\underline{c}_{\tilde{u}}(N) = Nz$. Since $\underline{c}_u(N)$ is continuous and strictly increasing in N from 0 to ∞ , so is $\underline{c}_{\tilde{u}}^{-1}(Nz)$. Thus, what Proposition 1 shows is that, for any cardinal preferences defined over consumption levels and product ranges that can be represented by a utility function satisfying the four conditions listed, the product continuum can be “re-measured” such that the preferences in question are representable by (6) (over c, N such that $c > Nz$).

Note also that the result of Proposition 1 does not require $\eta > 1$, though this is necessary for utility in consumption to be bounded above for any finite product range, as in the motivating case.

2.3 The irrelevance of inflation measurement

Assume utility function (4), and suppose N_t can grow across periods. As above, let us stipulate that the price level remains constant at 1, for each good type, at all times. 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

$$\tilde{c}_t \triangleq \frac{c_t}{\pi_t}. \tag{8}$$

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.

2.4 Endogenizing product development

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. [*Turn the next two paragraphs into a “Proposition”.*] 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

$$\int_t^\infty e^{-r(s-t)} \bar{c}_s ds = \frac{\bar{c}_t}{r + g_N - g_c}, \quad (9)$$

which grows at rate $g_{\bar{c}} = g_c - g_N$. (Note that, given $\delta > 0$ and $\eta > 1$, it follows from (20) 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

$$\frac{c_0}{r + g_N - g_c} e^{(g_c - g_N) \ln(i)/g_N} = i^\alpha \quad (10)$$

$$\implies i^{\frac{g_c - g_N}{g_N}} = i^\alpha \quad (11)$$

$$\implies g_N = \frac{g_c}{1 + \alpha}. \quad (12)$$

This will thus be the new product introduction rate observed in equilibrium. It follows from (23) that the interest rate, in turn, will equal

$$r = \delta + \eta(g_c - g_N) = \delta + \eta g_c \frac{\alpha}{1 + \alpha}. \quad (13)$$

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 oligopoly (rather than monopoly) or a production subsidy. But we will not explore this further here.

3 Long-term welfare

3.1 In total

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

$$u(c) = \frac{c^{1-\eta} - 1}{1 - \eta}, \quad \eta > 1. \quad (14)$$

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 $k > 1$:

$$\int_0^\infty \left(\frac{e^{kgt(1-\eta)}}{1-\eta} - \frac{e^{gt(1-\eta)}}{1-\eta} \right) dt = \frac{k-1}{kg(\eta-1)^2} < \infty. \quad (15)$$

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

$$\frac{P}{g(\eta-1)^2}. \quad (16)$$

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

$$u(c) = \frac{1}{4} \left(\left(\frac{1}{150} \right)^{-\frac{1}{4}} - c^{-\frac{1}{4}} \right). \quad (17)$$

Note that the added constant of

$$\frac{1}{4} \cdot \left(\frac{1}{150} \right)^{-\frac{1}{4}} \approx 0.87 \quad (18)$$

—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

$$\frac{\frac{1}{4} \left(\left(\frac{1}{150} \right)^{-\frac{1}{4}} - 1 \right)}{\frac{1}{4} \left(\frac{1}{150} \right)^{-\frac{1}{4}}} \approx 0.71. \quad (19)$$

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. [*This should all be replaced with something derived from the VSL, but still:*] 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 (16), is $1000P$. Since the “satiety level” here offers each individual approximately 0.87 units of welfare, by (18), 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 (15), 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.

3.2 On the margin

This would be where I explain the implications for Jones (2016), Aschenbrenner (2020), Jones (2023), etc.

4 Stylized facts and the equity premium

The following is essentially a simplification of a point made by Scanlon (2019) (as I learned since I first drafted this document).

To satisfy an individual's intertemporal Euler equation, in our framework, we must have

$$r = \delta + \left[\frac{d}{dt} \frac{\partial u(c_t, N_t)}{\partial c_t} \right] / \frac{\partial u(c_t, N_t)}{\partial c_t} = \delta + \eta g_{\bar{c}}, \quad (20)$$

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 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:

$$r = \delta + \tilde{\eta} g_c \quad \implies \quad \tilde{\eta} = \frac{r - \delta}{g_c}. \quad (21)$$

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

$$\tilde{\eta} < \eta = \frac{r - \delta}{g_c - g_N}. \quad (22)$$

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 individuals' utility functions from the risk preferences implied here yields very high values of η —sometimes even above 50. Estimating the curvature of individual 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.

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

$$g_N = g_c - \frac{r - \delta}{\eta}. \quad (23)$$

4.1 Relationship to existing resolutions to the equity premium puzzle

There are, broadly, two existing approaches to reconciling the observed discrepancy between η and $\tilde{\eta}$.

The first approach consists of positing that people have complex preference structures that intrinsically separate coefficients of relative risk aversion (here, η) from inverse elasticities of intertemporal substitution (here, $\tilde{\eta}$). (The first-introduced and most commonly used preference structure of this kind is that of Epstein and Zin (1989). A more general class of preference structures exhibiting the desired separation of η from $\tilde{\eta}$, to which Epstein-Zin preferences belong, are known as “recursive preferences”.) These preference structures fit the data at the cost of seeming somewhat *ad hoc*. In particular, they require us to abandon the simple, appealing assumption that individuals can be modeled as maximizing an expected sum of flow utility levels across periods.

The second approach does not abandon this assumption. Instead, it posits that, though people get richer over time (greatly driving down their marginal utility in consumption), some partially countervailing trend also unfolds over time whose effect is to raise people's marginal utility in consumption. One candidate for such an effect is “habit formation” (Fuhrer, 2000). As one's consumption rises, one's *own previous-period consumption* also rises. As a result of habituation to this previous-period consumption, one acts, in some sense, as if one's elevated consumption level is less than it is; one remains far from satiation. Another candidate for such an

effect is “keeping up with the Joneses” (Galí, 1994). The confounding trend in this case is that one’s own consumption rises over time alongside that of others. If one cares to some extent not only about one’s own consumption level but also about how it compares to others’, one’s desire for further consumption does not fall by as much following widespread economic growth as it does following a successful gamble offering an immediate spike in one’s own consumption.

The reconciliation introduced here belongs to the second approach. Consumption levels rise over time, but the introduction of new products over time increases marginal utility in consumption for any given (sufficiently high) level of consumption.

Note, however, that the habit formation and keeping up with the Joneses effects increase the marginal utility of future consumption by *depressing* the value of consumption growth. They leave individuals effectively poorer than they would be if their consumption grew in the absence of the effect in question. In a framework centered around the assumption of a fixed, concave utility function $u(c)$, this is inevitable: if $u'(\cdot)$ is observed to decline surprisingly slowly, it must be inferred that “ c ”, in some utility-relevant sense, has grown deceptively slowly. By contrast, new product introduction, as modeled here, increases the marginal utility of future consumption by making all consumption growth—including subsequent consumption growth—strictly more valuable.

Explanations for the equity premium puzzle (and related puzzles) are not exclusive, of course. Habits, preferences over relative consumption levels, and a desire to afford new products when they are introduced can all motivate people to save more than they would otherwise. Nevertheless, it may be encouraging that a model designed to capture the welfare implications of new product introduction more realistically turns out also to introduce a gap between the coefficient of relative risk aversion and the inverse elasticity of intertemporal substitution, in the observed direction.

5 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 then suggests that our forecasts of the long-term implications of economic growth are biased in at least two ways. First, our estimates of the long-term welfare benefits of economic growth are conventionally understated, and perhaps severely so. Second,

our estimates of the extent to which economic growth will motivate us to pursue goods like health and safety, rather than ever more consumption, are conventionally overstated. Finally, 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.

6 References

- Aschenbrenner, Leopold**, “Existential Risk and Growth,” September 2020. Global Priorities Institute (University of Oxford), Working Paper No. 6-2020.
- 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.
- Fuhrer, Jeffrey C.**, “Habit Formation in Consumption and Its Implications for Monetary-Policy Models,” *American Economic Review*, 2000, *90* (3), 367–390.
- Galí, Jordi**, “Keeping up with the Joneses: Consumption Externalities, Portfolio Choice, and Asset Prices,” *Journal of Money, Credit and Banking*, 1994, *26* (1), 1–8.
- Jones, Charles I.**, “Life and Growth,” *Journal of Political Economy*, 2016, *124* (2), 539–578.
- , “The A.I. Dilemma: Growth versus Existential Risk,” April 2023.
- 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.
- Romer, Paul**, “Endogenous Technological Change,” *Journal of Political Economy*, 1990, *98* (5), S71–S102.
- Scanlon, Paul**, “New Goods and Asset Prices,” *Journal of Financial Economics*, 2019, *132* (3), 140–157.

Appendices

A Proofs

A.1 Proof of Proposition 1

Consider a utility function $u(c, N)$ satisfying the conditions of Proposition 1. Since $u(\cdot)$ is continuously differentiable, $u^2(c, N)$ is continuous in c and N . $\underline{c}_u(N)$ is therefore continuous in N . It follows that $\underline{c}_u(\cdot)$ is a bijection on $\mathbb{R}_{>0}$.

Let

$$\begin{aligned} f_u(N) &\triangleq \underline{c}_u^{-1}(Nz), \\ \tilde{u}(c, N) &\triangleq u(c, f_u(N)), \end{aligned} \tag{24}$$

where

$$z \triangleq \begin{cases} \eta^{\frac{1}{\eta-1}}, & \eta \neq 1, \\ e, & \eta = 1 \end{cases} \tag{25}$$

and η is $u(\cdot)$'s RRA in c where $c > \underline{c}_u(N)$. Note that, since $\underline{c}_u(\cdot)$ is a bijection on $\mathbb{R}_{>0}$, $f_u(\cdot)$ is defined throughout $\mathbb{R}_{>0}$.

Also,

$$\begin{aligned} \underline{c}_{\tilde{u}}(N) &= \inf \{c : \tilde{u}^2(c, N) > 0\} \\ &= \inf \{c : u^2(c, f_u(N)) > 0\}, \end{aligned} \tag{26}$$

where the second equality holds because

$$\tilde{u}^2(c, N) = u^2(c, f_u(N))f'_u(N), \tag{27}$$

and

$$f'_u(N) = z\underline{c}_u^{-1'}(Nz) \tag{28}$$

is positive wherever it is defined, because $\underline{c}_u(\cdot)$ and therefore $\underline{c}_u^{-1}(\cdot)$ are strictly increasing. ($\underline{c}_u^{-1'}(\cdot)$ and therefore $\tilde{u}^2(\cdot)$ may sometimes be undefined, but only at a sparse set of points, by Lebesgue's theorem for the differentiability of monotone functions.) The infimum value of c at which (27) is positive is thus equal to (26). Finally, therefore,

$$\underline{c}_{\tilde{u}}(N) = \underline{c}_u(f_u(N)) = Nz. \tag{29}$$

Since $f_u(\cdot)$ is strictly increasing, and $u(\cdot)$ is strictly increasing in c and N where $c > \underline{c}_u(f_u(N))$, $\tilde{u}(\cdot)$ is strictly increasing in c and N where $c > \underline{c}_{\tilde{u}}(N)$. Likewise,

since $u(\cdot)$ exhibits the constant RRA η in c where $c > \underline{c}_u(f_u(N))$, so does $\tilde{u}(\cdot)$ where $c > \underline{c}_{\tilde{u}}(N)$. $\tilde{u}(\cdot)$ must therefore take the form

$$\tilde{u}(c, N) = \begin{cases} a(N) \frac{c^{1-\eta}-1}{1-\eta} + b(N), & \eta \neq 1, \\ a(N) \ln(c) + b(N), & \eta = 1 \end{cases} \quad (30)$$

for some functions $a(N) > 0$ and $b(N)$ and some $\eta > 0$, where $c > Nz$.

Observe that

$$\tilde{u}^1(\underline{c}_{\tilde{u}}(N), N) = u^1(Nz, f_u(N)). \quad (31)$$

Since $\underline{c}_u(f_u(N)) = Nz$, we know that $u^1(Nz, f_u(N))$ is independent of $f_u(N)$, because $u(\cdot)$ satisfies the ‘‘symmetry’’ condition. $\tilde{u}^1(\underline{c}_{\tilde{u}}(N), N)$ is thus independent of N . That is, $\tilde{u}(\cdot)$ also satisfies symmetry.

To find $a(N)$, therefore, we can differentiate (30) with respect to c and evaluate the result where $c = Nz$. By symmetry, we must have

$$a(N) (Nz)^{-\eta} \quad (32)$$

independent of N . This implies

$$a(N) = AN^\eta \quad (33)$$

for some constant $A > 0$, and thus

$$\tilde{u}(c, N) = \begin{cases} AN^\eta \frac{c^{1-\eta}-1}{1-\eta} + b(N), & \eta \neq 1, \\ AN \ln(c) + b(N), & \eta = 1. \end{cases} \quad (34)$$

Finally, by definition of $\underline{c}_{\tilde{u}}(N)$, we must have

$$\lim_{c \downarrow N\eta^{\frac{1}{\eta-1}}} \tilde{u}^2(c, N) = 0. \quad (35)$$

This implies

$$\begin{aligned} \lim_{c \downarrow N\eta^{\frac{1}{\eta-1}}} A\eta N^{\eta-1} \frac{c^{1-\eta} - 1}{1-\eta} + b'(N) &= 0 \\ \implies b'(N) &= -A \frac{1}{1-\eta} + A \frac{\eta}{1-\eta} N^{\eta-1} \\ \implies b(N) &= -AN \frac{1}{1-\eta} + AN^\eta \frac{1}{1-\eta} + B \end{aligned} \quad (36)$$

if $\eta \neq 1$, and

$$\begin{aligned} \lim_{c \downarrow Ne} A \ln(c) + b'(N) &= 0 \\ \implies b'(N) &= -A \ln(N) - A \\ \implies b(N) &= -AN \ln(N) + B \end{aligned} \quad (37)$$

if $\eta = 1$, for some constant B .

Substituting (36) and (37) into (34) gives

$$\tilde{u}(c, N) = \begin{cases} AN \frac{\left(\frac{c}{N}\right)^{1-\eta} - 1}{1-\eta} + B, & \eta \neq 1, \\ AN \ln \left(\frac{c}{N}\right) + B, & \eta = 1, \end{cases} \quad (38)$$

which is course an affine transformation of

$$\begin{aligned} N \frac{\left(\frac{c}{N}\right)^{1-\eta} - 1}{1-\eta}, & \quad \eta \neq 1, \\ N \ln \left(\frac{c}{N}\right), & \quad \eta = 1. \end{aligned} \quad (39)$$