

The Costs of Economic Growth

Charles I. Jones*

Stanford GSB and NBER

October 9, 2009 – Version 1.0

Abstract

The benefits of economic growth are widely touted in the literature. But what about the costs? Pollution, nuclear accidents, global warming, the rapid global transmission of disease, and bioengineered viruses are just some of the dangers created by technological change. How should these be weighed against the benefits, and in particular, how does the recognition of these costs affect the theory of economic growth? This paper shows that taking these costs into account has first-order consequences. Under standard preferences, the value of life may rise faster than consumption, leading society to value safety over economic growth. As a result, the optimal rate of growth may be substantially lower than what is feasible, in some cases falling all the way to zero.

*I am grateful to Raj Chetty, Pete Klenow, Omar Licandro, David Romer, Michele Tertilt, Martin Weitzman and participants in seminars at Berkeley, Stanford, and UCSD for comments, and to the National Science Foundation for financial support. Arthur Chiang, William Vijverberg and Mu-Jeung Yang provided excellent research assistance; Arthur deserves special thanks for his assistance with the proofs in the paper.

Certain events quite within the realm of possibility, such as a major asteroid collision, global bioterrorism, abrupt global warming — even certain lab accidents— could have unimaginably terrible consequences up to and including the extinction of the human race... I am not a Green, an alarmist, an apocalyptic visionary, a catastrophist, a Chicken Little, a Luddite, an anticapitalist, or even a pessimist. But... I have come to believe that what I shall be calling the “catastrophic risks” are real and growing...

— Richard A. Posner (2004, p. v)

1. Introduction

In October 1962, the Cuban missile crisis brought the world to the brink of a nuclear holocaust. President John F. Kennedy put the chance of nuclear war at “somewhere between one out of three and even.” The historian Arthur Schlesinger, Jr., at the time an adviser of the President, later called this “the most dangerous moment in human history.”¹ What if a substantial fraction of the world’s population had been killed in a nuclear holocaust in the 1960s? In some sense, the overall cost of the technological innovations of the preceding 30 years would then seem to have outweighed the benefits.

While nuclear devastation represents a vivid example of the potential costs of technological change, it is by no means unique. The benefits from the internal combustion engine must be weighed against the costs associated with pollution and global warming. Biomedical advances have improved health substantially but made possible weaponized anthrax and lab-enhanced viruses. The potential benefits of nanotechnology stand beside the threat that a self-replicating machine could someday spin out of control. Experimental physics has brought us x-ray lithography techniques and superconductor technologies but also the remote possibility of devastating accidents as we smash particles together at ever higher energies. These and other technological dangers are detailed in a small but growing literature on so-called “existential risks”; Posner (2004) is likely the most familiar of these references, but see also Bostrom (2002), Joy (2000), Overbye (2008), and Rees (2003).

¹For these quotations, see (Rees, 2003, p. 26).

Technologies need not pose risks to the existence of humanity in order to have costs worth considering. New technologies come with risks as well as benefits. A new pesticide may turn out to be harmful to children. New drugs may have unforeseen side effects. Marie Curie's discovery of the new element radium led to many uses of the glow-in-the-dark material, including a medicinal additive to drinks and baths for supposed health benefits, wristwatches with luminous dials, and as makeup — at least until the dire health consequences of radioactivity were better understood. Other examples of new products that were initially thought to be safe or even healthy include thalidomide, lead paint, asbestos, and cigarettes.

The benefits of economic growth are truly amazing and have made enormous contributions to welfare. However, this does not mean there are not also costs. How does this recognition affect the theory of economic growth?

This paper explores what might be called a “Russian roulette” theory of economic growth. Suppose the overwhelming majority of new ideas are beneficial and lead to growth in consumption. However, there is a tiny chance that a new idea will be particularly dangerous and cause massive loss of life. Do discovery and economic growth continue forever in such a framework, or should society eventually decide that consumption is high enough and stop playing the game of Russian roulette?

While this existential risk interpretation is one possible view of the model developed here, there is another more mundane interpretation. The mortality risk can apply to each person in the economy independently: the risk may be idiosyncratic rather than aggregate. In this case, an idea with a small probability of death kills a small but deterministic fraction of the population each period — reminiscent of the lead paint or asbestos examples.

What happens to growth when such costs are taken into account? The answer turns out to depend on preferences. For a large class of conventional specifications, including log utility, safety eventually trumps economic growth. The optimal rate of growth may be substantially lower than what is feasible, in some cases falling all the way to zero.

This paper is most closely related to the literature on sustainable growth and the environment; for example, see Gradus and Smulders (1993), Stokey (1998), and Brock and Taylor (2005). Those papers show that when pollution and the environment di-

rectly enter utility or the production function, a “growth drag” may result. Here, the key concern — the loss of life associated with potentially dangerous technologies — is quite different. Nevertheless, there are interesting links with this literature that will be discussed later.

Section 2 presents a simple model that illustrates the main results of the paper. The advantage of this initial framework is its simplicity, which allows the basic intuition of the results to shine through. The disadvantage is that the tradeoff between growth and safety is a black box. Sections 3 and 4 then develop a rich idea-based endogenous growth model that permits a careful study of the mechanisms highlighted by the simple model. Section 5 provides some numerical examples that help quantify the slowdown in growth that results from safety. Finally, Section 6 discusses a range of empirical evidence that is helpful in judging the relevance of these results.

2. A Simple Model

At some level, this paper is about speed limits. You can drive your car slowly and safely, or fast and recklessly. Similarly, a key decision the economy must make is to set a safety threshold: researchers can introduce many new ideas without regard to safety, or they can select a very tight safety threshold and introduce fewer ideas each year, potentially slowing growth.

To develop this basic tradeoff, we begin with a simple two period OLG model. Suppose an individual’s expected lifetime utility is

$$U = u(c_0) + e^{-\delta(g)}u(c), \quad c = c_0(1 + g) \tag{1}$$

where c denotes consumption, g is the rate of consumption growth, and $\delta(g)$ is the mortality rate so $e^{-\delta(g)}$ is the probability an individual is alive in the second period. A new cohort of young people is born each period, and everyone alive at a point in time has the same consumption — this generation’s $c_0(1 + g)$ is the next generation’s c_0 .

To capture the “slow and safely or fast and recklessly” insight, assume $\delta(g)$ is an increasing function of the underlying rate of economic growth. Faster growth raises the mortality rate. In the richer model in the next section, this “black box” linking growth and mortality will be developed with much more care. For the moment, let’s just ex-

plore its implications. (This is not to deny that other aspects of growth — for example improvements in health technology — could reduce mortality. We are focusing on the cost side of things for now, however.)

Each generation when young chooses the growth rate for the economy to maximize their expected utility in equation (1). The privately optimal growth rate balances the concerns for safety with the gains from higher consumption. The first order condition for this maximization problem can be expressed as

$$u'(c)c_0 = \delta'(g)u(c). \quad (2)$$

At the optimum, the marginal benefit from higher consumption growth, the left hand side, equals the marginal cost associated with a shorter life, the right hand side. This condition can be usefully rewritten as

$$1 + g = \frac{\eta_{u,c}}{\delta'(g)} \quad (3)$$

where $\eta_{u,c}$ is the elasticity of $u(c)$ with respect to c .

To make more progress, assume the following functional forms:

$$\delta(g) = \delta g \quad (4)$$

$$u(c) = \bar{u} + \frac{c^{1-\gamma}}{1-\gamma}. \quad (5)$$

where all parameter values are positive. Utility takes the familiar form that features a constant elasticity of marginal utility; the important role of the constant \bar{u} will be discussed momentarily.

2.1. Exponential Growth: $0 < \gamma < 1$

To begin, let's assume $\gamma < 1$ and set $\bar{u} = 0$. In this case, the elasticity of utility with respect to consumption is $\eta_{u,c} = 1 - \gamma$, so the solution for growth in (3) is

$$g^* = \frac{1 - \gamma}{\delta} - 1. \quad (6)$$

As long as δ is not too large, the model yields sustained positive growth over time. For example, if $\gamma = 1/2$ and $\delta = 1/10$, then $g^* = 4$ and $1 + g^* = 5$: consumption increases by a factor of 5 across each generation. This comes at the cost of a life expectancy that is less than the maximum, but such is the tradeoff inherent in this model.

One can check that this conclusion is robust to letting $\bar{u} \neq 0$. In general, that will simply introduce transition dynamics into the model with $\gamma < 1$. The key elasticity $\eta_{u,c}$ then converges to $1 - \gamma$ as consumption gets large, leading to balanced growth as an asymptotic result.

2.2. The End of Growth: $\gamma > 1$

What comes next may seem a bit surprising. We've already seen that this simple model can generate sustained rapid growth for a conventional form of preferences. What we show now is that in the case where γ is larger than one, the model does not lead to sustained growth. Instead, concerns about safety lead growth to slow all the way to zero, at least eventually.

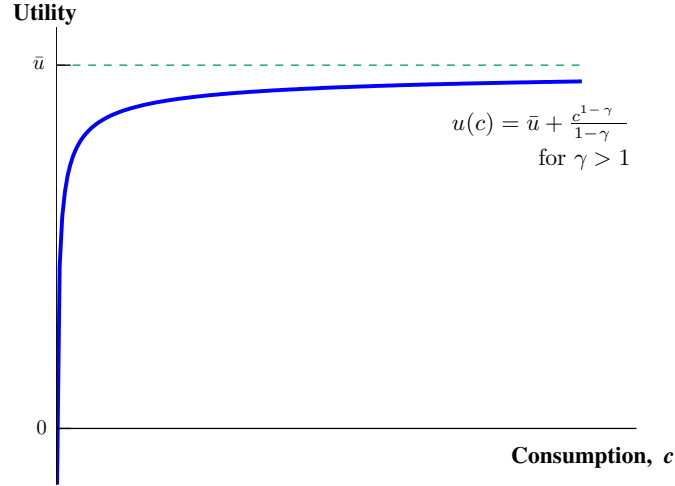
In this case, the constant \bar{u} plays an essential role. In particular, notice that we've implicitly normalized the utility associated with "death" to be zero. For example, in (1), the individual gets $u(c)$ if she lives and gets zero if she dies. But this means that $u(c)$ must be greater than zero for life to be worth living. Otherwise, death is the optimal choice at each point in time. With $\gamma > 1$, however, $\frac{c^{1-\gamma}}{1-\gamma}$ is less than zero. For example, this flow is $-1/c$ for $\gamma = 2$. Therefore for our problem to be interesting, we must add a positive constant to flow utility. In this case, the flow utility function is shown in Figure 1. Notice that flow utility is bounded, and the value of \bar{u} provides the upper bound.²

Assuming $\gamma > 1$ and $\bar{u} > 0$, the first order condition in (3) can be written as

$$(1 + g) \left(\bar{u} c_0^{\gamma-1} (1 + g)^{\gamma-1} + \frac{1}{1-\gamma} \right) = \frac{1}{\delta}. \quad (7)$$

Notice that the left-hand side of this expression is increasing in both c_0 and in g . As the

²As the figure illustrates, there exists a value of consumption below which flow utility is still negative. Below this level, individuals would prefer death to life, so they would randomize between zero consumption and some higher value; see Rosen (1988). This level is very low for plausible parameter values and can be ignored here.

Figure 1: Flow Utility $u(c)$ for $\gamma > 1$ 

Note: Flow utility is bounded for $\gamma > 1$. If $\bar{u} = 0$, then flow utility is negative and dying is preferred to living.

economy gets richer over time and c_0 rises, then, it must be the case that g declines in order to satisfy this first order condition. The optimal rate of economic growth slows along the transition path.

In fact, one can see from this equation that consumption converges to a steady state with zero growth. According to the original first order condition in (3), the steady state must be characterized by $\eta_{u,c}^* = \delta$ — that is, the point where the elasticity of the utility function with respect to consumption equals the mortality parameter. More explicitly, setting $g = 0$ in (7) reveals that the steady state value of consumption is given by

$$c^* = \left(\frac{1}{\bar{u}} \left(\frac{1}{\delta} + \frac{1}{\gamma - 1} \right) \right)^{\frac{1}{\gamma-1}}. \quad (8)$$

Because growth falls all the way to zero, mortality declines to zero as well and life expectancy is maximized.

To see the intuition for this result, recall the first order condition for growth: $1 + g = \eta_{u,c}/\delta$. When $\gamma > 1$ (or when flow utility is any bounded function), the marginal utility of consumption declines rapidly as the economy gets richer — that is, $\eta_{u,c}$ declines. This leads the optimal rate of growth to decline and the economy to converge to a steady

state level of consumption.

A crucial implication of the bound on utility is that the marginal utility of consumption declines to zero very rapidly. Consumption on any given day runs into sharp diminishing returns: think about the benefit of eating sushi for breakfast when you are already having it for lunch, dinner, and your midnight snack. Instead, obtaining extra days of life on which to enjoy your high consumption is a better way to increase utility.

This point can also be made with algebra. Consider the following expression:

$$\frac{u(c_t)}{u'(c_t)c_t} = \frac{1}{\eta_{u,c}} = \bar{u}c_t^{\gamma-1} + \frac{1}{1-\gamma}. \quad (9)$$

The left side of this equation is based on the flow value of an additional period of life, $u(c)$. We divide by the marginal utility of consumption to value this flow in units of consumption rather than in utils, so $u(c)/u'(c)$ is something like the value of a period of life in dollars. Then, we consider this value of life as a ratio to actual consumption.

The right side of this equation shows the value of life as a ratio to consumption under the assumed functional form for utility. Crucially, for $\gamma > 1$, the value of life rises faster than consumption. As the economy gets richer, concerns about safety become more important than consumption itself. This is the essential mechanism that leads the economy to tilt its allocation away from consumption growth and toward preserving life in the model.

Interestingly, this same result obtains with log utility. That is, the $\gamma = 1$ case of $u(c) = \bar{u} + \log c$ also leads growth to slow to zero even though utility is unbounded. In this case, $\eta_{u,c} = 1/u(c)$, so the elasticity of the utility function declines as consumption rises. Or, in terms of equation (9), the value of life as a ratio to consumption is just $u(c)$ itself, which grows with consumption.

2.3. Summary of the Simple Model

This simple model is slightly more flexible than the “Russian roulette” example given in the introduction. Rather than choosing between stagnation and a fixed rate of growth with a small probability of death, the economy can vary the growth rate and the associated death rate smoothly. This death rate can be given two different interpretations. It may apply independently to each person in the population, so that $e^{-\delta(g)}$ is the fraction

of the population that survives to old age in each cohort. Alternatively, it may represent an existential risk that applies to the entire economy (think of a nuclear holocaust or some unimagined disaster associated with the Large Hadron Collider).

With $\gamma < 1$, the optimal tradeoff between growth and mortality leads to sustained exponential growth, albeit with some positive death rate. In the idiosyncratic interpretation of the death rate, life expectancy is simply less than its maximum but the economy continues forever. In the existential risk interpretation, the economy grows exponentially until, with probability one, the existential risk is realized and the economy comes to an end.

A very different result occurs when $\gamma \geq 1$. In this case, the marginal utility of consumption in any period falls rapidly as individuals get richer. In contrast, each additional year of life delivers a positive and growing amount of utility. The result is an income effect that favors safety over growth. Growth eventually ceases, consumption settles to a constant, and life expectancy rises to its maximum. In the existential interpretation, the economy stops playing Russian roulette and, assuming it did not get unlucky before reaching the steady state, goes on forever.

3. Safety and Ideas in an Endogenous Growth Model

The simple model in the previous section is elegant and delivers clean results for the interaction between safety and growth. However, the way in which faster growth raises mortality is mechanical, and it is simply assumed that the economy can pick whatever growth rate it desires.

In this section, we address these concerns by adding safety considerations to a standard growth model based on the discovery of new ideas. The result deepens our understanding of the interactions between safety and growth. For example, in this richer model, concerns for safety can slow the rate of exponential growth from 4% to 1%, for example, but will never lead to a steady-state level of consumption. Alternatively, depending on the nature of technological danger, it could be optimal to slow the growth rate all the way to zero, but only as consumption rises to infinity. While supporting the basic spirit of the simple model, then, the richer model illustrates some important ways in which the simple model may be misleading.

The model below is a standard idea-based growth model, along the lines of Romer (1990) and Jones (1995). Researchers introduce new varieties of intermediate goods, and the economy's productivity is increasing in the number of varieties. The key change relative to standard models is that each variety i also comes with a danger level, z^i . Some ideas are especially dangerous (nuclear weapons or lead paint) and have a high value of z^i , while other ideas are relatively harmless and have a low z^i . A consumer's mortality rate then depends on the values of the z^i that are consumed as well as on the amount consumed.

3.1. The Economic Environment

The economy features three types of goods: consumption goods (which come in a range of varieties), ideas, and people. People and ideas are the two key factors of production, combining to produce the consumption goods and new ideas.

At any point in time, a variety of consumption goods indexed by i on the interval $[0, A_t]$ are available for purchase. We could define utility directly over this variety of goods, but for the usual reasons, it is easier to handle the aggregation on the production side. Hence, we assume these varieties combine in a CES fashion to produce a single aggregate consumption good:

$$C_t = \left(\int_0^{A_t} X_{it}^\theta di \right)^{1/\theta}, \quad 0 < \theta < 1. \quad (10)$$

New varieties (ideas) are produced by researchers. If L_{at} units of labor are used in research with a current stock of knowledge A_t , then research leads to the discovery of $\alpha L_{at}^\lambda A_t^\phi$ new varieties. This technology for producing new ideas is similar to Jones (1995).

What's novel here is that each new variety i is also associated with a danger level, $z^i > 0$. This danger level is drawn from a distribution with cdf $F(z)$ and is observed as soon as the variety is discovered.³ Researchers decide whether or not to complete the development of a new variety after observing its danger level. Given that varieties are otherwise symmetric, this leads to a cutoff level z_t : varieties with a danger level below

³We assume $F(0) = 0$, $F(z) > 0$ for $z > 0$, and that the first moment of $F(z)$ exists. These assumptions essentially mean that all ideas entail some danger, and the only way to have no technological danger is to set the threshold so that no new ideas are used.

z_t get implemented, whereas more dangerous varieties do not. z_t is a key endogenous variable determined within the model. The fraction $F(z_t) = \Pr[z^i \leq z_t]$ of candidate varieties get implemented, so the additional number of new varieties introduced at any point in time is

$$\dot{A}_t = \alpha F(z_t) L_{at}^\lambda A_t^\phi, \quad A_0 \text{ given.} \quad (11)$$

One unit of labor can produce one unit of any existing variety, and labor used for different purposes cannot exceed the total amount available in the economy, N_t :

$$\int_0^{A_t} X_{it} di + L_{at} \leq N_t. \quad (12)$$

This total population is assumed to grow over time according to

$$\dot{N}_t = (\bar{n} - \delta_t) N_t, \quad N_0 \text{ given.} \quad (13)$$

The parameter \bar{n} captures exogenous fertility net of any other sources of mortality unrelated to technological change.

Mortality from technological danger is denoted δ_t . In principle, it should depend on the amount of each variety consumed and the danger associated with each variety, and it could even be stochastic (nuclear weapons are a problem only if they are used). There could also be timing issues: the use of fossil fuels today creates global warming that may be a problem in the future.

These issues are interesting and could be considered in future work. To keep the present model tractable, however, we make some simplifying assumptions in determining mortality. In particular, all of the deaths associated with any new technology occur immediately when that technology is implemented, and the death rate depends on average consumption across all varieties. That is,

$$\delta_t = \bar{\delta} \dot{A}_t x_t \Gamma(z_t), \quad (14)$$

where $x_t \equiv \int_0^{A_t} X_{it} / N_t di / A_t$ is the average amount consumed of each variety and $\Gamma(z_t)$ is the average mortality rate of the varieties below the cutoff level z_t . That is, $\Gamma(z_t) \equiv E[z^i | z^i \leq z_t] = \int_0^{z_t} z dF(z) / F(z_t)$ is the average mortality rate associated with those new varieties that are actually consumed. The mortality rate δ_t , then, is the product of the

per capita quantity of new varieties consumed, $\dot{A}_t x_t$, and their average mortality rate.

Individuals care about expected utility, where the expectation is taken with respect to mortality. Let S_t denote the probability a person survives until date t conditional on being alive at date 0. Expected utility is given by

$$U = \int_0^{\infty} e^{-\rho t} u(c_t) S_t dt, \quad (15)$$

where

$$\dot{S}_t = -\delta_t S_t, \quad S_0 = 1. \quad (16)$$

Finally, we assume that flow utility $u(c)$ is

$$u(c_t) = \bar{u} + \frac{c_t^{1-\gamma}}{1-\gamma}, \quad c_t \equiv C_t/N_t, \quad \bar{u} > 0.^4 \quad (17)$$

3.2. A Rule of Thumb Allocation

Given the symmetry of X_{it} , there are two nontrivial allocative decisions that have to be made in this economy at each date. First is the allocation of labor between L_{at} and X_{it} . Second is the key tradeoff underlying this paper, the choice of the safety threshold z_t . A high cutoff for z_t implies that more new ideas are introduced in each period but it also means a higher mortality rate. This is the model's analog to driving fast and recklessly instead of slowly and safely.

In the next main section, we will let markets allocate resources and study an equilibrium allocation. To get a sense for how the model works, however, it is convenient to begin first with a simple rule of thumb allocation. For this example, we assume the economy puts a constant fraction \bar{s} of its labor in research and allocates the remainder symmetrically to the production of the consumption goods. In addition, we assume the safety cutoff is constant over time at \bar{z} .

Let g_x denote the growth rate of some variable x along a balanced growth path. Then, we have the following result (proofs for this and other propositions are given in Appendix A):

⁴In the case of $\gamma \leq 1$, $\bar{u} \leq 0$ is permissible.

Proposition 1 (BGP under the Rule of Thumb Allocation): *Under the rule of thumb allocation with $0 < \bar{s} < 1$ and $\bar{z} > 0$, there exists a balanced growth path such that $g_c = \sigma g_A$ where $\sigma \equiv (1 - \theta)/\theta$ and*

$$\delta^* = \bar{\delta} g_A (1 - \bar{s}) \Gamma(\bar{z}) \quad (18)$$

$$g_N = \bar{n} - \delta^* \quad (19)$$

$$g_A = \frac{\lambda(\bar{n} - \delta^*)}{1 - \phi} = \frac{\lambda \bar{n}}{1 - \phi + \lambda \bar{\delta} (1 - \bar{s}) \Gamma(\bar{z})}. \quad (20)$$


Along the balanced growth path, the mortality rate is constant and depends on (a) how fast the economy grows, (b) the intensity of consumption, and (c) the danger threshold. As in Jones (1995), the steady-state growth rate is proportional to the rate of population growth. However, the population growth rate is now an endogenous variable because of endogenous mortality. For example, an increase in research intensity \bar{s} will reduce the steady-state mortality rate (a lower consumption intensity) and therefore increase the long-run growth rate.

The effect of changing the danger threshold \bar{z} is more subtle and is shown graphically in Figure 2. As emphasized earlier, there is indeed a basic tradeoff in this model between growth and safety. Over the first 300+ years in the example, the safer choice of \bar{z} leads to slower growth as researchers introduce fewer new varieties. However, this tradeoff disappears in the long run because the growth rate itself depends on population growth. A safer technology choice reduces the mortality rate, raises the population growth rate, and therefore raises consumption growth in the long run.

4. A Competitive Equilibrium with Patent Buyouts

The rule of thumb allocation suggests that this model will deliver a balanced growth path with an interesting distinction between the medium-run and long-run tradeoffs between growth and safety. Moreover, the model features endogenous growth in the strong sense that changes in policy can affect the long-run growth rate. Somewhat surprisingly, neither of these results will hold in the competitive equilibrium, and our rule of thumb allocation turns out not to be a particularly good guide to the dynamics of the competitive equilibrium. Moreover, the features of the equilibrium allocation turn out

Figure 2: Growth under the Rule of Thumb Allocation



Note: There is a medium-run tradeoff between growth and technological danger, but no long-run tradeoff. In the long run, safer choices lead to faster net population growth and therefore faster consumption growth.

to hold in the optimal allocation as well, so the rule of thumb allocation is the one that proves to be misleading.

4.1. An Overview of the Equilibrium

A perfectly competitive equilibrium will not exist in this model because of the nonrivalry of ideas (Romer, 1990). Instead of following Romer and introducing imperfect competition, we use a mechanism advocated by Kremer (1998). That is, we consider an equilibrium in which research is funded entirely by “patent buyouts”: the government in our model purchases new ideas at a price P_{at} and makes the designs publicly available at no charge. The motivation for this approach is largely technical: it simplifies the model so it is easier to understand. However, there is probably some interest in studying this institution in its own right.

The other novel feature of this equilibrium is that we introduce a competitive market for mortality: idea producers pay a price v_t for every person they kill, and households “sell” their mortality as if survival were a consumer durable. This market bears some resemblance to one that emerges in practice through the legal system of torts and

liabilities.

In equilibrium, these two institutions determine the key allocations. Patent buyouts pin down the equilibrium amount of research and the mortality market pins down the danger cutoff.

4.2. Optimization Problems

The equilibrium introduces three prices: a wage w_t , the price of mortality v_t , and the price of new ideas P_{at} . The equilibrium then depends on three optimization problems.

First, a representative household supplies a unit of labor, chooses how much of her life to sell in the mortality market, pays a lump-sum tax τ_t , and eats the proceeds. Our timing assumption is that mortality is realized at the end of the period, after consumption occurs.

HH Problem: Given $\{w_t, v_t, \tau_t\}$, the representative household solves

$$\max_{\{\delta_t^h\}} \int_0^\infty e^{-\rho t} u(c_t) S_t dt$$

s.t. $c_t = w_t + v_t \delta_t^h - \tau_t$ and $\dot{S}_t = -\delta_t^h S_t$.

Next, a representative firm in the perfectly competitive market for the final good (FG) solves the following profit maximization problem:

FG Problem: At each date t , given w_t and A_t ,

$$\max_{\{X_{it}\}} \left(\int_0^{A_t} X_{it}^\theta di \right)^{1/\theta} - w_t \int_0^{A_t} X_{it} di.$$

Finally, a representative research firm produces ideas in the perfectly competitive idea sector and chooses a cutoff danger level z_t based on the price of mortality. The research firm sees constant returns to idea production at productivity α_t , so any effects associated with $\lambda < 1$ and $\phi \neq 0$ are external:

R&D Problem: At each date t , given $P_{at}, w_t, v_t, x_t, \alpha_t$,

$$\max_{L_{at}, z_t} P_{at} \alpha_t F(z_t) L_{at} - w_t L_{at} - v_t \delta_t N_t \quad \text{s.t.} \quad \delta_t = \bar{\delta} x_t \alpha_t F(z_t) L_{at} \Gamma(z_t).$$

4.3. Defining the Competitive Equilibrium

The competitive equilibrium in this economy solves the optimization problems given in the previous section and the relevant markets clear. The only remaining issue to discuss is the government purchase of ideas. We've already assumed the government pays a price P_{at} for any idea and releases the design into the public domain. We assume this is the only option for researchers — there is no way to keep ideas secret and earn a temporary monopoly profit. As discussed above, the reason for this is to keep the model simple; nothing would change qualitatively if we introduced monopolistic competition, either through secrecy or patents.

In addition, we assume the idea purchases are financed with lump-sum taxes on households and that the government's budget balances in each period. Moreover, we assume the government sets the price at which ideas are purchased so that total purchases of ideas are a constant proportion β of aggregate consumption; we will relax this assumption later.

The formal definition of the equilibrium allocation follows:

Definition A CE with patent buyouts for R&D consists of quantities $\{c_t, \delta_t^h, X_{it}, A_t, L_{at}, N_t, \tau_t, \delta_t, z_t, \alpha_t, x_t\}$ and prices $\{w_t, P_{at}, v_t\}$ such that

1. $\{c_t, \delta_t^h\}$ solve the HH Problem.
2. X_{it} solve the FG Problem.
3. $L_{at}, z_t, \delta_t, A_t$ solve the R&D Problem.
4. w_t clears the labor market: $\int_0^{A_t} X_{it} di + L_{at} = N_t$.
5. v_t clears the mortality market: $\delta_t^h = \delta_t$.
6. The government buys ideas: $P_{at} \dot{A}_t = \beta c_t N_t$.
7. Lump sum taxes τ_t balance the budget: $\tau_t = P_{at} \dot{A}_t / N_t$.
8. Other conditions: $\dot{N}_t = (\bar{n} - \delta_t) N_t$, $\alpha_t = \alpha L_{at}^{\lambda-1} A_t^\phi$, and $x_t \equiv \frac{1}{A_t} \int_0^{A_t} X_{it} di / N_t$.

4.4. The Benchmark Case

The equilibrium behavior of the economy depends in important ways on a few parameters. We specify a benchmark case that will be studied in detail, and then in subse-

quent sections consider the effect of deviating from this benchmark. In specifying the benchmark, it is helpful to note that in equilibrium $c_t = A_t^\sigma(1 - s_t)$, where $\sigma \equiv \frac{1-\theta}{\theta}$ is the elasticity of consumption with respect to the stock of ideas. The benchmark case is then given by

Assumption A. (Benchmark Case) Let $\eta \equiv \lim_{z \rightarrow 0} F'(z)z/F(z)$. Assume

A1. Finite elasticity of $F(z)$ as $z \rightarrow 0$: $\eta \in (0, \infty)$

A2. Rapidly declining marginal utility of consumption: $\gamma > 1$

A3. Knowledge spillovers are not too strong: $\phi < 1 + \eta\sigma(\gamma - 1)$.

We will discuss the nature and role of each of these assumptions in more detail as we develop the results. The least familiar assumption is A1, but note that both the exponential and the Weibull distributions have this property. In contrast, the lognormal and Fréchet distributions have an infinite elasticity in the limit as z goes to zero. This has interesting implications that we will explore in detail.

4.5. The Equilibrium Balanced Growth Path

In the benchmark case, the equilibrium allocation in this growth model with dangerous technologies exhibits a balanced growth path:

Proposition 2 (Equilibrium Balanced Growth): *Under Assumption A, the competitive equilibrium exhibits an asymptotic balanced growth path as $t \rightarrow \infty$ such that*

$$s_t \rightarrow \frac{\beta}{1 + \beta + \eta} \quad (21)$$

$$z_t \rightarrow 0 \text{ (and therefore } \delta_t \rightarrow 0) \quad (22)$$

$$\dot{z}_t/z_t \rightarrow g_z \equiv -(\gamma - 1)g_c, \quad \dot{\delta}_t/\delta_t \rightarrow g_z \quad (23)$$

$$\dot{c}_t/c_t \rightarrow g_c = \sigma g_A \equiv \frac{\lambda\sigma\bar{n}}{1 - \phi + \eta\sigma(\gamma - 1)} \quad (24)$$

$$u'(c_t)v_t \rightarrow \frac{\bar{u}}{\rho}. \quad (25)$$

The somewhat surprising result that emerges from the equilibrium under Assumption A is that mortality and the danger threshold, rather than being constant in steady state,

decline at constant exponential rates. Technological change becomes increasingly conservative over time, as an ever-rising fraction of possible new ideas are rejected because they are too dangerous.

The consequence of this conservative bias in technological change is no less surprising: it leads to a slowdown in steady-state growth. There are several senses in which this is true, and these will be explored as the paper goes on. But two are evident now. First, the negative growth rate of z_t introduces a negative trend in TFP growth for the idea sector, other things equal. Recall that the idea production function is $\dot{A}_t = \alpha F(z_t) L_{at}^\lambda A_t^\phi$. z_t declines at the rate $(\gamma - 1)g_c$, ultimately getting arbitrarily close to zero. By Assumption A1, the elasticity of the distribution $F(z)$ at zero is finite and given by η , so $F(z_t)$ declines at rate $\eta(\gamma - 1)g_c = \eta(\gamma - 1)\sigma g_A$.

The second way to see how this bias slows down growth is to focus directly on consumption growth itself. The steady-state rate of consumption growth is

$$g_c = \frac{\lambda\sigma\bar{n}}{1 - \phi + \eta\sigma(\gamma - 1)}. \quad (26)$$

The last term in the denominator directly reflects the negative TFP growth in the idea production function resulting from the tightening of the danger threshold.

That this slows growth can be seen by considering the following thought experiment. A feasible allocation in this economy is to follow the equilibrium path until z_t is arbitrarily small and then keep it constant at this value. This results in a mortality rate that is arbitrarily close to zero, and, as suggested by Proposition 1, the growth rate in this case will be arbitrarily close to $\lambda\sigma\bar{n}/(1 - \phi)$, which is clearly greater than the equilibrium growth rate. Rather than keep z constant at a small level, the equilibrium continues to reduce the danger cutoff, slowing growth. Some numerical examples at the end of this paper suggest that this slowdown can be substantial.

Of course, this raises a natural question: *Why* does the equilibrium allocation lead the danger threshold to fall to zero? To see the answer, first consider the economic interpretation of the mortality price v_t . This is the price at which firms must compensate households per unit of mortality that their inventions inflict. In the terminology of the health and risk literatures, it is therefore equal to the value of a statistical life (VSL).⁵

⁵Suppose the mortality rate is $\delta_t = .001$ and $v_t = \$1$ million. In this example, each person receives \$1000 ($= v_t\delta_t$) for the mortality risk they face. For every thousand people in the economy, one person will

Along the balanced growth path, the value of life satisfies equation (25):

$$u'(c_t)v_t \rightarrow \frac{\bar{u}}{\rho}.$$

This equation says that the value of life measured in utils is asymptotically equal to the present discounted value of utility: as consumption goes to infinity, flow utility converges to \bar{u} , so lifetime utility is just \bar{u}/ρ .

Viewed in another way, this equation implies that the value of life grows faster than consumption. Given our functional form assumption for preferences, $u'(c_t) = c_t^{-\gamma}$. So $c_t^{-\gamma}v_t$ converges to a constant, which means that $g_v \rightarrow \gamma g_c$. Because γ is larger than one, marginal utility falls rapidly, and the value of life rises faster than consumption.

With this key piece of information, we can turn to the first-order condition for the choice of z_t in the R&D Problem. That first-order condition is

$$z_t = \frac{P_{at}}{v_t N_t \bar{\delta} x_t} = \frac{\beta c_t}{v_t \bar{\delta} g_{At} (1 - s_t)}. \quad (27)$$

The first part of this equation says that the danger threshold z_t equals the ratio of two terms. The numerator is related to the marginal benefit of allowing more dangerous technologies to be used, which is proportional to the price at which the additional ideas could be sold. The denominator is related to the marginal cost, which depends on the value of the additional lives that would be lost.

The second equation in (27) uses the fact that $P_{at} \dot{A}_t = \beta C_t$ to eliminate the price of ideas. This last expression illustrates the key role played by the value of life. In particular, we saw above that v_t/c_t grows over time since $\gamma > 1$; the value of life rises faster than consumption. Because both g_{At} and $1 - s_t$ are constant along the balanced growth path, the rapidly rising value of life leads to the exponential decline in z_t . More exactly, v_t/c_t grows at rate $(\gamma - 1)g_c$, so this is the rate at which z_t declines, as seen in equation (23).

What is the economic intuition? Because γ exceeds 1, flow utility $u(c)$ is bounded and the marginal utility of additional consumption falls very rapidly. This leads the value of life to rise faster than consumption. The benefit of using more dangerous technologies is that the economy gets more consumption. The cost is that more people die. Because the marginal utility of consumption falls so quickly, the costs of people dying

die, and the total compensation paid out for this death will equal \$1 million.

exceed the benefit of increasing consumption and the equilibrium delivers a declining threshold for technological danger. Safety trumps economic growth.

4.6. Growth Consequences

From the standpoint of growth theory, there are some interesting implications of this model. First, as we saw in the context of the rule-of-thumb allocation, this model is potentially a fully endogenous growth model. Growth is proportional to population growth, but because the mortality rate is endogenous, policy changes can affect the mortality rate and therefore affect long-run growth, at least potentially.

Interestingly, however, that is not the case in the equilibrium allocation. Instead, the mortality rate trends to zero and is unaffected by policy changes in the long run. Hence, the equilibrium allocation features semi-endogenous growth, where policy changes have long-run level effects but not growth effects. In particular, notice that the patent buyout parameter β , which influences the long-run share of labor going to research, is not a determinant of the long-run growth rate in (26). Moreover, the invariance of long-run growth to policy is true even though a key preference parameter, γ , influences the long-run growth rate.

Finally, it is interesting to consider the special case of $\phi = 1$, so the idea production function resembles that in Romer (1990). In this case,

$$\frac{\dot{A}_t}{A_t} = \alpha F(z_t) L_{at}^\lambda.$$

Here, growth does not explode even in the presence of population growth, as can be seen in equation (26). Instead, the negative trend in z_t and the fact that the mortality rate depends on the growth rate conspire to keep growth finite.

5. Extensions

The crucial assumptions driving the results have been collected together and labeled as Assumption A. In this section, we illustrate how things change when these assumptions are relaxed. Briefly, there are two main findings. First, when we consider distributions with an infinite elasticity at $z = 0$, the concern for safety becomes even more extreme:

equilibrium growth slows all the way to zero asymptotically. Second, we highlight the role played by $\gamma > 1$: if instead $\gamma < 1$, then the equilibrium allocation looks like the rule of thumb allocation, selecting a constant danger threshold in steady state.

5.1. Relaxing A1: Letting $\eta = \infty$

Recall that η is the elasticity of the danger distribution $F(z)$ in the limit as $z \rightarrow 0$. Intuitively, this parameter plays an important role in the model because $F(z)$ is the fraction of new ideas that are used in the economy, and z is trending exponentially to zero. The term $\eta g_z = -\eta(\gamma - 1)g_c$ therefore plays a key role in determining the growth rate of ideas along the balanced growth path:

$$g_A = \frac{\lambda \bar{n}}{1 - \phi + \eta \sigma (\gamma - 1)}. \quad (28)$$

Assumption A1 says that η is finite. This is true for a number of distributions, including the exponential ($\eta = 1$), the Weibull, and the gamma distributions. However, it is not true for a number of other distributions. Both the lognormal and the Fréchet distributions have an infinite elasticity at zero, for example. Given that we have no prior over which of these distributions is most relevant to our problem, it is essential to consider carefully the case of $\eta = \infty$.

In fact, it is easy to get a sense for what will happen by considering equation (28). As η rises in this equation, the steady-state growth rate of the economy declines. Intuitively, a 1% reduction in z has a larger and larger effect on $F(z)$: an increasing fraction of ideas that were previously viewed as safe are now rejected as too dangerous. This reasoning suggests that as η gets large, the steady-state growth rate falls to zero, and this intuition is confirmed in the following proposition:

Proposition 3 (Equilibrium Growth with $\eta = \infty$): *Let Assumptions A2 and A3 hold, but instead of A1, assume $\eta = \infty$. In the competitive equilibrium, as $t \rightarrow \infty$*

1. *The growth rate of consumption falls to zero: $\dot{c}_t/c_t \rightarrow 0$*
2. *The level of consumption goes to infinity: $c_t \rightarrow \infty$*
3. *The technology cutoff, the mortality rate, and the share of labor devoted to research all go to zero: $z_t \rightarrow 0, \delta_t \rightarrow 0, s_t \rightarrow 0$.*

When $\eta = \infty$, the increasingly conservative bias of technological change slows the exponential growth rate all the way to zero. However, this does not mean that growth ceases entirely. Instead, the level of consumption still rises to infinity, albeit at a slower and slower rate.

5.2. Relaxing A2: Assume $\gamma < 1$

The most important assumption driving the results in this paper is that marginal utility diminishes quickly, in the sense that $\gamma > 1$. For example, the value of a year of life in year t as a ratio to consumption is

$$\frac{u(c_t)}{u'(c_t)c_t} = \bar{u}c_t^{\gamma-1} + \frac{1}{1-\gamma}.$$

For $\gamma > 1$, this rises to infinity as consumption grows. But for $\gamma < 1$, it converges to $1/(1-\gamma)$: the value of life is proportional to consumption. In this case, the elasticity of utility with respect to consumption remains positive rather than falling to zero, which keeps the value of life and consumption on equal footing. The result is that the trend toward rising safety disappears: the economy features exponential growth in consumption with a constant danger cutoff and a constant, positive mortality rate:

Proposition 4 (Equilibrium Growth with $\gamma < 1$): *Let Assumption A1 hold, strengthen A3 to $\phi < 1$, but instead of A2, assume $\gamma < 1$. The competitive equilibrium exhibits an asymptotic balanced growth path as $t \rightarrow \infty$ such that*

$$\begin{aligned} z_t &\rightarrow z^* \in (0, \infty), \quad \delta_t \rightarrow \delta^* \in (0, \infty) \\ s_t &\rightarrow \frac{\beta(1 - \frac{\Gamma(z^*)}{z^*})}{1 + \beta(1 - \frac{\Gamma(z^*)}{z^*})} \\ \frac{\dot{c}_t}{c_t} \rightarrow g_c &= \frac{\lambda\sigma(\bar{n} - \delta^*)}{1 - \phi} = \frac{\lambda\sigma\bar{n}}{1 - \phi + \lambda\bar{\delta}(1 - s^*)\Gamma(z^*)} \\ \frac{v_t}{c_t} &\rightarrow \frac{1}{1-\gamma} \cdot \frac{1}{\rho + \delta^* - (1-\gamma)g_c} \end{aligned}$$

For $\gamma < 1$, the economy looks very similar to the rule of thumb allocation; for example, compare the growth rate to that in Proposition 1. The economy features a constant

danger cutoff as well as endogenous growth: an increase in idea purchases by the government (a higher β) will shift more labor into research, lower the mortality rate, and increase the long-run growth rate.⁶

5.3. Optimality

In this section, we study the allocation of resources that maximizes a social welfare function. There are three reasons for this. First, it is important to verify that the declining danger threshold we uncovered in the equilibrium allocation is not a perverse feature of our equilibrium. Second, in the equilibrium allocation, individuals put no weight on the welfare of future generations; it is a purely selfish equilibrium. It is interesting to study the effect of deviations from this benchmark. Finally, our equilibrium allocation employed a particular institution for funding research: patent buyouts where spending on new ideas is in constant proportion to consumption. This institution is surely special (and not generally optimal), so it is important to confirm that it is not driving the results. The bottom line of this extension to an optimal allocation is that all of the equilibrium results hold up well.

In this environment with multiple generations, there is no indisputable social welfare function. However, a reasonably natural choice that serves our purposes is to treat flows of utility from different people symmetrically and to discount flows across time at rate ρ . This leads to the following definition of an optimal allocation:

Definition An *optimal allocation* in this economy is a time path for $\{s_t, z_t\}$ that solves

$$\begin{aligned} \max_{\{s_t, z_t\}} \int_0^{\infty} e^{-\rho t} N_t u(c_t) dt \quad \text{subject to} \\ c_t = A_t^\sigma (1 - s_t) \\ \dot{A}_t = \alpha F(z_t) s_t^\lambda N_t^\lambda A_t^\phi \\ \dot{N}_t = (\bar{n} - \delta_t) N_t \\ \delta_t = \bar{\delta} \alpha s_t^\lambda N_t^\lambda A_t^{\phi-1} (1 - s_t) \int_0^{z_t} z f(z) dz. \end{aligned}$$

⁶The intermediate case of log utility ($\gamma = 1$) requires separate consideration. In this case, the technology cutoff z_t still declines to zero, but this decline is slower than exponential. The long-run growth rate is then precisely back to the semi-endogenous growth case for $\eta < \infty$: $g_A = \lambda \bar{n} / (1 - \phi)$.

The optimal allocation can then be characterized as follows.

Proposition 5 (Optimal Balanced Growth): *Under Assumption A and the optimal allocation, the economy exhibits an asymptotic balanced growth path as $t \rightarrow \infty$ such that*

$$\frac{s_t}{1 - s_t} \rightarrow \frac{\lambda \sigma g_A}{(1 + \eta)(\rho - \bar{n} + (\gamma - 1)g_c) + (1 - \phi)g_A - \eta \sigma g_A}$$

$$z_t \rightarrow 0 \text{ (and therefore } \delta_t \rightarrow 0)$$

$$\dot{z}_t/z_t \rightarrow -(\gamma - 1)g_c, \quad \dot{\delta}_t/\delta_t \rightarrow -(\gamma - 1)g_c$$

$$\dot{c}_t/c_t \rightarrow g_c,$$

where g_c and g_A are the same as in the competitive equilibrium.

The key properties of the competitive equilibrium carry over into the optimal allocation. In particular, the danger threshold declines exponentially to zero at the rate $(\gamma - 1)g_c$, and this technological bias slows the growth rate of the economy. The long-run growth rate is the same as in the equilibrium allocation.

5.4. Comparing the Simple and Full Models

The simple toy model at the start of the paper and the richer model developed subsequently lead to slightly different conclusions. In the simple model, consumption converges to a constant value if $\gamma > 1$, while in the richer model the worst that happens is that the growth rate slows to zero as consumption rises to infinity. Why the difference?

The answer turns on functional forms and modeling choices about which we have very little information. We have already seen an example of this in the question about the magnitude of η . Another example goes back to the simple model. If instead of assuming the mortality function $\delta(g)$ is linear, we instead assume $\delta(g) = \delta g^\theta$, then the simple model leads consumption to rise to infinity as the growth rate slows to zero when $\theta > 1$.⁷

⁷The first order condition analogous to equation (3) becomes

$$g^{\theta-1}(1 + g) = \frac{\eta_{u,c}}{\delta\theta}$$

which implies that $g \rightarrow 0$ only occurs as $\eta_{u,c} \rightarrow 0$ when $\gamma > 1$.

Absent empirical evidence on the details of functional forms, the general result of this paper is that concerns for safety can lead to slowing growth, with the precise nature of the slowdown depending on modeling details.

6. Numerical Examples

We now report a couple of numerical examples to illustrate how this economy behaves along the transition path. We make some attempt to choose plausible parameter values and to produce simulations that have a “realistic” look to them. However, the model abstracts from a number of important forces shaping economic growth and mortality, so the examples should not be taken too literally. Mainly, they will illustrate the extent to which growth can be slowed by concerns about the dangers of certain technologies.

The first example features sustained exponential growth ($\eta < \infty$). The second assumes a distribution $F(z)$ with an elasticity that rises to infinity as z falls to zero. According to Proposition 3, this second example exhibits a growth rate that declines to zero, even though consumption itself rises indefinitely.

6.1. Benchmark Example

The basic parameterization of the benchmark case is described in Table 1. For the curvature of marginal utility, we choose $\gamma = 1.5$; large literatures on intertemporal choice (Hall 1988), asset pricing (Lucas 1994), and labor supply (Chetty 2006) suggest that this is a reasonable value. For $F(z)$, we assume an exponential distribution so that $\eta = 1$; we also assume this distribution has a mean of one. We set $\beta = .02$: government spending on ideas equals 2% of aggregate consumption. For the idea production function, we choose $\lambda = 1$ and $\phi = 1/2$, implying that in the absence of declines in z_t , the idea production function itself exhibits productivity growth. Finally, we assume a constant population growth rate of 1% per year. The remaining parameter values are relatively unimportant and are shown in the table. Other reasonable choices for parameter values will yield similar results qualitatively. The model is solved using a reverse shooting technique, discussed in more detail in Appendix B.


Figure 3 shows an example of the equilibrium dynamics that occur in this economy for the benchmark case. The economy features a steady-state growth rate of per capita

Table 1: Benchmark Parameter Values

Parameter Value	Comment
$\gamma = 1.5$	Slightly more curvature than log utility
$\eta = 1$	$F(z)$ is an exponential distribution
$\beta = 0.02$	Government spends 2% of consumption on ideas
$\lambda = 1, \phi = 1/2$	Idea production function: $\dot{A}_t = \alpha F(z_t) A_t^{1/2} L_{at}$
$\bar{n} = .01$	Long-run population growth rate
$\sigma = 2$	Elasticity of consumption wrt ideas
$\rho = .05$	Rate of time preference
$\bar{\delta} = 50$	Mortality rate intercept

Note: These are the baseline parameter values for the numerical examples.

Figure 3: Equilibrium Dynamics: Benchmark Case



Note: Simulation results for the competitive equilibrium using the parameter values from Table 1. Consumption growth settles down to a constant positive rate, substantially lower than what is feasible. The danger threshold and mortality rate converge to zero.

consumption of 1.33%. This constant growth occurs while the danger threshold and the mortality rate decline exponentially to zero; both z_t and δ_t grow at -0.67%.


Several other features of the growth dynamics are worth noting. First, the particular initial conditions we've chosen have the growth rate of consumption declining along the transition path; a different choice could generate a rising growth rate, although declining growth appears to be more consistent with the value of life in the model (more on this below).

Second, consider total factor productivity for the idea production function. With $\lambda = 1$, TFP is $\alpha F(z_t) A_t^\phi$. Because we've assumed $\phi = 1/2$, this production function has the potential to exhibit positive TFP growth as knowledge spillovers rise over time. However, a declining danger threshold can offset this. In steady state, TFP growth for the idea production function is $\phi g_A + \eta g_z = -0.33\%$. That is, even though a given number of researchers are generating more and more candidate ideas over time, the number that get implemented is actually declining because of safety considerations.

Finally, the steady-state growth rate of 1.33% can be compared to an alternative path. It is feasible in this economy to let the technology-induced mortality rate fall to some arbitrarily low level — such as 1 death per billion people — and then to keep it constant at that rate forever by maintaining a constant technology cutoff \bar{z} . As this constant cutoff gets arbitrarily small, the steady state growth rate of the economy converges to $\lambda \sigma \bar{n} / (1 - \phi)$ — that is, to the rule-of-thumb growth rate from Proposition 1. For our choice of parameter values, this feasible steady-state growth rate is 4.0% per year. That is, concerns for safety make it optimal in this environment to slow growth considerably relative to what is possible in the steady state.

The reason for this, of course, is the rising value of life, shown for this example in Figure 4. The value of life begins in period 0 at about 100 times annual consumption; if we think of per capita consumption as \$30,000 per year, this corresponds to a value of life of \$3 million, very much in the range considered in the literature (Viscusi and Aldy 2003; Ashenfelter and Greenstone 2004; Murphy and Topel 2005). Over time, the value of life relative to consumption rises exponentially at a rate that converges to 0.67%, the same rate at which mortality declines.

Figure 4: The Value of Life: Benchmark Case



Note: Simulation results for the competitive equilibrium using the parameter values from Table 1. The value of life rises faster than consumption.


6.2. Numerical Example When $\eta = \infty$

One element of the model that is especially hard to calibrate is the distribution from which technological danger is drawn, $F(z)$. The previous example assumed an exponential distribution so that $\eta = 1$; in particular, the elasticity of the distribution as z approaches zero is finite. However, this need not be the case. Both the Fréchet and the lognormal distributions feature an infinite elasticity. In Proposition 3, we showed that this leads the growth rate of consumption to converge to zero asymptotically. For this example, we consider the Fréchet distribution to illustrate this result: $F(z) = e^{-z^{-\psi}}$ and we set $\psi = 1.1$.⁸ Other parameter values are unchanged from the benchmark case shown in Table 1, except we now set $\bar{\delta} = 1$, which is needed to put the value of life in the right ballpark.

Figure 5 shows the dynamics of the economy for this example. The growth rate of consumption now converges to zero as $\eta(z)$ gets larger and larger, meaning that a given decline in the danger threshold eliminates more and more potential ideas. Interest-

⁸Assuming $\psi > 1$ ensures that the mean exists. The elasticity of this cdf is $\eta(z) = \psi z^{-\psi}$, so a small value of ψ leads the elasticity to rise to infinity relatively slowly.

Figure 5: Equilibrium Dynamics: Fréchet Case



Note: Dynamics when $F(z)$ is Fréchet, so $\eta = \infty$: growth slows to zero asymptotically. See notes to Figure 3.


ingly, this rising elasticity means that the danger threshold itself declines much more gradually in this example.

Figure 6 — with its logarithmic scale — suggests that this declining consumption growth rate occurs as consumption gets arbitrarily high. The value of life still rises faster than consumption, but the increase is no longer exponential.

7. Discussion and Evidence

A key result from both the simple model and the richer endogenous growth model is that safety has a large impact on the nature of economic growth if the marginal utility of consumption declines rapidly, that is if $\gamma > 1$. In this section, we discuss a range of evidence on γ . Is $\gamma > 1$ or $\gamma < 1$ likely to be the more relevant case? To summarize the findings, the evidence appears quite mixed. The case for $\gamma > 1$ may certainly be the relevant one, but it may not be. More work will be needed to decide.

Figure 6: Consumption and the Value of Life: Fréchet Case



Note: Dynamics when $F(z)$ is Fréchet, so $\eta = \infty$: consumption still rises to infinity. See notes to Figure 4.

7.1. Risk Aversion and the Elasticity of Intertemporal Substitution

In the most common way of specifying preferences for macro applications, the coefficient of relative risk aversion, γ in our notation, equals the inverse of the elasticity of intertemporal substitution. Large literatures on asset pricing (Lucas 1994) and labor supply (Chetty 2006) suggest that $\gamma > 1$ is a reasonable value.

Evidence on the elasticity of intertemporal substitution, $1/\gamma$ in our notation, is more mixed. The traditional view, supported by Hall (1988), is that this elasticity is less than one, consistent with the case of $\gamma > 1$. This view is supported by a range of careful microeconomic work, including Attanasio and Weber (1995), Barsky, Juster, Kimball and Shapiro (1997), and Guvenen (2006); see Hall (2009) for a survey of this evidence. On the other hand, Vissing-Jorgensen and Attanasio (2003) and Gruber (2006) find evidence that the elasticity of intertemporal substitution is greater than one, suggesting that $\gamma < 1$ could be appropriate.

7.2. Empirical Evidence on the Value of Life

Direct evidence on how the value of life changes with income is surprisingly difficult to come by. Most of the empirical work in this literature is cross-sectional in nature and focus on getting a single measure of the value of life (or perhaps a value by age); see Ashenfelter and Greenstone (2004), for example. There are a few studies that contain important information on the income elasticity, however. Viscusi and Aldy (2003) conduct a meta-analysis and find that across studies, the value of life exhibits an income elasticity below one. On the other hand, Costa and Kahn (2004) and Hammitt, Liu and Liu (2000) consider explicitly how the value of life changes over time. These studies find that the value of life rises roughly twice as fast as income, supporting the basic mechanism in this paper.

7.3. Evidence from Health Spending

The key mechanism at work in this paper is that the marginal utility of consumption falls quickly if $\gamma > 1$, leading the value of life to rise faster than consumption. This tilts the allocation in the economy away from consumption growth and toward preserving lives. Exactly this same mechanism is at work in Hall and Jones (2007), which studies health spending. In that paper, $\gamma > 1$ leads to an income effect: as the economy gets richer over time (exogenously), it is optimal to spend an increasing fraction of income on health care in an effort to reduce mortality. The same force is at work here in a very different context. Economic growth combines with sharply diminishing marginal utility to make the preservation of life a luxury good. The novel finding is that this force has first-order effects on the determination of economic growth itself.

What evidence is there for an income elasticity of health spending larger than one? Figure 7 shows two pieces of evidence. Panel (a) documents that health spending as a share of GDP is rising in many countries of the world, not only in the United States. Indeed, for the 19 OECD countries reporting data in both 1970 and 2006 (many not shown), all experienced a rising health share.

Panel (b) shows another well-documented fact: the share of GDP an economy spends on health rises with GDP per capita. That is, rich countries spend a higher fraction of their income on health than poor countries, although this result is more nuanced when

other controls are included; see Newhouse (1977) Reinhardt, Hussey and Anderson (2002), and the extensive survey in Gerdtham and Jonsson (2000). The cross-sectional evidence shown in Figure 7 for the year 2006 suggests a simple income elasticity of 1.145. Of course, this ordinary least squares estimate of the elasticity is biased because a country's per capita income is surely a function of its health spending, at least to some extent (Fogel (1994); Becker et al. (2005); Weil (2007); Acemoglu and Johnson (2007)).

Acemoglu, Finkelstein and Notowidigdo (2009a) estimate an elasticity of hospital spending with respect to transitory income of 0.7, less than one, using oil price movements to instrument local income changes at the county level in the southern part of the United States. While useful, it is not entirely clear that this bears on the key parameter here, as that paper considers income changes that are temporary (and hence might reasonably be smoothed and not have a large effect on health spending) and local (and hence might not alter the limited selection of health insurance contracts that are available).


7.4. Evidence from Accident Rates

Less direct evidence may be obtained by considering our changing concerns regarding safety. It is a common observation that parents today are much more careful about the safety of their children than parents a generation ago. Perhaps that is because the world is a more dangerous place, but perhaps it is in part our sensitivity to that danger which has changed.

Starr (1969) is an early example of a study that looked at accident rates for different technologies in a cost-benefit fashion. A trend toward increasing safety is revealed in Figure 8, which shows injury and fatality rates for motor vehicle, airplanes, and accidental drownings. (The choice of these categories was guided by a concern to minimize the role of progress in health care interventions, though this cannot be completely achieved).

Safety standards also appear to differ significantly across countries, in a way that is naturally explained by the model. While more formal data is clearly desirable, different standards of safety in China and the United States have been vividly highlighted by recent events in the news. Eighty-one deaths in the United States have been linked to the contamination of the drug heparin in Chinese factories (Mundy 2008). In the

Figure 8: The Trend toward Increasing Safety in the United States



Note: Motor vehicle injury rates are from the National Highway Transportation Safety Administration, “Traffic Safety Facts” DOT HS 811 172, June 2009. Aviation data are from the National Transportation Safety Board and various issues of *Injury Facts*. Drowning data are from various issues of the National Center for Health Statistics, Vital Statistics Data.

summer of 2007, 1.5 million toys manufactured for Mattel by a Chinese supplier were recalled because they were believed to contain lead paint (Spencer and Ye 2008). And in an article on the tragic health consequences for workers producing toxic cadmium batteries in China, the *Wall Street Journal* reports

As the U.S. and other Western nations tightened their regulation of cadmium, production of nickel-cadmium batteries moved to less-developed countries, most of it eventually winding up in China. “Everything was transferred to China because no one wanted to deal with the waste from cadmium,” says Josef Daniel-Ivad, vice president for research and development at Pure Energy Visions, an Ontario battery company. (Casey and Zamiska 2007)

A limitation of this evidence on safety, of course, is that the differences over time and across countries could reflect transition dynamics in a model with $\gamma < 1$.

7.5. Sustainable Growth and the Environment

This paper is most closely related to the literature on sustainable growth and the environment; for example, see Solow (1974), Stiglitz (1974), Gradus and Smulders (1993),

and Aghion and Howitt (1998, Ch. 5). Particularly relevant are Stokey (1998) and Brock and Taylor (2005), who study the environmental Kuznets curve in which pollution first rises and then falls with economic development. In these papers, pollution enters the utility function as a cost in an additively separable fashion from consumption. These models feature an income effect for $\gamma > 1$ because the utility from growing consumption is bounded. This leads to a “growth drag” from the environment: growth is slower than it would otherwise be because of environmental concerns. While the key issues here are very distinct — the utility costs of pollution in one case versus the loss of life associated with dangerous technologies in the other — it is interesting that the curvature of marginal utility plays a central role in both and can slow growth.

One of the ways in which pollution has been mitigated in the United States is through the development of new, cleaner technologies. Examples include scrubbers that remove harmful particulates from industrial exhaust and catalytic converters that reduce automobile emissions. Researchers can spend their time making existing technologies safer or inventing new technologies. Rising concerns for safety lead them to divert effort away from new technologies, which may slow growth. Acemoglu, Aghion, Bursztyn and Hemous (2009b) explore this kind of directed technical change in a model of growth and the environment.

8. Conclusion

Technological progress involves risks as well as benefits. Considering the risks posed to life itself leads potentially to first-order changes in the theory of economic growth. This paper explores these possibilities, first in a simple “Russian roulette” style model and then in a richer model in which growth explicitly depends on the discovery of new ideas. The results depend somewhat on the details of the model and, crucially, on how rapidly the marginal utility of consumption declines. It may be optimal for growth to continue exponentially despite the presence of existential risks, or it may be optimal for growth to slow to zero, even potentially leading to a steady-state level of consumption.

The intuition for the possible end to exponential growth turns out to be straightforward. For a large class of standard preferences, safety is a luxury good. The marginal utility associated with more consumption on a given day runs into sharp diminishing

returns, and ensuring additional days of life on which to consume is a natural, welfare-enhancing response. When the value of life rises faster than consumption, economic growth leads to a disproportionate concern for safety. This concern can be so strong that it is desirable that growth slow down.

The framework studied here clearly omits other factors that may be important. Health technologies can help to extend life, possibly offsetting some of the concerns here. Even dangerous technologies like nuclear weapons could have a life-saving use — for example if they helped to divert an asteroid that might otherwise hit the earth.

This paper suggests a number of different directions for future research on the economics of safety. It would clearly be desirable to have precise estimates of the value of life and how this has changed over time; in particular, does it indeed rise faster than income and consumption? More empirical work on how safety standards have changed over time — and estimates of their impacts on economic growth — would also be valuable. Finally, the basic mechanism at work in this paper over time also applies across countries. Countries at different levels of income may have very different values of life and therefore different safety standards. This may have interesting implications for international trade, standards for pollution and global warming, and international relations more generally.

A Appendix: Proofs of the Propositions

This appendix contains outlines of the proofs of the propositions reported in the paper.

Proof of Proposition 1. BGP under the Rule of Thumb Allocation

Equations (18) and (19) follow immediately from the setup. The growth rate of ideas then comes from taking logs and derivatives of both sides of the following equation, evaluated along a balanced growth path, and using the fact that $\delta^* = \bar{\delta}g_A(1 - \bar{s})\Gamma(\bar{z})$:

$$\frac{\dot{A}_t}{A_t} = \alpha F(\bar{z}) \frac{\bar{s}^\lambda N_t^\lambda}{A_t^{1-\phi}}.$$

QED.

Proof of Proposition 2. Equilibrium Balanced Growth

Solving the optimization problems that help define the equilibrium and making some substitutions leads to the following seven key equations that pin down the equilibrium values of $\{c_t, s_t, v_t, z_t, A_t, \delta_t, N_t\}$:

$$c_t = A_t^\sigma (1 - s_t) \quad (29)$$

$$\frac{s_t}{1 - s_t} = \beta - \frac{v_t \delta_t}{c_t} = \beta \left(1 - \frac{\Gamma(z_t)}{z_t}\right) \quad (30)$$

$$v_t = \frac{u(c_t)/u'(c_t)}{\rho + \delta_t + \gamma g_{ct} - g_{vt}} \quad (31)$$

$$z_t = \frac{\beta c_t}{v_t \bar{\delta} g_{At} (1 - s_t)} \quad (32)$$

$$\frac{\dot{A}_t}{A_t} = \alpha F(z_t) \frac{s_t^\lambda N_t^\lambda}{A_t^{1-\phi}} \quad (33)$$

$$\delta_t = \bar{\delta} g_{At} (1 - s_t) \Gamma(z_t) \quad (34)$$

$$\frac{\dot{N}_t}{N_t} = \bar{n} - \delta_t \quad (35)$$

We start with some preliminaries. Two technical assumptions are needed in the proof. First, we assume the first moment of $F(z)$ exists (this assumption can be weakened). Second, we assume that $F(z) > 0$ for $z > 0$. It is straightforward to show that $\Gamma(z) \leq z$ so that from (30) $s_t \in [0, \beta/(1 + \beta)]$ ensuring $s_t < 1$. Moreover, $s_t > 0$ if $z_t > 0$.

With these preliminaries out of the way, we begin by proving that $c_t \rightarrow \infty$. Especially in light of the results from the simple model at the start of the paper, this is an important step. This proof proceeds by contradiction. Suppose not. That is, suppose $c_t \rightarrow c_* \in (0, \infty)$. The contradiction arises because the model has a strong force for idea growth. Because $c_t = A_t^\sigma (1 - s_t)$ and $s_t \leq \frac{\beta}{1 + \beta}$, $\lim A_t \leq (1 + \beta) c_*^{1/\sigma}$, which means that $g_{At} \rightarrow 0$. Because (31) implies that $v_t \rightarrow v_*$, (32) requires

$$z_t = \frac{\beta c_t}{v_t \bar{\delta} g_{At} (1 - s_t)} \rightarrow \infty.$$

But this suggests that the economy is using all the new ideas that are created, which will produce a contradiction in that $g_A > 0$ will arise. To see this, consider the idea

production function in (33):

$$\lim_{t \rightarrow \infty} g_{At} = \lim_{t \rightarrow \infty} \alpha F(z_t) \frac{s_t^\lambda N_t^\lambda}{A_t^{1-\phi}} = \alpha \frac{s_*^\lambda N_*^\lambda}{A_*^{1-\phi}}$$

where we've used the fact that $F(\infty) = 1$ and $s_* > 0$ because $\lim z_t = \infty > 0$. The only way $g_{At} \rightarrow 0$, then, is if $N_* = 0$. However, this not the case. Instead, the death rate is given by $\delta_t = \bar{\delta} g_{At} (1-s_t) \Gamma(z_t) \rightarrow 0$ since $g_{At} \rightarrow 0$ while $\Gamma(\infty) < \infty$ due to the assumption that the first moment of $F(z)$ exists. But this means that $g_N \rightarrow \bar{n}$ and therefore $N_* \neq 0$, which is a contradiction, proving that it must be the case that $c_t \rightarrow \infty$.

Next, we show that $\delta_t \rightarrow 0$ and $z_t \rightarrow 0$. Consider equation (31):

$$v_t = \frac{u(c_t)/u'(c_t)}{\rho + \delta_t + \gamma g_{ct} - g_{vt}}.$$

With our specification of utility, $u(c_t)/u'(c_t) = \bar{u}c_t^\gamma + c_t/(1-\gamma)$. With $c_t \rightarrow \infty$ and $\gamma > 1$, $v_t/c_t \rightarrow \infty$, as the value of life rises faster than consumption. The presence of δ_t in the denominator above turns out not to be a problem, as one can use the fact that $\Gamma(z) \leq z$ to show that

$$\delta_t = \bar{\delta} g_{At} (1-s_t) \Gamma(z_t) \leq \bar{\delta} g_{At} (1-s_t) z_t = \beta \cdot \frac{c_t}{v_t},$$

so that $\delta_t \rightarrow 0$.

Equation (32) then implies that

$$z_t g_{At} = \frac{\beta c_t}{v_t \bar{\delta} (1-s_t)}. \quad (36)$$

But then $v_t/c_t \rightarrow \infty$ implies that $z_t g_{At} \rightarrow 0$.

We show that this requires z_t to go to zero. Why? Suppose not. That is, suppose instead that $z_t \rightarrow z_* > 0$ and $g_{At} \rightarrow 0$. Because $\delta_t \rightarrow 0$, the population grows at rate \bar{n} eventually and the law of motion for ideas (33) would lead to exponential growth in A_t , which is a contradiction. Therefore z_t has to go to zero.

Now consider growth rates along a balanced growth path. First return to the value of life in equation (31). In the long run, v_t is proportional to $u(c_t)/u'(c_t) = \bar{u}c_t^\gamma + c_t/(1-\gamma)$. Since $\gamma > 1$, the growth rate of the value of life must be equal to γg_c along an (asymptotic) balanced growth path. Equation (31) then implies the last result in the proposition, namely that $u'(c_t)v_t \rightarrow \bar{u}/\rho$. The fact that $g_v \rightarrow \gamma g_c$ immediately implies from (32)

that $g_z = -(\gamma - 1)g_c$.

To get the growth rate of ideas, recall that

$$\frac{\dot{A}_t}{A_t} = \alpha F(z_t) \frac{\bar{s}^\lambda N_t^\lambda}{A_t^{1-\phi}}.$$

Taking logs and derivatives of this equation along a balanced growth path, and using the fact that $\eta \equiv \lim_{z \rightarrow 0} F'(z)z/F(z)$ is finite from Assumption A, we have

$$(1 - \phi)g_A = \eta g_z + \lambda \bar{n}.$$

The growth rate results of the proposition then follow quickly, after we note that $g_c = \sigma g_A$ and $g_z = -(\gamma - 1)g_c$. For example

$$g_A = \frac{\lambda \bar{n}}{1 - \phi + \eta \sigma (\gamma - 1)}.$$

Finally, the share of labor devoted to research in steady state can be found from equation (30). Using L'Hopital's rule, one can show that

$$\lim_{z_t \rightarrow 0} \frac{\Gamma(z_t)}{z_t} = \lim_{z_t \rightarrow 0} \frac{\int_0^{z_t} z dF(z)}{z_t F(z_t)} = \lim_{z_t \rightarrow 0} \frac{z_t F'(z_t)}{F(z_t) + z_t F'(z_t)} = \lim_{z_t \rightarrow 0} \frac{\frac{z_t F'(z_t)}{F(z_t)}}{1 + \frac{z_t F'(z_t)}{F(z_t)}} = \frac{\eta}{1 + \eta}. \quad (37)$$

Substituting this into (30) yields the asymptotic value for s . QED.

Proof of Proposition 3. Equilibrium Growth with $\eta = \infty$

The proof of this proposition closely follows the proof of Proposition 2 above; after all, the main difference is the value of η . In particular, the proof that $c_t \rightarrow \infty$ and $z_t \rightarrow 0$ is identical.

We depart from the previous proof just after we showed that $z_t \rightarrow 0$, following equation (36). There we found that $z_t g_{At} \rightarrow 0$ and $z_t \rightarrow 0$. In Proposition 2, we then showed that $g_{At} \rightarrow g_A > 0$. Now instead we show that g_{At} must also go to zero.

Rewrite (33) as

$$\frac{g_{At} A_t^{1-\phi}}{N_t^\lambda} = \alpha F(z_t) s_t^\lambda.$$

Because $z_t \rightarrow 0$, the righthand side of this equation goes to zero. This means the left-

hand side must also, and this can happen in one of two ways. First, $g_{At} \rightarrow 0$, which is what we'd like to show. The alternative is that $g_{At} \rightarrow \hat{g} > 0$ such that $(1 - \phi)\hat{g} < \lambda\bar{n}$. But we can show that this alternative case leads to a contradiction. If g_{At} settles down to a positive constant rate, then many of the arguments used above to prove Proposition 2 hold. In particular, we can take logs and derivatives of the idea production function in (33) to see that

$$(1 - \phi)\hat{g} = \lim_{t \rightarrow \infty} \eta(z_t)g_z + \lambda\bar{n}.$$

Moreover, from (32) it must then be true that $g_z = -(\gamma - 1)g_c = -(\gamma - 1)\sigma\hat{g}$. However, our key assumption for this proposition is that the elasticity of $F(z)$ goes to infinity as $z \rightarrow 0$ implying that $\lim_{z \rightarrow 0} \eta(z) = \infty$. The negative growth rate of z_t then means that this equation cannot yield a positive solution for \hat{g} . Therefore the first case above must hold and $g_{At} \rightarrow 0$.

The fact that $s_t \rightarrow 0$ comes from (30) because $\lim_{z \rightarrow 0} \Gamma(z)/z = 1$ when $\eta = \infty$, as can be seen by the argument above in equation (37). Finally, since $c_t = A_t^\sigma(1 - s_t)$ and $g_{At} \rightarrow 0$, it must also be true that $g_{ct} \rightarrow 0$. QED.

Proof of Proposition 4. Equilibrium Growth with $\gamma < 1$

The equilibrium with $\gamma < 1$ is characterized by the same seven equations listed above in the proof of Proposition 2, equations (29) through (35). The proof begins in the same way, by studying the value of life in equation (31). With our specification of utility, $u(c_t)/u'(c_t) = \bar{u}c_t^\gamma + c_t/(1 - \gamma)$. Since $\gamma < 1$, the constant term disappears asymptotically and the growth rate of the value of life equals g_c along an (asymptotic) balanced growth path. Equation (31) then implies the last result in the proposition, giving the constant ratio of the value of life to consumption.

The fact that v_t/c_t converges to a constant means that z_t converges to a nonzero value, according to equation (32). Similarly, δ_t does as well, according to equation (34). The solution for the growth rate and the research share are found in ways similar to those in the proof of Proposition 2.

The solution for z^* involves a nonlinear equation with no explicit solution. For ex-

ample, one way to write this solution is

$$\frac{z^*}{\Gamma(z^*)} = \beta(1 - \gamma)(1 + \rho/\delta^*)$$

where $\delta^* = \bar{\delta}g_A^*(1 - s^*)\Gamma(z^*)$ and the equations stating g_A^* and s^* as functions of z^* are given in the statement of the proposition.

These equations can be combined to yield a single equation in z^* :

$$\frac{z^*}{\Gamma(z^*)} = \beta(1 - \gamma) \left(1 + \frac{\rho}{\bar{n}} \left[\frac{1 - \phi}{\lambda \bar{\delta} \Gamma(z^*)} \left(1 + \beta \left(1 - \frac{\Gamma(z^*)}{z^*} \right) \right) + 1 \right] \right). \quad (38)$$

I do not currently have a set of conditions that guarantee the existence or uniqueness of a solution to this equation, though some progress can certainly be made with the equation above (for example, β sufficiently small goes a long way toward getting existence and uniqueness). Numerical simulation for the case of the exponential distribution leads to a unique solution. At the moment, then, a technical condition for the proposition to hold is that a unique nonnegative solution z^* solves equation (38). QED.

Proof of Proposition 5. Optimal Balanced Growth

The Hamiltonian for the optimal growth problem is

$$H = N_t u(c_t) + \mu_{1t} \alpha F(z_t) s_t^\lambda N_t^\lambda A_t^\phi + \mu_{2t} N_t \left(\bar{n} - \bar{\delta} \alpha F(z_t) s_t^\lambda N_t^\lambda A_t^{\phi-1} (1 - s_t) \int_0^{z_t} z f(z) dz \right).$$

Applying the Maximum Principle, the first order necessary conditions are

$$H_s = 0: \quad N_t u'(c_t) A_t^\sigma = \mu_{1t} \lambda \frac{\dot{A}_t}{s_t} - \mu_{2t} N_t \left(\lambda \frac{\delta_t}{s_t} - \frac{\delta_t}{1 - s_t} \right)$$

$$H_z = 0: \quad \mu_{1t} = \mu_{2t} N_t \bar{\delta} z_t \cdot \frac{1 - s_t}{A_t}$$

$$\text{Arbitrage}(A_t): \quad \rho = \frac{\dot{\mu}_{1t}}{\mu_{1t}} + \frac{1}{\mu_{1t}} \left(N_t u'(c_t) \sigma \frac{c_t}{A_t} + \mu_{1t} \phi \frac{\dot{A}_t}{A_t} - \mu_{2t} N_t (\phi - 1) \frac{\delta_t}{A_t} \right)$$

$$\text{Arbitrage}(N_t): \quad \rho = \frac{\dot{\mu}_{2t}}{\mu_{2t}} + \frac{1}{\mu_{2t}} \left(u(c_t) + \mu_{1t} \lambda \frac{\dot{A}_t}{N_t} + \mu_{2t} (\bar{n} - \delta_t) - \mu_{2t} N_t \lambda \frac{\delta_t}{N_t} \right)$$

together with two transversality conditions: $\lim_{t \rightarrow \infty} \mu_{1t} A_t e^{-\rho t} = 0$ and $\lim_{t \rightarrow \infty} \mu_{2t} N_t e^{-\rho t} = 0$.

Combining these first order conditions (use the first and second to get an expression

Table 2: Key “State-Like” Variables for Studying Transition Dynamics

Variable	Meaning	Steady-State Value
$\ell_t \equiv \frac{v_t \delta_t}{c_t}$	Value of life \times mortality	$\ell^* = \beta \cdot \frac{\eta}{1+\eta}$
$m_t \equiv g_{A_t}$	Growth rate of A_t	$m^* = g_A$
δ_t	Mortality rate	$\delta^* = 0$
$w_t \equiv \frac{u(c_t)}{u'(c_t)c_t} \cdot \frac{c_t}{v_t}$	Value of a year of life relative to mortality price	$w^* = \rho$

for μ_{2t} and substitute this into the arbitrage equation for N_t) and rearranging yields:

$$\frac{\rho - g_{\mu 2t} - (\bar{n} - \delta_t) + \lambda \delta_t - \lambda g_{A_t} \bar{\delta} (1 - s_t) z_t}{\lambda \bar{\delta} g_{A_t} z_t \frac{1-s_t}{s_t} - \frac{\delta_t}{1-s_t} (\lambda \frac{1-s_t}{s_t} - 1)} = \frac{u(c_t)}{u'(c_t)c_t} \cdot (1 - s_t). \quad (39)$$

This is the key equation for determining the asymptotic behavior of z_t . In particular, along a balanced growth path, the right-hand-side goes to infinity for $\gamma > 1$. This requires that $z_t \rightarrow 0$ so that the denominator of the left side goes to zero (since $\delta_t \rightarrow 0$ as well). Moreover, with some effort, one can show that the denominator on the left side grows at the same rate as z_t along the balanced growth path, which implies that $g_z = -(\gamma - 1)g_c$ from the usual value-of-life argument used earlier.

The result for the growth rate of A_t and c_t follows by the same arguments as in the proof of Proposition 2. Finally, one can combine the first order conditions to solve for the allocation of research. QED.

B Appendix: Solving the Model Numerically

The transition dynamics of the equilibrium allocation can be studied as a system of four differential equations in four “state-like” variables that converge to constant values: ℓ_t , m_t , δ_t , and w_t . These variables, their meaning, and their steady-state values are displayed in Table 2.

Letting a “hat” denote a growth rate, the laws of motion for these state-like variables

are

$$\hat{\ell}_t = \frac{\rho + \delta_t - w_t + (\gamma - 1)\sigma m_t + \lambda(\bar{n} - \delta_t) - (1 - \phi)m_t}{1 + k_t \left(\frac{\lambda}{\beta - \ell_t} - \gamma \right) + \frac{\eta(z_t) + \theta(z_t)}{1 - \theta(z_t)}} \quad (40)$$

$$\hat{m}_t = - \left(\frac{\eta(z_t)}{1 - \theta(z_t)} + \frac{\lambda k_t}{\beta - \ell_t} \right) \hat{\ell}_t + \lambda(\bar{n} - \delta_t) - (1 - \phi)m_t \quad (41)$$

$$\hat{\delta}_t = \hat{m}_t + \left(k_t - \frac{\theta(z_t)}{1 - \theta(z_t)} \right) \hat{\ell}_t \quad (42)$$

$$\hat{w}_t = \hat{\delta}_t - \hat{\ell}_t + \left(\gamma - 1 + \frac{\delta_t}{w_t \ell_t} \right) (\sigma m_t + k_t \hat{\ell}_t) \quad (43)$$

where $\theta(z) \equiv z\Gamma'(z)/\Gamma(z) = \eta(z)(z/\Gamma(z) - 1)$ is the elasticity of the conditional expectation function and $k_t \equiv \ell_t/(1 + \beta - \ell_t)$. The only other variable that must be obtained in order to solve these differential equations is z_t , and it can be gotten as follows. First, $s_t = (\beta - \ell_t)/(1 + \beta - \ell_t)$. With s_t in hand, $\Gamma(z_t)$ can be recovered from the state-like variables using the mortality rate: $\delta_t = \bar{\delta}m_t(1 - s_t)\Gamma(z_t)$. Finally, $z_t = \beta\Gamma(z_t)/\ell_t$.

We solve the system of differential equations using “reverse shooting”; see Judd (1998, p. 355). That is, we start from the steady state, consider a small departure, and then run time backwards. For the results using the exponential distribution, we set $T = 600$; for the results using the Fréchet distribution, we set $T = 12150$.

An interesting feature of the numerical results is that $\hat{\ell}_t \approx 0$ holds even far away from the steady state. The reason is that $\lim_{z \rightarrow 0} \theta(z) = 1$: if z changes by a small percent starting close to zero, the conditional expectation changes by this same percentage. But this means that $\hat{\ell}_t \approx 0$ since there is a $1/(1 - \theta(z_t))$ term in the denominator. But since $z_t/\Gamma(z_t) = \beta/\ell_t$, if ℓ_t does not change by much, then z_t will not change by much either. QED.

References

- Acemoglu, Daron, Amy Finkelstein, and Matthew J. Notowidigdo, “Income and Health Spending: Evidence from Oil Price Shocks,” 2009. MIT working paper.
- and Simon Johnson, “Disease and development: the Effect of Life Expectancy on Economic Growth,” *Journal of Political Economy*, 2007, 115 (6), 925–985.
- , Philippe Aghion, Leonardo Bursztyn, and David Hemous, “The Environment and Directed Technological Change,” 2009. MIT working paper.

Aghion, Philippe and Peter Howitt, *Endogenous Growth Theory*, Cambridge, MA: MIT Press, 1998.

Ashenfelter, Orley and Michael Greenstone, "Using Mandated Speed Limits to Measure the Value of a Statistical Life," *Journal of Political Economy*, February 2004, 112 (1), S226–S267. Part 2.

Attanasio, Orazio P. and Guglielmo Weber, "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 1995, 103 (6), 1121–1157.

Barsky, R.B., F.T. Juster, M.S. Kimball, and M.D. Shapiro, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study*," *Quarterly Journal of Economics*, 1997, 112 (2), 537–579.

Becker, Gary S., Tomas J. Philipson, and Rodrigo R. Soares, "The Quantity and Quality of Life and the Evolution of World Inequality," *American Economic Review*, March 2005, 95 (1), 277–291.

Bostrom, Nick, "Existential Risks," *Journal of Evolution and Technology*, March 2002, 9.

Brock, William and M. Scott Taylor, "Economic Growth Theory and the Environment: A Review of Theory and Empirics," *Handbook of Economic Growth*, 2005, pp. 1749–1821.

Casey, Nicholas and Nicholas Zamiska, "Chinese Factory Is Identified In Tainted-Toy Recall Mattel's Reluctance Was Seen by Critics As Regulatory Flaw," *The Wall Street Journal Online*, August 8 2007.

Chetty, Raj, "A New Method of Estimating Risk Aversion," *American Economic Review*, December 2006, 96 (5), 1821–1834.

Costa, Dora L. and Matthew E. Kahn, "Changes in the Value of Life, 1940–1980," *Journal of Risk and Uncertainty*, September 2004, 29 (2), 159–180.

Fogel, Robert W., "Economic Growth, Population Theory, and Physiology: The Bearing of Long-Term Processes on the Making of Economic Policy," *American Economic Review*, June 1994, 84 (3), 369–395.

Gerdtham, Ulf-G. and Bengt Jonsson, "International Comparisons of Health Expenditure: Theory, Data and Econometric Analysis," in Anthony J. Culyer and Joseph P. Newhouse, eds., *Handbook of Health Economics*, North Holland, 2000.

- Gradus, Raymond and Sjak Smulders, "The Trade-off between Environmental Care and Long-Term Growth Pollution in Three Prototype Growth Models," *Journal of Economics*, 1993, 58 (1), 25–51.
- Gruber, Jonathan, "A Tax-Based Estimate of the Elasticity of Intertemporal Substitution," 2006. NBER Working Paper 11945.
- Güvenen, Fatih, "Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective," *Journal of Monetary Economics*, 2006, 53 (7), 1451–1472.
- Hall, R.E., "Reconciling Cyclical Movements in the Marginal Value of Time and the Marginal Product of Labor," *Journal of Political Economy*, 2009, 117 (2), 281–323.
- Hall, Robert E., "Intertemporal Substitution in Consumption," *Journal of Political Economy*, April 1988, 96 (2), 339–357.
- and Charles I. Jones, "The Value of Life and the Rise in Health Spending," *Quarterly Journal of Economics*, February 2007, 122 (1), 39–72.
- Hammit, James K., Jin-Tan Liu, and Jin-Long Liu, "Survival is a Luxury Good: The Increasing Value of a Statistical Life," July 2000. Harvard University mimeo.
- Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy*, August 1995, 103 (4), 759–784.
- Joy, Bill, "Why the World Doesn't Need Us," *Wired Magazine*, April 2000, 8 (4).
- Judd, Kenneth L., *Numerical Methods in Economics*, Cambridge, MA: MIT Press, 1998.
- Kremer, Michael, "Patent Buyouts: A Mechanism for Encouraging Innovation," *Quarterly Journal of Economics*, November 1998, 113, 1137–1167.
- Lucas, Deborah, "Asset Pricing with Undiversifiable Risk and Short Sales Constraints: Deepening the Equity Premium Puzzle," *Journal of Monetary Economics*, 1994, 34 (3), 325–342.
- Mundy, Alicia, "FDA Withholds List of Chinese Heparin Suppliers From Probe," *The Wall Street Journal*, May 10 2008, p. A4.
- Murphy, Kevin M. and Robert Topel, "The Value of Health and Longevity," June 2005. NBER Working Paper 11405.
- Newhouse, Joseph P., "Medical-Care Expenditure: A Cross-National Survey," *Journal of Human Resources*, 1977, 12 (1), 115–125.

- Overbye, Dennis, "Gauging a Collider's Odds of Creating a Black Hole," *New York Times*, April 2008.
- Posner, Richard A., *Catastrophe: Risk and Response*, Oxford University Press, 2004.
- Rees, Martin, *Our Final Century*, London: William Heinemann, 2003.
- Reinhardt, Uwe E., Peter S. Hussey, and Gerard F. Anderson, "Cross-National Comparisons of Health Systems Using OECD Data, 1999," *Health Affairs*, 2002, 21 (3), 169–181.
- Romer, Paul M., "Endogenous Technological Change," *Journal of Political Economy*, October 1990, 98 (5), S71–S102.
- Rosen, Sherwin, "The Value of Changes in Life Expectancy," *Journal of Risk and Uncertainty*, 1988, 1, 285–304.
- Solow, Robert M., "Intergenerational Equity and Exhaustible Resources," *Review of Economic Studies*, 1974, pp. 29–45.
- Spencer, Jane and Juliet Ye, "Toxic Factories Take Toll On China's Labor Force," *The Wall Street Journal*, January 15 2008, p. A1.
- Starr, Chauncey, "Social Benefit versus Technological Risk," *Science*, 1969, 165 (3899), 1232–1238.
- Stiglitz, Joseph, "Growth with Exhaustible Natural Resources: Efficient and Optimal Growth Paths," *Review of Economic Studies*, 1974, pp. 123–137.
- Stokey, Nancy L., "Are there Limits to Growth?," *International Economic Review*, 1998, pp. 1–31.
- Viscusi, W. Kip and Joseph E. Aldy, "The Value of a Statistical Life: A Critical Review of Market Estimates throughout the World," *Journal of Risk and Uncertainty*, 2003, 27, 5–76.
- Vissing-Jorgensen, Annette and Orazio P. Attanasio, "Stock-Market Participation, Intertemporal Substitution, and Risk-Aversion," *American Economic Association Papers and Proceedings*, May 2003, 93 (2), 383–391.
- Weil, David N., "Accounting for The Effect of Health on Economic Growth," *Quarterly Journal of Economics*, 2007, 122 (3), 1265–1306.