Growth Effects and the Cost of Business Cycles*

Gadi Barlevy

Department of Economics
Northwestern University
2003 Sheridan Road
Evanston, IL 60208

e-mail:barlevy@nwu.edu

February 1, 2001

Abstract

In his famous monograph, Lucas (1987) put forth an argument that the welfare gains from reducing the volatility of aggregate consumption are negligible. Subsequent work that has revisited Lucas' calculation has continued to find only small benefits from reducing the volatility of consumption, further reinforcing the perception that business cycles don't matter. This paper argues instead that fluctuations could affect the growth process, which could have much larger effects than consumption volatility. I present an argument for why stabilization could increase growth without a reduction in current consumption, which could imply substantial welfare effects as Lucas (1987) already observed in his calculation. Empirical evidence and calibration exercises suggest that the welfare effects can be quite substantial, possibly as much as two orders of magnitude greater than Lucas' original estimates.

*This paper was prepared for a conference to honor Assaf Razin. It is based on related work which was helped by the comments from Marco Bassetto, Larry Christiano, Elhanan Helpman, Sam Kortum, Robert Lucas, Alex Monge, Assaf Razin, Helene Rey, and various seminar participants.

1. Introduction

In a famous monograph, Lucas (1987) put forth a compelling argument that business cycles in the post-War U.S. involve only negligible welfare losses, thereby challenging the presumption that macroeconomic stabilization is highly desirable. His argument can be stated as follows. Consider a representative consumer with a conventional time-separable constant-relative risk aversion (CRRA) utility function

$$\sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma} - 1}{1-\gamma}$$

where $\gamma \geq 0$. Suppose this consumer is given a consumption stream C_t defined by

$$C_t = \lambda^t (1 + \varepsilon_t) C_0 \tag{1.1}$$

That is, consumption is initially equal to C_0 and then on average grows at a constant rate λ, so that after t periods average consumption is equal to $\lambda^t C_0$. Actual consumption each period is allowed to deviate from this average by a random factor $1 + \varepsilon_t$, where ε_t is an i.i.d. random variable with mean 0. Using data on per-capita consumption growth from the post-War period, we can estimate λ and the volatility of ε_t. To determine the costs of aggregate fluctuations, Lucas asked what fraction of initial consumption C_0 the consumer would be willing to sacrifice in order to avoid aggregate fluctuations, i.e. to replace ε_t in each period with its mean. For reasonable estimates of risk aversion γ, the answer turns out to be astonishingly small, less than one-tenth of one percent. By contrast, Lucas calculates the consumer would be willing to sacrifice a much larger fraction of initial consumption, about 20% when $\gamma = 1$, in order to increase the average growth rate λ by one percentage point. This leads him to conclude that growth is very important for welfare, but aggregate fluctuations are not.

Despite the flurry of papers that sought to dispute Lucas' insight, his essential claim that consumption risk at business cycle frequencies is associated with minor welfare costs appears to have survived intact. Various authors have modified key assumptions that are implicit in Lucas' calculation. These include calibrating consumption to individual income streams rather than to per-capita consumption and allowing for only self-insurance, as in Imrohoroglu (1989), Atkeson and Phalen (1994), and Krusell and Smith (1999); allowing for more persistent shocks to consumption, as in Obstfeld (1994); and allowing for non-expected utility classes of preferences that are arguably better at capturing attitude towards risk than the CRRA utility, as in Obstfeld

(1994), Pemberton (1996) and Dolmas (1998). However, all of these exercises have continued to yield costs of aggregate fluctuations that rarely exceed 1% for reasonable parameterizations. In order to convincingly argue that individuals would be substantially better off if the highs and lows of post-War business cycles were eliminated, then, it would appear that we must shift attention to costs of economic fluctuations that stem from something other than consumption risk per se.[1]

While consumption risk is the most obvious cost of aggregate fluctuations, Lucas' results that changes in growth have a substantial impact on welfare suggest that another potentially important cost of fluctuations could operate through their effect on the rate of economic growth. Specifically, suppose that the level of economic activity affects the incentives for agents to undertake investment that enhances future production, as is typical in many models of endogenous growth. Fluctuations in the level of economic activity could then conceivably affect the average rate at which consumption and income grow over time. Since Lucas demonstrated that even small changes in average growth rates have a dramatic impact on welfare, this channel could lead to a large cost of business cycles even for relatively modest growth effects. This paper seeks to formalize this intuition. That is, it examines whether eliminating fluctuations is likely to generate an increase in growth that would lead to a dramatic rise in welfare along the lines implied by Lucas' calculations. Although this question has already been tackled in Barlevy (2000), that paper focuses on a model in which R&D acts as the engine of growth. This paper instead focuses on a model of growth in which capital accumulation serves as the engine for growth. As such, it is closer to models that have been used by others in the literature to examine endogenous growth under uncertainty, and demonstrate more clearly why they fail to generate substantial costs of business cycles. In addition, since the model allows for long-run growth without external effects, it helps clarify certain issues that are obfuscated by the presence of externalities in the R&D model. Finally, while my previous paper argued that empirical evidence on R&D was consistent with growth effects that involve substantial gains from stabilization, this paper illustrates that similar conclusions can be reached from evidence on capital accumulation and investment.

[1] Detractors from this view include Campbell and Cochrane (1995) and Tallarini (1999). They argue that the large premium we observe on equity illustrates individuals are willing to pay substantial amounts to avoid fluctuations in consumption. However, one has to be careful drawing this conclusion from the equity premium. First, the equity premium could be due to market frictions rather than attitude towards risk. Second, as Alvarez and Jermann (1999) point out, the equity premium and the cost of consumption fluctuations are distinct concepts. They estimate a factor model for the marginal utility of consumption using financial data and put an upper bound on the costs of fluctuations at business cycle frequencies of 0.3%.

To preview my analysis, the model implies a consumption stream that is consistent with (1.1), but where the growth rate of consumption λ is no longer constant. Instead, the growth rate λ_t each period can be expressed as a composite function of the level of economic activity, i.e. $\lambda_t = \phi(i(\varepsilon_t))$, where $i(\cdot)$ denotes the investment rate for a given level of economic activity ε_t and $\phi(\cdot)$ denotes the growth rate of consumption for given level of investment. Removing fluctuations in ε_t will affect the average growth rate$E[\lambda] = E[\phi(i(\varepsilon_t))]$ in two different ways. First, depending on whether $i(\cdot)$ is concave or convex, stabilization can affect average growth by changing the average *level* of investment i. Second, depending on whether $\phi(\cdot)$ is concave or convex, stabilization can affect average growth by reducing the *volatility* of investment i. Previous work has assumed $\phi(\cdot)$ is linear, which allows for growth effects that only operate through the average level of investment. As is now well appreciated, this relationship is inherently ambiguous, i.e. $i(\cdot)$ can be either concave or convex, depending on the nature of preferences and technology. Regardless of whether stabilization leads to higher or lower investment, though, the implied welfare effects are likely to be far smaller in magnitude than those suggested by Lucas' calculation. This is because growth effects that are due to changes in average investment are inherently different from the growth effects Lucas studies. His welfare numbers capture the effect of increasing the growth rate of consumption starting from the same initial level of consumption, as illustrated by the solid and dashed lines in the first panel of Figure 1. As long as the agent does not discount future consumption too heavily, he will be vastly better off with a higher growth rate. However, an increase in growth that is due to a higher average investment leaves the agent with fewer resources to consume initially, since these must be allocated to investment. Thus, the resulting consumption path begins at a lower initial level, as illustrated in the bottom panel of Figure 1. Whether the agent prefers this new consumption stream depends on how the agent trades off present and future consumption, which explains why agents may voluntarily choose to lower the growth rate of consumption in response to the elimination of aggregate shocks. But more generally, since trading off present and future consumption is inherently different from increasing the growth rate of consumption from a given initial value, there is no reason to expect that the welfare gains from changing average investment in either direction will be anywhere near the magnitude of the huge gains from faster growth holding initial consumption fixed. Since previous authors focused exclusively on this type of growth effects, it is not surprising that they find only small welfare gains from growth effects.

By contrast, concavity in $\phi(\cdot)$ could generate an increase in growth without reducing average initial consumption. Concavity in $\phi(\cdot)$ implies diminishing returns to investment, i.e. an

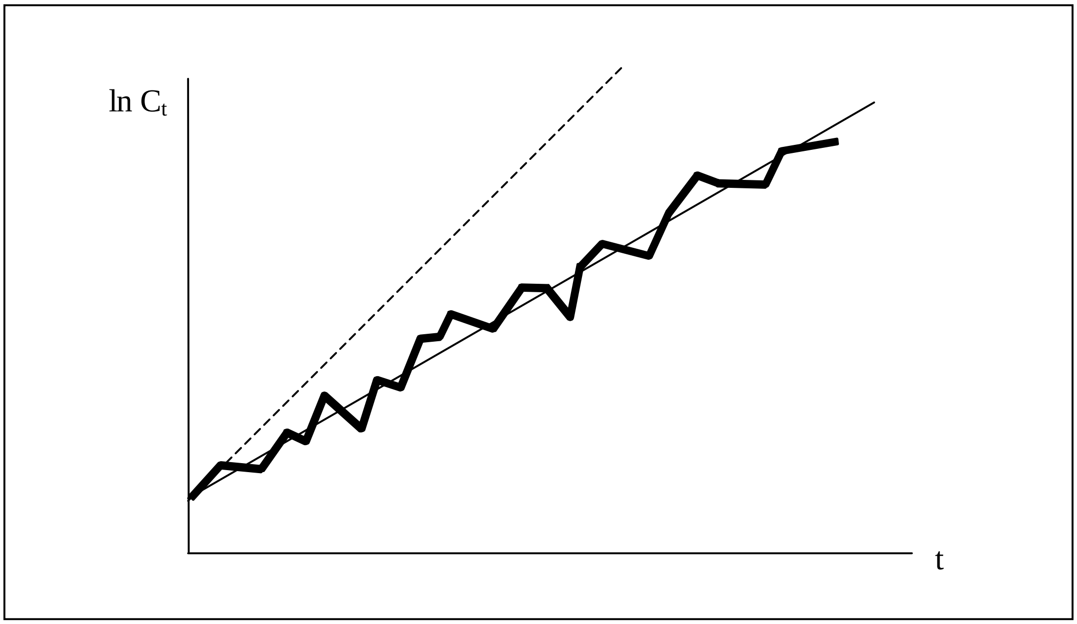

(a) Increased Growth from reduced volatility of investment

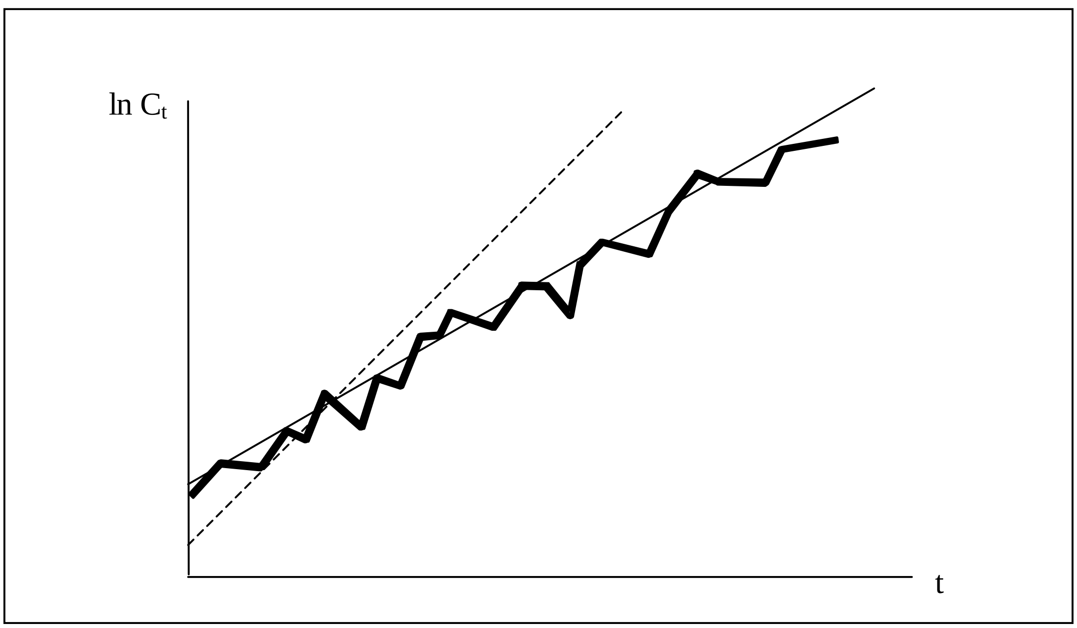

(b) Increased Growth from higher average investment

Figure 1: Consumption Paths under Endogenous Growth

additional unit of investment contributes less to the capital stock when investment is already high. With diminishing returns, keeping average investment unchanged but shifting resources from periods of high investment to those of low investment would lead to a faster accumulation rate of capital without requiring average investment – and thus average initial consumption – to fall. The welfare numbers Lucas computes for the value of additional growth therefore become relevant, which raises the possibility of substantial welfare costs from aggregate volatility even if growth effects are relatively modest. With sufficient concavity in the investment function $\phi(\cdot)$, a condition that is at the very least not contradicted in the data, one can obtain costs of aggregate fluctuations that are far larger than those that have been reported in previous work. As a rough benchmark, various pieces of evidence suggest that when we incorporate the growth effects of business cycles that occur because of diminishing returns to investment, individuals might be willing to sacrifice as much as 10% of their consumption to eliminate all fluctuations under reasonable specifications of utility.

The paper is organized as follows. Section 2 develops the model that can be used to distinguish between the role of the level of investment and the volatility of investment. Section 3 explores the quantitative implications of the growth effects. Section 4 concludes.

2. A Model of Endogenous Growth

To study the effects of economic fluctuations on growth and consequently on welfare, we need a model in which the growth rate is endogenous. Towards this end, I use a stochastic AK growth model. This specification has become a staple for modeling endogenous growth under uncertainty, and using it helps to relate the insights of this paper with previous findings. The first to analyze this model were Levhari and Srinivasan (1969), who used it to study savings decisions under uncertainty. They in turn solved an infinite horizon version of a problem that was originally studied by Phelps (1962). Leland (1974) subsequently reinterpreted this model in terms of long-run economic growth. Many authors have since used variations of this basic model to study environments with endogenous growth and aggregate uncertainty; recent examples include de Hek (1999) and Jones, Manuelli, and Stacchetti (1999).

The economy consists of a single representative agent who derives utility only from consumption. Time is discrete, and the agent discounts the future at a rate β. Following Lucas, I impose that the per-period utility function of the agent exhibits constant relative risk aversion, with a coefficient of risk aversion given by γ. Thus, given a consumption stream $\{C_t\}_{t=0}^{\infty}$, the utility of the agent is given by

$$\sum_{t=0}^{\infty} \beta^t U(C_t) = \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma} - 1}{1-\gamma}$$

Following Lucas, I will calibrate γ to equal 1. Since my analysis focuses on growth effects, the relevant interpretation for γ is as the inverse elasticity of intertemporal substitution. An estimate of 1 falls within the range of estimates for this measure in the empirical literature, e.g. Epstein and Zinn (1991).

The agent has direct access to a production technology. Although there is an equivalent decentralized representation of this economy in which agents own inputs which they sell to producers rather than having households engage directly in production, it is easier expositionally to assume the agent carries out production than to introduce factor markets. The only input in the economy is capital. Moreover, production is linear in capital, where the number of units that can be produced from a given unit of capital fluctuates stochastically over time. In particular,

$$Y_t = A_t K_t$$

where A_t is an i.i.d. random variable, and for simplicity I suppose there are only two possible realizations of A_t, i.e. $A_t \in \{A_0, A_1\}$, $A_1 > A_0$.[2] Although it is somewhat contrived to assume that capital is the sole factor of production, with additional notation one can modify the model to allow for labor and human capital, as illustrated by Manuelli, Jones, and Stacchetti (1999). They show that if the production function of output is homogeneous of degree 1 in physical and human capital, there exists a unique equilibrium in which both physical and human capital are accumulated at the same rate, making the two-factor model essentially equivalent to the one-factor model considered here.

Since capital is the sole factor of production, the trend path of output (and consequently consumption) depends on the evolution of the capital stock K_t. At date 0, the agent is endowed with some initial amount of capital K_0. Beyond this date, the level of capital depends on the

[2] Thus, the only source of fluctuations in this economy are productivity shocks. However, following Eaton (1981), we can reinterpret technology shocks in this model as government policy shocks. That is, suppose productivity were constant over time, i.e. $Y_t = AK_t$. Allowing for a government to collect an i.i.d random fraction τ_t of the income of the representative household to finance contemporaneous government purchases leaves the representative agent with an income of $Y_t = (1-\tau_t) AK_t \equiv A_t K_t$.

endogenous decisions of the agent. If an agent begins the period with K_t units, a fraction δ of the capital depreciates during the period, so that by the end of the period only $(1-\delta) K_t$ units remain. The agent can add to the stock of capital that is left by setting aside some of the output from the current period and converting it into capital. The technology for producing new capital from output is characterized by a function $\Phi(I_t, K_t)$ which depends on the existing stock of capital K_t and the amount of output set aside for producing capital I_t, where $\Phi(\cdot,\cdot)$ is assumed to be homogenous of degree 1 and increasing in its first argument. Given the homogeneity of degree 1, we can rewrite this production function as

$$\Phi(I_t, K_t) = \phi\left(\frac{I_t}{K_t}\right) K_t$$

where $\phi'(\cdot) > 0$. The stock of capital that is available for production at the beginning of period $t+1$ is thus given by

$$\begin{aligned} K_{t+1} &= \Phi(I_t, K_t) + (1-\delta) K_t \\ &= \left[\phi\left(\frac{I_t}{K_t}\right) + 1 - \delta\right] K_t \end{aligned} \tag{2.1}$$

By repeated substitution of the above equation to express the capital stock at date t as a function of the initial capital stock K_0:

$$K_t = \left[\prod_{s=0}^{t} \phi\left(\frac{I_s}{K_s}\right) + 1 - \delta\right] K_0$$

This specification corresponds to the framework originally studied by Levhari and Srinivasan (1969), except that they, and most subsequent authors who use this model, impose two additional restrictions. First, they impose full depreciation, i.e. $\delta = 1$. This assumption is not very appealing when we interpret K_t as capital, but it yields a closed-form solution for the agent's maximization problem. Second, they assume $\phi(\cdot)$ is linear, and more specifically is equal to the identity function. This assumption is natural under Levhari and Srinivasan's interpretation where K_t reflects wealth that is invested at a given interest rate. However, if we interpret K_t as capital, this assumption implies output can be converted one-for-one into capital. Following Uzawa (1969), it is also quite natural to consider concave specifications for $\phi(\cdot)$, since these allow for increasing marginal installation costs. In particular, concavity in $\phi(\cdot)$ assumes that at larger levels of investment, an increasing amount of capital is required just to install the new capital (or alternatively is eaten up in the process of putting the capital in place), a condition which appears to accord well with empirical evidence on investment.

To summarize, output in period t depends on the amount of capital available for production at the beginning of the period, i.e. $Y_t = A_t K_t$. Out of this output, the agent consumes an amount C_t, and uses the remainder $I_t = Y_t - C_t$ to invest in capital for the next period. It will prove convenient to define $c_t = \frac{C_t}{Y_t}$ as the fraction of output the agent consumes, and $i_t = \frac{I_t}{Y_t} = 1 - c_t$ as the fraction of output that the agent sets aside for investment. Using this notation, we can rewrite the consumption stream the agent chooses in a form that is reminiscent of Lucas' original specification:

$$\begin{aligned} C_t &= c_t A_t K_t \\ &= c_t A_t \left[\prod_{s=0}^{t} \phi(i_s A_s) + 1 - \delta \right] K_0 \\ &\equiv \left[\prod_{s=0}^{t} \lambda_s \right] (1 + \varepsilon_t) C_0 \end{aligned} \tag{2.2}$$

where $\lambda_s = \phi(i_s A_s) + 1 - \delta$ is the growth rate of the capital stock, $\varepsilon_t = \frac{c_t A_t}{c_0 A_0} - 1$ represents the level of consumption for a given level of capital, and C_0 is the initial level of consumption. Note that if the growth rate of capital λ_s were constant over time, the consumption stream the agent chooses would simplify to $\lambda^t (1 + \varepsilon_t) C_0$, which is exactly the form Lucas posited. The model instead yields a stochastic trend $\left[\prod_{s=0}^{t} \lambda_s \right] C_0$ in which the permanent component of consumption growth fluctuates over time.

Since the consumption stream above is endogenous, we can use this model to study how eliminating aggregate fluctuations should affect consumption choices and consequently the welfare of the representative agent, and thus revisit the question as to whether business cycles can involve substantial welfare losses. In the present context, the most natural notion for stabilization involves replacing the stochastic process for A_t with a constant process where productivity is equal to its average value $\overline{A} = \frac{1}{2}(A_0 + A_1)$. Before I proceed with the analysis, though, a few comments are in order. If we interpret A_t as exogenous productivity shocks, then unless government policy can directly affect technology, there is no way for government intervention to avoid the cost of aggregate fluctuations. Even though the government can replicate the effects of productivity shocks through taxes, the equilibrium is Pareto optimal, any such scheme that distorts the incentives of the agent to fool him into different consumption and savings choices will only result in lower welfare. Hence, the cost of business cycles represents a purely hypothetical one which cannot be avoided through active stabilization. By contrast, if we view A_t

as spurious shocks to fiscal policy as per footnote 2, we can interpret the cost of fluctuations as one that can be avoided by acting to eliminate arbitrary volatility in policy. However, it is important that the fluctuations that are eliminated are arbitrary rather than an optimal response to some other underlying shock. For example, one cannot use the logic of the model to argue governments should avoid seasonal fluctuations in spending such as snow removal in the winter. In that case, smoothing government spending over the year would eliminate volatility, but it would also prevent resources from being allocated to address underlying weather shocks that are seasonal in and of themselves. The harm caused by the latter may offset any of the benefits from eliminating fluctuations, negating any of the benefits from stabilization. Thus, to the extent the model implies a large cost of aggregate fluctuations, it can be used to justify stabilization policies only if underlying aggregate fluctuations are arbitrary and can be eliminated through government action – such as capricious policymaking or sunspots – rather than for any underlying source of macroeconomic volatility.

The consumption stream above is determined by the agent, who solves the maximization problem

$$V(K_0, A_0) \equiv \max_{C_t} E_0 \left[\sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma} \right] \tag{2.3}$$

subject to

$$K_{t+1} = \left[\phi \left(\frac{A_t K_t - C_t}{K_t} \right) + 1 - \delta \right] K_t$$

It is easier to analyze this problem using a recursive formulation to solve for the value function $V(K_t, A_t)$, namely

$$\begin{aligned} V(K_t, A_t) &= \max_{C_t} \left\{ \frac{C_t^{1-\gamma}}{1-\gamma} + \beta E\left[V(K_{t+1}, A_{t+1}) \right] \right\} \\ &= \max_{c_t} \left\{ \frac{(c_t A_t K_t)^{1-\gamma}}{1-\gamma} + \beta E\left[V\left(\left[\phi(i_t A_t) + 1 - \delta \right] K_t, A_{t+1} \right) \right] \right\} \end{aligned} \tag{2.4}$$

We guess that the value function $V(K_t, A_t)$ assumes the form

$$V(K_t, A_t) = \frac{a(A_t) K_t^{1-\gamma}}{1-\gamma}$$

in which case we can rewrite the above Bellman equation as

$$a(A_t) = \max_{c_t \in [0,1]} (c_t A_t)^{1-\gamma} + \beta E\left[a(A_{t+1}) \right] \left[\phi\left((1-c_t) A_t \right) + 1 - \delta \right]^{1-\gamma} \tag{2.5}$$

We can use this equation to solve for c_t (and thus $i_t = 1 - c_t$), allowing us to solve for the consumption stream the agent would choose when A_t fluctuates and when it is fixed. I begin analyzing this equation for the special case where $\phi(\cdot)$ is linear, i.e. $\phi(x) = x$, and then turn to the case where $\phi(\cdot)$ is concave.

2.1. Growth Effects through the Level of Investment

As noted earlier, previous authors have focused on the special case where $\phi(x) = x$. To allow for a closed form solution for c_t and i_t in this environment, I need to assume $\delta = 1$. The Bellman equation (2.4) above simplifies to

$$a(A_t) = \max_{c_t \in [0,1]} (c_t A_t)^{1-\gamma} + \beta E[a(A)]((1-c_t)A_t)^{1-\gamma}$$

$>$From the first order condition, we have

$$c_t^{-\gamma} = \beta E[a(A)](1-c_t)^{-\gamma}$$

which implies c_t is independent of A_t, i.e. the agent chooses to consume a constant fraction c of his income regardless of the current level of productivity. To solve for c, average the Bellman equation over the realizations of A to solve for $E[a(A)]$:

$$E[a(A)] = \frac{c^{1-\gamma} E\left[A^{1-\gamma}\right]}{1 - \beta(1-c)^{1-\gamma} E\left[A^{1-\gamma}\right]}$$

Substituting this back into the first order condition yields

$$c = 1 - \left(\beta E\left[A^{1-\gamma}\right]\right)^{\frac{1}{\gamma}} \tag{2.6}$$

$$i = \left(\beta E\left[A^{1-\gamma}\right]\right)^{\frac{1}{\gamma}} \tag{2.7}$$

To ensure an interior solution, we need to impose that $\beta E\left[A^{1-\gamma}\right] < 1$.

The implied growth rate of the capital stock is given by

$$\begin{aligned} \lambda_t &= \phi(iA_t) + 1 - \delta \\ &= iA_t \end{aligned}$$

It follows that the model admits growth effects only through changes in the average level of investment i. In particular, since i is constant even when A_t fluctuates, the average growth rate is given by

$$\begin{aligned} E[\lambda_t] &= E[iA_t] \\ &= iE[A_t] \\ &= i\overline{A} \end{aligned}$$

However, when aggregate productivity is stabilized so that $A_t = \overline{A}$ for all t, the implied growth rate is given by $i^*\overline{A}$, where i^* denotes the investment rate when aggregate productivity is constant (and equal to $\overline{A}$). Hence, the average growth rate of the capital stock $E[\lambda]$ will change if and only if $i^* \neq i$, i.e. if the average investment rate changes. Thus, when $\phi(\cdot)$ is linear, growth effects operate only if the elimination of aggregate shocks induces changes in the average level of investment. From the expression for i above, the fact that $E\left[A^{1-\gamma}\right] \gtreqless \overline{A}^{1-\gamma}$ for $\gamma \lesseqgtr 1$ implies that stabilizing productivity to its average value $\overline{A}$ will increase i if $\gamma < 1$ but decrease it if $\gamma > 1$. Thus, the growth rate of capital – and consequently consumption – can either increase or decrease when aggregate shocks are stabilized, depending on the underlying preferences of the agent. The observation that whether agents save more in the absence of fluctuations depends on whether the coefficient of relative risk aversion is less than 1 was first established by Phelps (1962), and is by now widely appreciated.

Given the above solutions for i and c, we can now formally study how stabilization affects the path of consumption (2.2) the agent chooses. Stabilization induces three distinct changes in the consumption stream. First, setting A_t equal to its average value serves to eliminate all variation around trend consumption. Formally, since A_t is the same for all t, the deviation from trend $\varepsilon_t = \dfrac{c_t A_t}{c_0 A_0} - 1 = 0$. Second, for a given investment rate i, stabilizing aggregate productivity A_t replaces the stochastic trend in consumption $\left[\prod_{s=0}^{t} iA_s\right] C_0$ with a deterministic trend $\left(i\overline{A}\right)^t C_0$ that have the same expectation. Thus, stabilization eliminates fluctuations *in* trend consumption as well as fluctuations *around* trend consumption. Finally, eliminating volatility affects the incentives of the agent to save, inducing a change in the investment rate from i to i^*. This will change the deterministic trend in consumption $\left(i\overline{A}\right)^t$ and thus the average growth rate of consumption.

Each of the three effects above raises the welfare of the agent; the first two effects eliminate temporary and persistent fluctuations in consumption, while the last effect reflects the gain from allowing the agent to adjust how the fraction of income he consumes in response to changes in the underlying environment he faces. However, as previous authors have shown, for reasonable specifications of utility, all three involve negligible changes in welfare when calibrated to U.S. consumption data. Lucas' original calculation showed that eliminating fluctuations in consumption around a trend will increase welfare by less than 0.1% even if we attribute all of the volatility in per-capita consumption to deviations from trend. Obstfeld (1994) evaluated the welfare gains of replacing a stochastic trend with a deterministic trend with the same mean. Even when he attributes all of the fluctuations in consumption to permanent shocks to trend consumption, he computes a welfare cost of no more than 0.3% when $\gamma \approx 1-2$. This just echoes the claim made in the Introduction that consumption risk at business cycle frequencies is of negligible importance. Moreover, growth effects associated with changes in average investment also yield only minimal welfare gains for reasonable parameter values. Epaulard and Pommeret (2000) and Matheron and Maury (2000) both compute these welfare gains, and find that for $\gamma \approx 1-2$, the gain from changes in the growth rate that are due to changes in average investment amounts to less than 0.1% of consumption. For $\gamma = 1$, this is not entirely surprising, since we just argued above that eliminating fluctuations has no effect on average investment and thus on growth when the coefficient of relative risk aversion is equal to unity. But even for plausible values of γ that are different from 1, the welfare gains from changes in i continue to be small.

The fact that changes in the growth rate involve only negligible welfare gains might seem at first to contradict the intuition in Lucas' original treatise in which he argued that even modest growth effects have huge welfare consequences. However, Lucas' result – that an agent would sacrifice up to 20% of his initial consumption to increase the growth rate by just one percentage point – pertains to an increase in growth starting from a given initial level of consumption C_0, as illustrated in the bottom panel of Figure 1. By contrast, growth effects that are due to changes in the average level of investment are conceptually quite different. In particular, if the growth rate of consumption rises because average investment i rises, it must be also true that the fraction of income the agent consumes, $c = 1 - i$, will fall. Hence, a higher average growth rate will be associated with a lower expected initial consumption $C_0 = c\overline{A}K_0$, as illustrated in the bottom panel of Figure 1. This fall in initial consumption wipes out most of the gains from faster growth that are inherent in Lucas' calculation. In order to generate substantial welfare gains along the magnitudes Lucas reports, one must argue that stabilization results in higher growth for a given initial consumption C_0, or alternatively, for a given level of average investment i.[3]

[3] It should be noted that welfare effects that are due to changes in the level of investment might be particularly small in the AK model, where the original growth rate (in the presence of shocks) is efficient. In a model with

2.2. Growth Effects through Investment Volatility

The last remark above hints at a role for concavity in the investment function $\phi(\cdot)$ in generating large welfare costs from aggregate fluctuations. In particular, suppose that we stabilize A and force the agent to leave average investment to capital ratio unchanged, i.e. we force the agent to choose a savings rate i^* such that

$$i^*\overline{A} = E[iA]$$

It follows that $c^* = 1 - i^*$ satisfies

$$\begin{aligned} c^*\overline{A} &= (1-i^*)\overline{A} \\ &= \overline{A} - E[iA] \\ & \quad E[(1-i)A] \\ &= E[cA] \end{aligned}$$

which insures average initial consumption $E[C_0] = E[cA]K_0$ remains unchanged by the elimination of fluctuations. However, as long as $i_0A_0 \neq i_1A_1$, concavity would imply that $E[\phi(iA)] < \phi(E[iA]) = \phi\left(i^*\overline{A}\right)$, so average growth is higher once shocks are eliminated even though average initial consumption remains unchanged. Intuitively, if the investment to capital ratio $\dfrac{I}{K}$ fluctuates over time, diminishing returns to investment imply that resources could be reallocated to achieve a higher average growth rate of the capital stock. By shifting some investment from periods of high investment, when the return to additional investment is fairly low since it involves high installation costs, to periods of low investment, when the return to additional investment is fairly high. Stabilization essentially achieves this by removing the incentives for the agent to change his investment to capital ratio over time, thus allowing the agent to attain a steeper consumption path without forcing him to give up initial consumption.

To establish the argument formally, I first need to argue that if A_t fluctuates over time, so will the equilibrium investment rate $\dfrac{I_t}{K_t} = i_tA_t$. However, I first need to impose some regularity conditions on $\phi(\cdot)$. Specifically, suppose $\phi(\cdot)$ is strictly concave, where $\lim_{x\to 0}\phi'(x) = \infty$ and $\lim_{x\to\infty}\phi'(x) = 0$. Under these assumptions, we can establish the following:

---

external effects where the equilibrium growth rate can be supoptimal, e.g. the R&D model examined in Barlevy (2000), a change in average investment may make agents significantly better or worse off by moving towards or away from the optimal growth rate. However, these welfare effects would still remain far smaller than those computed by Lucas, since shifting consumption from the present to the future or vice-versa affects welfare less than changing the slope of the consumption profile holding initial consumption fixed.

Proposition 1: Suppose $A_1 > A_0$. Then $\frac{I}{K}$ is increasing in A, i.e. the investment to capital ratio is strictly higher when aggregate productivity is higher.

Proof: From the first order condition of the Bellman equation, we have

$$(c_t A_t)^{-\gamma} = \beta E\left[a\left(A_{t+1}\right)\right]\left[\phi\left(\left(1-c_t\right)A_t\right)+1-\delta\right]^{-\gamma}\phi'\left(\left(1-c_t\right)A_t\right)$$

We can rearrange this equation to obtain

$$\phi'\left(i_t A_t\right) = \frac{1}{\beta E\left[a\left(A_{t+1}\right)\right]}\left(\frac{\phi\left(i_t A_t\right)}{A_t\left(1-i_t\right)}\right)^{\gamma} \tag{2.8}$$

Let $x = iA$. Then we can rewrite the first order condition as

$$\phi'\left(x\right) = k\left(\frac{\phi\left(x\right)+1-\delta}{A-x}\right)^{\gamma}$$

where k is a positive constant. Since $\phi'\left(x\right)$ is decreasing in x and $\left(\frac{\phi\left(x\right)}{A_t-x}\right)^{\gamma}$ is increasing in x, there exists at most one x which solves this equation. Existence then follows from the limit conditions $\lim_{x\to 0}\phi'\left(x\right)=\infty$ and $\lim_{x\to\infty}\phi'\left(x\right)=0$. If we rewrite this equilibrium condition as $f\left(x,A\right) \equiv \phi'\left(x\right) - \frac{1}{\beta E\left[a\left(A_{t+1}\right)\right]}\left(\frac{\phi\left(x\right)}{A-x}\right)^{\gamma} = 0$, the fact that $f_x < 0$ and $f_A > 0$ imply that as A rises, x must also rise to maintain $f\left(x,A\right)=0$, which establishes the claim follows. ■

The above Proposition establishes that if we stabilize A_t at its average value but force the agent to keep average investment to capital ratio unchanged, we can attain a higher average growth starting from the same initial consumption on average. Hence, Lucas' welfare calculations apply: if $\gamma = 1$, then for each one point increase in growth that is due to reduced volatility in investment, the agent would be willing to sacrifice approximately 20% of his initial consumption. This suggests that by stabilizing aggregate shocks in A_t and forcing the agent to keep the average investment to capital ratio constant, we can make the agent significantly better. Furthermore, by letting the agent choose a different investment rate $i \neq i^*$, we will only make the agent better off, since recall that the equilibrium of this economy is Pareto optimal. Thus, the growth effects that are due to investment volatility offer a lower bound on the welfare gains from eliminating aggregate fluctuations. More generally, if externalities imply the growth rate is inefficient, changes in average investment that are induced by the elimination of aggregate fluctuations could either increase or decrease welfare. However, as already noted above, the welfare implications of these changes are likely to be much smaller than those computed by Lucas. Thus, growth effects that stem from reduced volatility in investment have the potential to generate fairly substantial welfare gains from the elimination of aggregate volatility, much more than the estimates reported above. How large these welfare effects are likely to be depends on the extent of diminishing returns on investment, which must be determined based on the empirical evidence.

Before turning to the relevant empirical evidence, I close with a remark about the intuition behind my results. Given the potentially large gains from smoothing the investment to capital ratio, it seems natural to ask why the agent would willingly induce volatility in investment given that it has such a dramatic impact on welfare. To obtain some insights on this, note that the first order condition in (2.8) can be rewritten as

$$\phi'(i_t A_t) = \left[\frac{E\left[\beta V_K\left(K_{t+1}, A_{t+1}\right)\right]}{U'(C_t)}\right]^{-1}$$

This expression inside the brackets is just marginal q, i.e. it is the ratio of the marginal value of a unit of capital relative to the price of investment (which is just the price of output, here normalized to 1). Thus, the first-order condition can be rewritten as

$$\phi'(i_t A_t) = \frac{1}{q_t}$$

As long as q_t fluctuates over time, the agent will find it optimal to change his investment rate in response. From the proof of the proposition above, it follows that holding investment fixed, a positive productivity shock would tend to raise q_t, and with it the incentive to increase investment. Thus, the agent will find it optimal to increase investment in response to a positive productivity shock rather than keeping investment constant. In other words, the cost of aggregate fluctuations comes not from the fact that investment is volatile *per se*, but because the underlying environment makes it optimal for the agent to choose volatile investment. Eliminating fluctuations removes the incentives of the agents to change his investment rate over time, which makes him better off. But in the presence of aggregate fluctuations, forcing the agent to keep a constant investment to capital ratio would make him strictly worse off. Just because the agent chooses a volatile path for investment does not deny that he could be made significantly better off if the shocks that caused him to behave in this way were eliminated.

3. Quantitative Analysis

The preceding discussion demonstrates that in evaluating the effect of eliminating aggregate fluctuations on welfare, we should take into account both the reduction in consumption volatility as well as changes in the long-run average growth rate of consumption. As noted in the Introduction, previous authors have already established that consumption volatility appears to involve only negligible welfare losses. The question, then, is whether growth effects can lead to more significant welfare costs of aggregate fluctuations. One way to address this question is to adopt a reduced-form approach of estimating how the average growth rate depends on the underlying volatility present in the economy. The discussion suggests that in carrying out such an exercise, it is important to distinguish between changes in the growth rate $E[\lambda]$ holding average investment fixed from changes in $E[\lambda]$ that are stem from changes in the average level of investment, since the two are associated with very different welfare implications. In particular, an increase in growth for a given level of average investment is likely to generate much larger welfare effects, since it allows for more rapid consumption growth without a drop in average initial consumption. We can gain a sense of the magnitude of this growth effect by estimating average growth as a function of the volatility of investment, holding average investment fixed, i.e.

$$E[\lambda \mid \text{average } i] = f(\sigma_i) \tag{3.1}$$

Using the estimated $f(\cdot)$, we can predict the level of growth that would prevail if we eliminated volatility but maintained average investment at the same level. For the model developed above, the welfare gain from moving to this new consumption stream would establish a lower bound on the welfare gains from the elimination of aggregate fluctuations, since any additional changes in the investment rate that reallocate consumption between the present and the future only make the agent better off.

Fortunately, estimates of (3.1) already exist in the literature. In particular, Ramey and Ramey (1995) estimate a similar equation using cross-country data, both for a large set of countries as well as a sample that includes only OECD countries. However, they regress average growth on the volatility of output growth σ_λ rather than the volatility of the investment σ_i, which the model above identifies as the key factor in determining average growth. This is not much of a problem, though, since stabilization would eliminate volatility in both series, and so we can use either measure of volatility to infer the implied growth rate when the volatility of the underlying shock is set to 0.[4] Ramey and Ramey find that holding the average investment share

[4] In addition, Ramey and Ramey control for average investment using the investment share of output, i.e. the ratio of investment to output, rather than the investment to capital ratio as implied by the model. Given the difficulty of assembling reliable data on capital, the output share i is much easier to measure than the investment

of output fixed, a one percentage point reduction in the standard deviation of output growth is associated with an increased growth rate of 0.2%. Since the standard deviation of output growth in the U.S. is 2.5%, eliminating aggregate shocks altogether should increase the growth rate from 2.0% to 2.5%. Using the same dataset as Ramey and Ramey, I verified that estimating (3.1) using the standard deviation of the log investment share σ_i rather than σ_y yields the same predicted increase in average growth holding average investment fixed. Applying Lucas' estimate that an agent would sacrifice 20% of consumption for a 1 point increase in growth when $\gamma = 1$, the implied welfare gain from eliminating fluctuations in A_t is approximately 10% of initial consumption, two orders of magnitude greater than Lucas' original estimate. Thus, even though consumption risk represents only a minor burden of aggregate fluctuations, cross-country data seems to suggest a far more substantial burden on agents coming from the effect of aggregate fluctuations on long-run growth.

The above calculation reports the effects on growth if we forced average investment to remain unchanged. In general, though, stabilization could also affect the average level of investment, which would generate additional growth effects. In the model described above, allowing the agent to change the fraction of income he allocates to investment would only make him better off, and so any additional changes in growth would only make the agent even better off than if we restrict average investment to remain unchanged. But in models where the growth rate is inefficient, such changes in average investment could make the agent strictly worse off, in which case the above calculation will overstate the true cost of business cycles. This issue is to some extent ameliorated by another finding document by Ramey and Ramey, namely that the average level of i across countries appears to be uncorrelated with underlying economic volatility σ_y. Ramey and Ramey find this puzzling, relying on the intuition from previous work in which $\phi(\cdot)$ is assumed to be linear and where growth effects can only occur through changes in average investment. However, the model above delineates between growth effects that depend on the level of investment and those that depend on the volatility of investment, and the latter imply average growth could be higher in the absence of fluctuations even if average investment remains unchanged. The fact that virtually all of the growth effects Ramey and Ramey find do not appear to operate through investment is crucial for generating substantial welfare costs from aggregate fluctuations, since these are associated with more substantial welfare gains for the agent.

to capital ratio iA. However, as long as productivity shocks to A are relatively small, the investment share of GDP should serve as a useful proxy for I/K.

Although the cross-country data yields estimates of growth effects that are associated with a large cost of business cycles, we should proceed with some caution in reading this evidence. After all, differences in growth rates across countries could be due to a variety of differences across countries, a possibility that is underscored by the fact that some of Ramey and Ramey's estimates change dramatically with the addition of certain explanatory variables. While their point estimate for the coefficient on σ_λ tends to be clustered around 0.2%, their estimates range between 0.1 and 0.9%. In addition, although the negative relationship between the volatility of growth and the average rate of growth is statistically significant, it is not estimated with great precision. This suggests looking more deeply at the source of diminishing returns to investment and gauge whether it could plausibly generate an increase in growth from 2.0% to 2.5% as suggested by the cross-country evidence. In particular, the model suggests that the key factor in generating growth effects is curvature in the production of capital. Thus, as a complement to the reduced-form approach, we can ask whether the production function for investment goods $\phi(\cdot)$ exhibits the requisite concavity to generate substantial growth effects.

To address this question, consider first the case where $\phi(\cdot)$ is isoelastic, i.e.

$$\phi\left(\frac{I}{K}\right) = \left(\frac{I}{K}\right)^\phi$$

Diminishing and positive returns to investment require a value of $\phi \in (0,1)$. To determine what degree of diminishing returns is necessary to generate an increase in growth along the lines suggested by the cross-country evidence, I follow Barlevy (2000) in fitting consumption data to a two regime stochastic process that satisfies (2.2). Using post-war data, that paper estimates that the average growth rate λ is equal to 2.0%, with a standard deviation of 1.8%. For the isoelastic function to generate growth fluctuations that range between $2.0 \pm 1.8\%$, the investment rates $i_0 A_0$ and $i_1 A_1$ must satisfy

$$\begin{aligned} (i_0 A_0)^\phi + 1 - \delta &= 1.002 \\ (i_1 A_1)^\phi + 1 - \delta &= 1.038 \end{aligned}$$

Assuming a standard depreciation rate δ of 9% per year, we can rewrite the investment rates iA as functions of ϕ, i.e. $i_0 A_0 = (0.092)^{\frac{1}{\phi}}$ and $i_1 A_1 = (0.128)^{\frac{1}{\phi}}$. This allows us to compute the growth rate when investment is set equal to the average of $i_0 A_0$ and $i_1 A_1$ as a function of ϕ:

$$\phi\left(\frac{1}{2}\left((.092)^{\frac{1}{\phi}} + (.128)^{\frac{1}{\phi}}\right)\right) + 1 - \delta = \left(\frac{1}{2}\left((.092)^{\frac{1}{\phi}} + (.128)^{\frac{1}{\phi}}\right)\right)^\phi + .91$$

To generate an increase in average growth from 2.0% to 2.5% as is implied by the cross-country evidence, it is necessary that $\phi = 0.21$. To check whether this degree of concavity is reasonable, we can turn to empirical evidence that relates investment rates to q. In particular, for the isoelastic functional form, the first order condition (2.8) becomes

$$\ln\left(\frac{I}{K}\right) = \frac{1}{1-\phi}\ln q$$

Thus, for the isoelastic specification, an increase in growth of half a percentage requires an elasticity of investment with respect to q of $\frac{1}{1-0.21} = 1.26$. An even lower elasticity would imply even larger growth effects than those implied by the cross-country data. Turning to the literature on empirical investment equations, the estimated elasticity of investment with respect to q is typically lower than this estimate. For example, Eberly (1997) explicitly estimates an isoelastic specification using U.S. data, and obtains an elasticity of 1.22. Abel and Eberly (1995) also estimate an isoelastic specification using a different sample of firms, but obtain even smaller point estimates. Turning to other work on empirical investment equations, most researchers have tended to estimate the relationship between investment and q in levels rather than logs as would be implied by the isoelastic specification. However, we can still use these estimates to compute an elasticity at the sample mean. The reported elasticities of investment with respect to q are still lower than 1.26. For example, Abel's (1980) estimates for this the elasticity fall between 0.5 and 1.1 (p74). More recent work on investment regressions, such as Cummins, Hassett, Oliner (1998), who argue that their estimates find a much stronger response of investment to q than conventional studies, produce point estimates that cluster around unity. Thus, at a first pass, generating an increase in the growth rate of a half percentage point does not require implausibly large degrees of curvature in the investment function, at least when compared with available estimates of this curvature in the literature.

While the above discussion suggests modest growth effects that nonetheless yield large welfare costs of aggregate fluctuations are consistent with the empirical evidence on investment decisions, there are reasons to remain skeptical about estimates that are based on the elasticity of investment with respect to q. First, these estimates tend to be quite noisy, with some estimates well below unity. Such estimates are inconsistent with $\phi \in (0, 1)$, which is necessary for concavity for the particular specification above. However, very low estimates for this elasticity could simply be due measurement error; typically, instrumenting for q in these regressions yields higher elasticities that often exceed unity, while lower estimates are more prominent when measurement error is not properly accounted for. But even if we treat the higher point estimates as reliable, the isoelastic specification requires fairly volatile swings in q to induce agents to alter their investment decisions in a way that would lead to fluctuations in growth between 0.2% and 3.8%. From the first order condition above, we can compute the ratio of q_1 to q_0 that is necessary to induce the fluctuations in growth rate that we observe:

$$\frac{q_1}{q_0} = \frac{(i_0 A_0)^{\phi-1}}{(i_1 A_1)^{\phi-1}} = \left(\frac{.092}{.128}\right)^{\frac{.21-1}{.21}} = 3.46$$

This would be associated with a standard deviation for q of roughly 55%. This is somewhat more volatile that typical series for q. For example, Summers (1981) provides time series for both q and tax-adjusted q between 1931 and 1978. The standard deviations for the series he reports are 28% and 40%, respectively. Likewise, both investment and consumption turn out to be quite volatile under the isoelastic specification, with investment in peak periods almost four times as large as investment in low periods. A natural concern, then, is whether the mechanism of diminishing returns requires volatility in both investment and incentives for investment that exceed what we observe in the data.

To address this last question, I now turn to the question of how much of an increase in growth we can anticipate if we require that q exhibit reasonable fluctuations. Formally, I ask if for a given level of $\frac{q_1}{q_0}$ there exists a concave function $\phi(\cdot)$ that is consistent with the first-order condition and which implies an increase in growth of half a percentage point. This approach essentially places an upper bound on the amount of growth that is consistent with the observed volatility in q. While does not allow me to determine whether growth effects of the magnitude documented by Ramey and Ramey are *likely*, it does allow me to infer the more modest question of whether such growth effects are possible. A more precise assessment of the potential for cycles to affect average growth requires a better understanding of diminishing returns to investment at the aggregate level than can be gleaned from currently available literature.

The first order condition $\phi'(iA) = q^{-1}$ establishes a natural bound on the potential increase in growth as a function of the underlying volatility in q. The reason is that for any function $\phi'(\cdot)$, the first order conditions imply that

$$\frac{\phi'(i_0 A_0)}{\phi'(i_1 A_1)} = \frac{q_1}{q_0}$$

It follows that a restriction on the standard deviation of q, which effectively restricts the ratio

$\frac{q_1}{q_0}$, imposes limits on how much the first derivative $\phi'(\cdot)$ can change, and thus restricts $\phi''(\cdot)$.[5] This restriction is illustrated graphically in Figure 2. As the figure illustrates, we know that the growth rate $\lambda = \phi(\cdot) + 1 - \delta$ assumes a value of 1.002 when the investment is low, and a value of 1.038 when the investment is high. Moreover, we know that the slope of the growth rate with respect to investment at the two regimes is given by $\lambda'(i_0A_0) = \phi'(i_0A_0) = \frac{1}{q_0}$ and $\lambda'(i_1A_1) = \phi'(i_1A_1) = \frac{1}{q_1}$, respectively. Given particular values for the investment to capital ratio i_0A_0 and i_1A_1, it follows that any concave function $\phi(\cdot)$, the growth rate $\lambda(x) = \phi(x) + 1 - \delta$ which satisfies these properties is bounded above by the function

$$\widehat{\lambda}(x) = \min\left[1.002 + \frac{1}{q_0}(x - i_0A_0), 1.038 - \frac{1}{q_1}(i_1A_1 - x)\right] + 1 - \delta$$

The function $\widehat{\lambda}(x)$ therefore yields the maximum increase in growth that is possible given i_0A_0 and i_1A_1. Since $\phi(i_1A_1) - \phi(i_0A_0) = 0.036$ and the slope of $\phi(\cdot)$ for any concave function is confined to the interval $\left(\frac{1}{q_1}, \frac{1}{q_0}\right)$, the gap $\Delta iA = i_1A_1 - i_0A_0$ must lie in the interval $(.036q_0, .036q_1)$ for changes in investment to account observed fluctuations in output growth. Since q is approximately 1 on average in the data, this implies fairly reasonable fluctuations in investment to match up with observed fluctuations in the growth rate over time, particularly if q_0 is well below 1. As can be seen from Figure 2, the potential increase in growth depends on the value of ΔiA. Maximizing over all possible values of $\Delta iA \in (.036q_0, .036q_1)$,we obtain the following result:

Proposition 2: Stabilizing the average investment to capital ratio iA to a constant $i^*\overline{A} = E[iA]$ yields an average growth rate that is bounded above by

$$\frac{1}{q_1/q_0 + 1}\left[\phi(i_0A_0) + 1 - \delta\right] + \frac{q_1/q_0}{q_1/q_0 + 1}\left[\phi(i_1A_1) + 1 - \delta\right] \tag{3.2}$$

Proof: The maximization problem can be expressed as

$$\max_{\Delta iA}\left[\min\left(1.002 + \frac{1}{q_0}\left(\frac{\Delta iA}{2}\right), 1.038 - \frac{1}{q_1}\left(\frac{\Delta iA}{2}\right)\right)\right]$$

[5] Note that this approach is not robust to the presence of constraints on investment, e.g. financing constraints. If we account for those, $\phi'(iA)$ will not equal q^{-1} but the sum of this and the multiplier on the relevant constraint.

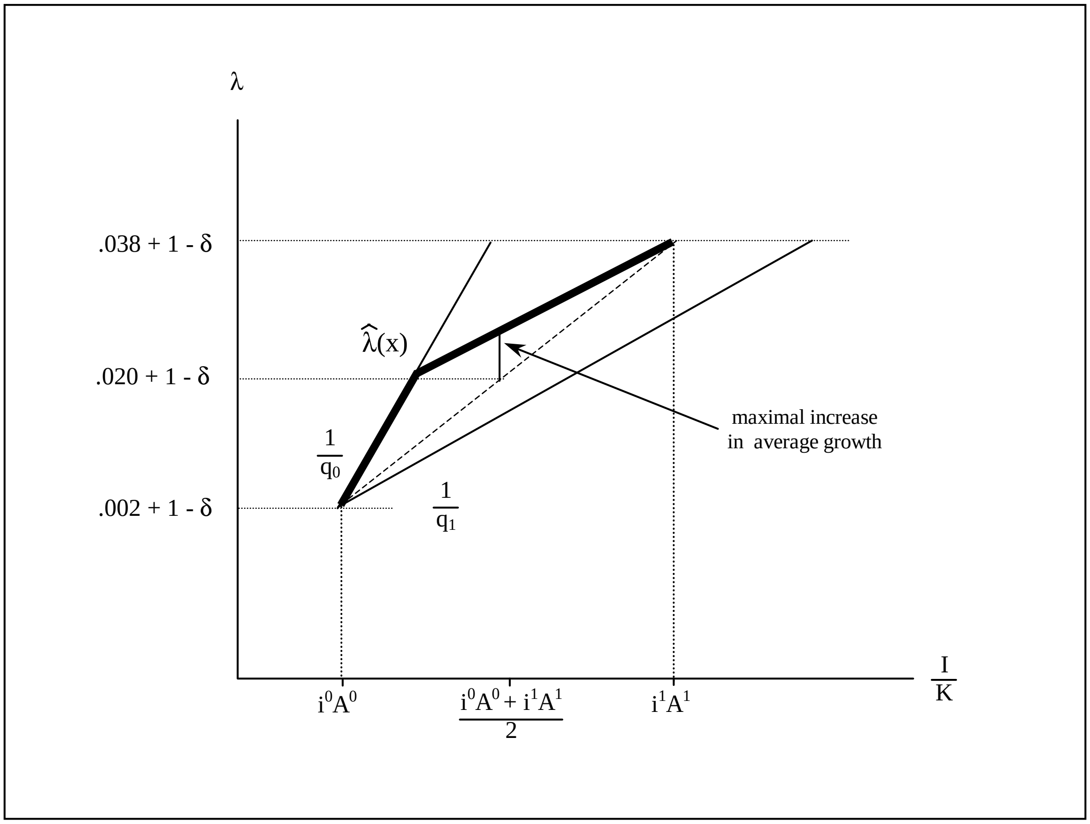

Figure 2: Bounds on Growth Effect

At the maximum, the two expressions must be equal, since if they are not equal, it would always be possible to increase this expression either by increasing ΔiA or decreasing ΔiA, depending on which expression is larger. Solving for ΔiA at the point of equality yields $\Delta iA = 0.072\frac{q_1 q_0}{q_1 + q_0}$, and substituting back in yields the desired result. ■

Note that the bound in Proposition 2 is sharp, i.e. there exists a function $\phi(\cdot)$ for which stabilization will yield the growth rate given in (3.2). This function is piecewise linear, and can be viewed as a linear approximation of the true function $\phi(\cdot)$ for a given value of ΔiA. Using the estimates of the volatility of q based on the data from Summers (1981), we obtain

$$\phi(E[iA]) + 1 - \delta \leq \frac{1}{2.78}1.002 + \frac{1.78}{2.78}1.038 = 1.025$$

using a standard deviation of 28% for q and

$$\frac{1}{3.33}1.002 + \frac{2.33}{3.33}1.038 = 1.027$$

for a standard deviation of 40%. Thus, the observed volatility of q does not preclude a growth rate of the magnitude suggested by the cross-country evidence. Of course, whether this growth rate is actually attained depends on the size of actual fluctuations in investment, i.e. ΔiA, and on the degree to which $\phi(\cdot)$ is well approximated by a piecewise linear function. Without a better understanding of diminishing returns to investment, the most we can say based on the available evidence is that notwithstanding the small costs of consumption risk computed by others, there remains a scope for fairly large welfare costs of economic fluctuations due to growth effects, with some evidence pointing to such large welfare effects as being quite likely.

4. Conclusion

This paper considers the potential cost of aggregate fluctuations that stem not from consumption volatility but from the effects of aggregate fluctuations on economic growth. In a sense, it closes a circle that began with Lucas (1987), who argued that growth matters for welfare while business cycles do not. By demonstrating that business cycles can affect the rate of economic growth, this paper argues that business cycles could matter precisely because they affect growth. It demonstrates that a plausible case can be made that for the U.S., stabilizing fluctuations can lead to an increase in the growth rate of half of a percentage point without affecting average initial consumption. This produces a cost of business cycles that is over 100 times larger than what Lucas computed based on the costs of consumption risk alone.

I close with a few caveats about the conclusions of the paper. First, given the inherent weakness in the relevant data, it is probably safest to interpret the argument in this paper as demonstrating that growth effects *could* generate a substantial cost of business cycles, not that it cannot. Since Lucas established a small lower bound on the costs of aggregate fluctuations, his argument appeared to close the book on the notion that stabilization could somehow achieve substantial welfare gains. The most compelling argument presented here establish a higher upper bound by allowing for growth effects. But without more concrete evidence on the nature of diminishing returns in the production of investment goods, it cannot offer a lower bound or a reliable estimate of how large these effects could be. Thus, its contribution lies in reopening the debate on whether business cycles are costly and whether stabilization is desirable, not in providing a final word in the debate. Second, even if growth effects are as substantial as this data implies, it is not obvious that one can recover these costs through implementation of aggregate stabilization policies. The model presented here introduces spurious volatility and assumes policy is able to shut it down and allow the economy to operate on a smooth track. Although the basic notion that stabilization can affect average long-run growth extends to other settings as well, if underlying fluctuations are not merely spurious, intervention may have offsetting costs that negate the benefits described here.

References

[1] Abel, Andrew, 1980. "Empirical Investment Equations: An Integrative Framework" in Carnegie-Rochester Conference Series on Public Policy, 12, p39-91.

[2] Abel, Andrew and Janice Eberly, 1995. "Investment and q with Fixed Costs: an Empirical Analysis" Mimeo, Wharton.

[3] Alvarez, Fernando and Urban Jermann, 1999. "Using Asset Prices to Measure the Cost of Business Cycles" Mimeo, Wharton.

[4] Atkeson, Andrew and Christopher Phelan "Reconsidering the Costs of Business Cycles with Incomplete Markets" *NBER Macroeconomics Annual*, edited by Stanley Fischer and Julio Rotemberg. Cambridge and London: MIT Press, p187-207.

[5] Barlevy, Gadi, 2000. "Evaluating the Costs of Business Cycles in Models of Endogenous Growth" Mimeo, Northwestern University.

[6] Campbell, John and John Cochrane, 1995. "By Force of Habit: a Consumption-Based Explanation of Aggregate Stock Market Behavior. NBER Working Paper #4995, January.

[7] Cummins, Jason, Kevin Hassett, and Stephen Oliner, 1999. "Investment Behavior, Observable Expectations, and Internal Funds" Federal Reserve System Discussion Paper 99/27.

[8] de Hek, Paul, 1999. "On Endogenous Growth Under Uncertainty" *International Economic Review*, 40(3), August, p727-44.

[9] Dolmas, Jim, 1998. "Risk Preferences and the Welfare Cost of Business Cycles" *Review of Economic Dynamics*, 1(3), July, p646-76.

[10] Eaton, Jonathan, 1981. "Fiscal Policy, Inflation, and the Accumulation of Risky Capital" *Review of Economic Studies*, 48(3), July, p435-45.

[11] Eberly, Janice, 1997. "International Evidence on Investment and Fundamentals" European Economic Review, 41(6), June, p1055-78.

[12] Epaulard, Anne and Aude Pommeret, 2000. "Recursive Utility, Growth, and the Welfare Cost of Volatility" Mimeo, ENSAE.

[13] Epstein, Larry and Stanley Zin, 1991. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis" *Journal of Political Economy*, 99(2), April, p263-86.

[14] Imrohoroglu, Ayse, 1989. "Cost of Business Cycles with Indivisibilities and Liquidity Constraints" *Journal of Political Economy*, 97(6), December, p1364-83.

[15] Jones, Larry, Rodolfo Manuelli, and Ennio Stacchetti, 1999. "Technology (and Policy) Shocks in Models of Endogenous Growth" Mimeo, Northwestern University.

[16] Krusell, Per and Anthony Smith, 1999. "On the Welfare Effects of Eliminating Business Cycles" *Review of Economic Dynamics*, 2(1), January, p245-272.

[17] Leland, Hayne, 1974. "Optimal Growth in a Stochastic Environment" *Review of Economic Studies*, 41(1), January, p75-86.

[18] Levhari, David and T.N. Srinivasan, 1969. "Optimal Savings Under Uncertainty" *Review of Economic Studies*, 36(1), April, p153-63.

[19] Lucas, Robert, 1987. *Models of Business Cycles.* Oxford: Basil Blackwell.

[20] Matheron Julien and Tristan-Pierre Maury. "The Welfare Cost of Fluctuations in AK Growth Models" Mimeo, University of Paris.

[21] Obstfeld, Maurice, 1994. "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability" *European Economic Review*, 38(7), August, p1471-86.

[22] Pemberton, James, 1996. "Growth Trends, Cyclical Fluctuations, and Welfare with Non-Expected Utility Preferences" *Economics Letters*, 50(3), March, p387-92.

[23] Phelps, Edmund, 1962. "The Accumulation of Risky Capital: a Sequential Utility Analysis" *Econometrica*, 30(4), October, p729-743.

[24] Ramey, Garey and Valerie Ramey, 1995. "Cross-Country Evidence on the Link Between Volatility and Growth" *American Economic Review*, 85(5), December, p1138-51.

[25] Summers, Larry, 1981. "Taxation and Corporate Investment: A q-Theory Approach" *Brookings Papers on Economic Activity*, 1, p67-121.

[26] Tallarini, Thomas, 1999. "Risk Sensitive Business Cycles" Mimeo, Carnegie Mellon University.

[27] Uzawa, Hirofumi, 1969. "Time Preference and the Penrose Effect in a Two-Class Model of Economic Growth" *Journal of Political Economy*, 77, p628-52.