

Evaluating the Costs of Business Cycles in Models of Endogenous Growth^a

Gadi Barlevy

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

e-mail: barlevy@nwu.edu

March 23, 2000

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.

^aThe author acknowledges comments from Marco Bassetti, Sam Kortum, Robert Lucas, Alex Monge, Helene Rey, and seminar participants at Northwestern, Purdue, U.C. Berkeley, and USC.

Introduction

In his famous monograph, Lucas (1987) put forth an elegant argument that the welfare effects of business cycles in the United States are negligible. The logic of his argument is as follows. Consider a representative consumer with a conventional time separable constant relative risk aversion (CRRA) utility function

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

where $\beta > 0$. This consumer is given a consumption stream C_t defined by

$$C_t = (1 + \bar{\Delta})^t (1 + \epsilon_t) C_0$$

where $\bar{\Delta}$ reflects average consumption growth over time and ϵ_t is an i.i.d. random variable with mean zero and variance $\frac{1}{\epsilon^2}$ that captures deviations of consumption from its trend growth rate. The parameters $\bar{\Delta}$ and $\frac{1}{\epsilon^2}$ can be estimated from data on log per capita consumption for the United States over the post-World War II period. To determine the costs of aggregate fluctuations, Lucas asks what fraction of initial consumption C_0 this consumer would be willing to sacrifice in order to stabilize his consumption stream, i.e. to replace ϵ_t with its mean $E(\epsilon_t) = 0$. For reasonable estimates of risk aversion γ , the answer turns out to be astonishingly small: less than 0.1%. By contrast, consumers would be willing to sacrifice a much larger fraction of initial consumption, about 20% when $\gamma = 1$, in order to increase the growth rate $\bar{\Delta}$ by one percentage point.

Various authors have since revisited Lucas' calculation, but for the most part have continued to find only small effects from stabilization, further reinforcing the perception that business cycles have only a negligible impact on welfare. One line of attack has focused on the shock process ϵ_t . The first to make this argument was Imrhoroglu (1989). She criticizes Lucas' calculation for its implicit assumption that agents have access to full insurance, which guarantees them the average level of consumption each period. If aggregate shocks do not affect agents uniformly and consumption insurance is imperfect, the volatility of each individual's consumption stream would be greater than the volatility of per capita consumption $\frac{1}{\epsilon^2}$. But when Imrhoroglu calibrates income streams to micro data, she finds a cost of aggregate fluctuations that is not much larger than Lucas' original estimate: only 0.3% when $\gamma = 1.5$. Moreover, subsequent papers by Atkeson and Phelan (1994) and Krusell and Smith (1999) argue that her notion of

stabilization removes too much idiosyncratic risk faced by agents along with aggregate risk. They find zero and even negative benefits to stabilization when aggregate risk is stabilized in a way that leaves the idiosyncratic risk faced by agents unaffected.¹ Another criticism focuses on the fact that the π_t process assumed in Lucas' calculation is not sufficiently persistent. If agents had to bear the consequences of aggregate shocks for extended periods, they would be more averse to fluctuations. The calculations when π_t follows a random walk with drift are reported in Obstfeld (1994a). For preference parameters that are similar to those used by Lucas, he finds costs of business cycles that are again about 0.3%.²

A second line of attack has focused on preferences, arguing CRRA utility is too restrictive and could underestimate the benefits of stabilization. Obstfeld (1994a) considers non-expected utility preferences that separate between risk-aversion and intertemporal substitution. But once again, by Obstfeld's own characterization, the welfare effect rises from a microscopic level to merely small, and the estimated cost remains below 1%.³ Pemberton (1996b) and Dumas (1998) consider a different class of non-expected utility preferences that exhibit 'first order' risk aversion, arguing that such preferences capture observed attitudes towards large and small bets that are inconsistent with expected utility. Some of their estimates are very large, but only under the assumption that π_t is close to a unit root. Even when π_t is autocorrelated with $\rho = .98$, Dumas reports that the costs of fluctuations do not exceed 1% for reasonable parameterizations. Extensions of Lucas' calculation therefore seem to confirm his original claim: the cost of consumption volatility from cycles calibrated to postwar data is trivial.

¹ In a parallel debate, Clark, Leslie, and Symons (1994) estimate the volatility of individual income from panel data and estimate a cost of cycles around 0.9% for the United Kingdom. This estimate is similarly criticized by Pemberton (1996a) for counting the reduction of purely idiosyncratic risk as a benefit of stabilization.

² Beaudry and Pagés (1999) combine the two criticisms by studying highly persistent idiosyncratic shocks. They obtain costs of at least 1%. This approach seems a promising way to generate more costly business cycles, especially since they introduce stabilization in a way that does not eliminate idiosyncratic risk. However, their model is vulnerable to Lucas' observation that if the costs of fluctuations stem from incomplete insurance, policymakers might be better off promoting insurance rather than stabilization.

³ Tallarini (1999) argues that Obstfeld's parameters are inconsistent with the equity premium. He generates a large cost of business cycles using a coefficient of risk aversion that is orders of magnitude greater than Obstfeld. A similar point is made in Campbell and Cochrane (1995) who use non-time-separable preferences to explain the equity premium, although Otkrok (1999) argues that a more "disciplined" calibration of non-separable preferences generates only small welfare effects. These arguments notwithstanding, the equity premium does not necessarily imply large costs of business cycles. First, the equity premium could be due to market frictions. Second, as Alvarez and Jermann (1999) point out, the equity premium and the cost of consumption fluctuations are distinct. They estimate a factor model for the marginal utility of consumption using financial data and put an upper bound on the costs of business cycles of 0.3%.

This paper pursues a different approach that could potentially generate much larger costs of business cycles than the previous work cited above. It is motivated by Lucas' original observation that even small changes in the growth rate \dot{A} have large implications for welfare. As such, if aggregate fluctuations somehow affected average long run growth, the costs of business cycles could potentially be much larger than previously estimated. This possibility is ruled out by Lucas' thought experiment; he treats the growth rate \dot{A} as an exogenous parameter that is unaffected by changes in Y . But standard models of endogenous growth predict the incentives of agents to engage in growth-enhancing activities depend on the level of economic activity, so that $\dot{A} = \dot{A}(Y)$. If this growth rate $\dot{A}(Y)$ is concave, stabilization will increase the long run growth rate of consumption. Since even small changes in \dot{A} have large welfare consequences, this could conceivably generate costs of business cycles that eclipse those described in previous work. The intuition behind this argument is illustrated graphically in Figure 1. Previous calculations measure the costs of business cycles by comparing utility from the observed volatile consumption stream represented by the heavy line with the utility from a consumption stream set to the trend rate of the original consumption stream, as represented by the thin solid line. By contrast, if stabilization affects average growth, the cost of business cycles will be the difference in utility from the original consumption path and from the consumption path represented by the dashed line. Since the latter provides more consumption at earlier dates, the implied costs of business cycles could be substantial as long as agents do not discount the future too heavily. According to this logic, the major cost of business cycles is not that consumption is volatile over time, as has been stressed in previous work, but that fluctuations impede the process of growth. While this idea has been occasionally discussed in previous work, it has yet to be incorporated into the original Lucas framework in a satisfactory way.⁴ This is unfortunate, because the notion that stabilization can increase growth and welfare is far from obvious. What assumptions imply that the growth rate \dot{A} is concave in the level of economic activity? How much additional growth should we expect from the elimination of cyclical fluctuations? Will such an increase in growth

⁴There is a rather diverse literature which argues that stabilization yields benefits other than reduced consumption volatility. For example, De Long and Summers (1988) and Ramey and Ramey (1991) argue stabilization increases the level of output, a point also raised in Chatterjee and Cochrane (1999). This is discounted by Romer (1996), who cites evidence that stabilization is unlikely to affect average output. More recent work studies the effects of stabilization on growth using models of factor accumulation with linear production technologies. This includes Aizenman and Marion (1993), Hopenhayn and Muniagurria (1996), de Heer (1999), and Jones, Manuelli, and Tacchetti (1999), as well as related work by Obstfeld (1994b) on the growth benefits of globalization. Except for Obstfeld, these papers are concerned with growth per se rather than welfare. Moreover, these models do not mesh well with the intuition behind Lucas' welfare calculations, since faster growth in this class of models does not necessarily imply higher welfare. Finally, Ramey and Ramey (1995) argue that volatility and growth are negatively related based on empirical evidence from a cross country evidence. But without a model, they cannot interpret the parameters they estimate structurally for welfare calculations.




Figure 1: Consumption Paths under Endogenous Growth

necessarily translate into higher welfare? If so, is it enough to generate larger welfare effects than in previous work?

To address these questions, I develop a model of endogenous growth in which shocks affect the level of economic activity. The model generates a reduced form growth rate $\dot{A} = \dot{A}(n, \sigma)$, where n is the amount of resources allocated to innovation activity and \dot{A} is the rate at which better technologies for producing goods are discovered for a given level of resources devoted to innovation. Setting σ to its expected value therefore affects average growth through two channels. First, stabilization can increase growth by increasing the level of innovation n , which will be the case if n is concave. Second, stabilization can increase growth by reducing the volatility of n , which will be the case if \dot{A} is concave. Previous work on growth under uncertainty has only considered the first effect, even though it is well known that the relationship between investment and uncertainty is ambiguous, i.e. n can be either concave or convex depending on somewhat arbitrary functional form assumptions. Even if stabilization increased average investment, the implications for welfare remain ambiguous, since shifting resources from production to innovation need not make individuals better off. The second channel, by contrast, is more equivocal: diminishing returns to innovation require that \dot{A} be concave, so stabilization will increase growth from a given level of resources n . Furthermore, the increase in growth will act to increase welfare.

The above discussion raises the question of how much growth could be generated from diminishing returns, and whether it can involve substantially larger welfare effects than those that can be traced to consumption volatility. Reduced form estimates on the relationship between average growth and the volatility of growth based on cross-country comparisons suggests removing aggregate fluctuations in the U.S. would raise per capita consumption growth from 2.0% to 2.5%. This increase implies a substantial cost of business cycles: when $\sigma = 1$, an individual would be willing to sacrifice 10% of his initial consumption to stabilize aggregate fluctuations and increase long run growth more than 100 times as much as Lucas originally estimated. Using empirical evidence on diminishing returns to R & D from micro studies confirms that even conservative estimates for diminishing returns in the production of new ideas can generate an increase in the growth rate of 0.5 percentage points. However, further scrutiny reveals that these estimates are based on functional forms that require implausible volatility in investment and equity values. Still, reasonable restrictions on the extent of volatility place an upper bound on how much stabilization can increase growth that is not much lower than

these original estimates – roughly 2.3%, which is valued at 5.7% of initial consumption. This suggests that the dramatic welfare implications based on reduced form estimates are consistent with plausible amounts of diminishing returns.

The paper is organized as follows. Section 1 develops the basic model of R & D with diminishing returns to innovation. Section 2 relates this model to previous work on stabilization and endogenous growth and distinguishes between effects on the level of investment and the volatility of investment. Section 3 estimates the size of the growth effects that can be attributed to diminishing returns and their implications for welfare. Section 4 concludes.

1. A Model of Diminishing Returns to Innovation

To study the effects of business cycles on growth, I need a model of endogenous growth that admits fluctuations in the level of economic activity. Models of technological innovation such as Grossman and Helpman (1991) and Aghion and Howitt (1992) satisfy this criterion. These models have two important features. First, they assume a monopolistically competitive framework which allows for different equilibrium levels of economic activity. Second, changes in economic activity affect the incentives of agents to innovate by changing the size of the market a monopolist can capture if he succeeds in developing a superior production technique. As such, aggregate fluctuations will have an impact on long run growth. I introduce fluctuations through shocks to the composition of government spending. This choice is motivated to minimize the number of auxiliary modelling assumptions that generate my results, not because I view the composition of government spending as an important source of aggregate fluctuations. Ultimately, the welfare effects I examine occur because stabilization reduces investment volatility. The same effects could arise in any environment where innovation fluctuates over time and the underlying production function for new ideas is concave; in fact, the next section illustrates similar effects in a model where the source of fluctuations can be interpreted as technology shocks. Since the source of aggregate fluctuations does not play a role in my subsequent attempts at quantifying the welfare effects of stabilization, the precise way in which I model it is unimportant.⁵ The model is presented in two parts. The...rst part describes its main features,

⁵ Although the source of aggregate fluctuations has no bearing on the question of whether business cycles can have...rst order welfare costs, it remains crucial for determining the feasibility and cost of stabilization, which ultimately determines whether stabilization is desirable. But the point of this paper is to challenge the notion that business cycles have no...rst order welfare effects, not to argue in favor of stabilization.

while the second characterizes its equilibrium.⁶

1.1. Setup

The economy contains three agents:

1. A representative agent, who consumes goods and supplies labor
2. A government, which taxes the agent and spends the revenue it collects
3. Entrepreneurs, who hire labor to produce goods and to develop new production methods

It will be helpful to preview the respective roles of the various parties. The agent plays an important but passive role: he consumes output, which provides us with a welfare measure to compare different economic environments. The government, as noted above, acts as the source of fluctuations, and can be used to introduce stabilization into the model. The role of entrepreneurs, aside from producing output, is to develop better production methods. This, in turn, determines the growth rate for the economy. The ultimate goal of the model is to gauge how business cycles affect the decisions of entrepreneurs to carry out innovation, and the consequences this has on the welfare of the representative agent.

I now describe the three agents in more detail. First, the representative agent has standard CRRA preferences over a consumption aggregate, i.e. his utility at a date t is given by

$$U(C_t) = \frac{C_t^{1-\gamma}}{1-\gamma}$$

where $\gamma > 0$ and C_t is a Cobb-Douglas aggregator over consumption goods c_{jt} for $j \in [0, 1]$:

$$C_t = \exp \left(\int_0^1 \ln c_{jt} dj \right)$$

Time is continuous, and the agent cares about his expected discounted utility

$$\int_0^\infty U(C_t) e^{-\rho t} dt$$

⁶Since writing this paper, I have become aware of a related model by Fatas (1998). Aside from differing in certain details from this model, he focuses on different questions.

where $\beta > 0$ is his discount rate. I choose CRRA utility despite its drawbacks which are discussed in the Introduction; these concern attitude towards risk, which is not important for my analysis. For this reason, β should be viewed as a measure of intertemporal substitution rather than a measure of risk-aversion.

The agent is endowed with L units of labor each instant. Leisure does not enter his utility function, so he supplies all of his labor whenever the wage is positive. The agent also owns all claims on the profits of entrepreneurs in this economy, positive and negative. Following Lucas, I initially abstract from savings; the agent must consume all of his disposable income, which is the sum of labor income and profits net of tax liabilities. However, this assumption is not important, as is demonstrated in Section 2 where I derive similar results for a model with savings. In the absence of saving the only choice the agent makes is how to allocate his income across different goods. Since C_t is a Cobb-Douglas aggregator, it is optimal to spend an equal amount on each good $j \in [0; 1]$. Integrating over all goods implies the spending rate on each good j is the same as disposable income,

$$p_j c_j = \int_0^1 p_j c_j dj = \frac{1}{2} + W - T$$

where $\frac{1}{2}$ denotes aggregate profits, W denotes labor income, and T denotes tax liabilities. Let $Y = \frac{1}{2} + W$ denote gross nominal income.

Next, I turn to the government sector. Government taxes the labor income of the agent and uses this revenue to finance various expenditures. The tax rate is a constant $\tau \in [0; 1]$. Since labor supply is inelastic, this tax is equivalent to a lump-sum tax. However, the fact that the tax is proportional to labor income is important, since it allows government revenue to grow as labor resources become more productive. Without this assumption, government share of output would eventually vanish. In what follows, I use labor as the numeraire good. The tax revenue in each instant is therefore $T = \tau L$. I assume the budget is balanced at each instant, i.e. the entire revenue T is always spent once collected. This revenue pays for goods the government consumes, and for workers the government hires. Hence,

$$T = G_t + N_t$$

where G_t denotes government spending on goods and N_t denotes the total wage bill of government employees. Given the normalization that labor is the numeraire, the latter is also the number of workers employed by the government. This notation might be a little confusing

at ...rst, especially since G_t is a nominal quantity while C_t is a real quantity. However, it is convenient to work with nominal quantities in solving the model, even though we ultimately care about real consumption. I denote nominal consumption by $P_t C_t = \sum_{j=0}^J p_{jt} C_{jt}$.

Government expenditures on goods are assumed to be allocated equally across the different commodities $j \in [0; 1]$, i.e. spending on each good j at date t is the same amount g regardless of the price of that good. It follows that aggregate government spending on goods is equal to g , since

$$G_t = \int_0^1 g dj = g$$

Building on Matsuyama (1995), I assume the government maintains a constant tax collection T but periodically shifts the composition of spending between goods and labor. Specifically, aggregate government expenditures G_t can assume one of two values, $G_1 > G_0$, and switch between these two values at a rate λ per unit of time. Denote the associated number of government workers employed in each regime by $N_0 = T - G_0$ and $N_1 = T - G_1$, respectively. Budget balance implies $N_0 > N_1$, i.e. more workers can work in the private sector when government shifts its spending towards goods.

Shifts in the composition of government spending generate aggregate fluctuations in the economy. This is because increased government spending on goods shifts labor resources out of the government sector and into the public sector. Because of monopolistic competition, workers generate more value working in production than working for the government: in the latter, they only generate labor income, while in the former they generate both labor income and profits. Hence, a shift in the composition of government spending towards goods raises the value of labor resources in this economy, which in turn affects the scale of production in the economy.

Finally, I turn to entrepreneurs. They produce the various goods $j \in [0; 1]$, as well as searching for better ways of producing these goods. In both capacities, they are driven by profit maximization. The nature of production in this economy is as follows. Each good j is associated with a number m_j that reflects the highest generation of technology available for producing that good. Put another way, m_j is the number of times the technology for producing good j has been improved upon since date $t = 0$. Each generation converts labor into output at a linear rate, but successive generations are more productive. Specifically, the m -th generation technology allows one unit of labor to produce λ^m units of output, where $\lambda > 1$ is the rate of

progress associated with each improvement. Technologies are protected by indefinite patents, so only the creator of the m -th generation technology can use it to produce goods. Still, the development of the m -th generation allows other entrepreneurs to develop the $m + 1$ -th generation, so discovering a new production technique helps both the innovator who discovers it and his competitors.

I first describe entrepreneurs in their capacity as producers. Each entrepreneur will produce using the highest generation he has access to. His only real decision is what price to charge for his good. Since both consumer and government spending on each good are independent of the price, demand for each good is unit elastic. Each monopolist will therefore want to set as high a price as possible; this way, he earns the same revenue but can produce fewer goods. But he cannot post a price that is too high; if his price exceeds the cost of producing a single unit of his next most efficient competitor, the latter will post a slightly lower price and steal away all of his business. Hence, producers quote a price equal to the marginal cost of their most efficient competitor, and the entrepreneur with the highest generation of technology will be the one supplying goods to the market. This implies that an entrepreneur with the m -th generation technology will set his price p_j to $\omega^{-(m-1)}$; his next most efficient competitor needs $\omega^{-(m-1)}$ workers to produce one unit of output, and each worker is paid a wage normalized to 1. At this price, the number of units the monopolist will sell is given by

$$\frac{Y - \omega T + G}{p_j} = \omega^{m-1} (Y - \omega T + G) \quad (1.1)$$

To produce each of these units, he needs to hire ω^{-m} units of labor. Hence, his total labor requirement is given by $\omega^{-1} (Y - \omega T + G)$. With labor as the numeraire, this is also the total cost of producing the quantity in (1.1). His profits will then equal his revenue net of costs, or

$$\begin{aligned} \pi_j &= p_j q_j - \frac{1}{\omega} (Y - \omega T + G) \\ &= \frac{\omega - 1}{\omega} (Y - \omega T + G) \end{aligned} \quad (1.2)$$

In what follows, I restrict attention to Markov equilibria in which nominal income Y depends only on the level of government spending G ; it is an open issue as to whether other equilibria also exist. In such an equilibrium, profits π_j will depend only on the level of government spending. Let π_{j0} denote profits when $G = G_0$ and π_{j1} when $G = G_1$.

Finally, I describe entrepreneurs in their capacity as innovators. I assume there are only two entrepreneurs in sector j at any point in time: one who owns the patent to the best available

technology and produces output, and the other who can work on developing a better technology for producing this good. Previous work has avoided having to make such assumptions by imposing constant returns to scale in research along with free entry. In this case, the number of potential entrants does not matter, and an incumbent monopolist will never choose to engage in research in equilibrium. To allow for diminishing returns to the research efforts of individual entrepreneurs, I need to impose these assumptions explicitly.⁷ If the entrepreneur in a given sector who engages in research employs n workers, he discovers the next generation technology $(m_j + 1)$ at a rate $\dot{A}(n)$ per unit of time. The function $\dot{A}(n)$ is strictly increasing and concave, and satisfies the usual boundary conditions $\lim_{n \rightarrow 0} \dot{A}'(n) = 1$ and $\lim_{n \rightarrow \infty} \dot{A}'(n) = 0$. Concavity in $\dot{A}(n)$ is associated with diminishing returns in research and development: the contribution of the marginal worker to the probability of success decreases with each additional worker. Denoting the value of a successful innovation by v , the innovator's problem is to choose n to maximize $\dot{A}(n)v - n$, so

$$\dot{A}'(n)v = 1 \quad (1.3)$$

1.2. Equilibrium

With the description of the economy complete, I can characterize its equilibrium. This is just a set of prices and quantities at each instant such that (1) agents choose prices and quantities optimally, (2) the government budget is balanced, and (3) output and labor markets clear. As noted before, I focus on Markov equilibria in which nominal variables vary only with the level of government spending. It will help to point out at this stage those features of an equilibrium which have a counterpart to Lucas' specification for consumption. First, an equilibrium is associated with a level of household consumption at each instant. Given a history up to date t , household consumption at date t takes on different values depending on the realization G_t . Consumption can therefore be above or below its expected value, which is analogous to the "term in Lucas' specification. Second, an equilibrium is associated with a level of employment in the innovation sector, which determines the rate of technological progress at that instant.

⁷An alternative approach is to allow for diminishing returns in aggregate R & D but for constant returns in the research efforts of individual entrepreneurs; in this case, no restrictions are necessary on the number of researchers or who engages in research. This approach is pursued, for example, in Stokey (1995). Under this scenario, condition (1.3) below is replaced by $n^{-1}\phi(n) v_t = 1$. One of the qualitative predictions of the model would be changed. In addition, for the constant elasticity function $\phi(n) = \Phi n^\xi$ Stokey considers, $\phi'(n)$ is proportional to $n^{-1}\phi(n)$, in which case the calibration exercise on the size of growth effects reported in Section 3 would also be unchanged.

This rate is analogous to the \dot{A} term in Lucas' specification.

The main step in solving for equilibrium is to solve for employment in the innovation sector n . To do this, I make use of the first order condition (1.3), which relates n to the value of a patent v . To solve for n , I express v in terms of n to obtain an equilibrium condition strictly in terms of n . Here, I use the approach of Lucas (1978). That is, suppose this economy had a market for patents. Since the representative agent must own all patents in equilibrium, the value of a patent must be such that the agent is indifferent between buying one and selling one. The expected utility from buying a patent is the marginal value of consumption one could afford with the profits it yields. This has an expected utility value of

$$E_t \int_0^{\infty} U'(C_{t+s}) \frac{dC_{t+s}}{dY_{t+s}} Y_{t+s} e^{-\rho(t+s)} ds$$

This has to be the same as the utility from selling the patent. In that case, the consumer can increase his income by v , which allows him to increase current consumption. The additional utility from this option is given by

$$U'(C_t) \frac{dC_t}{dY_t} v_t$$

Indifference between the two yields the value of a successful innovation:

$$v_t = E_t \int_0^{\infty} \frac{U'(C_{t+s})}{U'(C_t)} \frac{dC_{t+s}}{dC_t = dY_t} Y_{t+s} e^{-\rho(t+s)} ds$$

The following lemma establishes that v_t is well-defined when $\rho > 1$, and is linear in profits under the two regimes, Y_0 and Y_1 . Its proof, as well as those of all other claims in the paper, is delegated to the Appendix.

Lemma 1: Suppose $\rho > 1$. In a Markov stationary equilibrium, the value of a successful innovation is given by

$$v_t = \begin{cases} \frac{1}{\Gamma(n_0) \Gamma(n_1)^{1/2}} (\Gamma(n_0) Y_0 + \Gamma(n_1) Y_1) & \text{if } G_t = G_0 \\ \frac{1}{\Gamma(n_0) \Gamma(n_1)^{1/2}} (\Gamma(n_1) Y_1 + \Gamma(n_0) Y_0) & \text{if } G_t = G_1 \end{cases} \quad (1.4)$$

where

$$\Gamma(n) = 1 + \frac{1}{2} + \dot{A}(n) [1 + (\rho - 1) \ln_s]$$

If $n_0 = n_1$, the value of a patent v is higher in whichever regime offers higher profits π_i ; intuitively, a patent is more valuable if it pays out high dividends today rather than in the discounted future. But if $n_0 \neq n_1$, the patent could be more valuable when profits are relatively low, since v depends not only on the timing of payouts but also on the innovation rate $\dot{A}(n)$. This rate affects the value of a patent in two distinct ways. First, $\dot{A}(n)$ reflects the rate at which a patent is rendered obsolete by the arrival of a superior technology; a higher $\dot{A}(n)$ therefore reduces expected future profits, making the patent less valuable. Second, $\dot{A}(n)$ reflects the growth rate of the overall economy, which could make a patent either more or less valuable. With more rapid growth in productivity, real income grows more rapidly, raising demand for each good. This increases future expected profits, making patents more valuable. On the other hand, the more rapidly income grows, the more uneven is the anticipated future profile of consumption. This makes current consumption more valuable, reducing the value of the patent. Which of these effects dominates depends on how much the consumer is willing to substitute intertemporally. When $\gamma > 1$, so the agent is relatively less tolerant towards uneven consumption streams, the value of a patent unambiguously decreases with $\dot{A}(n)$.⁸

To express v solely in terms of n , I need to express profits π_i in terms of n . Here, I use the fact that aggregate profits π_i represent cumulative profits from all sectors net of the costs of innovation. That is,

$$\pi_i = \sum_{j=0}^{i-1} \pi_j g_i = \sum_{j=0}^{i-1} n_j g_i = \pi_i n$$

Combining this equation with (1.2) gives profits π_i as a function of n ,

$$\pi_i = (1 - \gamma) (1 - n_i / N)$$

which allows us to rewrite (1.3) as a system of two equations with n_0 and n_1 as the only endogenous variables:

$$\dot{A}'(n) \frac{1}{(n_1)^{1-\gamma} (n_0)^{\gamma-1/2}} [(1 - (n_j) + \gamma) L_i - (n_j) (n_i + N_i) i^{-1} (n_j + N_j)] = 1 \quad (1.5)$$

While obtaining a closed form solution for n_0 and n_1 is not generally possible, I can still characterize how employment, growth, profits, patent values, and gross income vary across the two regimes.

⁸ The focus on values of γ greater than 1 is also warranted by empirical evidence on the intertemporal elasticity of substitution. See, for example, Epstein and Zinn (1991).

Proposition 1: Suppose $\sigma > 1$. Then in equilibrium,

1. Employment in the innovation sector is increasing in G , i.e. $n_1 > n_0$. Hence, the growth rate $\dot{A}(n) \ln_s$ is increasing in G .
2. The value of patents is increasing in G , i.e. $v_1 > v_0$.
3. Nominal profits are increasing in G , i.e. $\pi_1 > \pi_0$.
4. Nominal income Y and consumption $P C$ could be either increasing or decreasing in G ; however, both are related to G in the same way.

Proposition 1 implies that a shift in government spending towards goods accelerates innovation, increases profits, and raises patent values. Intuitively, a shift in government spending towards goods increases demand, and consequently profits, for every operating monopolist. When profits are higher, there is greater incentive to capture the market, so n increases. This raises the growth rate of labor productivity, which is equal to $\dot{A}(n) \ln_s$, and which is also the growth rate of real output $\frac{\dot{Y}}{Y}$ and real consumption $\frac{\dot{C}}{C}$.

To relate Proposition 1 to more familiar notions of business cycles, it will be useful to return to the original Lucas setup. Recall that he posits the representative agent receives a consumption stream of the form

$$C_t = \prod_{s=0}^t (1 + \dot{A}_s) \prod_{t=0}^s (1 + \pi_t) C_0$$

where the growth rate \dot{A}_s is assumed constant. The term $\prod_{s=0}^t (1 + \dot{A}_s) C_0$ denotes expected consumption given cumulative growth up to period t , while actual consumption differs from this expectation by a stochastic term $\prod_{t=0}^s (1 + \pi_t)$ that reflects cyclical fluctuations. As anticipated by my previous remarks, this model admits an analogous continuous-time representation for the consumption aggregate

$$C_t = \exp \left(\int_0^t \dot{A}(n_s) \ln_s ds \right) \exp \left(\int_0^t \pi(G_s) ds \right) C_0 \quad (1.4)$$

The virtue of using labor as the numeraire good is that nominal quantities correspond to detrended real quantities. This is because aggregate real variables grow at the same rate as labor

productivity. Expressing the value of quantities in terms of units of labor effectively removes the growth term $\exp \int_0^t \dot{A}_s \ln s_s$ in (1.6), so that nominal consumption will be proportional to $1 + \pi_t$. This allows us to recover deviations from trend consumption using nominal consumption, and similarly for other real macroeconomic series in the model.

Armed with this observation, we can interpret Proposition 1 as a statement about the cyclical properties of various variables in this economy. It is natural to define the cycle in terms of real gross income, i.e. the economy is said to be in a boom if income is above its expected growth rate. This occurs when nominal gross income Y assumes its higher value. Proposition 1 implies consumption is procyclical, i.e. it attains its higher value when nominal income attains its higher value. Moreover, since the inequalities in Proposition 1 are strict, employment in the innovation sector, profits, equity values, and growth all fluctuate over time, although the direction of these series over the cycle is ambiguous. The fact that growth could be either procyclical or countercyclical should not be entirely surprising. Aghion and Saint Paul (1998) previously showed growth could be either procyclical or countercyclical, depending on the specification for R & D. This model makes a similar prediction, but for a different reason: with monopoly power in the production of goods, diverting workers from production to innovation reduces current income. If the increase in G diverts enough resources to innovation, gross income could fall even as each producer's profits rise. This is similar to effect in Helpman and Trajtenberg (1998) where the arrival of a more productive technology causes income to fall as resources are diverted to innovation. Empirically, profits, equity values, and R & D expenditures are all procyclical, which points to procyclical growth. This is also plausible given that innovation represents only a small share of the aggregate economy. Fortunately, though, whether growth is procyclical or countercyclical is irrelevant for how stabilization affects growth. The latter depends on whether $\bar{A}(\cdot)$ is concave in π , not on whether $\dot{A}(\cdot)$ is increasing or decreasing in π . It is to this issue that I now turn.

While it is useful heuristically to represent stabilization as favorable for growth whenever the growth rate \dot{A} is concave in the level of economic activity¹⁰, recall that the growth rate and the level of economic activity are determined simultaneously in equilibrium. Hence, the equilibrium growth rate is not technically a function of the scale of economic activity. To determine if stabilization increases long run growth, then, we must introduce a stabilization policy directly into the model and compare its equilibrium growth rate with the growth rate that prevails when fluctuations are allowed. Since government spending is the source of fluctuations, and

since its level is set by policymakers, it seems natural to define a stabilization policy as one which sets government spending on goods constant at the average of the two regimes, i.e. $G_t = \frac{1}{2} (G_0 + G_1) \triangleq \bar{G}$ for all t . Let $n(G)$ denote the number of workers employed in innovation when government spending is constant and equal to G , and $n^*(G)$ denote the number of workers employed in innovation when government spending fluctuates but is currently equal to G . Stabilization will increase average long run growth if and only if

$$\dot{A} \Big|_{n=\bar{n}} \Big|_{\bar{G}} > \frac{1}{2} [\dot{A} (n^*(G_0)) \Big|_{\bar{G}} + \dot{A} (n^*(G_1)) \Big|_{\bar{G}}] \quad (1.7)$$

The notation is meant to suggest that stabilizing government spending in this economy will increase growth if the growth rate $\dot{A} (n(\Phi)) \Big|_{\bar{G}}$ is “effectively concave” in government spending on goods. This is a modified notion of concavity, since (1.7) compares the average realized growth rate under two different stochastic processes for G_t , i.e. it compares the function $\dot{A} (n(\Phi))$ at \bar{G} with the average of a different function $\dot{A} (n^*(\Phi))$. The real concavity of $\dot{A} (\Phi)$ establishes the following sufficient (but not necessary) condition for stabilization to increase growth:

Proposition 2: (1.7) is satisfied if n is effectively concave in G , i.e. if

$$n \Big|_{\bar{G}} > \frac{1}{2} [n^*(G_0) + n^*(G_1)] \quad (1.8)$$

In other words, stabilization increases growth if it increases average investment n . Importantly, though, stabilization increases growth even if (1.8) holds with equality, i.e. when stabilization has no effect on average investment. This is because with diminishing returns to innovation, the fact that stabilization reduces the volatility of innovation activity acts to increase growth. Intuitively, stabilization shifts resources from peak periods of innovation, when their marginal product is low, to periods of less intensive innovation, when their marginal product is high. Of course, the resulting increase in growth could be offset if stabilization also reduces average employment in the innovation sector, i.e. if condition (1.8) is violated. The model does not allow us to pin down whether (1.8) is in fact satisfied, since this depends on the third derivative of $\dot{A} (\Phi)$. To see why, recall that n is determined by the first order condition $\dot{A}' (n)v = 1$. In order to establish whether n depends on v in a concave or convex manner, we differentiate this first order condition twice, which involves the third derivative \dot{A}''' whose sign and magnitude is ambiguous. This ambiguity reflects an indeterminacy in the relationship between investment and uncertainty that has already been well appreciated in the literature depending on the

production function and the investment cost function, investment can either increase or decrease in response to a mean preserving spread in productivity.⁹

To summarize, stabilization has an ambiguous effect on growth; but under diminishing returns, stabilization unambiguously increases growth for a given level of average investment. To better understand the role of the latter channel, I focus on the special case in which stabilization has no effect on the average investment, i.e. where n is effectively linear:

$$n^i \bar{G}^C = \frac{1}{2} [n^*(G_0) + n^*(G_1)] \quad (1.9)$$

Aside from theoretical interest, there is empirical justification for condition (1.9) in the recent work of Ramey and Ramey (1995), who study the relationship between growth and volatility for a sample of 92 countries and a subset of 240 ECD countries. They find increased volatility in output growth is associated with lower average output growth, but not with significant differences in investment rates. Ramey and Ramey find this puzzling given conventional wisdom that equates growth effects with changes in investment. However, this pattern is consistent with the notion that it is diminishing returns to innovation rather than changes in the level of innovation that lead stabilization to affect growth. Thus, their findings suggest we can ignore changes in average investment in studying the effects of stabilization on growth.¹⁰ Under this restriction, the model unambiguously predicts stabilization increases growth. More importantly, this increase in growth raises the welfare of the representative agent.

The last statement above underscores an important observation: even though both diminishing returns and changes in average investment allow stabilization to affect long run growth, the two channels have different implications for welfare. If stabilization raises growth by increasing average investment, the effect on welfare is ambiguous: shifting resources from production into innovation could make the agent better off or worse off, depending on his preferences between current and future consumption. But if stabilization increases growth by reducing

⁹ For example Hartman (1972) and Abel (1983) develop models where investment increases under uncertainty, while Bernanke (1983) and Dixit and Pindyck (1994) develop models where investment decreases under uncertainty. Caballero (1991) presents a unified framework in which investment can either increase or decrease with uncertainty, depending on the parameters of the production function and the investment cost function.

¹⁰ I should note that Ramey and Ramey use a broad measure of investment, while investment in this model only reflects expenditures on R & D. One should therefore interpret the implications of their results for this model with some caution. Still, the next section illustrates a factor accumulation model with similar implications and which admits a broader interpretation of investment.

the volatility of investment, the agent will be unambiguously better off. This follows from the welfare calculations reported in Lucas (1987). Recall that he finds substantial welfare gains from increasing the growth rate λ while holding average initial consumption C_0 fixed. The next lemma confirms that (1.9) insures setting G_t at its average value of \bar{G} also sets consumption equal to its average value

Lemma 2: Stabilizing government expenditures on goods also stabilizes consumption if and only if (1.9) holds, i.e.

$$P C'(\bar{G}) = \frac{1}{2} [P C'(G_0) + P C'(G_1)] \quad , \quad n'(\bar{G}) = \frac{1}{2} [n'(G_0) + n'(G_1)]$$

Diminishing returns therefore capture the intuition conveyed in Figure 1, i.e. stabilization results in a consumption stream that originates at the same initial level but grows more rapidly thereafter. By contrast, if stabilization increases growth by shifting resources into innovation, the resulting consumption stream would start at a lower average consumption level than is illustrated in the Figure. An increase in growth due to higher average investment will therefore not generate the same welfare gains as a similar increase in growth from diminishing returns, which is another reason why I focus on the latter channel. Hence, if stabilization affects growth only through diminishing returns, the growth effects of business cycles involve unambiguous welfare costs. The key question is whether these costs are larger than those computed in previous work. The answer to this question hinges on two issues. First, we need to assess by how much business cycles retard growth for plausible degrees of diminishing returns, an issue I address in Section 3. Second, although (1.9) rules out changes in average investment by assumption, we need to make sure that changes in average investment do not have additional implications for welfare that wipe out the costs due to diminishing returns. Judging by the evidence in Ramey and Ramey, these effects are likely to be small. In addition, the next section develops a model with different underlying assumptions that insure any changes in equilibrium investment, regardless of their impact on growth, increase welfare. In this case, the effects of stabilization holding the level of investment fixed necessarily provide a lower bound on the costs of business cycles.

A salient remark, note that Lemma 2 implies stabilizing government expenditures on goods also stabilizes consumption if and only if average employment is unchanged. This reflects a theme that has often been ignored in the literature, namely that there is no one correct notion of

stabilization. Lucas, for example, considers a policy setting consumption to its average value. This model considers a policy of setting government expenditures on goods to its average value. The two policies are only equivalent under strong linearity conditions, as represented by condition (1.9). Without linearity it would be impossible to simultaneously set several macroeconomic series equal to their average. But policies that stabilize different macroeconomic variables are conceptually quite different, and one has to be careful in comparing results based on different underlying notions of stabilization. To that end, imposing (1.9) is theoretically appealing because the model corresponds to exactly the same notion of stabilization that Lucas used in arguing business cycles involve relatively minor costs.

2. Alternative Models of Endogenous Growth

As noted in the Introduction, various authors have already investigated the implications of volatility in models of endogenous growth. However, even though this literature studies the effects of stabilization on growth, it has not discussed the welfare implications associated with such growth effects, nor has it examined how these models relate to Lucas' welfare calculations. This section revisits this literature and discusses its relationship to the model of the previous section. There are two reasons for doing this. First, the growth effects described in previous work reflect changes in average investment rather than diminishing returns. As such, these models do not necessarily imply faster growth is beneficial for the agent. Moreover, they continue to imply only negligible costs of business cycles for reasonable parameter values, which highlights the importance of diminishing returns. Second, previous authors have relied on models of factor accumulation with linear production technologies to generate endogenous growth rather than technological innovation. This section shows that this difference is irrelevant, and the results of the previous section could be obtained in models where factor accumulation is the engine of growth. Since the model differs in several dimensions from the one in the previous section, it clarifies that some of the simplifying assumptions in the previous section are not responsible for its results.

Consider the following model, taken from Jones, Manuelli, and Stacchetti (1999).¹¹ Time is

¹¹A side from Jones, Manuelli, and Stacchetti, various authors have used this model to study endogenous growth under uncertainty. It was first developed as a model of saving under uncertainty by Levhari and Srinivasan (1968). They solve an infinite horizon version of a problem first studied by Phelps (1962). Leland (1974) subsequently reinterpreted their model in terms of growth, and Eaton (1981) demonstrated how to introduce government

discrete. There are now two agents: the representative agent and the government. The agent has no labor resources, but is endowed with an initial amount K_0 of capital. He also has access to a linear production technology that converts a unit of capital into $1 + \lambda$ units of output, where $\lambda > 0$. Each period, the agent uses all of his capital to produce output. He then decides how much of this output to consume and how much to leave as capital for the subsequent production. This is in contrast with the previous section, where the agent was unable to save his income for future consumption. Production depreciates away all of the capital, so the only capital available for production in period $t+1$ is that which was set aside in period t . The agent has a conventional CRRA utility function over consumption in each period, and discounts the future at a rate β .

Just as before, I use government policy as the source of fluctuations in the model, although as previous authors have pointed out, policy shocks in this framework are isomorphic to technology shocks. That is, one can reinterpret the model as one where technology shocks drive fluctuations and tax policy is an instrument to stabilize against them. To remain consistent with the previous section, I will stick to the interpretation of shocks originating from government policy. Since there is no labor, government spending is devoted entirely to purchasing output. To finance its expenditures, the government taxes the income of the agent, so it takes a fraction τ_t of the output produced. As before, the budget is balanced every period, i.e. $G_t = \tau_t Y_t$. In contrast to the previous section, G_t now denotes real government spending. Once again, the share of government spending can assume only two values, which I assume is i.i.d. over time.¹² From budget balance, this implies the tax rate τ_t can take on two values $\tau_1 < \tau_0$, each equally likely. This leaves the agent with $(1 - \tau_t)(1 + \lambda)K_t = (1 + \lambda_t)K_t$ units of output each period.¹³ The problem facing the agent is given by

$$\max_{C_t} E_0 \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma}$$

subject to

$$K_{t+1} = (1 + \lambda_t)K_t - C_t$$

policy into this framework.

¹²By contrast, the previous section assumed G_t was Poisson, implying persistent shocks. Jones, Manuelli, and Stacchetti (1999) also examine a version of the model with persistent shocks.

¹³I restrict τ_0 and τ_1 so that $\beta E[(1 - \tau_t)(1 + \lambda)]^{1-\gamma} < 1$, which insures an interior solution.

where K_0 is given and $0 < C_t < (1 + s_t)K_t$. Applying a perturbation argument to the first order conditions of the agent, it can be shown that the solution to this problem is given by

$$\begin{aligned} C_t &= c(1 + s_t)K_t \\ K_{t+1} &= i(1 + s_t)K_t \end{aligned}$$

where i and c are constants, i is given by

$$i = \frac{\beta}{1 - E} \frac{h}{(1 + s_t)^{1-\gamma}} i^{\frac{1}{\gamma}} \quad (2.1)$$

and $c = 1 - i$. The growth rate of consumption is given by

$$\frac{C_{t+1}}{C_t} = i(1 + s_{t+1})i^{-1} \quad (2.2)$$

which in equilibrium is also the growth rate of output. The growth rate is proportional to the investment rate i , and thus depends on the behavior of the agent.

Just as in the previous section, we can gauge the effects of stabilization by setting the share of government spending on output, and thus the tax rate τ_t , to its expected value. This amounts to setting s_t equal to its average value. The only way this will affect the average growth rate in (2.2) is by affecting the propensity to save i . From (2.1), stabilization increases average growth if and only if $\sigma < 1$. Since estimates of σ are 1 or greater, stabilization either has no effect on growth or causes it to fall, in contrast with the previous model. But the fact that stabilization decreases growth does not imply the agent is worse off. After all, stabilization induces a lower growth rate because it leads the agent chooses to voluntarily consume a greater fraction of his output, i.e. to increase c and reduce $i = 1 - c$. Revealed preference implies that the agent must be better off under lower growth. More generally, because the equilibrium in this economy is Pareto optimal, any change in the equilibrium level of investment must make the agent better off, since the central planner cannot make the agent better off by reallocating resources into or out of the innovation sector. By comparison, the equilibrium in the previous section was not Pareto optimal given monopoly distortions and the inability of agents to borrow or lend, which is why changes in average investment could have reduced welfare in that framework.

Since stabilization can affect the savings decisions of the agent, this model also suggests growth effects which imply that Lucas' calculation underestimates the welfare under stabilization. In terms of Figure 1, the welfare of the household after stabilization does not correspond to

the consumption path set to the trend rate of the original consumption stream; instead, the consumption stream under stabilization could rotate depending to the optimizing decision of the agent. But the new consumption stream under stabilization will not be the same as the dashed line in Figure 1; for the consumption stream to have a steeper profile, i must be higher. But this implies that the initial level of consumption $C_0 = cK_0$ is lower, since $c = 1 - i$. Thus, the model above generates an additional cost of business cycles that reflects the effects of the cycle on the time profile of consumption, but which relies on changes in average investment and is therefore different from the channel of diminishing returns outlined in the previous section. Moreover, the model suggests this effect is likely to be quite small for reasonable estimates of σ . For the case where $\sigma = 1$, there are no growth effects and the model reduces to the one considered in Obstfeld (1994a), which involves a welfare gain from stabilization of only 0.3%. Using the estimated volatility in growth reported in the next section, I also computed the effects of stabilization for the case where $\sigma = 2$. For this value, the average growth rate declines from 2.00% to 1.98%. The implied welfare gains are once again almost identical to those reported by Obstfeld, which systematically fall well short of 1%.

While the above discussion suggests existing models of endogenous growth under uncertainty imply only small costs of business cycles, we could easily introduce diminishing returns in investment into these models and recover the potential for much larger effects. Following Uzawa (1963), suppose that the capital stock in period $t+1$ is given by

$$K_{t+1} = \phi(I_t; K_t)$$

where $I_t = (1 - \delta_t)K_t + C_t$ is the part of disposable income that is not consumed. The function $\phi(\cdot)$ is increasing and concave in I . This functional form implies that only ϕ percent of investment actually turns into capital, while the remainder of investment is eaten up in the process of installing the capital. I further assume the installation function exhibits constant returns to scale, so

$$\phi(I_t; K_t) = A \frac{I_t^\mu}{K_t^{1-\mu}}$$

The growth rate of the capital stock conditional on date t is given by

$$\frac{\dot{K}_t}{K_t} = A \frac{I_t^\mu}{K_t^{1-\mu}} - (1 - \delta_t) = A [i^*(\delta_t)(1 - \delta_t)]^\mu - (1 - \delta_t)$$

where i^* (Φ) is the fraction of after-tax income devoted to investment when s_t fluctuates over time. This growth rate is not the same as the conditional growth rate of consumption, but the two growth rates have the same unconditional expectation, and it is this expectation which is relevant for welfare. Stabilizing s_t continues to have ambiguous effects on expected growth which depend on ϕ and the third derivative \hat{A}''' . However, just as before, diminishing returns imply that stabilization will increase average growth for a given ratio of investment to capital. We can isolate the effects of diminishing returns by restricting attention to the case where the investment to capital ratio is effectively linear, i.e. where

$$i_{s_t} = \Phi i_{1+s_t} = \frac{\Phi}{2} [i^*(s_0)(1+s_0) + i^*(s_1)(1+s_1)]$$

The implications would then be identical as those of the model presented in Section 2. The only difference is that unlike that model, changes in average investment regardless of direction will necessarily make the agent better off. Thus, the welfare gains from diminishing returns in this model represent a lower bound on the costs of business cycles, and it will be enough to show diminishing returns on their own generate large welfare effects to establish that business cycles are costly.

3. Quantitative Analysis

The preceding discussion has argued that business cycles involve welfare losses that operate through endogenous growth channels, particularly because of diminishing returns in investment. But the most interesting question still remains: are these growth effects large enough to generate costs of business cycles that exceed those based only on consumption volatility? One way of addressing this question is to abstract from the source of diminishing returns and focus instead on the reduced form implications of the model for the relationship between volatility and average growth. This approach only requires information on the relationship between volatility and growth, either over time or from a cross section of countries. Such estimates are readily available from the work of Ramey and Ramey (1995). Using cross-country evidence, they found that controlling for the average investment share, 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%, this implies stabilization should increase the growth rate from 2.0% to 2.5%. Applying Lucas' estimate that an agent would sacrifice 20% of consumption for a 1 point increase in growth when $\phi = 1$ implies a welfare gain of 10% of initial consumption, two orders of magnitude greater than Lucas' original

estimate. Moreover, since Ramey and Ramey found no evidence of systematic changes in average investment rates associated with different levels of volatility, there should be no offsetting effects from changes in average investment that could lower welfare in a second best world.

While these estimates imply an enormous cost of business cycles, we should proceed with some caution before we accept them as evidence of costly business cycles. After all, differences in growth rates across countries could be due to unmeasured heterogeneity, 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 volatility 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 assess whether it could plausibly generate an increase in growth from 2.0% to 2.5%.

For stabilization to generate higher growth from a given level of innovation, two conditions are necessary: the amount of innovation n has to be volatile over the cycle, and the function $\hat{A}(n)$ that maps resources employed in innovation into growth has to be concave. Standard measures of R & D inputs compiled by the National Science Foundation, including industry expenditures on R & D and the number of workers employed as scientists or engineers, certainly exhibit volatility over time; log changes in both series have a standard deviation of about 5%. Moreover, work by Fatas (1994) shows that changes in R & D expenditures are related to the business cycle, although that relationship seems to have weakened in the 1990s. But these measures may fail to adequately capture the volatility in all of the inputs that enter into R & D. For example, volatility in innovation might be better reflected in the time spent by scientists and engineers on innovation projects rather than in employment numbers. This is partly addressed by using R & D expenditures, but this measure too is likely to be problematic. For instance, the physical implementation of new technologies is likely to be counted as investment rather than R & D, even though resources devoted to implementation should be counted as inputs into the innovation process. Hence, instead of calibrating the model to the volatility of data on R & D expenditures, I calibrate it to the volatility of consumption growth. In other words, rather than matching the volatility of inputs that go into producing innovation, I match the volatility of the outputs of innovation. Since the ultimate goal is to predict the growth rate in the absence of cycles, it seems only reasonable to calibrate the function $\hat{A}(n)$ to the growth rates we do

observe over the business cycle.

To measure volatility in growth in a way that is consistent with the theoretical framework above, I discretize (1.6) and estimate its underlying parameters from annual consumption data. The model implies consumption between period $t-1$ and t will change for two reasons. First, technological innovation improves productivity and allows more consumption goods to be produced from a given amount of resources. This growth rate is equal to $\hat{A}(n_t) \ln s_t$, which depends on the underlying growth regime in the current period. Second, regime changes trigger changes in the amount of resources employed in production and thus the level of consumption. The change in log consumption between two periods is therefore given by

$$\ln C_t = \hat{A}(n_t) \ln s_t + (\ln s_t - \ln s_{t-1}) \quad (3.1)$$

In fitting this stochastic process to actual consumption data, I maintain the implicit assumption of two regimes, which I denote by $s_t \in \{0, 1\}$.¹⁴ Without loss of generality, I designate $s_t = 0$ as the low productivity growth regime, i.e. $\hat{A}_0 < \hat{A}_1$. Let $\ln s_t - \ln s_{t-1} = \ln s_t - \ln s_{t-1}$ denote the difference in the level of consumption across the two regimes. With this notation, we can rewrite (3.1) as $\hat{A}_0(1 - s_t) + \hat{A}_1 s_t + (\ln s_t - \ln s_{t-1})$. I then estimate \hat{A}_0 , \hat{A}_1 , and $\ln s_t - \ln s_{t-1}$ from annual data by minimizing the mean square error over all possible sequences $\{s_t\}$:

$$\min_{\phi_0, \phi_1, \epsilon, \{s_t\}_{t=1951}^{1998}} \sum_{t=1951}^{1998} (\ln C_t - [\hat{A}_0(1 - s_t) + \hat{A}_1 s_t + (\ln s_t - \ln s_{t-1})])^2 \quad (3.2)$$

We can think of this minimization problem in two stages. First, for each sequence $\{s_t\}$, I look for the vector $(\hat{A}_0; \hat{A}_1; \ln s_t - \ln s_{t-1})$ which minimizes mean square error. I then look for the sequence of realizations s_t for which this minimum mean square error is lowest. The problem with solving (3.2) is that the number of possible sequences involves $2^{49} = 5.6 \times 10^{14}$ combinations. To get at the minimum, I follow a routine that is similar to simulated annealing. That is, I guess an initial sequence of s_t and then allow for small stochastic perturbations around the original sequence. If a lower value is achieved, I use the new minimum as my initial sequence. I repeated this procedure from several different initial conditions and also allowed for larger perturbations to check against local optima. The estimates at the minimum were $\hat{A}_0 = 0.02$, $\hat{A}_1 = 0.38$, and

¹⁴Maintaining the Markov structure is essential for identification. To see why, consider the general equation $\Delta \ln C_t = \phi_t + \Delta \varepsilon_t$. The problem is that since innovation depends on the scale of economic activity, $\phi(\varepsilon_t)$ and $\Delta \varepsilon_t$ are correlated. If we allow ε_t to follow a generalized AR process, it will be impossible to identify $\phi(\varepsilon_t)$ from $\Delta \varepsilon_t$, since both reduce to lag functions of ε_t . Even if we impose that ε_t is i.i.d., we could not separately identify $\rho(\phi_t, \Delta \varepsilon_t)$ and $\text{var}(\phi_t)$. However, we can still estimate a range for $\text{var}(\phi_t)$ by looking across all values of ρ . The standard deviation of ϕ_t lies between 0.6 and 3.0%, which includes my estimate of 1.8% below.

$\sigma = 0.07$. This implies that the growth rate in consumption because of technological progress fluctuates between 0.2% and 3.8% per year, for an average of 2% per year, just as Lucas finds. The fact that the point estimate for σ is negative suggests consumption growth is countercyclical, i.e. consumption grows more rapidly if it is below its expected level. This finding is somewhat disturbing since I previously argued procyclical growth is more consistent with the empirical evidence on procyclical profits and equity values. However, this is not important for calculating the effects of stabilization, and, as I discuss further below, it is not very robust.

The final step in computing the growth effects associated with stabilization is to specify an innovation rate $\dot{A}(n)$ and calibrate it so that it generates growth rates of 0.2% and 3.8%. Previous models of technological progress that allow for diminishing returns have tended to focus on the constant elasticity function $\dot{A}(n) = \phi n^\phi$; for example, this is the specification Stokey (1995) uses to study the implications for diminishing returns on equilibrium growth. The main advantage of the CES function is that it corresponds to the functional form empirical researchers have estimated from micro data. In particular, there is now a sizable literature that estimates production functions for new ideas by using R & D expenditures as the input and patents as output. This specification accords with a narrow interpretation of innovation in the model, since the model treats each patent as a successful innovation. More generally, not all patents contribute equally to growth, but estimates of diminishing returns from patent data should provide an upper bound on the true ϕ as long as the production of successful innovations exhibits diminishing returns in the number of patents. Griliches and Pakes (1984) and Hall, Griliches, and Hausman (1986) both estimate that the elasticity of patents with respect to total R & D expenditures at the firm level is 0.6, while the elasticity with respect to current R & D expenditures is only 0.3. These findings suggest fairly strong diminishing returns, but at the level of the firm; external spillovers at the aggregate level might cause these estimates to exaggerate the extent of diminishing returns in aggregate R & D. However, Kortum (1992) estimates R & D production functions at the industry level and finds even greater diminishing returns, with the estimates of the elasticity in the range between 0.1 and 0.25. Doubling the total resources devoted to R & D in an industry leads to a smaller increase in successful patenting than a doubling of resources devoted to R & D in a particular firm would, suggesting that diminishing returns carry over to more aggregate levels as well.

For this range of estimates for ϕ , we can estimate the growth rate $\dot{A}^{\frac{1}{1-\phi}}(n_0 + n_1) \ln s$ at average employment given that $\dot{A}(n_0) \ln s = 0.2$ and $\dot{A}(n_1) \ln s = 3.8$. Although this question might

appear to require more information in order to identify n_0 , n_1 , ϕ , and $\ln s$, the CES function allows us to determine the growth rate strictly as a function of η :

$$\phi^{\frac{1}{\eta}} (n_0 + n_1) = \frac{1}{2} (0.002)^{\frac{1}{\eta}} + (0.038)^{\frac{1}{\eta}}$$

Using the most conservative estimate of $\eta = 0.6$, setting R & D at its average level implies a growth rate of 2.52%, exactly the same as the reduced form estimates from Ramey and Ramey. Other estimates would imply even larger growth effects. Thus, estimates of diminishing returns from micro data are consistent with the increase in growth that we see using reduced form regressions on cross country comparisons.¹⁵

Although the consistency in estimates from two distinct empirical approaches might be reassuring at first, the CES specification above hides several discomforting features that should cast doubt on its relevance. The CES function can match the volatility in consumption growth rates only if R & D resources are incredibly volatile; the implied ratio of R & D levels between peak and trough levels is given by

$$\frac{n_1}{n_0} = \frac{0.038}{0.002} \eta^{\frac{1}{0.6}} = 135:29$$

which is implausibly large for even a broad definition of innovation activity. Likewise, from the first order condition (1.3), we know that the ratio of equity values is proportional to the marginal productivity of innovation, i.e.

$$\frac{v_1}{v_0} = \frac{A'(n_0)}{A'(n_1)} = \frac{0.038}{0.002} \eta^{\frac{0.4}{0.6}} = 7:12$$

The problem lies in the restrictive functional form imposed by the CES function; the only way to generate the low level of growth associated with the low productivity regime is to have very few resources devoted to innovation. This is inconsistent with the fact that consumption growth can be almost zero even as aggregate R & D expenditures remain high. Thus, the constant elasticity function provides a poor approximation for characterizing the aggregate production of innovation.

Unfortunately, the literature offers little guidance as to what a more reasonable innovation function might look like. Still, the model provides some restrictions on the shape of the function

¹⁵ It should be noted that the CES function $\phi(n) = \Phi n^\eta$ is not consistent with (1.9). Numerical simulation of the model suggested that average employment in innovation declines under stabilization, although the extent of the change depends on additional parameter values.

$\hat{A}(n)$ that allow us to gauge whether a function that implies more plausible volatility can generate an increase in growth that is consistent with the estimates in Ramsey and Ramsey. Recall that the first order condition (1.3) restricts the rate at which $\hat{A}'(n)$ changes between n_0 and n_1 , since

$$\frac{\hat{A}'(n_0)}{\hat{A}'(n_1)} = \frac{v_1}{v_0} \quad (3.3)$$

This restriction turns out to be useful in determining whether more plausible levels of diminishing returns can generate significant increases in growth. The reason for this is illustrated graphically in Figure 2. We know that the growth rate $\hat{A}(n) \ln s$ assumes a value of 0.2 at some employment level n_0 and a value of 3.8 at some other employment level n_1 . Suppose the slope of the growth rate with respect to n evaluated at n_1 , i.e. $\hat{A}'(n_1) \ln s$, is equal to some real number a . Consistency imposes the following limit on the possible values of a

$$\frac{v_0}{v_1} \frac{\hat{A}(n_1) - \hat{A}(n_0)}{n_1 - n_0} \ln s < a < \frac{\hat{A}(n_1) - \hat{A}(n_0)}{n_1 - n_0} \ln s \quad (3.4)$$

From (3.3), we know that at $n = n_0$, we have $\hat{A}'(n_0) \ln s = \frac{v_1}{v_0} a$. The two lines originating at n_0 and n_1 with these respective slopes form an upper envelope for any possible function $\hat{A}(n) \ln s$, as illustrated in the Figure. Maximizing $\hat{A}^{\frac{1}{2}}[n_0 + n_1] \ln s$ over all possible values of a we can derive an upper bound on how much additional growth could be generated in a model which generates reasonable volatility in equity values. Proposition 3 establishes this bound as a function of the ratio $\frac{v_1}{v_0}$:

Proposition 3: If $\hat{A}'(n_0) = \frac{v_1}{v_0} \hat{A}'(n_1)$, then

$$\hat{A}^{\frac{1}{2}}[n_0 + n_1] \ln s \leq \frac{1}{v_1 = v_0 + 1} \hat{A}(n_0) \ln s + \frac{v_1 = v_0}{v_1 = v_0 + 1} \hat{A}(n_1) \ln s \quad (3.5)$$

The bound in Proposition 3 is tight, since there exists a concave piecewise linear function $\hat{A}(n)$ for which the growth rate at average employment is exactly equal to the bound. Thus, Proposition 3 is a precise estimate of the maximum possible increase in growth that can be obtained through diminishing returns for a given level of volatility implied by $\frac{v_1}{v_0}$. At the same time, the restriction implied by (3.3) does not impose a lower bounds on how much growth diminishing returns could generate, since values of a near either of the endpoints in (3.4) imply the growth rate at average employment $\hat{A}^{\frac{1}{2}}[n_0 + n_1] \ln s$ is arbitrarily close to 2.0%.




Figure 2: Computing Upper Bounds on the Effect of Stabilization

Proposition 3 therefore tells us only whether the growth effects estimated in Ramey and Ramey could arise under plausible formulations of diminishing returns, not whether they in fact do.

In estimating the ratio $\frac{V_1}{V_0}$, a natural benchmark is stock market data which reflects the value of...rms over the cycle. Conventional estimates, such as Christiano and Fisher (1998), suggest the standard deviation of the S&P 500 is on the order of 10%. However, since new...rms tend to have more volatile equity values, and new...rms are listed in exchanges and included in stock indices such as the S&P only with a lag, this is likely to be an underestimate of true volatility in profit opportunities over the cycle. As an alternative measure, recall that v is a weighted discounted flow of profits across the two regimes. Hence, the ratio of profits in the two regimes serves as an upper bound on the ratio $\frac{V_1}{V_0}$. Data on after-tax corporate profits in the U.S. suggests a standard deviation for $\ln(\text{profits})$ of about 10 ; 15%.¹⁶ A standard deviation of 15% would imply that the ratio $\frac{V_1}{V_0} = \frac{1.15}{0.85} = 1.35$. If the standard deviation of v is allowed to be as large as 20%, $\frac{V_1}{V_0}$ will equal 1.5. Substituting this into (3.5) yields

$$\Delta \ln \left[\frac{1}{2} (n_0 + n_1) \right] \ln s \cdot \begin{matrix} \text{2:27} \\ \text{2:36} \end{matrix} \text{ if } \frac{V_1}{V_0} = \frac{1.15}{0.85} = 1.35 \quad \text{if } \frac{V_1}{V_0} = \frac{1.20}{0.80} = 1.50$$


Since this calculation is independent of n_0 and n_1 , the above estimates imply that diminishing returns that rely on reasonable volatility in both R & D and in equity values can still allow stabilization to increase growth from 2.0% to as much as 2.36%, well within the margin of error of the estimates provided by Ramey and Ramey. This implies an upper bound on the costs of business cycles of 0.27E20 = 5.4% of initial consumption when the standard deviation of v is set to 15%, and 7.2% when the standard deviation is equal to 20%, both of which are significantly larger than previous estimates for the costs of business cycles. While these bounds suggest that estimates in Ramey and Ramey are slightly overstated, they do not rule out significant growth effects as inconsistent with empirically plausible levels of diminishing returns.

I end my discussion with two remarks. First, although my estimates are identified by estimating a two state model, similar results obtain when I estimate a three state model that allows for


¹⁶ This covers the same period 1951 – 1998. This estimate accounts for detrending implied productivity growth using consumption growth. Detrending has only a minor impact on the estimate of the standard deviation, since $\Delta \ln(\text{profits})$ is far more volatile than $\Delta \ln(C)$.

three growth rates $\hat{A}_0, \hat{A}_1, \hat{A}_2$ and for two changes in levels $\ln c_2 - \ln c_1$ and $\ln c_1 - \ln c_0$. In estimating the three state model, I obtained a negative point estimate for the lowest growth rate \hat{A}_0 . Since technological regress is incompatible with the model, I repeated the minimization under the constraint that $\hat{A}_0 \geq 0$. The estimates then were $\hat{A}_0 = 0.000$; $\hat{A}_1 = 0.022$ and $\hat{A}_2 = 0.036$, $\ln c_2 - \ln c_1 = 0.08$, and $\ln c_1 - \ln c_0 = 0.001$ and $\ln c_2 - \ln c_1 = 0.09$. The difference between the lowest and highest growth rate is the same as in the two state model, so the upper bound on the increase in growth implied in Proposition 3 would be unchanged. But the three state model is still noteworthy in several respects. First, unlike the two state model, it does not predict countercyclical growth: $\ln c_2 - \ln c_1$ is positive, so that more rapid growth is associated with higher consumption levels. The estimate for $\ln c_1 - \ln c_0$ is negative, but is essentially indistinguishable from 0. Second, it provides some evidence of diminishing returns, since $j \ln c_2 - \ln c_1 > j \ln c_1 - \ln c_0$ while $j \hat{A}_2 - \hat{A}_1 < j \hat{A}_1 - \hat{A}_0$. Lastly, the sequence $\{s_t\}$ that achieves the smallest mean square error in the three state model accords quite well with conventional measures of the business cycle, in contrast with the two state model which requires a fair amount of volatility to match consumption growth, especially in the 1990s when annual consumption growth was consistently close to its historical average of 2.0%. The three state model, on the other hand, accords almost perfectly with NBER business cycle dates. The simulated consumption paths and changes in consumption for the two state and the three state model are illustrated in Figure 3.

Second, although I have focused on innovation as the engine of growth, the same basic argument could be applied to the factor accumulation model in Section 2 in calibrating an investment function $\hat{A} = \frac{1}{K}$. The reason I chose to focus on R & D is both because it seems to be a more relevant source of growth for developed countries such as the U.S., as well as the fact that estimates of curvature in investment based on q regressions appear to be noisy and less reliable. Still, whatever estimates there are point to substantial curvature in investment, since investment appears to respond so sluggishly to changes in q I did experiment with introducing adjustment costs to the model presented in Section 2, and I could generate growth effects on par with Ramey and Ramey using specifications for $\hat{A}(q)$ that match empirical estimates in the investment literature. As in the case with R & D, these estimates turned out to rely on implausible volatility in investment. However, we can once again establish an upper bound similar to that in Proposition 3. In the model developed in Section 2, the ratio of $\hat{A}'(q)$ turns out to depend on the volatility of the investment to consumption ratio. Since this series has a standard deviation of about 15%, this model again confirms that reasonable specifications for the underlying mechanism of diminishing returns can account for the growth effects documented



(a) Log Consumption - Actual vs. Simulated



(b) Δ Log Consumption - Actual vs. Simulated

Figure 3

by Ramey and Ramey, which in turn imply large welfare costs of business cycles.

4. Conclusion

This paper explores the costs of business cycles in a model of endogenous growth. Although previous work has explored the role of aggregate fluctuations in models of endogenous growth, and has even noted that growth effects could have implications for the costs of business cycles, the notion that business cycles involve substantial costs has yet to be analyzed formally. This paper identifies diminishing returns to innovation as an important channel in allowing business cycles to generate substantial welfare effects for plausible empirical conditions. Reduced form estimates from cross-country comparisons suggest that for the U.S., this channel would generate an increase in growth from 2.0% to 2.5%, which is associated with an increase in welfare for which an individual would be willing to sacrifice 10% of his initial consumption. Estimates of diminishing returns from micro data are consistent with these reduced form estimates, although they require implausibly volatile R & D and equity values to generate observed volatility in consumption growth. Still, restrictions on the innovation function that accord with volatility data allow an increase in growth from 2.0% to 2.36%, well within the range of estimates from cross-country data and which continue to deliver costs of business cycles that are far larger than reported in previous work.

Since there is already an extensive literature on the relationship between growth and cycles, it would seem fitting to end with a recap of what this paper contributes to this literature. First and foremost, it emphasizes the effect of volatility on growth through diminishing returns to investment, whereas previous work has emphasized the effect of stabilization on average investment. This explains why previous work has made conflicting predictions on the effects of stabilization on growth, reflecting a general ambiguity on the effects of uncertainty on investment. The model presented here is just as ambiguous about the theoretical implications of business cycles on average growth since the average level of investment can either rise or fall when fluctuations are reduced. However, it unambiguously predicts stabilization will increase growth at a fixed level of investment. Since empirical evidence finds little evidence that volatility affects average investment rates across countries, this appears to be the more relevant channel for studying the effects of business cycles on long run growth.

Second, this paper examines the welfare effects of stabilization rather than just the effects of

stabilization on the growth rate per se. Carrying out these welfare calculations explicitly shows that business cycles can matter. This conclusion stands in contrast to the large body of work which computes only small welfare gains from the volatility of consumption over the business cycle. As the discussion in the Introduction reveals, the initial observation by Lucas that consumption volatility from aggregate fluctuations has negligible welfare implications seems to have held up to subsequent scrutiny, except perhaps when such shocks are both very persistent and highly volatile. This is because fluctuations in aggregate consumption in the postwar era are sufficiently small that they fail to register any impact for reasonable aversion to fluctuations. But this does not imply that stabilization cannot produce first order welfare effects through other channels. In a sense, this paper closes the circle that Lucas originally began: he used his calculation to argue that because consumption volatility has negligible welfare effects, only growth matters. But if business cycles affect the process of long run growth in a way that is consistent with his thought experiment, which this paper argues, they will matter as well.

Appendix

Proof of Lemma 1: Recall that the value of a patent is given by

$$v_t = E_t \int_t^{\infty} \frac{U'(C_{t+s})}{U'(C_t)} \frac{dC_{t+s}}{dC_t} \frac{1}{\beta_{t+s}} e^{-\int_t^{t+s} \beta_s ds} ds = E_t \int_t^{\infty} \frac{1}{\beta_{t+s}} e^{-\int_t^{t+s} \beta_s ds} ds$$

I begin by characterizing the evolution of the discount factor β_{t+s} and of expected profits $E_t[\pi_{t+s}]$ assuming government spending remains constant between t and $t+s$, i.e. where for all $i \geq 2$ $[t, t+s]$ $G_i = G_1$:

1. From (1.1), we have

$$g_{t+s} = \frac{Y_{t+s} i}{\beta_{t+s}}$$

which implies

$$\begin{aligned} \int_0^{\infty} \ln g_{t+s} g ds &= \int_0^{\infty} \ln \frac{Y_{t+s} i}{\beta_{t+s}} \frac{1}{\beta_{t+s}} ds \\ &= \int_0^{\infty} \ln(Y_{t+s} i) \frac{1}{\beta_{t+s}} ds - \int_0^{\infty} \ln \beta_{t+s} \frac{1}{\beta_{t+s}} ds \\ &= \int_0^{\infty} \ln(Y_{t+s} i) \frac{1}{\beta_{t+s}} ds - \int_0^{\infty} \ln_{s^*} m_s(j)_{i-1} g ds \\ &= \int_0^{\infty} \ln(Y_{t+s} i) \frac{1}{\beta_{t+s}} ds - \int_0^{\infty} [m_s(j)_{i-1}] g ds \quad \text{where} \\ &= \int_0^{\infty} \ln(Y_{t+s} i) \frac{1}{\beta_{t+s}} ds - (\ln_{s^*}) (m_t + \dot{A}_i s_{i-1}) \end{aligned}$$

where

$$m_{t+s} = \int_0^{\infty} [m_s(j)_{i-1}] ds = E[m_s(j)]_{i-1}$$

The last step uses the fact that $(m_{t+s} - m_t) \gg \text{Poisson}(\dot{A}_i s)$. This implies as long as the same regime i prevails between dates t and $t+s$,

$$C_{t+s} = \exp \int_0^{\infty} \ln g_{t+s} g ds = e^{(\ln_{s^*})(m_{t+1} + \dot{A}_i s)} (Y_{t+s} i)$$

so that

$$\frac{dC_{t+s}}{dY_{t+s}} = e^{(\ln_{s^*})(m_{t+1} + \dot{A}_i s)}$$

2. Using the fact that $\frac{U'(C_{t+s})}{U'(C_t)} = \frac{C_{t+s}}{C_t} \mu_i$, I compute the ratio of consumption $\frac{C_{t+s}}{C_t}$. The growth of consumption at a given instant is given by

$$\begin{aligned} d \ln C_t &= \frac{Z_1}{Z_1} d(\ln g_t) \\ &= \frac{Z_1}{Z_1} [(1 - \dot{A}_i) dt + \dot{A}_i (\ln g_t) dt] \\ &= \frac{Z_1}{Z_1} [(1 - \dot{A}_i) dt + \dot{A}_i (\ln g_t) dt] \\ &= \dot{A}_i (\ln g_t) dt \end{aligned}$$

so that as long as regime i prevails between time t and $t+s$, we have

$$C_{t+s} = C_t e^{(\ln g_t) \dot{A}_i s}$$

3. Prob. μ_{t+s} are either equal to μ_i if a better technology was not invented and 0 if it was. For a constant level of government spending the probability that no new discovery is made between dates t and $t+s$ is given by $e^{-\dot{A}_i s}$. Hence

$$E_t[\mu_{t+s}] = \mu_i e^{-\dot{A}_i s}$$

Combining these three results, the profit flow at time $t+s$ evaluated according to the respective and conditional on the same regime prevailing between time t and $t+s$ is given by

$$\frac{U'(C_{t+s})}{U'(C_t)} \frac{dC_t}{dY_t} \mu_{t+s} = \begin{cases} \mu_i \exp(\ln g_t (m_{ti} + 1 + \dot{A}_i (1 - \mu_i) s)) & \text{if no new innovation arrives} \\ & \text{(a probability } e^{-\dot{A}_i s} \text{ event)} \\ 0 & \text{else} \end{cases}$$

A sufficient condition for v_i to be finite is for

$$\frac{1}{2} + \dot{A}_i + (\mu_i - 1) \ln g_t \dot{A}_i > 0$$

which is satisfied for $\mu_i > 1$.

Define $Z_t = e^{\ln g_t (m_{ti} + 1 + \dot{A}_i (1 - \mu_i) s)}$. If the regime changes at time $t+s$, then we update the value of m according to

$$m_{t+s} = m_t + \dot{A}_i (1 - \mu_i) s$$

It follows that the evolution of z_t exhibits the following law of motion within a given regime

$$\frac{z_t}{z_i} = (1 + \hat{A}_{it}(\epsilon_i - 1)) \ln_s$$

The value of owning a patent can therefore be characterized by the following asset equation

$$\begin{aligned} \frac{1}{2}v_i(z) &= z^{\frac{1}{2} + \hat{A}_i + 1} (v_i - v_j) + v^0(z) z^{\hat{A}_i} \hat{A}_i v_i \\ &= z^{\frac{1}{2} + \hat{A}_i + 1} (v_i - v_j) + z^{\hat{A}_i} (\epsilon_i - 1) (\ln_s) v_i^0(z) + \hat{A}_i v_i \end{aligned}$$

so we have a system of equation

$$\begin{aligned} (\frac{1}{2} + \hat{A}_0 + 1)v_0(z) + \hat{A}_0 v_0(z) + z^{\hat{A}_0} (\epsilon_0 - 1) (\ln_s) v_0^0(z) &= z^{\frac{1}{2} + \hat{A}_0 + 1} v_0 \\ (\frac{1}{2} + \hat{A}_1 + 1)v_1(z) + \hat{A}_1 v_1(z) + z^{\hat{A}_1} (\epsilon_1 - 1) (\ln_s) v_1^0(z) &= z^{\frac{1}{2} + \hat{A}_1 + 1} v_1 \end{aligned}$$

To solve the above system, we guess

$$\begin{aligned} v_0 &= \bar{v}_0 z \\ v_1 &= \bar{v}_1 z \end{aligned}$$

Matching coefficients yields the coefficients \bar{v}_0 and \bar{v}_1 reported in the text. \neq

Proof of Proposition 1: I begin by showing that $n_1 > n_0$. For suppose not, i.e. $n_1 \leq n_0$. Then the fact that $N_0 > N_1$ implies

$$L(n_1, N_1) > L(n_0, N_0)$$

so that $\frac{1}{2}v_1 > \frac{1}{2}v_0$. Since $n_1 > n_0$ and $\hat{A}^0(n) > 0$, it follows that $\hat{A}(n_1) > \hat{A}(n_0)$. In this case,

$$\begin{aligned} \frac{v_0}{v_1} &= \frac{[\frac{1}{2} + 1 + (\epsilon_0 - 1) \ln_s] \hat{A}(n_1) \frac{1}{2} v_0 + 1 \frac{1}{2} v_1}{[\frac{1}{2} + 1 + (\epsilon_1 - 1) \ln_s] \hat{A}(n_0) \frac{1}{2} v_1 + 1 \frac{1}{2} v_0} \\ &< \frac{[\frac{1}{2} + 1 + (\epsilon_0 - 1) \ln_s] \hat{A}(n_0) \frac{1}{2} v_0 + 1 \frac{1}{2} v_1}{[\frac{1}{2} + 1 + (\epsilon_1 - 1) \ln_s] \hat{A}(n_0) \frac{1}{2} v_1 + 1 \frac{1}{2} v_0} \\ &= \frac{1 (\frac{1}{2} v_0 + \frac{1}{2} v_1) + [\frac{1}{2} + (\epsilon_0 - 1) \ln_s] \hat{A}(n_0) \frac{1}{2} v_0}{1 (\frac{1}{2} v_0 + \frac{1}{2} v_1) + [\frac{1}{2} + (\epsilon_1 - 1) \ln_s] \hat{A}(n_0) \frac{1}{2} v_1} \\ &< 1 \end{aligned}$$

However, from the first order condition (1.3), $n_0 > n_1 \Rightarrow v_0 > v_1$, which is a contradiction. Hence, $n_1 > n_0$. This establishes part (1) of the Proposition. From the first order condition (1.3), this also implies $v_1 > v_0$, which establishes part (2).

To establish (3), suppose that $n_1 - n_0 > N_0 - N_1 > 0$, so

$$L(n_0, N_0) > L(n_1, N_1)$$

and $v_0 > v_1$. But then,

$$\begin{aligned} \frac{v_0}{v_1} &= \frac{(\frac{1}{2} + \frac{1}{\sigma} + (1 + \frac{1}{\sigma}) \ln_{\sigma} \bar{A}(\eta)) v_0 + \frac{1}{\sigma} v_1}{(\frac{1}{2} + \frac{1}{\sigma} + (1 + \frac{1}{\sigma}) \ln_{\sigma} \bar{A}(\eta)) v_1 + \frac{1}{\sigma} v_0} \\ &> \frac{(\frac{1}{2} + \frac{1}{\sigma} + (1 + \frac{1}{\sigma}) \ln_{\sigma} \bar{A}(\eta)) v_0 + \frac{1}{\sigma} v_1}{(\frac{1}{2} + \frac{1}{\sigma} + (1 + \frac{1}{\sigma}) \ln_{\sigma} \bar{A}(\eta)) v_1 + \frac{1}{\sigma} v_0} \\ &= \frac{\frac{1}{\sigma} (v_0 + v_1) + (\frac{1}{2} + (1 + \frac{1}{\sigma}) \ln_{\sigma} \bar{A}(\eta)) v_0}{\frac{1}{\sigma} (v_0 + v_1) + (\frac{1}{2} + (1 + \frac{1}{\sigma}) \ln_{\sigma} \bar{A}(\eta)) v_1} \\ &> 1 \end{aligned}$$

which contradicts the fact that $v_1 > v_0$. Hence, $v_1 > v_0$, establishing (3).

To show (4), note that nominal income is given by

$$Y = \frac{1}{\sigma} + W = \frac{1}{\sigma} + L_i \eta$$

so that $Y_1 - Y_0$ is given by

$$\begin{aligned} \frac{1}{\sigma} \eta_1 - \frac{1}{\sigma} \eta_0 + (\eta_1 - \eta_0) &= (\zeta_i - 1)(N_0 + \eta_0 - N_1 + \eta_1) + (\eta_1 - \eta_0) \\ &= \frac{1}{\sigma} (\eta_1 - \eta_0) + (\zeta_i - 1)(N_1 - N_0) \end{aligned}$$

whose sign is ambiguous since we can only establish that

$$N_0 - N_1 > \eta_1 - \eta_0 > 0$$

whereas $Y_1 > Y_0$ only if

$$N_0 - N_1 > \frac{\sigma}{\zeta_i - 1} (\eta_1 - \eta_0) > \eta_1 - \eta_0$$

It is possible to generate results whether the first inequality is and is not satisfied. The fact that $C = Y - T$ establishes (4). \square

Proof of Lemma 2: Applying my previous results, we have that for each regime $i \in \{1, 2\}$,

$$\begin{aligned} (PC)_i &= Y_i - T \\ &= \frac{1}{\sigma} + L_i \eta_i - T \\ &= (\zeta_i - 1)(\frac{1}{\sigma} - N_i + \eta_i) + L_i \eta_i - T \\ &= \frac{1}{\sigma} (\frac{1}{\sigma} - \eta_i) + (\zeta_i - 1)N_i - T \\ &= \frac{1}{\sigma} (\frac{1}{\sigma} - \eta_i) + (\zeta_i - 1)G_i - T \end{aligned}$$

Thus, comparing nominal consumption under stabilization is given by

$$\frac{1}{\sigma} L_i - \frac{1}{\sigma} T_i - \frac{1}{\sigma} \eta_i \bar{G} + (\zeta_i - 1) \bar{G}$$

while average nominal consumption in a world of stochastic shocks is given by

$$\bar{c} = \frac{1}{2} (c^0(G_0) + c^0(G_1)) + \frac{1}{2} (G_0 + G_1)$$

Rearranging these two equations and using the fact that $\bar{c} = \frac{1}{2} (G_0 + G_1)$, it follows that

$$\begin{aligned} p(c) \bar{c} &> \frac{p(c^0(G_0)) + p(c^0(G_1))}{2} \\ p(c) \bar{c} &= \frac{p(c^0(G_0)) + p(c^0(G_1))}{2} \\ p(c) \bar{c} &< \frac{p(c^0(G_0)) + p(c^0(G_1))}{2} \end{aligned} \quad , \quad \begin{aligned} n(\bar{c}) &< \frac{n^0(G_0) + n^0(G_1)}{2} \\ n(\bar{c}) &= \frac{n^0(G_0) + n^0(G_1)}{2} \\ n(\bar{c}) &> \frac{n^0(G_0) + n^0(G_1)}{2} \end{aligned}$$

establishing the claim. \square

Proof of Lemma 3: Simple differentiation yields

$$\frac{d}{dn} \frac{p(n)}{A(n)} = \frac{A'(n) - A''(n)n}{n^2}$$

By concavity we know that

$$A'(n) > A'(0) + A''(n)n$$

Since $A'(0) = 0$ by assumption, it follows that the derivative is positive. \square

Proof of Proposition 3: The maximization problem can be expressed as

$$\max_a \min \left\{ A(n) \ln a + \frac{v_1}{v_0} a \frac{n}{2}; A(n) \ln a + \frac{1}{2} a \frac{n}{2} \right\}$$

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 a or decreasing a , depending on which expression is larger. Solving for a at the point of equality and substituting back in yields the desired result. \square

References

- [1] Abel, Andrew 1983. "Optimal Investment under Uncertainty" *American Economic Review* 73(1), March, p228-33.
- [2] Aghion, Philippe and Peter Howitt, 1992. "A Model of Growth through Creative Destruction" *Econometrica* 60(2), March, p323-51.
- [3] Aghion, Philippe and Gilles Saint Paul, 1998. "Virtues of Bad Times: Interaction between Productivity Growth and Economic Fluctuations" *Macroeconomic Dynamics*, 2(3), September, p322-44.
- [4] Aizerman, Joshua and Nancy Marion, 1993. "Policy Uncertainty, Persistence, and Growth" *Review of International Economics*, 1(2), June, p145-63.
- [5] Alvarez, Fernando and Urban Jermann, 1999. "Using Asset Prices to Measure the Cost of Business Cycles" *imeo* Wharton.
- [6] Atkeson, Andrew and Christopher P helan "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.
- [7] Beaudry, Paul and Carmen Pages, 1999. "The Cost of Business Cycles and the Stabilization Value of Unemployment Insurance" *imeo* University of British Columbia
- [8] Bernanke, Ben, 1983. "Irreversibility, Uncertainty, and Cyclical Investment" *Quarterly Journal of Economics* 98(1), February, p85-106
- [9] Caballero, Ricardo, 1991. "On the Sign of the Investment-Uncertainty Relationship" *American Economic Review* 81(1), March, p279-88.
- [10] Campbell, John and John Cochrane, 1995. "By Force of Habit: a Consumption-Based Explanation of Aggregate Stock Market Behavior." *NBER Working Paper* # 4995, January.
- [11] Chatterjee, Satyajit and Dean Corbae, 1999. "A Welfare Comparison of Pre and Post WW II Business Cycles: Some Implications for the Role of Postwar Macroeconomic Policies" *imeo* University of Pittsburgh.
- [12] Christiano, Larry and Jonas Fisher. "Stock Market and Investment Good Prices: Implications for Macroeconomics" *imeo* Northwestern University, 1998.

- [13] Clark, Kenneth, Derek Leslie, and Elizabeth Symons, 1994. "The Costs of Recession" *Economic Journal*, 104(422), January, p20-36
- [14] deHek, Paul, 1999. "On Endogenous Growth Under Uncertainty" *International Economic Review*, 40(3), August, p727-44.
- [15] DeLong Bradford and Lawrence Summers, 1988. "How Does Macroeconomic Policy Affect Output?" *Brookings Papers on Economic Activity*, 2, p433-80.
- [16] Dixit, Avinash and Robert Pindyck, 1994. *Investment Under Uncertainty*. Princeton: Princeton University Press.
- [17] Dumas, Jim, 1998. "Risk Preferences and the Welfare Cost of Business Cycles" *Review of Economic Dynamics*, 1(3), July, p646-76
- [18] Eaton, Jonathan, 1981. "Fiscal Policy, Inflation, and the Accumulation of Risky Capital" *Review of Economic Studies*, 48 (3), July, p435-45.
- [19] 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
- [20] Fatas, Antonio, 1998. "Do Business Cycles Cast Long Shadows? Short-Run Persistence and Economic Growth" *mimeo, IN SEAD*.
- [21] Giliches, Zvi and Ariel Pakes, 1984. "Patents and R & D at the Firm Level: a First Look" in *R & D, Patents, and Productivity*, edited by Zvi Giliches. Chicago: University of Chicago Press, p55-72.
- [22] Grossman, Gene and Elhanan Helpman, 1991. *Innovation and Growth in the Global Economy*. Cambridge, MA: MIT Press.
- [23] Hall, Bronwyn, Zvi Giliches, and Jerry Hausman, 1986 "Patents and R and D: Is There a Lag?" *International Economic Review*, 27(2), June, p265-83.
- [24] Hartman, Richard, 1972. "The Effects of Price and Cost Uncertainty on Investment" *Journal of Economic Theory*, 5(2), October, p258-66
- [25] Helpman, Elhanan and Manuel Trajtenberg, 1998. "A Time to Sow and a Time to Reap: Growth Based on General Purpose Technologies" *General Purpose Technologies and Economic Growth*, edited by Elhanan Helpman. Cambridge, MA: MIT Press.

- [26] Hopenhayn, Hugo and Maria Muniaurria, 1996. "Policy Variability and Economic Growth" *Review of Economic Studies*, 63(4), October, p61-65.
- [27] Imrhoroglu, Ayse, 1989. "Cost of Business Cycles with Indivisibilities and Liquidity Constraints" *Journal of Political Economy*, 97(6), December, p1364-83.
- [28] Jones, Larry, Rodolfo Manuelli, and Ennio Stacchetti, 1999. "Technology (and Policy) Shocks in Models of Endogenous Growth" *mimeo Northwestern University*
- [29] Kortum, Samuel, 1992. "Inventions, R&D, and Industry Growth" PhD Thesis, Yale University.
- [30] Krusell, Per and Anthony Smith, 1999. "On the Welfare Effects of Eliminating Business Cycles" *Review of Economic Dynamics*, 2(1), January, p245-272.
- [31] Leland, Hayne, 1974. "Optimal Growth in a Stochastic Environment" *Review of Economic Studies*, 41(1), January, p75-86
- [32] Lelhari, David and T.N. Srinivasan, 1983. "Optimal Savings Under Uncertainty" *Review of Economic Studies*, 36(1), April, p153-63.
- [33] Lucas, Robert, 1978. "Asset Prices in an Exchange Economy" *Econometrica*, 46(6), November, p1429-45.
- [34] Lucas, Robert, 1987. *Models of Business Cycles*. Oxford: Basil Blackwell.
- [35] Matsuyama, Kiminori, 1995. "Complementarities and Cumulative Processes in Models of Monopolistic Competition" *Journal of Economic Literature*, 33(2), June, p701-29.
- [36] Obstfeld, Maurice, 1994a. "Evaluating Risky Consumption Paths: The Role of Intertemporal Substitutability" *European Economic Review*, 38(7), August, p1471-86
- [37] Obstfeld, Maurice, 1994b. "Risk-Taking, Global Diversification, and Growth" *American Economic Review*, 84(5), December, p1310-1329.
- [38] Otrck, Christopher, 1999. "On Measuring the Welfare Cost of Business Cycles" *mimeo University of Virginia*
- [39] Pemberton, James, 1996a. "The Costs of Recession: A Comment" *Economic Journal*, 106(437), July, p77-80.

- [40] Pemberton, James, 1996. "Growth Trends, Cyclical Fluctuations, and Welfare with Non-Expected Utility Preferences" *Economics Letters*, 50(3), March, pp37-92.
- [41] Ramsey, Gary and Valerie Ramsey, 1991. "Technology Commitment and the Cost of Economic Fluctuations" *NBER Working Paper* # 3755, June
- [42] Ramsey, Gary and Valerie Ramsey, 1995. "Cross-Country Evidence on the Link Between Volatility and Growth" *American Economic Review*, 85(5), December, pp1138-51.
- [43] Romer, David, 1996 *Advanced Macroeconomics*, McGraw-Hill.
- [44] Stokey, Nancy, 1995. "R & D and Economic Growth" *Review of Economic Studies*, 62(3), July, pp465-89.
- [45] Tallarini, Thomas, 1999. "Risk Sensitive Business Cycles" *mimeo, Carnegie Mellon University.*
- [46] Uzawa, Hirofumi, 1968. "Time Preference and the Penrose Effect in a Two-Class Model of Economic Growth" *Journal of Political Economy*, 77, pp48-52.