

On the Cyclicity of Research and Development*

Gadi Barlevy

Economic Research
Federal Reserve Bank of Chicago
230 South La Salle
Chicago, IL 60604

e-mail: gbarlevy@frbchi.org

September 12, 2006

Abstract

Economists have recently revived the notion that recessions play a useful role in fostering innovation and growth. Yet a major source of growth, R&D, is procyclical. This paper argues one reason why R&D might be procyclical is because of a dynamic externality inherent to R&D that makes entrepreneurs short-sighted and concentrate their innovation in booms even though it is optimal to concentrate it in recessions. Even if additional forces imply that procyclical R&D is desirable, equilibrium R&D in a decentralized environment is likely to be too procyclical, and it would be better if some R&D would be shifted to recessions.

*I am grateful to the editor and two anonymous referees who made extremely valuable suggestions in improving the paper. I would also like to thank Jeff Campbell, Marty Eichenbaum, Jonas Fisher, Huw Lloyd-Ellis, Merritt Lyon, Kiminori Matsuyama, Alex Monge, Joanne Roberts, and Fabrizio Zilibotti for their comments, as well as seminar participants at Northwestern, the Stockholm School of Economics, Queens' University, the London School of Economics, the NBER Summer Workshop, the Canadian Macro Study Group, and the Society of Economic Dynamics.

Introduction

In recent years, economists have resurrected the Schumpeterian idea that recessions promote activities which contribute to long-run productivity.¹ Modern reincarnations of this hypothesis emphasize the role of intertemporal substitution: since the opportunity cost of achieving productivity growth – the forgone output or sales that could have been achieved instead – is lower in recessions, there is greater incentive to undertake such activities in downturns. Recessions should therefore stimulate long-run productivity growth, and cyclical fluctuations might even contribute positively to welfare by allowing the economy to grow at a lower resource cost.

Although some growth-enhancing activities – such as schooling – are countercyclical, one of the major sources of long-run productivity growth – research and development (henceforth R&D) – appears to be procyclical. Yet R&D seems like an activity that should be concentrated in recessions: it is labor intensive, and while labor productivity in producing goods appears to decline in recessions, Griliches (1990) observes that productivity in innovation is acyclical. The purpose of this paper is to provide some insight on why R&D might be procyclical despite this apparent incentive for it to be countercyclical. In particular, is the procyclical pattern in R&D inefficient, necessitating policy intervention such as countercyclical R&D subsidies?

This paper argues that externalities inherent to R&D can induce a procyclical bias in R&D relative to the socially optimal path for R&D. Moreover, this bias can be quantitatively significant, enough to turn R&D procyclical in environments where it should be countercyclical. That said, there may be other forces that contribute to the procyclicality of R&D, some of which may make procyclical R&D desirable. Hence, it might not necessarily be optimal to concentrate R&D in recessions as suggested by the Schumpeterian view, but the bias still implies resources for R&D are underprovided in recessions. In other words, private entrepreneurs allocate too much of their R&D activity to booms, but the optimal path may still be procyclical.

The bias I describe relies on the notion that new ideas benefit not only the innovator who came up with it but others as well. For example, others might be able to improve on the original idea, incorporate some of its features, or adopt it in full once patents expire. It has been observed since at least Arrow (1962) that if innovators cannot appropriate the spillovers from their research, there will be too little R&D. But there is a less appreciated temporal

¹See, Hall (1991), Mortensen and Pissarides (1994), and Gomes, Greenwood, Rebelo (2001) on how recessions affect search; Cooper and Haltiwanger (1993), Caballero and Hammour (1994), Aghion and St. Paul (1998), and Canton and Uhlig (1999) on how they affect technical change; and DeJong and Ingram (2001), Dellas and Sakellaris (2003), and Barlevy and Tsiddon (2006) on how they affect human capital accumulation.

aspect to such spillovers, namely that the benefits which accrue to others typically occur later in time. For example, rival innovators are more likely to improve upon or imitate an idea as more time passes. Similarly, with limited-duration patents, rivals benefit from a new idea only after a patent expires. The incentive for an innovator to undertake R&D is thus driven by the short-term benefits of innovation which accrue principally to her. Whereas a benevolent planner would take into account that research undertaken in a recession yields benefits when economic conditions improve, private agents discount these longer-run benefits and fail to fully exploit downturns to carry out innovation at a lower cost. The more procyclical are profits, the stronger this distortion will be. If profits are sufficiently procyclical, innovation can fall enough in recessions to turn R&D procyclical on the basis of this distortion alone.

This bias is related to work by Shleifer (1986) which argues that entrepreneurs introduce new technologies *en masse* in booms in order to capture high profits. However, Shleifer examines when firms implement new technologies, not when they undertake R&D. When Francois and Lloyd-Ellis (2003) endogenize innovation in Shleifer's model, they find it is countercyclical: entrepreneurs engage in R&D in recessions, then wait to implement their ideas in booms. By contrast, I assume firms do not wait to implement new ideas, e.g. because they are impatient. If innovators must implement immediately, the higher profits in booms could potentially conflict with the incentives to engage in intertemporal substitution.² The contribution of this paper is to show that because of dynamic externalities, entrepreneurs tend to chase short term profits at the expense of intertemporal substitution, engaging in too little R&D during recessions.

To illustrate this point, I construct an equilibrium model of R&D that captures the essence of the Schumpeterian view: there is a fixed supply of resources that can be used for either production or innovation, and the relative productivity of the two uses fluctuates over time in a way that leads the opportunity cost of innovation to fall in recessions. At the same time, I incorporate a dynamic spillover whereby successful innovation today benefits future entrepreneurs. I solve for the equilibrium of this model and discuss its tendency towards too little R&D in recessions. I then calibrate the model to gauge whether the bias is large.

The paper is organized as follows. Section 1 presents evidence on the procyclicality of R&D. Section 2 sorts through potential sources of procyclical R&D to focus the subsequent theory. Section 3 presents the model and illustrates its inherent procyclical bias. Section 4 calibrates the model to assess the quantitative significance of the bias. Section 5 concludes.

²Another difference is that in Shleifer's model cycles occur because new ideas are implemented with delay. Thus, it is not possible to ask how entrepreneurs who implement immediately behave in his model during recessions, since without delay there is no cycle. I instead assume cycles are due to exogenous shocks.

1. Evidence on the Cyclicity of R&D

I begin my analysis by documenting the cyclical behavior of R&D activity. There are two main data sources on R&D activities in the U.S. The first is the National Science Foundation (NSF), which sponsors annual surveys on R&D activities conducted by private and public entities. Among private entities, questionnaires are administered to a subset of firms from the universe of companies tracked by the Census Bureau. Large companies that are known in advance to be active in R&D are included in this subset with certainty; additional firms are sampled each year from the remaining population of firms. Companies are questioned about their R&D expenditures, defined as activities whose “purpose is to do one or more of the following things: pursue a planned search for new knowledge, whether or not the research has reference to a specific application; apply existing knowledge to problems involved in the creation of a new product or process, including work required to evaluate possible uses; or apply existing knowledge to problems involved in the improvement of a present product or process.” Firms are also asked about the source of funding for these activities, how their R&D expenditures are allocated between wage and salary payments and other expenses such as materials, and the number of scientists and engineers employed in R&D.

The other source of data on R&D expenditures is the Standard & Poor’s Compustat database of publicly traded companies. As part of their 10-K filing, publicly traded companies must disclose their expenditures on research and development, defined as “planned search or critical investigation aimed at discovery of new knowledge” and “translation of research findings or other knowledge to an existing product or process whether intended for sale or use.” While publicly traded firms are a more selected sample than those surveyed by the NSF, the virtue of Compustat is that it allows us to relate R&D expenditures to firm and industry data.

In documenting how R&D activity evolves with macroeconomic conditions, I use data from both datasets and confirm that both yield similar patterns. I then consider evidence on potential explanations for these patterns using firm and industry level data from Compustat. Appendix A provides more details on how the data was assembled.

1.1. The Cyclicity of Research and Development

I begin with data from the NSF. Figure 1 plots the growth in real R&D expenditures performed and financed by private industry as reported by the NSF between 1958 and 2003. Expenditures are deflated using the implicit GDP deflator, as the NSF suggests. As evident from Figure 1,

the growth rate of R&D covaries positively with real GDP growth. The correlation between the two series is 0.39, and regressing growth in real R&D expenditures on real GDP growth yields a coefficient of 0.69 with a standard error of 0.25. It is also apparent from Figure 1 that growth in real R&D expenditures tends to fall in NBER recession years, although in the recessions of both the early 1980s and 1990s R&D growth remained steady during the recession and only declined afterwards. The fact that R&D expenditures as reported by the NSF tend to grow more rapidly in booms than in recessions has been noted by previous researchers. Fatas (2000) plots essentially the same two series (up through 1997), although his measure of R&D includes publicly-funded R&D. Griliches (1990) surveys the evidence on patents, and suggests both patents and R&D (the latter based on NSF data) are procyclical. More recently Comin and Gertler (2006) decompose the same R&D expenditure data used here into frequencies, and find that R&D expenditures are procyclical at business cycle frequencies and even more so at medium-run frequencies of between 8 and 50 years.

Although various authors have used real R&D expenditures as a measure of innovative activity, this measure has some drawbacks. First, since there is no long series on the price of R&D, deflating by GDP deflator may lead to spurious fluctuations in R&D that reflect mismeasurement of prices rather than true changes in R&D activity. Second, firms may not be consistent in what activities they consider to be R&D. I therefore turn to the growth rate in the number of full-time equivalent scientists and engineers employed in R&D in industry as reported by the NSF, narrowly defined to include “all persons engaged in scientific or engineering work at a level which requires a knowledge of physical or life sciences or engineering or mathematics” and whose “experience is equivalent to completion of a 4-year college course with a major in these fields, regardless of whether or not they actually hold a degree in this field.” Although scientists and engineers account for only a part of all R&D inputs, the virtue of this measure is that it does not depend on prices and is thus more likely to reflect true R&D activity. This growth rate is also depicted in Figure 1. As evident from the figure, the growth rate in employment of scientists and engineers closely tracks the growth rate in R&D expenditures; the correlation between the two series is 0.61 over the period in which both series are available. However, employment growth is less correlated with GDP growth; the correlation between the two is only 0.10, and regressing the growth in employment of scientists and engineers on the growth in real GDP does not yield a statistically significant coefficient. Still, the fact that employment growth is highly correlated with R&D expenditures suggests R&D expenditures reflect changes in real resources devoted to research and development activities rather than spurious mismeasurement.

I next turn to the Compustat data. In contrast to NSF data, Compustat only reports R&D expenditures, not employment. Also, it is not a representative sample; the firms in Compustat

are publicly traded and tend to be large (measured either by assets or employees). However, since a disproportionate share of R&D expenditures in the NSF is done by large firms, R&D patterns among these firms may capture quite well the behavior of total R&D expenditures. Using the firms in Compustat that engage in some R&D, I estimate average growth in R&D expenditures each year. That is, I took all domestically incorporated firms that report positive R&D in both the current and previous year for each year, then averaged the log growth in their R&D expenditures, weighting by the initial level of R&D expenditures of each firm. To minimize the effects of outliers on the average growth rate, I removed observations whose absolute change in R&D expenditures ranked in the top 5% of all observations. Figure 2 plots this series, together with the growth rate in R&D expenditures reported by the NSF depicted in Figure 1. The two series are consistent, both in level and in the way they evolve over time. Nevertheless, there are a few differences in the two series. Average growth of R&D expenditures among Compustat firms fell more sharply in the early 1980s than among all firms as reported by the NSF. In addition, R&D growth in the late 1990s appeared to be more erratic among Compustat firms, although both datasets show a sharp decline in real R&D expenditures around the 2001 recession. Average growth of R&D in Compustat is actually more correlated with real GDP growth than the NSF series; the correlation between the two series for the same period between 1958 and 2003 is 0.49, and regressing average growth in R&D on growth in real GDP yields a higher coefficient of 0.94 with a standard error of 0.25.

Since Compustat reports R&D data by firm, we can use this data to examine whether the procyclical pattern in R&D is also apparent at the industry level.³ That is, do firms tend to expand R&D at a faster rate when output in the industry they operate in is rising? Fortunately, Compustat assigns each firm a four-digit standard industry classification (SIC) code. For many firms it also reports the North American industry classification system (NAICS) industry code. Even when Compustat fails to assign a NAICS code, in many cases I was able to assign one based on the assigned SIC code using conversion tables provided by the Census Bureau.

Measuring output at the industry level is somewhat more problematic than at the aggregate level; in particular, which is a better indicator of a booming industry, gross output or value added? In the absence of a clear answer, I consider both measures. First, using the NBER Manufacturing and Productivity database compiled by Bartelsman, Becker, and Gray, I construct real gross output for 459 four-digit SIC industries from 1958 to 1996. This dataset provides highly disaggregated data and covers an extended period. However, it only encom-

³While the NSF does break down total R&D expenditures by SIC industry codes, many of these series contain missing values or are deliberately withheld for privacy reasons.

passes manufacturing. While this excludes certain R&D-intensive industries such as computer programming services and communications, R&D is heavily concentrated in manufacturing. According to both the NSF and Compustat, the share of R&D expenditures in manufacturing has consistently ranged between 70 – 80% of total R&D expenditures since 1990, and over 90% earlier. NSF data suggests manufacturing accounts for a similar share of employment of scientists and engineers. To incorporate data outside of manufacturing, I also used data from Bureau of Economic Analysis (BEA) on real gross output and real value added by industry. The BEA provides these estimates at the three-digit industry level rather than at the four digit level (and using the NAICS classification as opposed to SIC). This data also covers a shorter period than the NBER database: gross output data is available for 1988-2004, and value added is available at the three digit level for 1978-2004, and at the two digit level for 1947-2004.

Armed with this data, I assigned each firm in Compustat its respective industry output measure whenever available. I then regressed the growth of real R&D expenditures of each firm on its industry’s output growth along with firm fixed effects.⁴ This regression should capture the extent to which firms deviate from their long-run growth rate of R&D as industry conditions vary. The results are reported in the first row of Table 1. As evident from the table, the coefficient on industry output is positive and statistically significant for all three measures. To isolate the reaction of firms to industry-specific (as opposed to aggregate) output growth, I introduced year effects to absorb both cyclical effects and changes in patent policy or R&D subsidies that affect all industries. As can be seen in the table, the coefficients on output growth decline for all three measures, but they remain positive and statistically significant. Hence, even after controlling for aggregate conditions, industries where production is increasing tend to undertake more R&D, not less. Firms do not appear to substitute towards R&D when industry output is low or divert resources from R&D when industry output is high.

1.2. Evidence on Possible Sources of R&D Procyclicality

Since Compustat contains additional firm-level information, we can use it to examine potential explanations for the procyclical pattern in R&D activity. For example, could the procyclical pattern in R&D be picking up some other firm-level variation that is correlated with industry or aggregate output? Does the procyclicality of R&D vary systematically across industries, and

⁴Rafferty and Funk (2004) also regress R&D growth among Compustat firms against the growth in gross output using fixed effects, although they only used gross output from the NBER Manufacturing and Productivity database. But their regression includes additional explanatory variables, among them the firm’s own sales, which make it difficult to interpret the coefficient on growth in industry output as a pure measure of cyclical sensitivity.

if so can it be related to industry characteristics? I now explore such questions. My analysis is motivated by particular hypotheses regarding the procyclicality of R&D. These are discussed in more detail in the remainder of the paper, so I only briefly outline them here.

One possible explanation for the procyclicality of R&D is that firm earnings fall in recessions, making it more difficult to finance innovative activities. According to this view, the positive coefficient on industry output is really standing in for omitted measures of firm balance sheet positions, and including these should reduce the implied procyclicality of R&D. Since Compustat includes financial information on firms, we can test this prediction. Table 2 repeats the regressions from Table 1 but adding various balance sheet variables. Specifically, I add the current year cash flow (before R&D expenses) and the residual cash flow in the previous year after R&D expenses, as well as current and lagged values of firm assets, firm liabilities, plant, property, and equipment, short-term debt (debt due within one year) and long-term debt. All variables are deflated by the implicit GDP deflator. As evident in Table 2, these variables have almost no effect on the estimated coefficient on industry output. The increase in R&D expenditures in periods of rising industry output evidently does not reflect omitted changes in the balance sheet positions of firms. Adding second and third order terms of all of the above balance sheet variables left the coefficient on output growth similarly unchanged.

The fact that balance sheet information cannot explain any of the apparent procyclicality of R&D might seem surprising at first. However, this result should be interpreted with caution. In particular, it does not mean credit constraints are unimportant for R&D. In fact, although not reported in the table, the coefficient on lagged cash flow is significant in two of the specifications reported in Table 2, suggesting that similarly-sized firms (as measured by assets) with more cash do tend to have higher R&D growth. Rather, the results in Table 2 imply that the tendency of firms to accelerate their R&D when industry output rises does not reflect how relatively cash-rich or cash-poor they are. This result is consistent with Himmelberg and Petersen (1994) who argue R&D is sensitive to cash flow, but that this relationship is also hard to detect because firms tend to smooth R&D in response to temporary cash flow shocks by, among other things, curtailing their investment in physical capital. In other words, cash-rich firms undertake more R&D activity than cash-poor firms, but regardless of the scale of their R&D activity, both tend to minimize the disruptive effect of temporary fluctuations in cash flows on their R&D.

Another explanation for procyclical R&D argues that firms engage in more R&D in booms because this is when the present discounted value of the profits they expect to earn if successful is highest. According to this hypothesis, firms in industries where the present discounted value of profits rises more in booms should also have more procyclical R&D. In addition, for reasons

that will become clear once I present the model, industries where the present discounted value of profits is only modestly procyclical should have *countercyclical* R&D.

To test this prediction, I partitioned firms into two-digit SIC codes. For industries in which relatively few firms reported R&D (e.g. particular industries within agriculture, retail trade, and wholesale trade), I aggregated firms into their respective one-digit industries. Conversely, for industries within two-digit codes that are particularly R&D intensive, (e.g. the pharmaceutical, communication equipment, computers, and computer programming), I divided the original two-digit code into multiple industries. This left me with 44 industries. I first needed a measure of the procyclicality of present discounted profits for these industries. I used the stock market values of publicly traded firms in each industry as a stand-in for present discounted profits, since these should in theory reflect the discounted value of future dividends of these firms. Thus, I took all domestically incorporated firms in each industry where available and regressed the growth rate in the real stock prices of firms on the growth of real GDP and a constant. Let $\widehat{\beta}_{stock}$ denote the coefficient on real GDP growth from this regression, which represents how stock market values in a given industry covary with the business cycle. Next, for each industry I took all domestically incorporated R&D-performing firms and regressed the growth rate in real R&D expenditures of each firm on the growth rate of real GDP and a constant. Let $\widehat{\beta}_{R\&D}$ denote the coefficient on real GDP growth in this regression. Figure 3 plots $\widehat{\beta}_{stock}$ and $\widehat{\beta}_{R\&D}$ across industries, where the size of each point corresponds to the inverse of the standard error in estimating $\widehat{\beta}_{R\&D}$. The reason for weighting is that some sectors have few firms that report any R&D (e.g. transportation, utilities, insurance, and finance), and $\widehat{\beta}_{R\&D}$ is poorly estimated in those sectors. As evident from the figure, industries with more procyclical stock values tend to have more procyclical R&D growth. For example, computer programming and semiconductors, whose stock market values are far more procyclical than all remaining industries, also have highly procyclical R&D, while primary metals, petroleum, mining, and agriculture, whose stock market values are the least procyclical, display weakly countercyclical R&D growth (although in all three cases $\widehat{\beta}_{R\&D}$ is statistically indistinguishable from zero).⁵

Regressing $\widehat{\beta}_{R\&D}$ on $\widehat{\beta}_{stock}$ using the same set of weights yields the following relationship:

$$\widehat{\beta}_{R\&D} = -1.58 + 0.39 \times \widehat{\beta}_{stock}$$

(0.30) (0.04)

⁵As an aside, neither stock market values nor R&D are particularly procyclical for pharmaceuticals. This is probably because innovation is an exceptionally drawn out process due to regulation, insulating the industry from cyclical fluctuations. Although pharmaceuticals account for a significant share of total R&D, they do not contribute much to the cyclical patterns in R&D this paper focuses on. One should therefore avoid drawing inference from that sector for the phenomena studied here.

This confirms that the positive relationship in Figure 3 is statistically significant. Note that the constant in this regression is significantly negative, suggesting that industries where stock market values are only modestly procyclical will tend to exhibit countercyclical R&D. While this finding is consistent with the model I present below, one has to be careful in interpreting this finding. In particular, this intercept is extrapolated, since I estimate $\hat{\beta}_{stock}$ to be positive and significant in all 44 industry categories. Moreover, while I estimate negative values of $\hat{\beta}_{R\&D}$ in several industries, in none of the cases is this negative value statistically significant at the 10% level. Thus, the evidence that industries for which stock market values are only moderately procyclical tend to have countercyclical R&D is at best indirect.

2. Previous Explanations for the Procyclicality of R&D

Before attempting to model the procyclical bias in R&D described in the Introduction, I first review some of the theories advanced in the literature for the procyclicality of R&D. The purpose of this section is to help focus the subsequent theoretical model by determining what elements are and are not important for studying the *cyclical* behavior of R&D.

As already noted in the previous section, one common explanation for the procyclicality of R&D involves credit market frictions. This hypothesis is explicitly developed in Aghion, Angeletos, Banerjee, and Manova (2005). In their model, the opportunity cost of R&D is lower in recessions. But downturns also reduce the amount of internal funds available for firms to finance ongoing R&D projects. Aghion *et al* show that if firms had unlimited access to credit, they would concentrate growth-enhancing activities in recessions. But if firms anticipate they will be borrowing-constrained, they would instead focus their R&D efforts in booms for fear that they will not be able to bring projects initiated in recessions to fruition. While this explanation is intuitive, the evidence described in the previous section suggests it is unlikely to account for the procyclical pattern in R&D we observe, at least insofar as credit constraints are reflected in firm balance sheets. This conclusion is lent further credence by the observation that the lion's share of the cyclical fluctuations in R&D come from large, publicly traded firms account which are generally less likely to suffer from financing difficulties. I therefore chose not to incorporate credit constraints in modelling the evolution of R&D over the business cycle. However, this does not imply credit constraints are not important for understanding R&D more generally.

Another explanation for the procyclicality of R&D is that recessions do not create opportunities for intertemporal substitution as assumed under the Schumpeterian view. There are two different versions of this argument. According to one view, the specialized labor employed in

R&D is not readily substitutable for production workers at high frequencies. While there is some truth that scientists and engineers may not be easily shifted to production, NSF data suggests that on average 40% of wage payments in R&D are allocated to support staff, which are more easily reallocated into production. Moreover, an important input into R&D is managerial attention, which certainly could be alternatively devoted to production issues. A separate argument for why recessions might not encourage intertemporal substitution is that R&D relies on produced goods as an input. But if more goods can be produced during booms, there will be more resources available to use for R&D. Aghion and Saint Paul (1998) explore this possibility formally, and show that innovation will be procyclical if productivity-improving activities use final goods rather than factor inputs.⁶ But they then reject this view, arguing that data on how demand shocks affect long-run productivity is inconsistent with this assumption. More directly, as Griliches (1984) observes, the primary input into R&D is labor, not produced goods. Hence, although some R&D certainly involves produced goods, I will assume in my model that R&D only uses labor inputs that could alternatively be used for production.

A more compelling explanation for the procyclicality of R&D is that labor supply itself varies over the cycle. Specifically, if labor resources are procyclical, it may not be desirable to divert already scarce labor resources in recessions from production to R&D. Conversely, there may no need to draw down labor from R&D during booms given that labor resources are more plentiful. This is precisely how Fatas (2000) generates procyclical R&D in his paper. While I abstract from this aspect in my initial analysis, I do so largely for reasons of exposition. When I look at the quantitative aspects of the model, I will allow for cyclical variation in labor supply.

3. A Model of R&D and Analysis of a Special Case

To study the evolution of R&D over the business cycle, I adapt a standard real business cycle model to include an innovation sector as modeled by Grossman and Helpman (1991). Production in the model economy involves two factors of production, capital and labor. These can be used to produce a single final good via a two-step process. This final good in turn can be used for three purposes: consumption, production of new capital, and as an indirect input into the production of final goods (a feature sometimes referred to as roundabout production). Labor resources, in addition to being used for production, can be used to carry out R&D in order to augment labor productivity in producing this final good. Thus, production and R&D compete

⁶Following Rivera-Batiz and Romer (1991), this is known as the lab-equipment model. Comin and Gertler (2006) also assume R&D uses final goods, and note that this assumption helps to generate procyclical R&D.

for labor resources. This section lays out the model, and then focuses on a special case of the model that can be solved analytically to illustrate its key features. In the next section, I calibrate the model in order to analyze its quantitative properties.

I assume households in the economy can be represented by a single agent with constant-relative risk-aversion utility over the amount of the final good C_t consumed at date t :

$$U(C_t) = \frac{C_t^{1-\gamma} - 1}{1-\gamma} \quad (3.1)$$

Utility is discounted at rate ρ . The representative agent is endowed with a constant labor endowment L per unit time and an initial capital stock K_0 normalized to one. Since the agent is assumed to derive no utility from leisure, the entire labor endowment will be supplied at every date. However, as alluded to in the previous section, variable labor supply can be an important source of procyclical R&D. Although it will be more convenient to abstract from variation in labor supply for now, I will introduce such variation in the next section.

Turning to technology, consider first the production of capital goods. I assume a linear technology whereby one unit of the final good can be converted into q units of capital. Assuming capital depreciates at rate δ , the instantaneous rate of change in the capital stock is

$$\dot{K}_t = qI_t - \delta K_t \quad (3.2)$$

where I_t denotes the amount of final goods allocated to the production of capital.

Next, consider the production of the final good. I assume it is constructed in two stages. First, labor is converted into a continuum of intermediate goods indexed by $j \in [0, 1]$. Second, the intermediate goods are combined with capital to produce the final good.

Let us begin with the second stage. It will be convenient to proceed as if all intermediate goods must first be assembled into a composite good, whose quantity I denote by X_t , according to a Cobb-Douglas technology. That is, x_{jt} of each good j together yield

$$X_t = \exp \left[\int_0^1 \ln x_{jt} dj \right] \quad (3.3)$$

The production of the final good is itself Cobb-Douglas in the composite good and capital, i.e. X_t units of composite good and K_t units of capital yield Y_t units of the final good, where

$$Y_t = z_t K_t^\alpha X_t^{1-\alpha} \quad (3.4)$$

The coefficient z_t reflects productivity in the final goods sector. To capture the fact that labor productivity in the goods sector varies over the business cycle, I let z_t follow a Markov switching process between two states, $Z_1 \geq Z_0$, with a constant hazard rate μ . This is the only source of fluctuations in the model. I treat these fluctuations as exogenous, although one could potentially endogenize them.⁷

I next turn to the production of intermediate goods. For reasons that will become clear below, I assume the production of each intermediate good entails both a fixed cost and a variable cost. The fixed cost is assumed to be denominated in final goods, i.e. F_t units of final good have to be sacrificed per instant if a given intermediate good is to be produced. As the subscript suggests, this amount is time varying, in a way described in more detail below. Once this fixed cost is incurred, labor can be converted into intermediate goods according to a good-specific linear technology. Let L_{jt} denote the amount of labor employed in the production of good j . The amount x_{jt} of intermediate good j produced at date t is given by

$$x_{jt} = \lambda_{jt} L_{jt} \tag{3.5}$$

The coefficient λ_{jt} is assumed to have the following structure:

$$\lambda_{jt} = \lambda^{m_{jt}} \tag{3.6}$$

for some $\lambda > 1$, where m_{jt} is an integer that denotes the generation of technology used for producing good j at date t . Each good j starts out at generation m_{j0} , respectively, but agents can advance to higher-generation technologies by engaging in R&D. That is, starting with generation m_j , devoting R_j units of labor to R&D on good j gives rise to a hazard ϕR_j of discovering generation m_j+1 in the next instant, which given $\lambda > 1$ will be more productive than its predecessor. In line with the evidence on the acyclicity of productivity in the innovation sector in Griliches (1990), I assume ϕ does not vary with time.

The discovery of a new technology, in addition to increasing the productivity in good j by a factor of λ , leads to two additional outcomes. First, it allows innovators to begin working on the next generation technology. R&D is therefore incremental, in the sense that the m -th generation has to be created before the $m+1$ -th technology can be. Each new discovery immediately becomes public knowledge, allowing rivals to incorporate it as they work on subsequent generations. This feature introduces a dynamic externality, since whenever a researcher discovers a new generation she allows others to build on and profit from her work.

⁷For example, Benhabib and Farmer (1994) describe an economy with spillovers in which the scale of production, and thus the productivity of individual producers, fluctuate endogenously over time. Fluctuations in z_t could also reflect unmodelled changes in utilization rates.

Second, I assume each new discovery spawns inferior imitations. That is, new ideas can be copied, albeit imperfectly, in the sense that imitation versions are more costly to operate. Although the imitation technology will never actually be used, it plays an important role in the analysis. In particular, the technology available to rivals affects the profits the leading producer can earn in a decentralized market. Without imitations, rivals could use older generation technologies, which have higher variable costs but the *same* fixed cost as the leading producer. In the model, this feature generates counterfactual implications about the markup producers charge. I therefore assume discoveries spawn imitations which have a *lower* fixed cost than the leading version (but a higher variable cost), resulting in more empirically plausible markups. This assumption can be defended on the grounds that imitators do not incur the fixed costs of patent protection that a leading-edge producer would. To simplify the analysis, I assume the imitation technology involves zero fixed cost and has a marginal cost that is λ times the marginal cost of the leading technology, although this exact specification is not essential.

Let $M_t = \int_0^1 m_{jt} dj$ denote the average generation across intermediate goods, and let $R_t = \int_0^1 R_{jt} dj$ denote total employment in R&D aggregated over all intermediate goods. One can show that in any equilibrium (and along the optimal path), each sector will use the same amount of labor $L_{jt} = L - R_t$. Substituting in this expression, we have

$$Y_t = z_t K_t^\alpha X_t^{1-\alpha} = z_t [\lambda^{M_t} (L - R_t)]^{1-\alpha} K_t^\alpha \quad (3.7)$$

Output grows as labor productivity improves and more capital is accumulated. To avoid the situation where the economy outgrows the fixed cost, I scale this fixed cost in proportion to the rate at which output increases, namely λ^{M_t} . Specifically,

$$F_t = \lambda^{M_t} F \quad (3.8)$$

for some constant $F > 0$. The notion that fixed costs grow over time is plausible; for example, the cost of overhead labor should naturally rise with overall labor productivity.⁸

This describes the economic environment. In the remainder of this section, I focus on a special case of this model which is analytically tractable. In particular, I restrict $\gamma = 0$ so that the agent is risk-neutral, and set $q = \delta = 0$ so that the value of capital is always fixed at one

⁸This begs the question why I did not model the fixed cost directly in terms of labor. The reason is that the value of labor in a decentralized market would change one-for-one with z_t . As a result, the fixed cost F_t will vary greatly over the cycle. But, in practice, real wages tend to be somewhat rigid over the cycle, even though they do grow over longer time horizons. Fixed costs should therefore not vary much over the cycle, a feature I wanted the model to adhere to. Assuming the fixed cost is denominated in final goods captures rigidity in the salaries of overhead labor without unnecessarily complicating the model.

(i.e. capital can be viewed as a fixed factor, such as land). Finally, in the absence of capital accumulation, (3.8) must be modified to $F_t = \lambda^{(1-\alpha)M_t} F$, the rate at which output grows. None of the other parameters is assigned a value.

Consider first what a benevolent social planner would do in this economy. She would maximize the well-being of the representative household, choosing how much labor to allocate to R&D and each intermediate good (and if capital were accumulable, how much capital to produce) as well as which technology to use to produce each intermediate good. Clearly, the planner would always use the leading-edge technology for any good. It is also easy to show that she should allocate equal the same amount of labor, $L - R_t$, to each intermediate good. The planner's problem thus reduces to choosing a path for R&D:

$$V_i(M_0) = \max_{R_t} E \left[\int_0^\infty \lambda^{(1-\alpha)M_t} \left[z_t (L - R_t)^{1-\alpha} - F \right] e^{-\rho t} dt \mid z_0 = Z_i \right] \quad (3.9)$$

subject to the constraint that

$$\dot{M}_t = \phi R_t$$

i.e. the change in the average generation of technology corresponds to the fraction of successful entrepreneurs who discover the next generation in their respective line. For this problem to be well-defined, we need to ensure the planner cannot achieve infinite utility. This requires that the economy cannot grow faster than the discount rate ρ . Since the maximal growth rate occurs when all labor resources are used in R&D, this condition can be written as

$$\rho > (1 - \alpha) \phi L \ln \lambda \quad (3.10)$$

We can rewrite the planner's problem recursively as follows:

$$\rho V_i(M) = \max_{R \in [0, L]} \left\{ \lambda^{(1-\alpha)M} \left[Z_i (L - R)^{1-\alpha} - F \right] + \mu (V_{1-i}(M) - V_i(M)) + \frac{\partial V_i}{\partial M} \phi R \right\} \quad (3.11)$$

I now establish the following result:⁹

Proposition 1: In any interior path, optimal innovation is countercyclical along the optimal path, i.e. $R_0 > R_1$.

This claim formalizes the basic Schumpeterian argument: since the returns to production are low when z_t is low, it is preferable to shift resources towards R&D in these periods. However,

⁹ All propositions relate to interior solutions where R&D is always positive. An earlier version of this paper dealt with corner cases. One can typically guarantee an interior solution by assuming L is sufficiently large.

there are two aspects of this special case that make countercyclical R&D particularly desirable. First, labor L is assumed to be constant over time. If instead labor were more plentiful when z were high, it might not be necessary to divert resources from R&D to take advantage of temporarily high productivity in the goods sector. Second, the assumption of risk-neutrality, which I impose for analytical convenience, contributes to making countercyclical R&D more desirable. In particular, a drawback of shifting resources from production to innovation during recessions is that it makes output more volatile; the decline in the amount of labor used in production compounds the already low productivity. To the extent this makes consumption more volatile, risk-neutrality may overstate the desirability of countercyclical R&D. In the next section I explicitly take into account these two considerations.

I now show that in same environment for which Proposition 1 dictates a countercyclical path for R&D, equilibrium R&D in a decentralized market can be procyclical. I first need to specify how the decentralized economy is organized, and then solve for its equilibrium.

Both intermediate and final are produced by profit-maximizing firms. The technology for producing final goods is freely available, so profits in this sector will equal zero in equilibrium. By contrast, intermediate goods producers are assumed to enjoy some market power: the entrepreneur who discovers the m -th generation for producing good j earns a patent that grants him exclusive rights to use this technology; without this patent, no innovation would ever take place. In what follows, I focus on equilibria where R_{jt} is the same across all goods j , and where in addition the common value of R&D in all sectors, R_t , can be expressed as a function of aggregate productivity z_t at date t alone. This restriction is natural given these features are true for the optimal path. Formally, I restrict attention to symmetric Markov perfect equilibria. Let R_i denote the level of R&D common to all sectors when productivity $z_t = Z_i$ for $i \in \{0, 1\}$.

I normalize the wage per unit labor to 1. Let p_{jt} denote the equilibrium price of intermediate good j . Given the Cobb-Douglas aggregator X , the demand of final goods producers for each intermediate good j will be unit elastic. Thus, each intermediate-goods producer would want to charge as high a price as possible: his revenue will be constant regardless of the price he charges, but at higher prices he can produce fewer goods and lower his costs. However, if he were to charge too high of a price, his rivals could underprice him and steal away his business. Given my assumptions, the relevant threat involves the knock-off version of his own technology, which has no fixed cost but a marginal cost that is λ times as large. Hence, the price required to deter rival entry is the marginal cost of the rival technology, i.e. $p_{jt} = \lambda^{-(m_{jt}-1)}$. At this price, the producer would be charging a constant markup λ over his own marginal cost.

Let $e_{jt} = p_{jt}x_{jt}$ denote total expenditures by final goods producers on intermediate good j . Given the Cobb-Douglas specification for X , final goods producers will equalize expenditures across all intermediate goods, so $e_{jt} = e_t$ for all j . Each intermediate goods producer thus earns e_t in revenue, and given my normalization of the wage to 1, incurs variable costs equal to the number of workers he employs. The latter expression is equal to $\lambda^{-m_{jt}}x_{jt}$, or alternatively $\lambda^{-m_{jt}}e_t/p_{jt}$. Using the expression for p_{jt} derived above, profits are given by

$$\pi_{jt} = \left(1 - \frac{1}{\lambda}\right) e_t - P_t F_t \quad (3.12)$$

where P_t denotes the price of the final good. Profits are thus the same for all goods j , i.e. $\pi_{jt} = \pi_t$ for all j . Equation (3.12) helps to explain the reason for some of my earlier assumptions. It states that profits net of fixed costs are equal to total sales e_t times one minus the inverse markup. This equation holds more generally for any constant marginal cost technology,

$$\pi_t = \left(1 - \frac{1}{\lambda_t}\right) e_t - FC_t \quad (3.13)$$

where λ_t represents a potentially time-varying markup and FC_t represents the fixed cost at date t . In analyzing the cyclical behavior of profits, Ramey (1991) divides equation (3.13) through by sales e_t to arrive at an expression for the ratio of profits to sales. Since in practice profits are more procyclical than sales, she not surprisingly finds this ratio is highly procyclical. But as evident from (3.13), in the absence of fixed costs this ratio inherits the cyclical properties of the markup λ_t . Since markups are moderately countercyclical, we would need to allow for a fixed cost of production to have any hope of according with the data. At the same time, λ_t cannot be so countercyclical as to offset the role of the fixed cost. This explains the need for an imitation technology with a lower fixed cost than the leading technology. If the next most efficient producer had the same fixed cost as the leading producer, the resulting markup would be sufficiently countercyclical to make π_t/e_t acyclical. Getting the model to match the cyclical behavior of profits is important, since in the model it is profits that ultimately drive R&D.

I next derive expressions for e_t and P_t . Since there is no capital accumulation, in equilibrium the representative household spends all of its income on final goods. Thus, household expenditures on final goods $P_t(Y_t - F_t)$ equal household income, i.e. the sum of aggregate profits $\Pi_t = \int_0^1 \pi_t dj$ and factor payments:

$$P_t(Y_t - F_t) = \Pi_t + r_t K + L \quad (3.14)$$

where r_t denotes the rental rate of capital at date t . Given the Cobb-Douglas technology for final goods, expenditures on intermediate goods e_t should account for a fraction $1 - \alpha$ of the total cost of final good production. Since the market for final goods is perfectly competitive,

the latter equals the market value of final goods produced, i.e. $e_t = (1 - \alpha) P_t Y_t$. Similarly, expenditures on capital goods $r_t K$ equal $\alpha P_t Y_t$. Substituting into (3.14) and rearranging yields

$$\pi_t = (\lambda - 1)(L - R_t) - P_t F_t \quad (3.15)$$

As for the price of final goods P_t , competition in the market for final goods implies P_t equals the minimum cost to produce a single unit of the good, i.e.

$$P_t = \min_{x_{jt}, K_t} \left\{ \int_0^1 p_{jt} x_{jt} dj + r_t K_t \right\} \\ \text{s.t. } z_t K_t^\alpha \left(\exp \left[\int_0^1 \ln x_{jt} dj \right] \right)^{1-\alpha} = 1$$

Using the fact that $p_{jt} = \lambda^{-(m_{jt}-1)}$, $r_t K_t = \alpha P_t Y_t$, and $K_t = 1$ for all t , one can show that

$$P_t = \frac{\lambda(L - R_t)^\alpha}{(1 - \alpha) z_t \lambda^{(1-\alpha)M_t}} \quad (3.16)$$

Hence, profits can be expressed as

$$\pi_t = (\lambda - 1)(L - R_t) - \frac{\lambda(L - R_t)^\alpha F}{(1 - \alpha) z_t} \quad (3.17)$$

Finally, let v_j denote the value of a successful innovation of intermediate good j . Entrepreneurs who succeed in innovation earn profits (3.17) as long as their technology is the most advanced. Since profits π_{jt} are the same for all j and R_{jt} is assumed to be the same for all j , v_j will be the same for all j as well. Let v denote the common value of a successful innovation in all sectors, and let v_i denote this value if current productivity $z_t = Z_i$. Let \mathbb{I}_{jt} be an indicator that equals 1 if the leading-edge producer of good j at date 0 is still the leading edge producer at date t , and zero otherwise. Since the representative agent owns all claims in equilibrium, v_i must leave him indifferent to buying an additional claim. This indifference condition implies

$$v_i = E \left[\int_0^\infty \mathbb{I}_{jt} \cdot \frac{U'(C_t)/P_t}{U'(C_0)/P_0} \pi_t e^{-\rho t} dt \mid z_0 = Z_i \right] \\ = E \left[\int_0^\infty \mathbb{I}_{jt} \cdot \frac{P_0}{P_t} \pi_t e^{-\rho t} dt \mid z_0 = Z_i \right] \quad (3.18)$$

where the expectation above is taken over all possible paths for z_t and \mathbb{I}_{jt} . A firm trying to become the leading producer will choose R to maximize the expected value from a successful innovation net of R&D costs, $\phi R v - R$. It follows that in equilibrium $\phi v_i \leq 1$ for $i \in \{0, 1\}$, with strict equality if $R_i > 0$.

Substituting in for π_t and P_t into (3.18) yields

$$v_i = E \left[\int_0^\infty \mathbb{I}_t \cdot \frac{(L - R_0)^\alpha}{Z_i} \frac{\lambda^{(1-\alpha)M_t}}{\lambda^{(1-\alpha)M_0}} \left[(\lambda - 1) z_t (L - R_t)^{1-\alpha} - F \right] e^{-\rho t} dt \mid z_0 = Z_i \right] \quad (3.19)$$

As in (3.10), we once again need to assume that growth is not too rapid:

$$\ln \lambda < (1 - \alpha)^{-1} \quad (3.20)$$

I now establish the following result:

Proposition 2: Given (3.20), there exist two values F^* and \bar{F} such that if $F > F^*$ and $F < \bar{F}$, then there exists a pair $R_0 < R_1$ where $\phi v_0(R_0, R_1) = \phi v_1(R_0, R_1) = 1$. By contrast, when $F = 0$, any solution for the system $\phi v_i(R_i, R_{1-i}) = 1$ must satisfy $R_0 > R_1$.

Proposition 2 suggests that if fixed costs are large, so that profits are highly volatile, there can be an equilibrium with procyclical R&D. At the same time, fixed costs cannot be too large, or else profits will be so low that entrepreneurs have no incentive to try to achieve them. Conversely, if fixed costs are small, so that profits are only as volatile as sales, equilibrium R&D will be countercyclical. Intuitively, if profits are only modestly procyclical, the value of innovation in booms will rise, but by less than the cost of R&D. Procyclical R&D requires that profits be at least more procyclical than the cost of R&D to ensure that innovation is more profitable in booms. Recall that earlier I found that industries with only moderately procyclical stock prices are indeed predicted to have countercyclical R&D.

Why does the decentralized market fail to reproduce the countercyclical path for R&D the planner would like? The reason is that in a decentralized market, entrepreneurs care too much about the short term benefits from their R&D, correctly expecting that the benefits from their research in the more distant future are more likely to accrue to others. Formally, consider the ratio of the real value of a successful innovation in a boom to its value in a recession. This ratio is given by

$$\frac{v_1/P_1}{v_0/P_0} = \frac{E \left[\int_0^\infty \mathbb{I}_t \cdot \lambda^{(1-\alpha)M_t} \left[z_t (L - R_t)^{1-\alpha} - \frac{F}{\lambda - 1} \right] e^{-\rho t} dt \mid z_0 = Z_i \right]}{E \left[\int_0^\infty \mathbb{I}_t \cdot \lambda^{(1-\alpha)M_t} \left[z_t (L - R_t)^{1-\alpha} - \frac{F}{\lambda - 1} \right] e^{-\rho t} dt \mid z_0 = Z_i \right]} \quad (3.21)$$

The analogous expression for the benevolent planner is the ratio of the marginal value improving the technology, $\partial V/\partial M$, in a boom relative to a recession. This expression is given by

$$\frac{\partial V_1/\partial M}{\partial V_0/\partial M} = \frac{E \left[\int_0^\infty \lambda^{(1-\alpha)M_t} \left[z_t (L - R_t)^{1-\alpha} - F \right] e^{-\rho t} dt \mid z_0 = Z_1 \right]}{E \left[\int_0^\infty \lambda^{(1-\alpha)M_t} \left[z_t (L - R_t)^{1-\alpha} - F \right] e^{-\rho t} dt \mid z_0 = Z_0 \right]} \quad (3.22)$$

Comparing the two expressions reveals two differences. First, in the decentralized market the fixed cost is scaled by a factor of $(\lambda - 1)^{-1}$. This is because entrepreneurs care not about

the output they produce but the profits they earn. While they only earn a proportion of the output they produce, they bear all of the cost of the output they use to cover their fixed cost. Entrepreneurs therefore weight the output they produce and the output they use up differently, unlike the planner. This distortion is not relevant for our purposes, and can be nullified by setting $\lambda = 2$ so the markup is 100%. Second, and more importantly, the entrepreneur multiplies the flow value at each date by \mathbb{I}_t , an indicator of whether he will still be the leading producer at date t . Since the probability that $\mathbb{I}_t = 1$ decreases with t , the integrals in (3.21) assign more weight to values closer to date $t = 0$ than the integrals in (3.22), i.e. private agents care less about the long-run benefits of innovation than the planner. The ratio of the respective values of innovation will thus be distorted in the direction of the ratio of profits in booms to profits in recessions, since these are the expressions that receive the largest weight under (3.22). The bigger the ratio of profits in booms to their value in recessions, the more entrepreneurs value innovations in booms than in recessions, and the more biased R&D will be towards booms.

Although the special case of the model above is analytically convenient for demonstrating the bias, it is far too stylized to help us gauge if the bias can play a significant role in making equilibrium R&D procyclical. To address this question, I shall now return to the general model and analyze its quantitative implications.

4. Quantitative Analysis

The general model involves several complications that the special case analyzed earlier allowed me to side-step. I first discuss some of these issues, and then I proceed to solve the model numerically for specifically calibrated parameter values. The section concludes with several robustness exercises, including the introduction of variable labor supply.

4.1. Analysis of the General Model

In analyzing the general model, it will prove helpful to restrict γ to 1. This is the value I use in the subsequent calibration, and the case of log utility turns out to be analytically convenient. The planner's problem in this case reduces to

$$V_i(K_0, M_0) = \max_{R_t, I_t} E \left[\int_0^\infty \ln \left(z_t K_t^\alpha [\lambda^{M_t} (L - R_t)]^{1-\alpha} - \lambda^{M_t} F - I_t \right) e^{-\rho t} dt \mid z_0 = Z_i \right]$$

s.t. 1. $\dot{M}_t = \phi R_t$
 2. $\dot{K}_t = qI_t - \delta K_t$

Let $k = \lambda^{-M} K$ and $\iota = \lambda^{-M} I$. Using the law of motion for M , one can show that $V_i(K_0, M_0) = v_i(k) + M_0 \frac{\ln \lambda}{\rho}$, where $v_i(k)$ satisfies

$$\rho v_i(k) = \max_{\iota, R} \left\{ \begin{array}{l} \ln \left(Z_i k^\alpha (L - R)^{1-\alpha} - F - \iota \right) + \frac{\phi R \ln \lambda}{\rho} + \\ \frac{\partial v_i}{\partial k} (q\iota - (\delta + \phi R \ln \lambda) k) + \mu (v_{1-i}(k) - v_i(k)) \end{array} \right\} \quad (4.1)$$

The planner now has *two* relevant control variables, investment and R&D. The first-order conditions for the maximization problem with respect to each are given by

$$\begin{aligned} \frac{1}{Z_i k^\alpha (L - R)^{1-\alpha} - F - \iota} &= q \frac{\partial v_i}{\partial k} \\ \frac{(1 - \alpha) Z_i k^\alpha (L - R)^{-\alpha}}{Z_i k^\alpha (L - R)^{1-\alpha} - F - \iota} &= \left(\frac{1}{\rho} - k \frac{\partial v_i}{\partial k} \right) \phi \ln \lambda \end{aligned}$$

Substituting the first equation into the second yields the following formula for R_i , the value of R&D when productivity is equal to Z_i :

$$R_i = L - \left[\left(\frac{1}{\rho (\partial v_i / \partial k)} - k \right) \frac{\phi \ln \lambda}{q (1 - \alpha) Z_i k^\alpha} \right]^{-1/\alpha} \quad (4.2)$$

Rather than two numbers R_0 and R_1 , an optimal plan now corresponds to two functions $R_0(k)$ and $R_1(k)$. This raises a question of what it means for R&D to be procyclical, since changes in z_t affect not only how much R&D is desirable at a given k but also k itself. I shall refer to a policy as procyclical if $R_1(k) > R_0(k)$ for any k in the limiting set the economy settles down to in the long-run, i.e. for any value of k that occurs infinitely often with probability 1. Similarly, a policy is said to be countercyclical if $R_1(k) < R_0(k)$ for all such k .

Next, consider the decentralized equilibrium of the economy. The production side of the economy is essentially the same as in the special case analyzed earlier, except that the price of final goods P now explicitly depends on the level of capital. In particular, (3.16) generalizes to

$$P = \frac{\lambda^{1-M} (L - R)^\alpha}{z (1 - \alpha) k^\alpha} \quad (4.3)$$

The main difficulty in solving for an equilibrium is that once we move away from the assumption of risk-neutrality, evaluating the value of a successful innovation v in (3.18) necessitates an expression for consumption C_t . Without accumulable capital, consumption is equal to net output $Y_t - F_t$. When capital is accumulable, we instead need to explicitly solve the household's problem to derive C_t , i.e. we must determine how a household should optimally divide its wealth between consumption, capital, and claims to the profits of intermediate good producers.

Since the household must own all of the claims on profits in equilibrium, it will be convenient to proceed as if there were a mutual fund company that pooled all entrepreneurs into a single portfolio on behalf of the household. Arbitrage requires the value of this portfolio to be the same as the cost of buying up all firms, which is just $\int_0^1 v dj = v$. To insure the fund continues to own all incumbents, it must pay the research expenses of any potential innovator in exchange for the rights to the patent if the innovator is successful, i.e. the fund deducts an operating expense R out of dividends. Thus, as far as the household is concerned, it can either allocate its wealth to physical capital or to an asset whose price is v and which yields a dividend of $\Pi = \pi - R$ per unit time, where π is given by (3.15).

Let w denote the household's nominal wealth and σ denote the fraction of its wealth that it allocates to capital. Given the linear technology for producing capital goods, the price of a unit of capital, P_K , is equal to qP . If aggregate productivity remains constant, the nominal return per unit of capital is $r + \dot{P}_K$, and the number of units of capital it holds is $\sigma w/P_K$. Similarly, the nominal return per share of the mutual fund it owns is $\Pi + \dot{v}$, and the number of shares it owns in the mutual fund is $(1 - \sigma)w/v$. In equilibrium, however, $\dot{v} = 0$. In addition to the returns to its assets, the household also earns labor income and spends some of its resources on consumption. Hence, the evolution of nominal wealth w while z_t is constant is given by

$$\dot{w} = \left[\left(\frac{r}{P_K} + \frac{\dot{P}_K}{P_K} \right) \sigma + \frac{\Pi}{v} (1 - \sigma) \right] w + L - \lambda^M P c \quad (4.4)$$

where $c = \lambda^{-M} C$. If productivity z_t did change, the nominal value of the physical capital the household owns would jump together with P_K . The nominal value of wealth held in the mutual fund would not change, however, since the value of the mutual fund $v = \phi^{-1}$ independently of aggregate productivity. Hence, the wealth of the household will jump from w to w^* where

$$w^* = \left[\frac{P'_K}{P_K} \sigma + (1 - \sigma) \right] w \quad (4.5)$$

and P'_K is the price of capital under the new level of productivity. Let W denote the aggregate wealth of the economy. In equilibrium, of course, $w = W$. However, since individual households act as price takers, they treat the path of W as given and assume it determines the values of all relevant economic variables. Let $R_i(W)$ denote the equilibrium employment in R&D when $z_t = Z_i$ and aggregate wealth is W_i . We can then express k in terms of W , since $W = P_K K + v = q\lambda^M P k + \phi^{-1}$. Using the expression for P in (4.3), we have

$$k_i(W) = \left[\frac{(W - \phi^{-1}) Z_i (1 - \alpha)}{q\lambda (L - R_i(W))^\alpha} \right]^{\frac{1}{1-\alpha}}$$

We can similarly express the nominal quantities r , P , and Π as functions of W . This implies we can express the household problem recursively in terms of two state variables, w and W :

$$\rho V_i(w, W) = \max_{\sigma, c} \left\{ \begin{array}{l} \ln c + \frac{\phi R \ln \lambda}{\rho} + \frac{\partial V_i}{\partial w} \dot{w} + \frac{\partial V_i}{\partial W} \dot{W} + \\ \mu (V_{1-i}(w^*, W^*) - V_i(w, W)) \end{array} \right\} \quad (4.6)$$

subject to (4.4) and (4.5), the free entry condition $\phi v = 1$, and the laws of motion for W , i.e. if z_t remains constant over the next instant, then

$$\dot{W} = (r + \dot{P}_K) \lambda^M k + \Pi + L - \lambda^M P c(W, W) \quad (4.7)$$

while if z_t changes over the next instant, W will jump to W^* , where

$$W^* = q \lambda^M P_{1-i} k_i(W) + \phi^{-1}$$

The first order conditions for the household problem with respect to σ and c are given by

$$\left(\frac{r}{P_K} + \frac{\dot{P}_K}{P_K} \right) - \phi \Pi = \mu \frac{\partial V_{1-i}(w^*, W^*) / \partial w^*}{\partial V_i(w, W) / \partial w} \left[1 - \frac{P_{1-i}}{P_i} \right] \quad (4.8)$$

$$\frac{1}{P c(w, W)} = \frac{\partial V_i}{\partial w} \quad (4.9)$$

An equilibrium is a set of functions $w_i^*(w)$, $V_i(w, W)$, and $R_i(W)$ which satisfy the system of equations (4.5), (4.6), and (4.8).

4.2. Calibration and Results

The equations that define an equilibrium must be solved numerically, requiring me to assign particular values to the various parameters of the model. Since the model is essentially a standard real business cycle model with endogenous growth, many of the parameters I use have already been discussed in the real business cycle literature. The values I use are as follows:

Table 3

γ	1.00	α	0.33	λ	1.20
ρ	0.05	Z_0	0.94	ϕ	0.10
q	1.00	Z_1	1.06	F	3.6
δ	0.08	μ	0.20	L	30.8

The first two parameters correspond to utility terms. As already anticipated, I assume log utility by setting the coefficient of relative risk aversion γ to 1. Normalizing a unit of time in

the model to correspond to a year, I set the discount rate ρ to 5%. The next several parameters relate to production. First, I set $q = 1$ so consumption and investment goods trade one-for-one, a common assumption in real business cycle models. I set the depreciation rate of capital δ to 8% per year. The share of capital in the production of final goods α is set to one third. To match the 6% unconditional standard deviation of detrended productivity growth we observe in the data, I set Z_0 to 0.94 and Z_1 to 1.06. I set the transition rate μ so that a complete cycle is 10 years, slightly longer than the 8 year frequency often used to identify business cycle fluctuations. For λ , I follow Rotemberg and Woodford (1999) in calibrating the markup to 20%. The productivity term ϕ turns out to be a scaling parameter; I normalize it to 0.10.

The remaining two parameters, F and L , are chosen to match the growth rate of GDP per capita and the average GDP share of R&D. Empirically, the average growth rate of GDP per capita is 2% per year. The GDP share for total R&D (both private and public) has been roughly stable at about 2.5% for much of the post-War period. However, the relevant share for my purposes is private R&D, which has trended during this period from 1% prior to just over 2%. To err on the side of caution, I forced the model to match the higher R&D share of 2%. This implies to a lower value for F and hence *less* procyclical R&D. If I had instead calibrated F to match a smaller R&D share, I would have needed a higher value of F to drive down profits and make R&D sufficiently less attractive that fewer resources would be devoted to it. It is not obvious whether the appropriate output measure in the model is gross output or output net of fixed costs. However, at $F = 3.6$, R&D accounts for 2.0% of gross output and 2.2% of net output, so the distinction is relatively minor. Interestingly, for these parameter values the model generates reasonable time variation in R&D, even though it was not designed to: the standard deviation of log R&D share over time is 0.139 and 0.136 for gross and net output respectively, compared to 0.137 for the log share of total R&D between 1953 and 2002.

As a fraction of output, the parameters in Table 3 imply a fixed cost equal to 8.1% of gross output (and 8.8% of net output). By comparison, Ramey (1991) and Basu (1996) suggest non-production workers as a proxy for overhead labor. Non-production workers account for 20% of the labor force during the post-War period. Since labor accounts for two-thirds of output, this suggests an overhead cost of 13% of output. If anything, my estimate is overly conservative.

To solve the model, I use a collocation method in which I approximate the respective value functions with n -th order polynomials, where I choose the coefficients of the polynomial by requiring that the asset equations hold exactly at $n + 1$ points. Thus, to solve the planner's problem I approximate $v_i(k)$ using a polynomial in k . To solve for the decentralized equilibrium, I approximate $R_i(W)$ and $w_i^*(w)$ using polynomials in W and w , respectively, and the function

$V_i(w, W)$ with the polynomial $\sum_{k=0}^n \sum_{\ell=0}^{n-k} a_{k\ell} w^k W^\ell$. The coefficients of each respective polynomial are chosen so that either the equilibrium conditions or the planner's first order conditions hold exactly at particular values of k , w , and W , respectively.¹⁰ The results reported here are based on $n = 4$, although I confirmed that higher order polynomials yielded nearly identical results.

The purpose of solving for equilibrium is to determine whether at the calibrated parameter values the model can deliver procyclical R&D, and what role the bias plays in accounting for this pattern. However, since the model abstracts from variable labor supply, it is still premature at this stage to gauge the empirical plausibility of the R&D path the model generates.

Figure 4 plots the functions $R_i(k)$ along the equilibrium path over the limiting set for k . For the values in Table 3, equilibrium R&D is procyclical. The extent of this procyclicality is modest: R&D in a boom is about 3 – 8% higher for a given value of k than in a recession. Moreover, since an increase in z_t stimulates capital accumulation, which raises k and provides even further incentive to undertake R&D, the actual path for R&D will appear to be even more procyclical. By contrast, solving the planner's problem reveals that the optimal path for R&D is countercyclical, i.e. $R_0(k) > R_1(k)$ for all limiting values of k . This suggests the procyclical bias in R&D must be large at empirically plausible parameter values, enough to turn R&D procyclical all by itself in an environment where it should be countercyclical.

One caveat in comparing the equilibrium and optimal paths for R&D is that the two differ not only in their implied cyclicity of R&D but also in the average level of R&D. At the assigned values, optimal R&D is much higher than equilibrium R&D. One might therefore ask if the planner would still prefer a countercyclical path if she were restricted to the same average level of R&D as prevails in equilibrium. This question can be addressed using a perturbation argument. Suppose aggregate productivity fluctuates between $Z_0 = 1 - \varepsilon$ and $Z_1 = 1 + \varepsilon$ for a small value ε . Let R denote non-stochastic steady state equilibrium R&D when $Z_0 = Z_1 = 1$. For ε small, the planner's constrained optimum can be approximated by $R_0 = (1 - \zeta) R$ and $R_1 = (1 + \zeta) R$ for some ζ . For $\varepsilon = 0.01$, I find that the household's value function is decreasing in ζ in the neighborhood of $\zeta = 0$. Thus, for small shocks, the optimal policy that is constrained to keep average R&D unchanged would opt to concentrate R&D in recessions. I confirm that this remains true for $\varepsilon = 0.06$, i.e. starting from a constant path for R&D, a small countercyclical

¹⁰Following the recommendation of Judd (1998), these points correspond to the roots of the Chebyshev polynomials, adapted to the limiting interval for the relevant variable. For $V_i(w, W)$ I use the triangular array $\{w_i, W_j\}_{1 \leq i \leq j \leq n+1}$ where w_i and W_j represent the roots of Chebyshev polynomials adapted to the limiting interval for equilibrium wealth. Note that I need to approximate $V(w, W)$ both on and off the equilibrium path (in which $w = W$) to approximate both $\partial V_i / \partial w$ and $\partial V_i / \partial W$.

perturbation from a constant R&D path raises welfare, while a small procyclical perturbation lowers it. The model unambiguously views procyclical equilibrium R&D as inefficient.

4.3. Robustness

The calibration exercise above reveals that the procyclical bias in R&D can be strong enough to account for the procyclicality of R&D all by itself. It also suggests that optimal policy should actually reverse the timing of R&D. The remainder of this section explores the robustness of these two results. To preview the remainder of the section, I first argue that small changes in certain parameters can significantly weaken the magnitude of the bias, enough that equilibrium R&D will no longer be procyclical. This suggests the bias towards procyclical R&D might not be able to account for the procyclicality of R&D by itself. I then argue that variable labor supply, which I have emphasized throughout as a potentially important consideration, can restore the procyclicality of R&D in these cases. Interestingly, plausible degrees of variable labor supply cannot on their own generate procyclical labor, suggesting both elements are important. Finally, I argue that once we allow for variable labor supply, it may no longer be optimal to turn R&D countercyclical (although it will still be desirable to shift some R&D from booms to recessions).

I begin by examining the effects of changing some of the parameters in Table 3. One parameter of interest is the fixed cost F . Recall that the magnitude of the fixed cost determines how much profits vary with z_t . If I set F to zero (while adjusting L to keep average growth at 2%), equilibrium R&D turns countercyclical, just as predicted by Proposition 2 for the special case of the model analyzed earlier. So how large must the fixed cost be for equilibrium R&D to be procyclical? The answer turns out to be 3.3, which corresponds to 7.9% of gross output and 8.6% of net output at the implied parameter values. It would therefore not take a much smaller fixed cost than the one I originally calibrated to turn R&D countercyclical. That said, fixed costs are probably higher than my estimate, not lower. Aside from the observation that the share of non-production workers all by itself is roughly 13%, profit rates in the U.S. are quite small, while smaller values of F would imply large profit rates (as well as counterfactually large R&D shares as becoming a leading-edge producer turns more lucrative).

Another parameter of interest is the markup λ , especially given the difficulty in estimating marginal cost properly. Since markups determine profits, it is not surprising that this parameter has important implications for the magnitude of the procyclical bias in R&D. Again, as I varied λ , I also adjusted together with L and F to match a 2% average growth rate and a 2% share of R&D in gross output. Since a higher markup by itself implies higher profits, we need a

higher fixed cost (measured as a fraction of output) to lower profits back down and keep the equilibrium share of R&D at 2%. But since a higher fixed cost is associated with more volatile profits, a higher markup λ should lead to more procyclical R&D while a lower markup should lead to less procyclical R&D. I find that λ must be at least 1.15 for R&D to remain procyclical; for values of λ below this cutoff, there will be some values of k that prevail in the long-run for which $R_0(k) > R_1(k)$. When $\lambda = 1.10$, equilibrium R&D will be strictly higher in recessions for all limiting values of k . Thus, slightly smaller markups (in an absolute sense) dramatically lower the magnitude of the bias inherent to R&D.

Yet another important parameter is μ , which governs the persistence of fluctuations. If profits revert to their long-run average very quickly, firms will expect roughly similar profit streams regardless of whether they undertake R&D in recessions or in booms. This suggests that the bias will lessen as z_t becomes less persistent, i.e. at higher values of μ . In fact, μ does not need to be much higher for the procyclicality evident in Figure 4 to break down. If $\mu \geq .23$, which implies cycles on average last about 9 years rather than 10, there will be levels of k in the limit set for which $R_0(k) > R_1(k)$. A similar issue arises if we allow for asymmetric transition rates. Suppose the transition rate for z_t is given by μ_0 in recessions and μ_1 in booms. As an example, suppose we set $\mu_0 = .5$ and $\mu_1 = .125$, so that the average cycle still lasts 10 years on average, but recessions last only 2 years while booms last 8 years. To leave the unconditional standard deviation of productivity unchanged, I also recalibrate $Z_0 = .88$ and $Z_1 = 1.03$. Equilibrium R&D in this case turns out to be countercyclical. More generally, the bias appears to become weaker whenever at least one of the states is not very persistent. This pattern is consistent with the findings in Comin and Gertler (2006) that R&D is more procyclical at medium term frequencies than at high frequencies; shocks that cause profits to remain below-average for extended periods are more likely to depress R&D than high frequency shocks.

In sum, the finding that the procyclical bias can account for procyclical R&D appears somewhat fragile: a slightly lower fixed cost, a slightly lower markup, or a slightly less persistent process all lead to countercyclical equilibrium R&D. This does not deny the main point of the quantitative exercise, however, which is that the procyclical bias in R&D can be significant. Specifically, making profits more volatile by introducing fixed costs leads to large shifts in R&D activity towards booms. What the robustness analysis instead suggests is that there are probably additional forces that contribute to the procyclical pattern in R&D which I have not accounted for. As noted earlier, an important consideration I have ignored up to now is that labor supply varies over the cycle. I now modify the model to incorporate this possibility.

Rather than introduce leisure as a separate argument in the utility function and allowing

households to choose their labor supply, I follow Fatas (2000) in assuming labor varies exogenously over the cycle. Although this ignores certain welfare issues, this approach should still serve us well in gauging whether the model is able to account for cyclical patterns in R&D. In particular, if I did introduce a preference for leisure, I would have had to calibrate preferences so that the implied endogenous variation in hours will accord with the same volatility in hours I already match with exogenous variation. As long as variation in hours does not affect other aspects of the model such as the marginal utility of consumption, the implied equilibrium paths should be the same whether labor varies endogenously or exogenously. Empirically, the standard deviation of detrended hours is on the order of 6%, roughly the same as for total factor productivity. I therefore assume that when $z_t = Z_0$ then $L_t = 0.94 \cdot L$ and when $z_t = 1$ then $L_t = 1.06 \cdot L$, where L is the same as in Table 1.

Before examining whether this modification can salvage the ability of the model to generate procyclical R&D for slight perturbations of the parameters in Table 3, I first ask whether variable labor can generate procyclical R&D on its own. In particular, suppose I weaken the procyclical bias in R&D by setting the fixed cost F to zero. Could fluctuations in labor supply generate a procyclical pattern in R&D? The answer is no: when I set $F = 0$ and recalibrate L to match a steady-state growth rate of 2%, equilibrium R&D turns out to be countercyclical. Variable labor supply is compatible with procyclical R&D when $F = 0$ only when the standard deviation of labor is at least 11%, almost twice as much as labor productivity. Thus, when I minimize the extent of the procyclical bias in R&D in the model, it needs extraordinarily volatile labor supply to overcome the incentive to concentrate innovation activity in recessions. Although previous papers with variable labor have been able to generate procyclical equilibrium R&D without introducing fixed costs to make profits more volatile, e.g. Fatas (2000) and Comin and Gertler (2006), they also abstract from changes in the opportunity cost of R&D over the cycle. The findings here suggest that if this feature were incorporated, these models would have a difficult time generating procyclical R&D in equilibrium without further modifications.

Although variable labor supply cannot generate procyclical R&D in the absence of fixed costs, it can contribute to making equilibrium R&D procyclical when fixed costs lead to moderately volatile profits that are not quite enough to make equilibrium R&D procyclical. To gauge the ability of variable labor supply to amplify the procyclicality of R&D, I re-solve the model for the benchmark parameters in Table 3 allowing fluctuations in labor supply of 6%. Figure 5 compares the implied equilibrium path for R&D with the path in Figure 4 where labor is constant over time. With variable labor supply, the long-run range for k is wider, and for each k the variation in R&D dwarfs the variability when labor is assumed fixed. More precisely, whereas the standard deviation of the log share of R&D when labor is fixed over time is

0.139, with variable labor supply this standard deviation almost triples to 0.414. The empirical counterpart between 1953 and 2002, as noted earlier, is 0.137. This suggests that once we allow for variable labor supply, the model predicts overly volatile R&D. However, recall that fairly modest changes to some of the parameters can turn R&D countercyclical in the absence of labor supply shocks. For example, suppose we set $\mu_0 = 0.5$ and $\mu_1 = 0.125$, and in addition we lower the markup λ to 1.15. In the absence of variable labor supply, equilibrium R&D would be moderately countercyclical. However, with variable labor supply, equilibrium R&D is procyclical. The implied standard deviation of the log share of R&D is 0.193, much closer to its empirical counterpart. Allowing for slightly countercyclical markups can presumably bring down the volatility of R&D even lower so as to accord with what we observe in the data. Quantitatively, then, both the procyclical bias and variable labor supply play important roles in accounting for the procyclicality of R&D, since the only way for the model not to generate overly volatile R&D is if neither factor generates procyclical R&D by itself.

Finally, I turn to the question of the nature of optimal policy. When I assumed constant labor supply over time, I found that the optimal path for R&D was countercyclical. That is, it is desirable not just to reallocate some R&D from booms to recessions on the margin to offset the tendency of private agents to undertake too much R&D in booms – policymakers should actually reverse the timing of R&D, as the Schumpeterian view would suggest. However, when labor supply varies over time, it is no longer obvious that reversing R&D is still optimal. The intuition behind the Schumpeterian view is that we should allocate more resources to production in booms when the return to this activity is relatively high, and we should allocate more resources to R&D in recessions when the return to this activity is relatively high. When labor resources are fixed, this necessitates shifting labor resources between the two activities. But if labor resources are more abundant in booms, it might not be necessary to draw down resources from innovation to allocate more resources to production and take advantage of its temporarily high return. Similarly, if labor resources are more scarce in recessions, it might be too costly to take away resources from this activity and devote it to innovation. To gauge whether optimal policy remains countercyclical with variable labor supply, I once again resort to a perturbation argument. In particular, suppose aggregate productivity fluctuates between $Z_0 = 1 - \varepsilon$ and $Z_1 = 1 + \varepsilon$ and labor supply fluctuates between $L_0 = (1 - \varepsilon)L$ and $L_1 = (1 + \varepsilon)L$. Would a planner constrained to keep average R&D at the equilibrium level prefer to vary R&D with or against the cycle? Formally, if we represent the path for R&D by $R_0 = (1 - \zeta)R$ and $R_1 = (1 + \zeta)R$ for some ζ , would welfare be increasing in ζ or decreasing? For $\varepsilon = 0.01$, I find that the household's value function is now increasing in ζ in the neighborhood of $\zeta = 0$. Thus, for small shocks, a planner who is constrained to keep average R&D unchanged would opt to concentrate R&D in booms, not recessions. This result remains true for $\varepsilon = 0.06$. While

the optimal policy would still dictate subsidizing R&D in recessions to undo the tendency of short-sighted entrepreneurs to engage in too much R&D during these periods, it will no longer dictate concentrating R&D in recessions. Even though entrepreneurs fail to take full advantage of intertemporal substitution in R&D, procyclical R&D may not be inherently inefficient.

5. Conclusion

This paper examines why R&D activity is procyclical even though the Schumpeterian view implies recessions are the ideal time to undertake such activities. The main result of the paper is that because of dynamic externalities, there is a tendency in decentralized markets towards engaging in too much R&D in booms. Quantitatively, it appears that this bias could generate procyclical equilibrium R&D in an environment where the optimal path for R&D is countercyclical. Moreover, this bias implies that society would be better off reallocating some of its R&D from booms to recessions, precisely because it allows the economy to grow at a lower resource cost as the Schumpeterian view suggests. However, this conclusion does not necessarily imply that R&D should be countercyclical. In particular, certain considerations absent from the Schumpeterian view, such as variable labor supply, can make it desirable to undertake more R&D in booms, just not to the full extent it is in a decentralized market. The mere fact that R&D is procyclical is therefore not in itself *prima facie* evidence of inefficiency, although the bias described in the paper suggests that in practice there is probably too little intertemporal substitution of the type emphasized by the Schumpeterian view. In other words, in a world of cyclical fluctuations, growth is more costly than it needs to be. This highlights a welfare cost associated with business cycles that is distinct from the one described in Lucas (1987), or even the cost described in Barlevy (2004) which directly concerns the welfare consequences of business cycles through their effects on growth.

I close with a few remarks on some issues that the model presented here has ignored. One feature absent from the analysis is the possibility of a lag between when an initial discovery is made and when the idea it spawns can be put to practical use. In particular, firms in the model can expect with some probability to profit immediately from their R&D. If we think of R&D as a process of creating new ideas, this seems rather far-fetched. One could instead interpret R&D in the model as consisting mostly of actions involving the development of ideas as opposed to the creation of ideas; good ideas are always available, but entrepreneurs need to hire labor resources to make them practical. The increase in R&D during recessions should then be seen as an acceleration in the development of existing projects. This interpretation is reasonable, given that on average roughly 70% of R&D expenditures are at the development stage according

to NSF estimates. However, the notion that ideas are always readily available and are not the result of deliberate effort is unsatisfactory. If research activity declines in recessions, there shouldn't be as many ideas around for entrepreneurs to develop in booms. A more satisfactory approach would be to model research and development as separate stages, both of which require resources that could alternatively be used in production, and study their behavior over the cycle. Comin and Gertler (2006) make some progress in this direction. Although a model would be necessary to fully flesh out these issues, I suspect that both research and development would tend to be biased towards booms. Intuitively, even if an idea discovered at date t could only be implemented at date $t + T$, as long as shocks are persistent the occurrence of a recession today should increase the probability of lower profits T periods from now. The bias will be small for large T , and in industries such as pharmaceuticals where diffusion lags are on the order of 12-15 years the incentive to concentrate R&D in booms may be negligible. However, lags between discovery and implementation are considerably shorter in other sectors, such as software and computer equipment, and research in these may be markedly biased.

Another feature missing from the model is the possibility of strategically delaying using an idea that has already been developed. In particular, firms might choose to engage in R&D during recessions, when the opportunity cost of innovation is relatively low, but then wait to implement their ideas when profits are high. Shleifer (1986) and Francois and Lloyd Ellis (2003) both advocate this view. However, there is reason to believe that the possibility of strategic delay is not too important in practice. First, if entrepreneurs are sufficiently impatient, they may not want to wait to implement new technologies. Second, empirical evidence suggests firms are not very reluctant to release the results of their research. For example, Griliches (1990) reports that firms tend to take out patents – and thus publicize their new ideas – very soon after undertaking R&D efforts and long before they actually put their new ideas to use. In addition, strategic delay over the cycle would imply a mismatch between R&D activity and patenting over the business cycle, since R&D activity should peak in recessions while patents should peak in booms. However, Griliches reports that R&D and patents are highly synchronized over the business cycle. These findings suggest strategic delay is not widespread in practice. Nevertheless, given the emphasis in Shleifer (1986) on booms as periods of mass implementation, it is certainly worth incorporating this feature in future work.

Appendix A: Data Construction

NSF data

Data on real R&D expenditures performed and financed by industry are taken from National Patterns of Research Development Resources, 2003, Appendix B, Table B-10, column 23. Data on full-time equivalent employment from 1957-1999 were taken from Tables H-19 and B-25 in the NSF Industrial Research & Development Information System, which reports the number of employees for January of each year. I assign each January's observation to the previous year. Beyond 1999, I took data from various NSF reports. Data for January 2000 was taken from Table 1 in the NSF report "U.S. Industrial R&D Expenditures and R&D-to-Sales Ratio Reach Historical Highs in 2000." Data for January 2001 was taken from Table 1 in the NSF report "Largest Single-Year Decline in U.S. Industrial R&D Expenditures Reported for 2002." Data for January 2002 and January 2003 were taken from Table 3 in the NSF report "Increase in US Industrial R&D Expenditures Reported for 2003 Makes Up For Earlier Decline." All reports are available from the NSF website, www.nsf.gov.

Compustat

All Compustat variables were taken from the North American Industrial dataset. The particular data used correspond to the following original series in Compustat:

- R&D expenditures: item46
- Cash flow (after R&D expenditures): item14 + item18
- Cash flow (before R&D expenditures): item14 + item18 + item 46
- Assets: item6
- Liabilities: item181
- Short term debt (due within one year): item44
- Long term debt (due beyond one year): item 9
- Net value of capital stock: item8
- Market value of equity at year end: item24 \times item25 + item10

All series are deflated by the implicit GDP deflator to arrive at real values. To match the timing, I assigned the stock price to the year after it was reported, since it is quoted for the end of the year. To compute average growth in R&D, I took all firms that reported positive amounts of R&D at some point. From these, I excluded companies that were incorporated outside of the U.S., identified based on whether the variable FINC provided by Compustat is not equal to zero. In addition, I searched for duplicate entries for the same company (identified by the presence of PRE-FASB in their title) and eliminated those. Although only a few companies have duplicate records, some of these are large companies (e.g. Ford Motor Company, General Motors, General Electric) and allowing for duplicates can affect the weighted averages I compute. In computing the cyclicalities of stock prices, my selection criterion was analogous: I took all firms for which both a closing price and

Industry output

For the NBER manufacturing industry database compiled by Bartelsman, Becker, and Gray, I constructed real gross output as follows. For each industry, I compute nominal gross output as the sum of value added and material costs. To arrive at real gross output, this was divided by the shipments deflator provided for each industry.

Real gross output for certain industries is available from the Bureau of Economic Analysis from 1987 to 2004. The BEA breaks down industries into three-digit NAICS codes. However, some industries are available at a more

disaggregate level (for example, motor vehicles are distinguished from other transportation equipment, and legal services are distinguished from computer system design and miscellaneous professional and technical services) while some are only available at a more aggregated level (for example, real output is not reported for separate three digit industries for wholesale trade or retail trade). For all Compustat firms which I could assign a NAICS code to, I paired it with the relevant BEA classification that included that code and assigned it the correspond value for gross output. Value added is available from the Bureau of Economic Analysis from 1947 to 2004 and is classified for the same industry groups.

Appendix B: Proofs of Propositions

Proof of Proposition 1: For given values of $\{R_i\}_{i=0,1}$, the system given by (3.11) reduces to ordinary linear differential equations in $V(Z_i, M)$. Standard theorems ensure this system has a unique solution. Hence, starting with values for R_i , we can use the method of undetermined coefficients to find the unique value functions $V(Z_i, M)$ associated with a given pair (R_0, R_1) . I conjecture that the value function $V(\cdot, \cdot)$ takes the form

$$V(Z_i, M) = v_i \lambda^{M(1-\alpha)}$$

Differentiating this function with respect to M yields

$$\frac{\partial V}{\partial M} = (1 - \alpha) v_i \lambda^{M(1-\alpha)} \ln \lambda$$

which simplifies the differential equations above to a system of independent linear equations in the coefficients v_i :

$$\rho v_i = Z_i (L - R_i)^{1-\alpha} - F + \mu (v_{1-i} - v_i) + (1 - \alpha) v_i \phi R_i \ln \lambda$$

This yields a unique solution (v_0, v_1) as functions of (R_0, R_1) .

Since the RHS of (3.11) is strictly concave in R_i , the first order condition is both necessary and sufficient to characterize the optimal R_i . The first order condition is given by

$$-(1 - \alpha) Z_i \lambda^{M(1-\alpha)} (L - R_i)^{-\alpha} + \frac{\partial V}{\partial M} \phi \leq 0 \quad (5.1)$$

with equality if $R_i > 0$. Substituting the expression for $V(\cdot, \cdot)$, we obtain

$$R_i = \begin{cases} L - \left(\frac{Z_i}{v_i \phi \ln \lambda} \right)^{\frac{1}{\alpha}} & \text{if } v_i > \frac{Z_i}{\phi L^\alpha \ln \lambda} \\ 0 & \text{else} \end{cases} \quad (5.2)$$

If we substitute this expression into the asset equation (3.11), we obtain a pair of equations with v_{1-i} as a function of v_i that hold at the optimal R_i :

$$v_{1-i} = g_{1-i}(v_i) = \begin{cases} \frac{(\rho + \mu - (1 - \alpha) \phi L \ln \lambda)}{\mu} v_i - \frac{\alpha}{\mu} Z_i^{\frac{1}{\alpha}} (v_i \phi \ln \lambda)^{1 - \frac{1}{\alpha}} + \frac{F}{\mu} & \text{if } v_i > \frac{Z_i}{\phi L^\alpha \ln \lambda} \\ \frac{\rho + \mu}{\mu} v_i - \frac{\mu}{\mu} \frac{Z_i L^{1-\alpha}}{\mu} + \frac{F}{\mu} & \text{else} \end{cases}$$

The optimal program corresponds to any pair (v_0^*, v_1^*) which solves the equations

$$\begin{aligned} v_1^* &= g_1(v_0^*) \\ v_0^* &= g_0(v_1^*) \end{aligned}$$

The function $g_{1-i}(\cdot)$ is continuous and differentiable, since the left and right hand derivatives at $v_i = \frac{Z_i}{\phi L^\alpha \ln \lambda}$ are both equal to $\frac{\rho + \mu}{\mu}$. Since $\rho > (1 - \alpha) \phi L \ln \lambda$, it follows that $\frac{\partial g_{1-i}(v_i)}{\partial v_i} > 1$ for all v_i . To prove the system of

equations has a unique solution, note that since $\frac{dg_{1-i}}{dv_i} > 1 > 0$ for all v_i , the function $g_{1-i}(\cdot)$ must be invertible. An equilibrium therefore involves a value v_0^* such that $g_1(v_0^*) - g_0^{-1}(v_0^*) = 0$. Differentiating this condition with respect to v_0^* yields

$$\frac{d}{dx} [g_1(x) - g_0^{-1}(x)] = \frac{dg_1}{dx} - \left(\frac{dg_0}{dx}\right)^{-1} > 0$$

This monotonicity insures there is at most one value of v_0^* . To establish existence, note that $g_1(0) < 0$ while $g_0^{-1}(0) > 0$. Hence, $g_1(0) - g_0^{-1}(0) < 0$, and is finite. The fact that $\lim_{x \rightarrow \infty} \frac{dg_1}{dx} > 1 > \lim_{x \rightarrow \infty} \left(\frac{dg_0}{dx}\right)^{-1}$ implies $\frac{\partial}{\partial x} [g_1(x) - g_0^{-1}(x)]$ is strictly bounded away from 0, and so $g_1(x) - g_0^{-1}(x) \rightarrow \infty$ as $x \rightarrow \infty$. The existence of v_0^* follows from continuity.

Next, suppose that the optimal path dictates $R_i > 0$ for both i . I need to show $R_0 > R_1$. The proof proceeds in two steps. First, I argue that $v_1^* > v_0^*$. Since $R_i > 0$, the asset equations imply

$$\begin{aligned} v_{1-i}^* &= \frac{(\rho + \mu - (1 - \alpha)\phi L \ln \lambda)}{\mu} v_i^* - \frac{\alpha}{\mu} Z_i^{\frac{1}{\alpha}} (v_i^* \phi \ln \lambda)^{1 - \frac{1}{\alpha}} + \frac{F}{\mu} \\ &\equiv a v_i^* - b Z_i^{\frac{1}{\alpha}} (v_i^*)^{1 - \frac{1}{\alpha}} + \frac{F}{\mu} \end{aligned}$$

Consider the fixed point \hat{v}_i which solves

$$\hat{v}_i = a \hat{v}_i - b Z_i^{\frac{1}{\alpha}} (\hat{v}_i)^{1 - \frac{1}{\alpha}} + \frac{F}{\mu}$$

It is easy to show \hat{v}_i exists and is unique. Implicit differentiation implies

$$\frac{d\hat{v}_i}{dZ_i} = \frac{\frac{b}{\alpha} \left(\frac{\hat{v}_i}{Z_i}\right)^{1 - \frac{1}{\alpha}}}{(\alpha - 1) + \frac{1 - \alpha}{\alpha} b \left(\frac{Z_i}{\hat{v}_i}\right)^{\frac{1}{\alpha}}} > 0$$

so that $Z_0 < Z_1 \Rightarrow \hat{v}_0 < \hat{v}_1$. Since $\frac{dg_i^{-1}}{dx} < 1$, we know that for any $x < \hat{v}_1$, it follows that $x - g_0^{-1}(x) < 0$. Hence,

$$\begin{aligned} g_1(\hat{v}_0) - g_0^{-1}(\hat{v}_0) &= \hat{v}_0 - g_0^{-1}(\hat{v}_0) \\ &< 0 \end{aligned}$$

where the inequality uses the fact that $\hat{v}_0 > \hat{v}_1$. Since $g_1(v_0^*) - g_0^{-1}(v_0^*) = 0$ and $g_1(x) - g_0^{-1}(x)$ is increasing in x , it follows that $v_0^* > \hat{v}_0$. But since $\frac{dg_i^{-1}}{dx} > 1$, the fact that $g_1(\hat{v}_0) = \hat{v}_0$ implies $g_1(x) > x$ for any $x > \hat{v}_0$. Hence, $g_1(v_0^*) > v_0^*$. But since $v_1^* = g_1(v_0^*)$, it follows that $v_1^* > v_0^*$.

Next, I use the fact that $v_1^* > v_0^*$ to argue $\frac{v_1}{Z_1} < \frac{v_0}{Z_0}$, which is sufficient to establish $R_1 < R_0$ from the first-order condition above. Combining the equations $v_{1-i}^* = g_{1-i}(v_i^*)$ for both values yields the equation

$$a v_0^* - b Z_0^{\frac{1}{\alpha}} (v_0^*)^{1 - \frac{1}{\alpha}} - v_1^* = a v_1^* - b Z_1^{\frac{1}{\alpha}} (v_1^*)^{1 - \frac{1}{\alpha}} - v_0^*$$

which can be rearranged to yield

$$\frac{v_0^*}{v_1^*} = \frac{(a + 1) - b \left(\frac{Z_1}{v_1^*}\right)^{\frac{1}{\alpha}}}{(a + 1) - b \left(\frac{Z_0}{v_0^*}\right)^{\frac{1}{\alpha}}}$$

so that

$$v_1^* > v_0^* \Leftrightarrow \frac{v_1^*}{v_0^*} < \frac{Z_1}{Z_0}$$

But given the expression for R_i in (5.2), this implies $R_0 > R_1$. ■

Proof of Proposition 2: I begin by deriving analytical expressions for the value of a successful innovation v_i at each level of productivity Z_i . For any z_t -measurable function $X(\cdot)$, the integral

$$W_i(M_0) = E \left[\int_0^\infty \mathbb{I}_t \cdot \lambda^{(1-\alpha)M_t} X(z_t) e^{-\rho t} dt \mid z_0 = Z_i \right]$$

where $\dot{M}_t = \phi R$ and $R(\cdot)$ is z_t -measurable can be characterized by the recursive system of equations

$$(\rho + \mu) W_i(M) = \lambda^{(1-\alpha)M} X(Z_i) + \mu W_{1-i}(M) + \left[\frac{\partial W_i}{\partial M} - W_i(M) \right] \phi R_i$$

Using the method of undetermined coefficients, we can verify that $W_i(M) = w_i \lambda^{(1-\alpha)M}$ where

$$w_i = \frac{\omega(R_{1-i}) X(Z_i) + \mu X(Z_{1-i})}{\omega(R_i) \omega(R_{1-i}) - \mu^2}$$

and

$$\omega(R) = \rho + \mu + (1 - (1 - \alpha) \ln \lambda) \phi R$$

Using the expression for profits π_t and setting $y_i = L - R_i$, the value of a successful innovation v_i can be written as

$$v_i = \frac{(\lambda - 1) \left[\omega(L - y_{1-i}) y_i + \mu \frac{Z_{1-i}}{Z_i} y_{1-i}^{1-\alpha} y_i^\alpha \right] - [\omega(L - y_{1-i}) + \mu] \frac{\lambda y_i^\alpha F}{(1 - \alpha) Z_i}}{\omega(L - y_i) \omega(L - y_{1-i}) - \mu^2} \quad (5.3)$$

In an interior equilibrium, $v_0 = v_1 = \phi^{-1}$. We can easily rule out the case where $y_1 = 0$ in equilibrium, since this implies the marginal product of labor is infinite and hence must yield higher utility than using it in R&D. For $y_1 \neq 0$, we can rewrite v_0 and v_1 in terms of y_1 and $\xi = \frac{y_0}{y_1}$ as

$$v_0 = \frac{(\lambda - 1) \left[\omega(L - y_1) \xi + \mu \frac{Z_{1-i}}{Z_i} \xi^\alpha \right] y_1 - [\omega(L - y_1) + \mu] \frac{\lambda \xi^\alpha F y_1^\alpha}{(1 - \alpha) Z_0}}{\omega(L - y_1) \omega(L - \xi y_1) - \mu^2} \quad (5.4)$$

$$v_1 = \frac{(\lambda - 1) \left[\omega(L - \xi y_1) + \mu \frac{Z_0}{Z_1} \xi^{1-\alpha} \right] y_1 - [\omega(L - \xi y_1) + \mu] \frac{\lambda F y_1^\alpha}{(1 - \alpha) Z_1}}{\omega(L - y_1) \omega(L - \xi y_1) - \mu^2} \quad (5.5)$$

A necessary condition for equilibrium is that $v_0 - v_1 = 0$, which can be rearranged to yield the condition

$$\frac{\lambda F y_1^{\alpha-1}}{Z_1 (\lambda - 1) (1 - \alpha)} (A_0 - A_1(y_1) \xi - A_2(y_1) \xi^\alpha) - 1 + \xi + h(\xi) = 0 \quad (5.6)$$

where

$$\begin{aligned} A_0 &= \frac{\omega(L) + \mu}{\omega(L)} \\ A_1(y_1) &= \frac{(1 - (1 - \alpha) \ln \lambda) \phi y_1}{\omega(L)} \\ A_2(y_1) &= \frac{\omega(L - y_1) + \mu}{\omega(L)} \frac{Z_1}{Z_0} \end{aligned}$$

and

$$h(\xi) = \frac{\mu}{\omega(L)} \left[\frac{Z_1}{Z_0} \xi^\alpha - \frac{Z_0}{Z_1} \xi^{1-\alpha} \right]$$

For convenience, let us define

$$Q(\xi, y_1) \equiv \frac{\lambda F y_1^{\alpha-1}}{Z_1 (\lambda - 1) (1 - \alpha)} (A_0 - A_1(y_1) \xi - A_2(y_1) \xi^\alpha) - 1 + \xi + h(\xi)$$

so that we can now rewrite (5.6) more compactly as $Q(\xi, y_1) = 0$.

Step 1: I first show that there exist values F^* and \bar{F} such that if $F > F^*$ and $F < \bar{F}$, there exists an F^* such that for all $F > F^*$, there exists a pair (y_0, y_1) where $y_0 > y_1$ and $Q\left(\frac{y_0}{y_1}, y_1\right) = 0$. This implies the associated $R_1 > R_0$. I begin with the following lemma:

Lemma: Suppose $\ln \lambda < (1 - \alpha)^{-1}$. For any $F > 0$, there exists a unique $\hat{y}_F > 0$ such that $v_0(\hat{y}_F, \hat{y}_F) = v_1(\hat{y}_F, \hat{y}_F)$. Moreover, there exists an $F^* > 0$ such that $v_i(\hat{y}_F, \hat{y}_F) < \phi^{-1}$ for $F < F^*$ and $v_i(\hat{y}_F, \hat{y}_F) > \phi^{-1}$ for $F > F^*$.

Proof: Consider the equation $v_0(y, y) = v_1(y, y)$. Substituting in and rearranging yields

$$\mu(Z_1 + Z_0) \frac{y^{1-\alpha}}{F} + \frac{\lambda}{(\lambda - 1)(1 - \alpha)} (1 - (1 - \alpha) \ln \lambda) \phi y = \frac{\lambda(\omega(L) + \mu)}{(\lambda - 1)(1 - \alpha)} \quad (5.7)$$

Since $\ln \lambda < (1 - \alpha)^{-1}$, the LHS of this equation is monotonically increasing in y given (3.20) and ranges from 0 to ∞ as y ranges from 0 to ∞ . Since the RHS above is strictly positive, there exists a unique value \hat{y}_F for which the equation is satisfied. Moreover, this \hat{y}_F is monotonically increasing in F . Taking limits, $\hat{y}_F \rightarrow 0$ as $F \rightarrow 0$, while $\hat{y}_F \rightarrow L + \frac{\rho + 2\mu}{(1 - (1 - \alpha) \ln \lambda) \phi}$ as $F \rightarrow \infty$, at which point $\omega(L - \hat{y}_F) = -\mu$.

At $y_0 = y_1 = y$, the value of a successful innovation is given by

$$v_i(y, y) = (\lambda - 1) \frac{\left(\omega(L - y) + \mu \frac{Z_{1-i}}{Z_i} \right) y - \frac{(\omega(L - y) + \mu) \lambda F y^\alpha}{(\lambda - 1)(1 - \alpha) Z_i}}{\omega^2(L - y) - \mu^2}$$

Setting $y = \hat{y}_F$ and using (5.7), this expression reduces to

$$\begin{aligned} v_i(\hat{y}_F, \hat{y}_F) &= (\lambda - 1) \frac{\left(\omega(L - \hat{y}_F) + \mu \frac{Z_{1-i}}{Z_i} \right) \hat{y}_F - \frac{\mu(Z_1 + Z_0) \hat{y}_F}{Z_i}}{\omega^2(L - \hat{y}_F) - \mu^2} \\ &= \frac{(\lambda - 1) \hat{y}_F}{\omega(L - \hat{y}_F) + \mu} \end{aligned}$$

Hence, $v_i(\hat{y}_F, \hat{y}_F)$ is monotonically increasing in \hat{y}_F , which in turn is monotonically increasing in F . As noted above, $\hat{y}_F \in \left[0, L + \frac{\rho + 2\mu}{(1 - (1 - \alpha) \ln \lambda) \phi} \right)$, which implies v_i ranges between 0 and ∞ . The existence of F^* follows from continuity. ■

I now prove the claim using F^* as defined in the lemma. For any value of F , define the set

$$\Omega_F = \left\{ (y_0, y_1) \mid y_1 > 0, Q\left(\frac{y_0}{y_1}, y_1\right) = 0 \right\}$$

I will show that for $F > F^*$ as defined in the lemma, there exists an element $(y_0, y_1) \in \Omega_F$ such that $y_0 > y_1$ and $v_0(y_0, y_1) = v_1(y_0, y_1) = \phi^{-1}$. Take any value of $\xi \geq 1$, and consider the values of y_1 for which $(\xi y_1, y_1) \in \Omega_F$. For $\xi_1 = 1$, we know from the lemma that there exists a unique such value, namely $y_1 = \widehat{y}_F$. I now argue that for any $\xi > 1$, there exists a unique value \widehat{y}_ξ such that $(\xi \widehat{y}_\xi, \widehat{y}_\xi) \in \Omega_F$.

I begin by rearranging $Q(\xi, y_1)$ to get

$$Q(\xi, y_1) = \frac{\lambda F}{Z_1(1-\alpha)(\lambda-1)} \left(\left[\frac{Z_1}{Z_0} \xi^\alpha - \xi \right] \frac{(1-(1-\alpha)\ln\lambda)\phi}{\omega(L)} y_1^\alpha + \frac{\omega(L)+\mu}{\omega(L)} \left[1 - \frac{Z_1}{Z_0} \xi^\alpha \right] y_1^{\alpha-1} \right) - H(\xi)$$

where $H(\xi) = h(\xi) - 1 + \xi$ does not depend on y_1 . For $\xi \geq 1$, the coefficient on y_1^α is positive, while the coefficient on $y_1^{\alpha-1}$ is negative. Hence, for a fixed ξ , $Q(\xi, y_1)$ is monotonically increasing in y_1 . Next, if we take the limit as $y_1 \rightarrow 0$, the limit tends to $-\infty$ since $y_1^{\alpha-1}$ grows arbitrarily large and its coefficient is negative. Finally, if we take the limit as $y_1 \rightarrow \infty$, the limit tends to ∞ since y_1^α grows arbitrarily and its coefficient is positive. The existence and uniqueness of \widehat{y}_ξ follow from continuity. Continuity also implies that the path $(\xi \widehat{y}_\xi, \widehat{y}_\xi)$ for $\xi \geq 1$ forms a continuous path in (y_0, y_1) space.

As $\xi \rightarrow \infty$, the path $(\xi \widehat{y}_\xi, \widehat{y}_\xi)$ converges to $(y_0, 0)$ for some y_0 . This is illustrated graphically in Figure B1, where the heavy line depicts a segment of the set Ω_F . A fixed ξ corresponds to a ray in (y_0, y_1) space, and increasing ξ implies rotating this ray clockwise towards the y_0 axis. Using the expressions for $v_0(y_0, y_1)$ and $v_1(y_0, y_1)$ and substituting in $y_1 = 0$ implies $v_1 = 0$, and

$$v_0 = \frac{(\lambda-1)\omega(L)y_0 - \frac{\omega(L)+\mu\lambda F}{(1-\alpha)Z_0}y_0^\alpha}{\omega(L-y_0)\omega(L) - \mu^2}$$

There are two values of y_0 at which $v_0 = v_1$, namely $y_0 = 0$ and

$$y_0 = \widetilde{y}_0 \equiv \left[\frac{(\omega(L)+\mu)\lambda F}{\omega(L)(\lambda-1)(1-\alpha)Z_0} \right]^{\frac{1}{1-\alpha}}$$

In what follows, we will want to insure that $\omega(L - \widetilde{y}_0)\omega(L) - \mu^2 > 0$, i.e. the present discounted values of profits are well-defined at the point $(\widetilde{y}_0, 0)$. From the value of \widetilde{y}_0 above, this implies

$$F < (1-\alpha)Z_0 \frac{\lambda-1}{\lambda} \left[\frac{\omega(L)}{\omega(L)+\mu} \right]^\alpha \left[\frac{\omega(L)-\mu}{\phi(1-(1-\alpha)\ln\lambda)} \right]^{1-\alpha} \equiv \overline{F}$$

Next, I argue that $\lim_{\xi \rightarrow \infty} (\xi \widehat{y}_\xi, \widehat{y}_\xi) = (\widetilde{y}_0, 0)$. By continuity, the only two possible limits are $(0, 0)$ and $(\widetilde{y}_0, 0)$. To rule out $(0, 0)$, suppose this were the limit. Then by continuity, $Q(y_0, 0) \geq 0$ for all $y_0 > 0$, since for any ξ , $Q(\xi y_1, y_1) > 0$ for any $y_1 > \widehat{y}_\xi$. But since $Q(y_0, 0) < 0$ for all $y_1 \in (0, \widetilde{y}_0)$, this yields a contradiction.

Finally, I need to show that there exists a $\xi > 1$ for which $v_0(\xi \widehat{y}_\xi, \widehat{y}_\xi) = v_1(\xi \widehat{y}_\xi, \widehat{y}_\xi) = \phi^{-1}$. By definition, $v_0(\xi \widehat{y}_\xi, \widehat{y}_\xi) = v_1(\xi \widehat{y}_\xi, \widehat{y}_\xi)$ for all ξ . The issue is whether this joint value is equal to ϕ^{-1} . To determine this, note that at $\xi = 1$, $\widehat{y}_\xi = \widehat{y}_F$. If $F > F^*$, then $v_0(\widehat{y}_F, \widehat{y}_F) > \phi^{-1}$. Moreover, in the limit as $\xi \rightarrow \infty$, the value $v_0(\xi \widehat{y}_\xi, \widehat{y}_\xi) \rightarrow 0$. If $v_0(\xi \widehat{y}_\xi, \widehat{y}_\xi) < \infty$ for all $\xi \geq 1$, then by continuity there exists a value of $\xi > 1$ such that $v_0(\widehat{y}_F, \widehat{y}_F) = \phi^{-1}$. However, $v_0(\xi \widehat{y}_\xi, \widehat{y}_\xi)$ may not be bounded for all $\xi \geq 1$, and as we trace out ξ , the value of $v_0(\xi \widehat{y}_\xi, \widehat{y}_\xi)$ might rise to $+\infty$ and jump to $-\infty$ if the denominator in v_0 turns from being positive to negative. In this case, $v_0(\xi \widehat{y}_\xi, \widehat{y}_\xi)$ might reach 0 without ever equalling ϕ^{-1} . However, if that is the case, since $F < \overline{F}$ then by continuity there exists a $\widehat{\xi}$ such that $\omega(L - \xi \widehat{y}_\xi)\omega(L - \widehat{y}_\xi) - \mu^2 > 0$ for all $\xi > \widehat{\xi}$. Thus, $v_0(\xi \widehat{y}_\xi, \widehat{y}_\xi)$ must traverse from $+\infty$, the last point at which the denominator switched from being negative to being positive, down to 0, and so at some point by continuity it must equal ϕ^{-1} . ■

Step 2: I next show that if $F = 0$, then any solution to the system of equations $v_0(y_0, y_1) = v_1(y_0, y_1) = \phi^{-1}$ must have $y_0 < y_1$ and hence $R_0 > R_1$. Note that when $F = 0$, this equation reduces to

$$1 - \xi = h(\xi) \tag{5.8}$$

It will suffice to show that any solution ξ^* which solves (5.8) lies in the unit interval $(0, 1)$, since this implies $y_0 < y_1$, and hence that $R_0 > R_1$ if $F = 0$.

I begin by establishing there exists a $\xi^* \in (0, 1)$ which solves (5.8). If $\xi = 0$, we have

$$1 - \xi = 1 > 0 = h(\xi)$$

while if $\xi = 1$, we have

$$1 - \xi = 0 < \frac{\mu}{\omega(L)} \left[\frac{Z_1}{Z_0} - \frac{Z_0}{Z_1} \right] = h(\xi)$$

where we use the fact that $\omega(L) > 0$ given that $\rho > (1 - \alpha)\phi L \ln \lambda$. The existence of ξ^* follows from continuity.

To prove there is no solution $\xi^* > 0$, differentiate $h(\cdot)$ to obtain

$$h'(\xi) = \frac{\mu}{\omega(L)} \left[\alpha \frac{Z_1}{Z_0} \xi^{\alpha-1} - (1 - \alpha) \frac{Z_0}{Z_1} \xi^{-\alpha} \right]$$

For $\xi \geq 1$, we have

$$\begin{aligned} \alpha \frac{Z_1}{Z_0} \xi^{\alpha-1} - (1 - \alpha) \frac{Z_0}{Z_1} \xi^{-\alpha} &> -(1 - \alpha) \frac{Z_0}{Z_1} \xi^{-\alpha} \\ &> -1 \end{aligned}$$

Since at $\xi = 1$, $h(\xi) > 1 - \xi$, a necessary condition for there to exist a $\xi^* > 1$ such that $1 - \xi^* = h(\xi^*)$ is that there exists a $\xi > 1$ such that $h'(\xi) < -1$. So there can be no $\xi > 1$ for which $1 - \xi = h(\xi)$. Hence, $R_0 > R_1$ in any interior equilibrium. ■

Table 1: Correlation of R&D with industry output measures

Dependent variable: growth in real R&D

	4 digit SIC code manufacturing only gross output 1959-1996	3 digit NAICS code all industries gross output 1988-2004	3 digit NAICS code all industries value added 1978-2004*
output growth (no time fixed effects)	0.2102 0.0262	0.9229 0.0468	0.4140 0.4140
output growth (with time fixed effects)	0.1354 0.0292	0.3631 0.0649	0.1861 0.0365
Firm fixed effects	Yes	Yes	Yes
N	29,618	43,557	62,715
# of firms	3,454	6,160	7,719
# of industries	130	51	54

*Value added data is available from 1950 on for 2-digit NAICS codes. For industries where the 3-digit code accounted for virtually all of the value added in the respective 2-digit industry, I used the 2-digit output and tracked data back to 1950. Observations prior to 1978 account for less than 2% of my observations.

Table 2: Partial correlation of R&D with industry output measures after controlling for balance sheet variables*

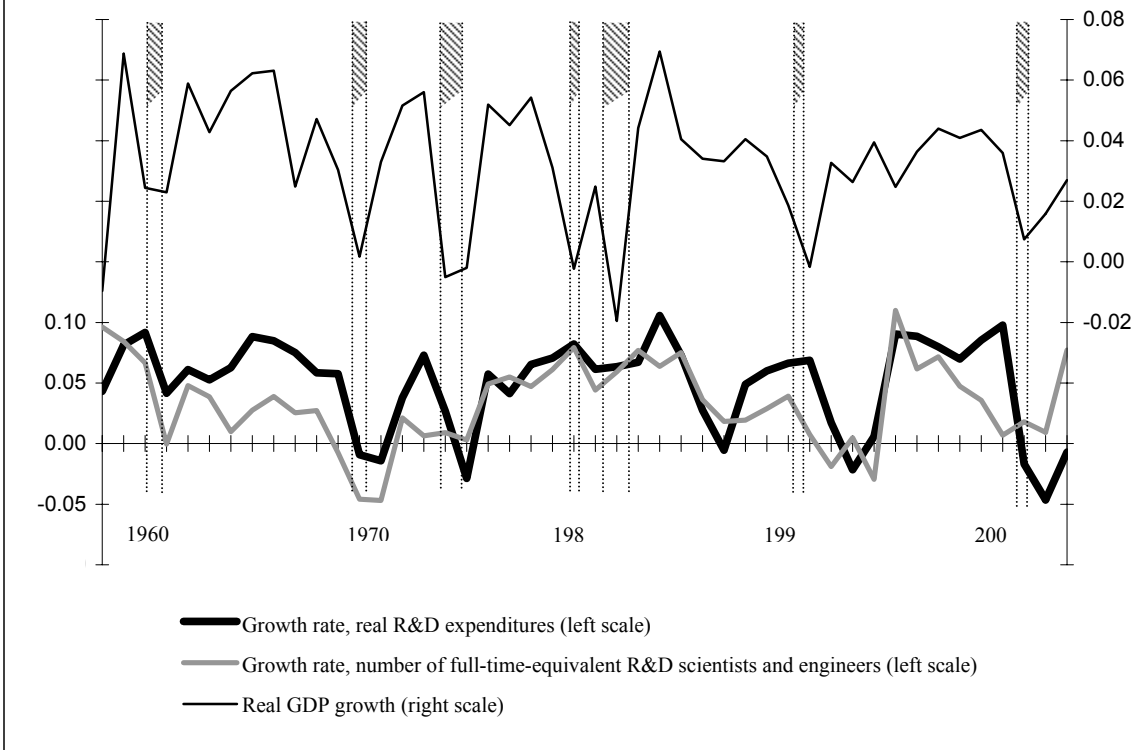
Dependent variable: growth in real R&D

	4 digit SIC code manufacturing only gross output 1959-1996	3 digit NAICS code all industries gross output 1988-2004	3 digit NAICS code all industries value added 1978-2004**
output growth (no time fixed effects)	0.2057 0.0270	0.9116 0.0475	0.4115 0.0307
output growth (with time fixed effects)	0.1331 0.0301	0.3538 0.0657	0.1875 0.0370
Firm fixed effects	Yes	Yes	Yes
N	28,389	42,598	61,336
# of firms	3,429	6,124	7,674
# of industries	129	51	55

*Balance sheet variables included in all regressions (but not reported in the table) are current year cash flow (before R&D expenditures), lagged cash flow (after R&D expenditures), and current and lagged assets, liabilities, plant property and equipment, short-term debt and long-term debt, all deflated by the implicit GDP deflator.

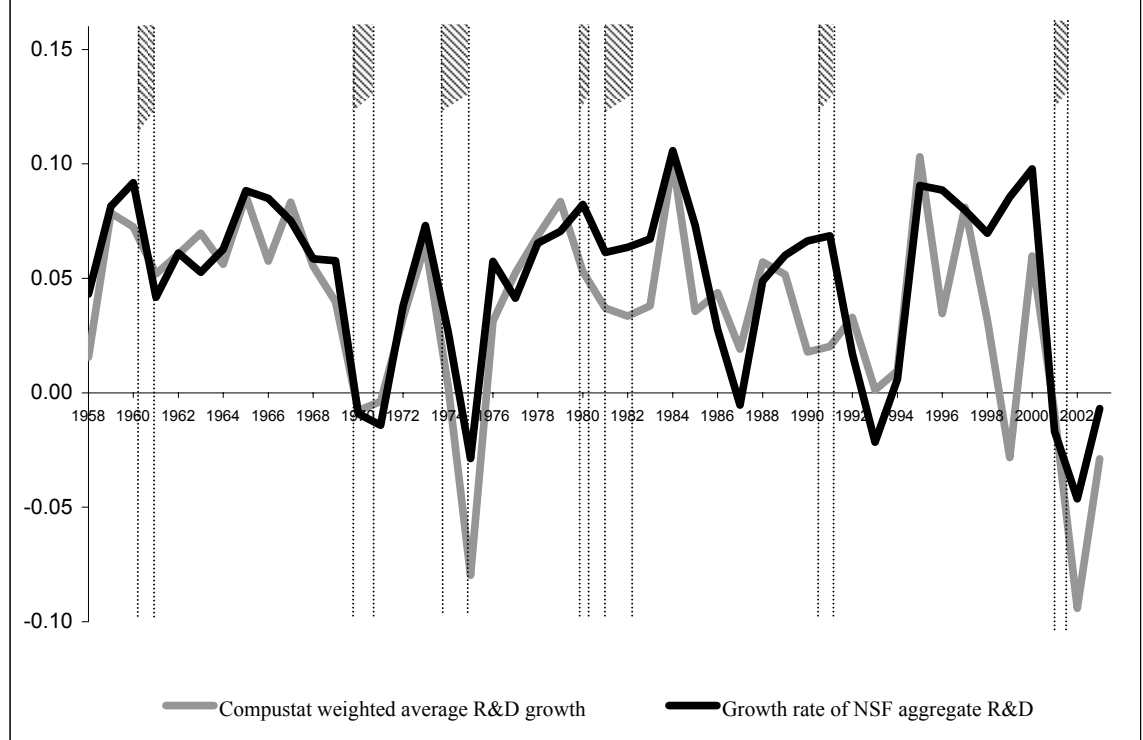
**Value added data is available from 1950 on for 2-digit NAICS codes. For industries where the 3-digit code accounted for virtually all of the value added in the respective 2-digit industry, I used the 2-digit output and tracked data back to 1950. Observations prior to 1978 account for

Figure 1: Measures of R&D over the Business Cycle




Source: NSF

Figure 2: R&D Growth, NSF vs. Compustat



$\beta_{R\&D}$ – Cyclicity
of R&D



β_{stock} – Cyclicity of stock prices

Figure 3: Cyclicity of stock prices (β_{stock}) and R&D ($\beta_{R\&D}$) across industries
(Observations weighted by the standard error of the estimate of R&D cyclicity in each industry)




Figure 4: Equilibrium path for R&D (as a function of k) for parameters in Table 3




Figure 5: Equilibrium R&D for parameters in Table 3, fixed and variable labor




Figure B1: The set $\Omega_F = \{(y_0, y_1) \mid Q(y_0/y_1, y_1) = 0\}$

References

- [1] Aghion Philippe, George-Marios Angeletos, Abhijit Banerjee, and Kalina Manova, 2005. "Volatility and Growth: Credit Constraints and Productivity-Enhancing Investment" NBER Working Paper #11349.
- [2] Aghion, Philippe and Gilles Saint Paul, 1998. "Virtues of Bad Times: Interaction between Productivity Growth and Economic Fluctuations" *Macroeconomic Dynamics*, September, 2(3), p322-44.
- [3] Arrow, Kenneth, 1962. "Economic Welfare and the Allocation of Resources for Invention" in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, edited by Richard Nelson, Princeton, NJ: Princeton University Press, p609-626.
- [4] Barlevy, Gadi, 2004. "The Cost of Business Cycles under Endogenous Growth" *American Economic Review*, September, 93(4), p964-90.
- [5] Barlevy, Gadi and Dani Tsiddon, 2006. "Earnings Inequality and the Business Cycle" *European Economic Review*, 50(1), January, p55-89.
- [6] Basu, Susanto, 1996. "Procyclical Productivity: Increasing Returns or Cyclical Utilization?" *Quarterly Journal of Economics*, 111(3), August, p719-51.
- [7] Benhabib, Jess and Roger Farmer, 1994. "Indeterminacy and Increasing Returns" *Journal of Economic Theory*, 63(1), June, p19-41.
- [8] Caballero, Ricardo and Mohamad Hammour, 1994. "The Cleansing Effect of Recessions" *American Economic Review*, December, 84(5), p1350-68.
- [9] Canton, Eric and Harald Uhlig, 1999. "Growth and the Cycle: Creative Destruction versus Entrenchment" *Journal of Economics*, 69(3), p239-66.
- [10] Comin, Diego and Mark Gertler, 2006. "Medium-Term Business Cycles" *American Economic Review*, September, 96(3), June, p523-51.
- [11] Cooper, Russell and John Haltiwanger, 1993. "The Aggregate Implications of Machine Replacement: Theory and Evidence" *American Economic Review*, June, 83(3), p181-186.
- [12] Dellas, Harris and Plutarchos Sakellaris, 2003. "On the Cyclical Behavior of Schooling: Theory and Evidence" *Oxford Economic Papers*, January, 55(1), p148-72.
- [13] DeJong, David and Beth Ingram, 2001. "The Cyclical Behavior of Skill Acquisition" *Review of Economic Dynamics*, July, 4(3), p536-61.

- [14] Fatas, Antonio, 2000. "Do Business Cycles Cast Long Shadows? Short-Run Persistence and Economic Growth" *Journal of Economic Growth*, 5(2), June, p147-62.
- [15] Francois, Patrick and Huw Lloyd-Ellis, 2003. "Animal Spirits Through Creative Destruction" *American Economic Review*, June, 93(3), p530-50.
- [16] Gomes, Joao, Jeremy Greenwood, and Sergio Rebelo, 2001. "Equilibrium Unemployment" *Journal of Monetary Economics*, 48(1), p109-52.
- [17] Griliches, Zvi, 1984. "Comment on R&D and Innovation: Some Empirical Findings" in *R&D, Patents, and Productivity*, edited by Zvi Griliches, Chicago: University of Chicago Press, p148-149
- [18] Griliches, Zvi, 1990. "Patent Statistics as Economic Indicators: a Survey" *Journal of Economic Literature*, December, 28(4), p1661-1707.
- [19] Grossman, Gene and Elhanan Helpman, 1991. "Quality Ladders in the Theory of Growth" *Review of Economic Studies*, 58(1), January, p43-61.
- [20] Hall, Robert 1991. "Labor Demand, Labor Supply, and Employment Volatility" *NBER Macroeconomics Annual*. Cambridge, Mass: MIT Press.
- [21] Himmelberg, Charles and Bruce Petersen, 1994. "R&D and Internal Finance: A Panel Study of Small Firms in High-Tech Industries" *Review of Economics and Statistics*, 76(1), February, p38-51.
- [22] Judd, Kenneth, 1998. *Numerical Methods in Economics*. Cambridge, MA: MIT Press.
- [23] Lucas, Robert, 1987. *Models of Business Cycles*. Oxford: Basil Blackwell.
- [24] Mortensen, Dale and Chris Pissarides, 1994. "Job Creation and Job Destruction in the Theory of Unemployment" *Review of Economic Studies*, 61(3), July, p397-415.
- [25] Ramey, Valerie, 1991. "Comment on Markups and the Business Cycle" *NBER Macroeconomics Annual*, vol 6, p134-9.
- [26] Rafferty, Matthew and Mark Funk "Demand Shocks and Firm-financed R&D Expenditures," *Applied Economics*, 36(14), August, p1529-1536.
- [27] Rivera-Batiz, Luis and Paul Romer, 1991. "Economic Integration and Endogenous Growth," *Quarterly Journal of Economics*, 106(2), May, p531-55.
- [28] Rotemberg, Julio and Michael Woodford, 1999. "The Cyclical Behavior of Prices and Costs" in *Handbook of Macroeconomics* (1B), vol. 15, edited by John Taylor and Michael Woodford. Amsterdam: Elsevier Science, p1051-1135.

- [29] Shleifer, Andrei, 1986. "Implementation Cycles" *Journal of Political Economy*, 94(6), December, p1163-90.