

MARKET STRUCTURE, ORGANIZATIONAL STRUCTURE, AND R&D DIVERSITY

September 26, 2002

Joseph Farrell

*Department of Economics
University of California at Berkeley
Berkeley, California 94720-3880
farrell@econ.berkeley.edu*

Richard J. Gilbert

*Department of Economics
University of California at Berkeley
Berkeley, California 94720-3880
gilbert@econ.berkeley.edu*

and

Michael L. Katz*

*Walter A. Haas School of Business
University of California at Berkeley
Berkeley, California 94720-1900
katz@haas.berkeley.edu*

Prepared for “Imperfect Economics: Essays in Honor of Joseph Stiglitz,” forthcoming,
MIT Press

* Katz is the Deputy Assistant Attorney General for Economic Analysis at the U.S. Department of Justice. The views expressed are those of the authors and do not necessarily represent the views of the U.S. Department of Justice. The essay does not draw on any confidential materials or information. Farrell thanks SIEPR for research support through a Cain Fellowship.

ABSTRACT

We examine the effects of market structure and the internal organization of firms on equilibrium R&D projects. We compare a monopolist's choice of R&D portfolio to that of a welfare maximizer. We next show that Sah and Stiglitz's finding that the market portfolio of R&D is independent of the number of firms under Bertrand competition extends to neither Cournot oligopoly nor a cartel. We also show that the ability of firms to pre-empt R&D by rivals along particular research paths can lead to socially excessive R&D diversification. Lastly, using Sah and Stiglitz's definition of hierarchy, we establish conditions under which larger hierarchies invest in smaller portfolios.

Keywords: Internal organization, market power, research and development.

JEL Numbers: L1, L12, and O31

I. INTRODUCTION

An extensive literature examines the relationship between market structure and research and development (R&D) activities, primarily comparing privately and socially optimal investment levels, or innovative *efforts*, along a single dimension. In practice, however, there are many different research paths that a firm might pursue. Hence, progress also hinges on the research *directions* chosen by firms and the extent to which firms *diversify* their approaches to R&D by pursuing multiple directions simultaneously.

Many people suspect that a diversity of approaches goes along with a diversity of approachers, and antitrust authorities have expressed concern over the effects of mergers, for instance, on research diversity.¹ But it is not immediately clear why this should be a concern. After all, a single organization can have incentives to pursue diverse approaches. Put simply, why can't a single organization do everything a group of firms can do, plus take advantage of coordination where beneficial? We answer this question by exploring how market structure and organizational structure affect the social portfolio of R&D approaches.

Joe Stiglitz is no stranger to these issues. Two lines of Joe's research are directly on point. One is his work on R&D competition when each firm can pursue multiple paths simultaneously. Sah and Stiglitz (1987) established conditions under which the total number of paths pursued in a market is independent of the number of firms. While provocative, this result applies to a limited set of market structures. Below, we examine several additional settings. We call the influence of market structure on diversity the role of "external factors."

¹ See Robinson (1999) and Rubinfeld and Hoven (2001) on the proposed merger of Lockheed-Northrop and Grumman.

Sah and Stiglitz (1985, 1986, and 1988) also provided a second line of directly relevant work, this time on organizational design. Sah and Stiglitz explored how the architecture of an economic organization—who collects information, with whom it is communicated, and how decisions are made—affects the quality of decision making. Below, we build on their model of hierarchical architectures to examine how firms’ choices of internal organization affect R&D diversity. We call this the role of “internal factors.”

The remainder of this chapter is organized as follows. Section II lays out several assumptions maintained throughout the analysis. Section III establishes a benchmark by examining a firm that is the sole potential innovator and is a unitary, profit-maximizing decision maker. Section IV examines external factors by considering the interaction of several unitary, profit-maximizing decision makers. Specifically, it examines whether highly concentrated industries will predictably give rise to different R&D portfolios than less concentrated industries. Section V turns to internal issues and examines how organizational choices interact with market structure to affect the equilibrium R&D portfolio.

II. OUR CHOICE OF RESEARCH PATH

Many factors influence firms’ R&D strategies, and to identify diversification incentives we begin by eliminating or holding constant potentially confounding influences. First, with product innovation, the value of R&D diversity as a response to uncertainty may become confounded with the value of product variety.² To avoid this problem, we restrict attention to process innovations and assume throughout that each firm has a product of fixed characteristics.

² There is a large literature on deterministic product selection, which establishes that a monopolist may choose greater or less variety than is socially optimal and than would a multi-firm market. Joe Stiglitz is a prominent contributor to this literature. See Dixit and Stiglitz (1977). See also Katz (1980).

Second, we restrict attention to situations in which R&D activities along different paths are substitutes for one another. We do this because complementarities could introduce economies of scope that can separately influence the choice of R&D portfolios. We define a project as a level of effort along a particular R&D path. We assume that each project, j , gives rise to a stand-alone cost level, c^j , and that a firm's unit cost is equal to the lowest realized cost level over the set of projects the firm has undertaken: firm i 's unit production costs are $c_i = \min \{ c^1, c^2, \dots, c^k \}$ when firm i undertakes projects 1 through k .

Third, we also want projects to be substitutes rather than complements in terms of the interaction of effort levels across projects. Throughout most of the analysis we assume that there are no technological spillovers within or across firms: the distribution of results from project j undertaken by firm i is independent of the efforts devoted to other projects by firm i or its rivals.³

Fourth, we want to distinguish the incentive to diversify R&D paths from incentives for firms to choose different types of R&D projects. The economics literature suggests that incumbent firms with market power have stronger incentives than new competitors to invest in incremental innovations.⁴ Furthermore, a large business strategy literature suggests that incumbent firms tend to look to innovate in areas they already know.⁵ While these biases are themselves of considerable interest, we focus on diversification narrowly defined in order clearly to identify various forces at work. Thus, we assume that, conditional on the level of effort, each

³ Notice that this assumption does not imply that the resulting cost levels of different projects are uncorrelated conditional on the effort levels.

⁴ See Gilbert and Newbery (1982) and Katz and Shapiro (1987).

⁵ Christensen (1997) argues that market leaders in the computer disk drive industry repeatedly failed to embrace the next technological revolution because they focused too much on meeting customer demands for incremental improvements to existing technologies. Henderson (1993) invokes organizational factors and economic incentives to explain why new entrants in semiconductor photolithography equipment often leapfrogged existing market leaders.

of the different substitute projects has the same cost distribution as any other.

III. UNITARY MONOPOLY

We begin by comparing the diversification incentives of a profit-maximizing monopolist with those of a total-surplus-maximizing decision maker. Consider a set of portfolios of R&D projects where each portfolio entails the same level of aggregate R&D expenditures. Each portfolio gives rise to a distribution function for the firm's cost level. A profit-maximizing decision maker chooses the portfolio that maximizes the expected value of $\pi(c)$, the monopoly profits earned with unit production cost level c . A welfare-maximizing decision maker chooses the portfolio that maximizes the expected value of $W(c)$, the sum of profits and consumer surplus when the firm chooses the monopoly price corresponding to marginal cost c . Note that both profits and welfare are decreasing functions of c . Moreover, profits are convex in c ; a profit maximizer has incentives to take risks with R&D. Does the firm do so to an efficient degree?

As is well known, a monopolist facing a downward-sloping demand function undertakes too little cost-reducing R&D because a fall in marginal costs leads to an equilibrium increase in consumer surplus. There is not a similarly general result for the monopolist's attitude towards risk. Define $x(p)$ as the quantity demanded at price p , $r(c)$ as the monopoly price given costs c , and $x^*(c) \equiv x(r(c))$. By the envelope theorem, $\pi'(c) = -x^*(c)$. The change in total surplus is

$$W'(c) = (r - c) \frac{dx^*}{dc} - x^*(c). \text{ Using the Lerner condition, } W'(c) = -x^*(c)(1 + \varphi(c)), \text{ where } \varphi(c) \equiv \frac{dr}{dc} \geq 0 \text{ is the pass-through rate. Suppose that } \varphi(c) \text{ is constant over } c, \text{ as is the case with linear,}$$

constant elasticity, or rectangular (*i.e.*, all consumers have the same reservation price) demand.

Integrating the expression for the derivatives of profits and welfare demonstrates that

$W(c) = \alpha + (1 + \varphi)\pi(c)$, where α is a constant. Therefore, when $\varphi(c)$ is constant, profit-maximizing and welfare-maximizing decision makers have identical preference orderings over risky portfolios that require equal R&D expenditures.

When the pass-through rate varies with c , profit-maximizing monopolist may have different attitudes toward risk in c than a welfare maximizer. One measure of preferences toward risk is the Arrow-Pratt coefficient of absolute risk aversion, $-u''(x)/u'(x)$, where $u(x)$ is a payoff function. Twice differentiating the expressions for welfare and profits yields

$$W''(c) = -\frac{dx}{dr} \varphi(c)(1 + \varphi(c)) - x^*(c)\varphi'(c) \quad \text{and} \quad \pi''(c) = -\frac{dx}{dr} \frac{dr}{dc} = \left[-\frac{r}{x} \frac{dx}{dr} \right] \frac{dr}{dc} \frac{x}{r} = \eta \varphi(c) \frac{x}{r} > 0,$$

where $\eta \equiv -\frac{r}{x} \frac{dx}{dr}$ is the elasticity of demand. Combining the calculations above,

$$-\frac{W''(c)}{W'(c)} = -\frac{\pi''(c)}{\pi'(c)} - \frac{\varphi'(c)}{1 + \varphi(c)}. \quad (1)$$

Consider a choice between one R&D portfolio that yields a non-degenerate distribution of cost levels and another requiring the same R&D expenditures that yields a particular cost level with certainty. When the pass-through rate everywhere increases with c , the coefficient of absolute risk aversion is everywhere higher for the profit-maximizing monopolist and thus, whenever the welfare maximizer weakly prefers the risky R&D portfolio, so does the profit maximizer.⁶ Similarly, if the pass-through rate decreases with c , then whenever the profit maximizer weakly prefers the risky R&D portfolio, so does the welfare maximizer.

⁶ This result is a variant of the following result, which applies to two decision makers with monotonically increasing, concave objective functions who are choosing between a lottery and a sure thing. (In our setting, two decision makers with monotonically decreasing, convex objective functions choose between a lottery and a sure thing.) If one decision maker everywhere has a higher coefficient of absolute risk aversion than the other, then if the decision maker with a higher coefficient of absolute risk aversion weakly prefers the lottery, so does the decision maker with the lower coefficient. See, for example, Kreps (1990 at 86).

As is well known, the Arrow-Pratt measure is of limited usefulness for analyzing choices between two risky portfolios. Ross (1981) has proposed a stronger measure for dealing with such situations. Modifying his definition to fit the present setting, $W(c)$ is said to be “strongly more risk loving” than $\pi(c)$ if and only if $\inf_c \frac{W''(c)}{\pi''(c)} \geq \sup_c \frac{W'(c)}{\pi'(c)}$. If one decision maker is strongly more risk loving than the other, then when the second one would choose the riskier of two portfolios, so would the first (Ross 1981, Application 1). Application of this measure generally requires consideration of specific demand functions. However, one can make the following observations, which limit the possible nature of any divergence between private and social attitudes toward risk in c . Using our earlier expressions for the derivatives of profits and welfare, $\frac{W''(c)}{\pi''(c)} = 1 + \varphi(c) - \frac{r}{\eta} \frac{\varphi'(c)}{\varphi(c)}$ and $\frac{W'(c)}{\pi'(c)} = 1 + \varphi(c)$. Hence, if $\varphi'(c) > 0$, the welfare maximizer cannot be strongly more risk loving than the profit maximizer. Similarly, if $\varphi'(c) < 0$, the profit maximizer cannot be strongly more risk loving than the welfare maximizer.

The analysis above compares private and social attitudes toward risk with respect to cost realizations. But to understand any divergence between the privately chosen degree of diversification and the social optimum, one must also understand the generally complex relationship between riskiness and diversification. As noted above, a profit-maximizing monopolist will—from a social perspective—tend to under-invest in cost-reducing R&D. To distinguish this effort bias from any diversification bias, we assume that the firm has a fixed total R&D expenditure of E and we examine the effects of spreading expenditure over additional projects. In choosing the degree of diversification, at least two factors come into play in addition to the firm’s attitudes toward risk: (1) different projects are substitutes, and (2) changing the level of effort devoted to an R&D project changes its distribution of returns.

Consider the first additional factor. Because the outputs of successful R&D projects are perfect substitutes, there is no incremental private or social value to adding a project whose distribution of resulting cost reductions mirrors that of a project already in the firm's portfolio.⁷ Thus, all else equal, there is value in pursuing negatively correlated projects even if the decision maker is risk loving.

Now focus on the second factor by supposing the outcomes of different projects are independent and identically distributed conditional on effort levels, with common distribution function $G(c|e)$, where e is the amount of effort devoted to that project. We normalize the price of effort at 1 and assume that there is an additional fixed cost of F per project. If the optimal allocation of effort across k active projects is uniform, then $E = k(e+F)$ and the lowest realized cost has density $kg(c|E/k - F)(1 - G(c|E/k - F))^{k-1}$, where g is the density associated with G .

A central question is whether differences in attitudes toward risk lead a profit-maximizing monopolist to choose a different value of k than would a total-surplus maximizer. Building on yet another line of Joe Stiglitz's work (Diamond and Stiglitz (1974) and Rothschild and Stiglitz (1970 and 1971)), one portfolio of projects is said to be riskier than the other if the distribution of costs associated with the first portfolio is equal to the distribution of the second plus a mean preserving spread. Unfortunately, characterizing the effects of k on a portfolio's distribution of cost realizations can be difficult because the distribution of returns from each project in a portfolio generally varies as total effort is distributed more thinly across projects.

There is one case in which the private and social portfolio decisions can readily be

⁷ More precisely, there is no incremental value to a project for which there is no state of the world that occurs with positive probability in which that project has a strictly lower cost realization than any other project in the firm's portfolio.

compared: models such as Sah and Stiglitz (1987), in which any project has only two possible outcomes, “success” and “failure.” Suppose that failure leaves a firm’s cost unchanged, while success lowers the firm’s marginal cost to c^* . Let $\rho(e)$ denote a project’s probability of success given that effort level e is devoted to that project. We assume throughout this essay that $\rho(0) = 0$, $\rho'(e) > 0$, and $\rho''(e) < 0$. For a fixed level of total R&D expenditure, both the social and private programming problems are to allocate the expenditure across projects to maximize the probability that at least one project succeeds. This common program can be expressed as

$$\min \prod_{i=1}^k (1 - \rho(e_i))$$

subject to $kF + \sum_{i=1}^k e_i \leq E$,

where k is the total number of projects receiving positive effort. It immediately follows that: (a) attitudes toward risk have no effect on the optimal choice of project diversity, and (b) any privately optimal allocation of R&D effort is also socially optimal.

For the analysis that follows, it is useful to characterize the optimal allocation of effort more fully. Forming the Lagrangian and differentiating yields first-order conditions

$$\rho'(e_i) \prod_{j \neq i} (1 - \rho(e_j)) = \mu \quad i = 1, 2, \dots, k, \quad (2)$$

where μ is the multiplier for the budget constraint. For any two projects i and j receiving positive effort, one must have $\rho'(e_i)(1 - \rho(e_j)) = \rho'(e_j)(1 - \rho(e_i))$. When $\frac{\rho'(e)}{1 - \rho(e)}$ is strictly decreasing

in e , all active projects must receive the same level of effort.⁸ In what follows, we assume that this condition is satisfied.

Now suppose that E is endogenous. The argument above implies that all active projects will receive a common level of effort, e . The firm chooses e and k to maximize expected profits. Let π be the expected incremental profit from a successful project. The optimal level of effort satisfies Equation (2) with $\mu = 1/\pi$ and, ignoring integer constraints, the marginal R&D project must just break even in equilibrium,

$$\pi\rho(e)(1 - \rho(e))^{k-1} = F + e .$$

These two equations imply

$$\frac{\partial\rho(e)/\partial e}{\rho(e)} = \frac{1}{F + e} . \quad (3)$$

Making sufficient assumptions about the curvature of $\rho(e)$, there is a unique optimal per-project effort level, e^* .⁹ Notice that e^* depends on neither k nor π . Finally, define $\rho \equiv \rho(e^*)$.

IV. EXTERNAL CONCERNS

In this section, we examine the interaction of multiple unitary decision makers. The R&D decisions of different firms interact in several ways. One is through product-market competition: successful R&D by one firm affects the returns to R&D that are enjoyed by product-market rivals. Other effects can arise when intellectual property rights, such as patents, enable an

⁸ If $\frac{\rho'(e)}{1-\rho(e)}$ is strictly increasing in e , the firm undertakes only one R&D project.

⁹ A sufficient condition is for the elasticity of $\rho(e)$ with respect to e to be non-increasing in e because then $\frac{e\rho'(e)}{\rho(e)}$ must cross $\frac{e}{F+e}$ once from above.

initial innovator to preempt later ones following a similar R&D path or when firms conducting R&D compete for scarce inputs, such as trained research personnel.

Our interest is in how market structure—acting through its influence on the nature of these interactions—shapes the market-wide portfolio of R&D projects.¹⁰ As Sah and Stiglitz (1987) observed, many models of R&D investment forcibly underestimate the R&D diversification that may arise in concentrated market structures by flatly assuming that each firm undertakes only a single R&D project. Sah and Stiglitz emphasized that such an assumption is unrealistic and provided a set of circumstances in which a highly concentrated industry undertakes the same total number of R&D projects as a more atomistic one. Below, we generalize their result and show how it depends on assumptions made about the interactions identified above. We find that equilibrium R&D diversity generally depends on market structure, but in complex ways not well captured by the conventional one-project-per-firm assumption.

A. Innovation Competition with Nonexclusive Intellectual Property Rights

Following Sah and Stiglitz (1987), assume that each of N producers of a homogeneous product can pursue one or more cost-reducing R&D projects. For simplicity, assume that all firms have the same constant marginal costs $c_i = c^0$, $i = 1, \dots, N$ before any discovery is made. The outcomes of the projects are stochastically independent, whether pursued by the same firm or by different ones. Firms draw from an infinite pool of projects, so that the chance of any two firms undertaking the same project is nil. The results of all undertaken projects become common knowledge before price or output decisions are made. As in our earlier example, an unsuccessful

¹⁰ By “structure” we mean the underlying tastes and technology. In much of what follows, we treat the number of firms as an exogenously given element of market structure. This should be viewed as a short hand for the endogenous determination of the number of firms as a result of tastes, technology, and possibly government policies, such as antitrust.

project leaves a firm's cost unchanged, while a successful project lowers the firm's marginal cost to $c^* < c^0$. The marginal cost reduction from an additional successful project is zero.

The intellectual property rights regime is such that successful R&D by one firm neither allows other firms to take advantage of the results of that R&D nor precludes other firms from making use of their own successful R&D. We have in mind an environment in which R&D projects are protected by trade secrets. Secrecy prevents a firm from appropriating a rival's successful R&D but does not prevent any firm from exploiting the results of its own R&D.¹¹

1. *Nash-Bertrand Competition*

Sah and Stiglitz (1987) showed that the total, market-wide number of R&D projects pursued in equilibrium is independent of the number of firms in the industry when firms are Nash-Bertrand product-market competitors. With Bertrand competition and constant marginal costs, all producers earn zero profits if more than one firm succeeds at R&D or if all of them fail. Firm i earns positive product-market profits, π , if and only if its R&D alone succeeds.

If firm j engages in k_j R&D projects, $j = 1, 2, \dots, N$, then firm i 's net expected profit is

$$\pi q(k_i) \prod_{j \neq i} (1 - \rho)^{k_j} - k_i m,$$

where $q(k_i) = 1 - (1 - \rho)^{k_i}$ is the probability that at least one of firm i 's projects succeeds,

$\prod_{j \neq i} (1 - \rho)^{k_j}$ is the probability that all projects by other firms fail, and $m = F + e^*$ is the cost of an

optimally scaled R&D project. With Bertrand competition and constant marginal costs, the

¹¹ This secrecy also makes licensing difficult. For a discussion of the difficulties of selling information, see Arrow (1962 at 614-16). For a discussion of licensing in the presence of the potential theft of information shown to the prospective buyer, see Anton and Yao (1994).

equilibrium total number of R&D projects, $K^* = \sum_{j=1}^N k_j$, satisfies

$$\pi \rho (1 - \rho)^{K^*} \leq m \leq \pi \rho (1 - \rho)^{K^*-1}.$$

The Sah and Stiglitz result that the extent of R&D diversity is independent of market structure under these conditions is evident from these inequalities: K^* is independent of the number of firms in the industry.¹²

Consider the productive and allocative efficiency properties of the equilibrium. An important property of Bertrand competition is that if at least one project is successful, all equilibrium production is by a firm with the lower cost level, c^* . Thus, production efficiency is independent of how the projects are spread across firms. This property does not hold for other forms of product-market competition and, as we will see below, this has important implications for the effect of R&D competition on industry costs and welfare. Turning to allocative efficiency, the equilibrium price is lower if two or more firms have successful R&D projects than if only one firm does. Hence, as long as demand is not perfectly inelastic, allocative efficiency depends on the distribution of projects across firms.¹³

Before discussing the Sah and Stiglitz invariance result further, we observe that, with independently and identically distributed R&D projects, K is a useful measure of R&D diversity. When the returns to different R&D projects are correlated to various degrees, the extent of diversity also depends on the extent of correlation among the projects firms choose to pursue. For the reasons discussed in Section III above, a firm has incentives to pursue a portfolio of R&D

¹² When industry demand is perfectly inelastic, one can extend this result to projects that have more than two possible outcomes. The reason is that a project's value depends only on how much better it is than the next-most-successful project, regardless of who owns it.

¹³ This point is made by Sah and Stiglitz (1987 at 103 and 104).

projects whose returns are negatively correlated. It is easy to see that, under Nash-Bertrand competition with undifferentiated products, there also is an incentive for a firm to choose projects with returns negatively correlated with those of rival firms because there is no value in being successful if another firm has been successful as well.

2. *A Perfect Cartel*

The reason for the invariance of the total number of R&D projects in the industry with respect to the total number of firms under homogeneous Bertrand competition is that the incremental private benefit of a successful R&D project: (a) is zero if there is at least one other successful project, *whether pursued by the same firm or by another*, and (b) is independent of the number of firms and the set of unsuccessful projects if it is the unique successful project. Thus, a firm contemplating an incremental project calculates the profitability of the project based on the total number of projects in the industry, not on their allocation among firms.

As we will illustrate with the models of this subsection and the next, the Sah and Stiglitz invariance result is not robust to the nature of product-market competition. Before considering these formal models, it is useful to understand intuitively where the result breaks down. As long as any two successful projects are perfect substitutes, one success for a firm makes additional successes worthless regardless of the nature of product-market competition. However, the extent to which success by one firm affects the value of success for its rivals does vary with the nature of competition. In particular, under many forms of competition, there is a positive prize associated with being one of several firms to have successfully innovated. Thus, property (a) of the Sah-Stiglitz model is not generally satisfied.¹⁴ Turning to property (b), this too does not

¹⁴ Sah and Stiglitz (1987 at 107) themselves make this point.

generally hold because other firms will affect the successful innovator's product-market output and price, and in many models a change in the number of rivals (with the same costs as one another) will change the value of being a unique successful innovator.

Consider the polar opposite case from Bertrand competition: firms collude perfectly on price. We continue to assume that firms compete in research and development. Specifically, assume that demand is inelastic at quantity D^0 up to the reservation value v and firms share industry revenues vD^0 equally. Firm i 's expected profit is

$$\frac{D^0}{N} \left[(v - c^0) + (c^0 - c^*) (1 - (1 - \rho)^{k_i}) \right] - k_i m$$

when it undertakes k_i R&D projects. Each firm's optimal number of R&D projects satisfies

$$\frac{D^0}{N} (c^0 - c^*) \rho (1 - \rho)^{k^*} \leq m \leq \frac{D^0}{N} (c^0 - c^*) \rho (1 - \rho)^{k^*-1} .$$

The inequalities imply that k^* is a non-increasing function of D^0 / N , which clearly falls as N rises. The total industry number of R&D projects is not, in general, invariant to the structure of the industry.¹⁵ With our assumed R&D technology, a second success by a given producer is worthless to that firm. However, the value of its first success is now independent of whether other firms have succeeded or not. Moreover, the value of success depends on the producer's share of total output, and each firm's optimal investment in R&D, k^* , is decreasing in the number of firms. Hence, neither property (a) nor (b) of the Sah and Stiglitz model now holds.

¹⁵ The equilibrium industry-wide total number of projects is $K \equiv Nk^*$. Ignoring integer constraints, the sign of dK/dN is equal to the sign of $\ln \left(\frac{q(k^*)}{1 - q(k^*)} \right)$ where $q(k) \equiv 1 - (1 - \rho)^k$. The sign is positive if k^* is sufficiently large and negative if k^* and ρ are sufficiently small.

Although the total number of R&D projects can increase with N , expected welfare in this example is a non-increasing function of the number of firms in the industry. Expected industry profits (which are equal to expected total surplus in this model) fall as N rises because there is less diffusion of a successful innovation. A firm with lower production costs sells only D^0 / N units of output, and thus expected production costs are higher when a given number of R&D projects are spread across more firms. Moreover, each firm's expected marginal cost of output rises with the number of firms because each firm undertakes fewer R&D projects. Thus, while more firms may contribute to greater R&D diversity as measured by the total number of projects, expected total surplus falls and expected production costs rise.

Because the firms do not compete in price with one another in this assumed cartel, licensing would be jointly profitable. Above, we implicitly assumed that informational asymmetries and the intellectual property rights regime make such licensing infeasible. However, under different informational and property right assumptions, the invariance result would reappear. Specifically, the number of R&D projects would be independent of the number of firms if: (1) a monopoly licensor could fully extract the value of its innovation from its licensees, and (2) two or more successful innovators would compete in the licensing market in Nash-Bertrand fashion, driving the equilibrium license fee to zero. When conditions (1) and (2) are satisfied, properties (a) and (b) of the Sah-Stiglitz model hold. Condition (1) is necessary because a firm fully internalizes the benefits of innovation for its own sales, which are a function of N , while the total benefits of industry-wide licensing are independent of N .

3. *Nash-Cournot Oligopoly*

Suppose the N firms are Nash-Cournot competitors in the product market. Unlike a perfect cartel, a firm's payoff from innovation depends on the number of other firms that

innovate successfully. Unlike Bertrand competition, the firm's payoff from innovation can be positive even if it shares the market with other successful innovators.

Assume that N firms sell a homogeneous product with linear inverse demand with intercept A and slope $-b$. Define the average marginal cost of all firms other than firm i :

$\tilde{c}_{-i} \equiv \frac{1}{N-1} \sum_{j \neq i} \tilde{c}_j$, where \tilde{c}_j is a random variable that takes the value c^* with probability

$q(k_j) = 1 - (1 - \rho)^{k_j}$ and c^0 with probability $1 - q(k_j)$. Standard analysis (see, e.g., Vives (1999)

and Yi (1999)) shows that firm i 's profit as a function of industry cost realizations is¹⁶

$$\pi_i(c_i, \tilde{c}_{-i}) = \frac{1}{b} \left[\frac{A - c_i + (N-1)(\tilde{c}_{-i} - c_i)}{N+1} \right]^2.$$

In a Cournot oligopoly, firm i 's profit depends on the average of all other active firms' costs and does not depend upon how those costs are distributed among its rivals. Let \bar{c}_{-i} be the expected value of \tilde{c}_{-i} . Firm i 's expected profit is

$$\begin{aligned} E\pi_i(c_i, \tilde{c}_{-i}) &= \frac{1}{b} \left[\frac{A - c_i + (N-1)(\bar{c}_{-i} - c_i)}{N+1} \right]^2 + \frac{(N-1)^2}{b(N+1)^2} \text{var}(\tilde{c}_{-i}) \\ &= \pi_i(c_i, \bar{c}_{-i}) + \frac{(N-1)^2}{b(N+1)^2} \text{var}(\tilde{c}_{-i}). \end{aligned}$$

If firm i undertakes k_i R&D projects, its expected profit net of R&D expenditures is

$$(1 - \rho)^{k_i} E\pi_i(c^0, \tilde{c}_{-i}) + (1 - (1 - \rho)^{k_i}) E\pi_i(c^*, \tilde{c}_{-i}) - k_i m$$

and the expected benefit to firm i from an additional R&D project is

$$\rho(1 - \rho)^{k_i} [\pi_i(c^*, \bar{c}_{-i}) - \pi_i(c^0, \bar{c}_{-i})] - m.$$

¹⁶ We are assuming that all firms are active producers, which is the case for $c^0 - c^*$ sufficiently small.

In a symmetric equilibrium, each of N firms invests in k R&D projects. Then $\bar{c}_j = c^* + (1 - \rho)^k (c^0 - c^*)$ for all $j = 1, \dots, N$. Figure 1 shows the total number of R&D projects in a symmetric N -firm oligopoly when $A = 80$, $b = 1$, $\rho = .1$, $c^0 = 10$, $c^* = 6$, and $m = 5$. This example ignores integer constraints and assumes projects are divisible. As Figure 1 shows, total investment in R&D has an inverted-U shape. That is, the equilibrium total number of R&D projects reaches a maximum at intermediate levels of market concentration.

Intuitively, there are two offsetting forces at play as the number of firms rises. One, as the number of firms rises, each firm's sales fall and, thus, so do the benefits of successful, unit-cost-reducing R&D. This effect leads each firm to do less R&D as the number of firms rises. Two, unlike the Nash-Bertrand case, each firm can benefit from successful R&D even if other firms succeed as well. This effect can raise the total industry incentives to conduct multiple projects. Up to some number of firms (four in this example), the total industry investment in R&D increases. It falls for larger numbers of firms.

In the present model, even if an increase in the number of firms leads to a larger total number of R&D projects, it leads to (weakly) fewer projects per firm. Moreover, unlike the Nash-Bertrand case, a Nash-Cournot competitor whose R&D succeeds does not serve the entire market. As a result, both firm and industry expected unit costs increase with the number of firms, even if the total number of R&D projects and, therefore, the probability of successful R&D also increase with the number of firms. Recall that the perfect cartel equilibrium exhibits a similar pattern. The total number of R&D projects can increase with the number of firms, but expected unit costs at the firm and industry levels weakly increase with more firms.

In contrast to the case of a perfect cartel facing inelastic demand, the beneficial impact on prices of an increase in the number of Cournot competitors can outweigh the negative impact on

industry costs. For the parameter values in the example reported in Figure 1, expected welfare increases with N up to six firms, and decreases for larger numbers of competitors.

B. Innovation Competition with Exclusive Success or Scarce R&D Inputs

When innovation is protected through intellectual property laws that grant the right to exclude others (*e.g.*, patent), if one firm succeeds along a particular path, then other firms may be unable to make use of the results of their R&D if they have followed the same research path.¹⁷ Thus, when there are finitely many potential paths, or some firms can choose their R&D paths after observing the choices of their rivals, the possibility of preemption arises. We will show how preemption can give rise to an incentive to diversify as a “spoiler” strategy.

Consider a model in which each firm allocates effort along two paths, A and B, each of which is a patent race. While innovations on the two paths are perfect substitutes for reducing production costs, a patent for an innovation on one path does not block use of an innovation on the other one. On each path, each firm’s probability of success is an increasing function of its own efforts and a decreasing function of its rival’s. Thus, firm 1 has success probability $a(e_1^A, e_2^A)$ along path A, where e_i^A is firm i ’s effort along path A. Firm 2 has a symmetric success probability, $a(e_2^A, e_1^A)$. We use analogous notation for path B. Unlike in the non-exclusive case, at most one firm can “succeed” on any path.

Private payoffs are as follows: if one firm succeeds along at least one path and its rival succeeds along neither, then gross payoffs are $(\pi^M, 0)$, where π^M denotes monopoly profits in the product market. If each firm succeeds along a different path, then gross payoffs are (π^D, π^D) ,

¹⁷ We are assuming that a successful innovator has no obligation to share the fruits of its R&D with product-market rivals and that rival firms cannot use independent invention as a defense to an infringement claim.

where π^D denotes the per-firm duopoly payoffs and $2\pi^D < \pi^M$. Firm 1's expected payoff is

$$a(e_1^A, e_2^A)[1 - b(e_2^B, e_1^B)]\pi^M + a(e_1^A, e_2^A)b(e_2^B, e_1^B)\pi^D + \\ b(e_1^B, e_2^B)[1 - a(e_2^A, e_1^A)]\pi^M + b(e_1^B, e_2^B)a(e_2^A, e_1^A)\pi^D - e_1^A - e_1^B .^{18}$$

A similar expression holds for firm 2. Hence, the marginal return to e_1^A is (using subscripts on $a(\cdot, \cdot)$ and $b(\cdot, \cdot)$ to denote partial derivatives):

$$a_1(e_1^A, e_2^A) [(1 - b(e_2^B, e_1^B))\pi^M + b(e_2^B, e_1^B)\pi^D] - a_2(e_2^A, e_1^A) b(e_1^B, e_2^B)[\pi^M - \pi^D] - 1 .$$

The term -1 is the direct cost of firm 1's marginal effort along path A. The first term (beginning with a_1) is the expected private value of increasing firm 1's own success along path A, while the term with a_2 (recall that $a_2 < 0$) captures the private value of reducing firm 2's probability of succeeding along path A, which is valuable to firm 1 in the event that firm 1 succeeds along B. The term with a_2 is a raising-rival's-cost effect: firm 1 tries to succeed along path A in part in order to stop firm 2 from doing so. This component of firm 1's private benefit is proportional to firm 1's probability of success along path B. Note that

$$\frac{\partial^2 \pi_1(e_1^A, e_1^B)}{\partial e_1^A \partial e_1^B} = -[a_1(e_1^A, e_2^A)b_2(e_2^B, e_1^B) + b_1(e_1^B, e_2^B)a_2(e_2^A, e_1^A)][\pi^M - \pi^D] > 0 .$$

Thus, there is a *private complementarity* between success on A and success on B. Intuitively, if a firm knows it is going to succeed along one path, it has strong private incentives to succeed along the other in order to protect the profits it can earn from being the unique successful innovator.

In contrast to the private incentives, there is a *social substitutability* between success along path A and success along B. If the firm succeeds along path A, there is no social value to the firm's succeeding along path B as well because there is no incremental cost reduction. The

¹⁸ For simplicity, in this subsection we assume $F = 0$.

private complementarity can create a private incentive to diversify along both paths even if the technology of R&D is characterized by increasing returns to effort along any one path so that specialization would be efficient. Moreover, with non-exclusive success, increasing returns to R&D would give rise to private incentives to specialize efficiently.

Similarly to preemptive patenting, a firm can have incentives to diversify its projects in order to raise rivals' R&D costs. If some inputs for R&D have upward sloping supply curves, then a firm may expand its R&D along a particular path to make R&D more expensive for its competitors, and thereby decrease their probabilities of success along that path.

V. INTERNAL CONSIDERATIONS

So far, we have assumed that firms behave as unitary decision makers acting to maximize profits subject to market competition. This is, at best, an approximation; a well-run firm solves a complex multi-layer principal-agent problem and can thus be viewed as maximizing profits subject to many constraints, only some of which come from product-market competition. In the present section, we briefly examine the effects of various private responses to the need to aggregate and exchange information within an organization in order to make decisions.

To focus on the influence of organizational design on the choice of R&D portfolios, we first consider a setting in which the equilibrium number of R&D projects is independent of the number of firms, holding organizational structure constant across firms, and we examine how changing the organizational structure affects the equilibrium number of R&D projects.

Similar to the first model of Section IV, consider a market in which firms are Nash-Bertrand competitors in the product market and choose among stochastically independent projects, each of which has only two possible outcomes, success and failure. In contrast to our earlier model, assume that there are two classes of R&D projects: "good" and "bad." A good

project has a probability of success, $\rho(e) > 0$ for all $e > 0$ and $\rho(0) = 0$, while a bad project has a zero probability of success for any level of effort.

Suppose the manager choosing whether to undertake an R&D project believes that a fraction $\omega \in (0,1)$ of the projects proposed by the organization's staff are good. Let π be the prize from having the sole successful project. As earlier, the prize is 0 if two or more firms have successful projects. The expected incremental value of a marginal project is equal to

$$\max_e \omega \rho(e) H \pi - e - F ,$$

where H is the probability that all other projects (of that firm or any other) fail. Inserting ω into the derivation in Section III, the optimal level of effort per project, e^* is given by Equation (3) and is independent of ω and π .

In a symmetric equilibrium with K projects, $H = (1 - \omega \rho)^{K-1}$, where as before $\rho \equiv \rho(e^*)$. When $\omega \rho \pi > F + e^*$, it is profitable for at least one firm to engage in R&D and, ignoring integer constraints, the equilibrium number of projects satisfies

$$\omega \rho (1 - \omega \rho)^{K-1} \pi = F + e^* = m, \quad (4)$$

As expected, the equilibrium number of projects is independent of N .

In what follows, it is useful to understand the relationship between the equilibrium number of R&D projects, K^* , and the probability that a project is good, ω . We can examine this relationship by taking the natural logarithm of the condition for the equilibrium number of projects and totally differentiating with respect to ω :

$$\left[\frac{1}{\omega} - \frac{(K^* - 1)\rho}{1 - \omega \rho} \right] d\omega + \ln(1 - \omega \rho) dK^* = 0 .$$

Hence,

$$\frac{dK^*}{d\omega} = -\frac{1 - \omega\rho K^*}{\omega(1 - \omega\rho)\ln(1 - \omega\rho)} \quad (5)$$

The denominator is negative, and the sign of $\frac{dK^*}{d\omega}$ is equal to the sign of $1 - \omega\rho K^*$.

By Equation (4), $K^* = 1 - \left[\frac{\ln \omega\rho\pi - \ln m}{\ln(1 - \omega\rho)} \right]$, and thus

$$1 - \omega\rho K^* = 1 - \omega\rho + \omega\rho \left[\frac{\ln \omega\rho\pi - \ln m}{\ln(1 - \omega\rho)} \right]. \quad (6)$$

Because $0 < \omega\rho < 1$, the first two terms on the right-hand side of Equation (6) sum to a positive number and the denominator of the third term is negative. By Equation (4), $m < \omega\rho\pi$ and thus the third term is negative.

If π is sufficiently large, the sum of the three terms is negative. Hence, given any positive value of $\omega\rho$, if π is sufficiently large, then a change in organizational design that increases ω will result in fewer equilibrium R&D projects industry wide. In other words, an increase in the effectiveness of project selection will reduce equilibrium R&D diversity. Conversely, for any admissible value of $\omega\rho$, if π is sufficiently close to $\frac{m}{\omega\rho}$ from above, the sum of the three terms is positive. Hence, for a given value of ω , if π is sufficiently close to $\frac{m}{\omega\rho}$ from above, an increase in the effectiveness of project selection will result in greater equilibrium R&D diversity.

While the equilibrium number of projects undertaken may rise or fall with ω , the equilibrium probability that at least one project will succeed always rises. The reason is as follows. From Equation (4), $(1 - \omega\rho)^{K^*-1}$ is equal to $\frac{m}{\omega\pi\rho}$, which falls as ω rises. The probability that all projects fail is $(1 - \omega\rho)^{K^*} = (1 - \omega\rho)(1 - \omega\rho)^{K^*-1}$. Because both factors on the right-hand side of this equality fall as ω rises, the probability that all projects fail must fall as ω rises.

Thus far, we have taken ω to be exogenous. However, various aspects of organizational design affect an enterprise's ability to pursue a favorable selection of R&D projects. One element is the extent to which internal reward structures align employee incentives with those of shareholders. Suppose, for example, that the R&D staff in an organization have very low-powered incentives. Then the technical staff may propose "bad" projects because these projects generate utility to the staff as interesting research problems even though they hold no commercial promise. Conversely, in an organization that has compensation and promotion schemes that align the incentives of R&D staff with those of stockholders, higher-level managers can correctly assume that a higher percentage of projects advanced by the technical staff are good projects.

A second factor is how the organization aggregates diverse information and views held by different members of the organization. The design of organizations to accomplish this task was the subject of an important line of research reported in Sah and Stiglitz (1985, 1986, and 1988). Here, we analyze a model building on this line of research. Firms choose from the same, infinitely large pool of potential projects. When a firm chooses a project, the firm can expend resources to evaluate the project before committing effort to it. Unlike in Sah and Stiglitz, projects are substitutes for one another and, thus, there are declining incremental social and private values of undertaking additional projects.

An organization chooses how many evaluations of a proposed R&D project to conduct. Each evaluation incurs a cost, s , to obtain a binary signal of whether the project is good or bad. We say that an evaluator "approves" a project when the signal indicates that the project is good. The probability that a single evaluator will approve a project is γ if the project is in fact good and β if the project is in fact bad. We assume that evaluations are informative (*i.e.*, $\gamma > \beta$) and that the evaluations are independent of one another conditional on the true type of the project.

As defined by Sah and Stiglitz, under an L -level *hierarchy*, a project is evaluated sequentially and a negative evaluation at any point leads to the project's being rejected without any further evaluation. Thus, a project is accepted if and only if all L levels of the hierarchy give it a positive evaluation. When ω_0 is the prior probability that a project is good, if the project has been approved by an L -level hierarchy, then the posterior probability that the project is good is

$$\omega(L) = \left[\frac{\omega_0 \gamma^L}{\omega_0 \gamma^L + (1 - \omega_0) \beta^L} \right].$$

Note that $\omega(L)$ goes to 1 as L goes to infinity, and $\omega(L)$ goes to ω_0 as L goes to 0.

On average, a proposed project will be evaluated

$$\zeta(L) \equiv \omega_0 \left[\frac{1 - \gamma^L}{1 - \gamma} \right] + (1 - \omega_0) \left[\frac{1 - \beta^L}{1 - \beta} \right]$$

times. The fraction of proposed projects that a L -level hierarchy will deem good and therefore eligible for investment is $\omega_0 \gamma^L + (1 - \omega_0) \beta^L$. Consequently, the evaluation cost per approved

project in an L -level hierarchy is $\frac{s\zeta(L)}{\omega_0 \gamma^L + (1 - \omega_0) \beta^L}$, which is an increasing function of L . The

expected evaluation cost per undertaken project increases with the number of levels in the hierarchy because each project has to be reviewed at each level and the organization has to sort through a larger number of projects to select one in which to invest.

Evaluation costs are in addition to per-project R&D costs, $F + e^*(L)$, where $e^*(L)$ satisfies

$$\frac{\partial \rho(e^*) / \partial e}{\rho(e^*)} = \frac{1}{F + e^*(L) + \frac{s\zeta(L)}{\omega_0 \gamma^L + (1 - \omega_0) \beta^L}} \quad .^{19}$$

By the concavity of $\rho(\cdot)$, $e^*(L)$ is increasing in L . Hence, the sum of the effort and evaluation costs per undertaken project, $m(L) \equiv F + e^*(L) + \frac{s\zeta(L)}{\omega_0 \gamma^L + (1 - \omega_0) \beta^L}$, is increasing in L . An additional level increases the probability that R&D projects approved for investment will succeed, but the cost per undertaken project also rises.²⁰

Suppose there are N hierarchies, each of which has L levels. Define $\rho^*(L) \equiv \rho(e^*(L))$. Generalizing our earlier discussion of the Bertrand case, if the firms undertake a total of K projects, no firm can increase its expected profits by screening one more or one fewer project if

$$\omega(L) \rho^*(L) (1 - \omega(L) \rho^*(L))^{K-1} \pi \geq m(L) \geq \omega(L) \rho^*(L) (1 - \omega(L) \rho^*(L))^K \pi .$$

Thus, ignoring integer constraints,

$$\omega(L) \rho^*(L) (1 - \omega(L) \rho^*(L))^{K-1} \pi = m(L) .$$

As before, the equilibrium number of R&D projects under Bertrand competition is independent of the number of firms in the industry holding L fixed. The number of projects does, however, depend on the size of the hierarchies, which may itself depend on the number of firms.

Our analysis above identifies several effects on R&D from increasing the number of layers in each hierarchy. The net result depends on technological parameters. Figure 2 reports the results of simulations with $\omega_0 = 0.2$, $\gamma = 0.6$, $\beta = 0.4$, and $\pi = 300$ under the assumption that

¹⁹ This result is another application of the derivation in Section III and implicitly assumes sufficient curvature of $\rho(e)$ to ensure that a unique solution exists.

²⁰ Moreover, if—unlike the present model—there is a limited number of good projects, then a firm's mistakenly rejecting a good project, which becomes more likely as L increases, will be costly.

the scale of R&D projects is technologically fixed such that $F + e^*(L) = 5$ and $\rho^*(L) = 0.1$ for all L . The figure illustrates the equilibrium outcome for various levels of the cost per evaluation, s . Given these parameters, for small values of L , π is sufficiently close to $\frac{m}{\omega(L)\rho^*(L)}$ from above that $\frac{dK^*}{d\omega}$ is positive. For larger values of L , however, $\frac{dK^*}{d\omega}$ is negative. When there are no per-project evaluation costs ($s = 0$), an increase in L increases the probability of success for each project actually undertaken and, from our earlier result, increases the total probability that R&D is successful for the market as a whole. Furthermore, for small hierarchies, the numerator in Equation (5) is positive and the total number of R&D projects in the market increases with the number of layers in each hierarchy. However, the results change dramatically when project evaluations are costly. For example, when $s = 0.8$, the total number of R&D projects and the total probability of successful R&D fall with the number of layers in each hierarchy, and firms do not do any R&D at all if the number of levels in the hierarchies exceeds 2.

Although we have treated the number of hierarchical levels as given, in practice each firm chooses the number of levels in its organization. We next briefly explore some of the forces at work. We show that, under the assumption that the scale of R&D projects is technologically fixed at e^* with success probability ρ , a larger reward for successful innovation leads firms to invest in (weakly) more layers of hierarchy.

Suppose there is a fixed cost of S per level of hierarchy within an organization in addition to the per-evaluation cost, s . A firm's choice of k and L can be broken into two steps. For any given probability of success, $1-\alpha$, the firm chooses k and L to

$$\begin{aligned} & \min km(L) + LS \\ & \text{subject to } [1 - \rho\omega(L)]^k \leq \alpha . \end{aligned} \tag{7}$$

The firm then chooses α to maximize $(1-\alpha)\pi - \Phi(\alpha)$, where $\Phi(\alpha)$ is the optimized value of the objective function in the sub-problem above and π is the prize associated with having at least one successful project.²¹ By standard revealed preference arguments, $\Phi(\cdot)$ is a non-increasing function and the firm's choice of α is non-increasing in π .

We next examine the comparative statics of L varying α exogenously.²² It will be convenient to define $\theta(L) \equiv \ln [1 - \rho\omega(L)]$ and write Inequality (7) as $k\theta(L) \leq \ln \alpha$.

The Kuhn-Tucker conditions for the choice of k and L include

$$k[m(L) + \vartheta\theta(L)] = 0 \quad (8)$$

and

$$L[S + km'(L) + k\vartheta\theta'(L)] = 0 \quad , \quad (9)$$

where ϑ is the multiplier for the constraint. If $k > 0$, then Equation (8) implies that $\vartheta = \frac{-m(L)}{\theta(L)}$.

Substituting this expression into Equation (9) yields the necessary condition

$$L \left[S + km(L) \left(\frac{m'(L)}{m(L)} - \frac{\theta'(L)}{\theta(L)} \right) \right] = 0. \quad (10)$$

By Equation (10), if $S > 0$, then $L > 0$ only if

$$\frac{m'(L)}{m(L)} < \frac{\theta'(L)}{\theta(L)}. \quad (11)$$

Now consider two values of the failure probability, $0 < \alpha_1 < \alpha_2$. Define $k_1' \equiv \frac{\ln \alpha_1}{\theta(L_2)}$ and

$k_2' \equiv \frac{\ln \alpha_2}{\theta(L_1)}$. By construction, (k_1', L_2) satisfies the constraint in the firm's program when the

²¹ The value of π depends both on the nature of product-market competition and the intensity of R&D competition that the firm faces. Here, our reduced form allows for arbitrary forms of product competition and assumes that the firm makes its R&D decisions holding its rivals' R&D strategies fixed.

failure probability is α_1 , and (k'_2, L_1) satisfies the constraint when the failure probability is α_2 .

The optimality of (k_1, L_1) and (k_2, L_2) implies

$$k'_1 m(L_2) + L_2 S \geq k_1 m(L_1) + L_1 S$$

and

$$k'_2 m(L_1) + L_1 S \geq k_2 m(L_2) + L_2 S \quad .$$

Adding these two inequalities and rearranging terms yields

$$(k'_1 - k_2) m(L_2) \geq (k_1 - k'_2) m(L_1) \quad . \quad (12)$$

Using the definitions of the k'_i and the fact that the constraint in the firm's program is satisfied with equality by each (k_i, L_i) pair, it follows that Inequality (12) is satisfied if and only if

$$[\ln \alpha_1 - \ln \alpha_2] \left[\frac{m(L_2)}{\theta(L_2)} - \frac{m(L_1)}{\theta(L_1)} \right] \geq 0 \quad . \quad (13)$$

By hypothesis, $0 < \alpha_1 < \alpha_2$. Thus, the first term in square brackets is negative. Hence, the second term in square brackets must be non-positive. By Inequality (11), $\frac{m(L)}{\theta(L)}$ is increasing in L (recall that $m(\cdot)$ is a positive, increasing function, while $\theta(\cdot)$ is a negative, decreasing function).²³ Therefore, Inequality (13) can be satisfied only if $L_1 \geq L_2$.

Now return to our earlier model of a perfect cartel of N firms with no licensing. In that model, π is a decreasing function of N and the analysis above establishes that the equilibrium value of L is non-increasing in N . Intuitively, as N rises, each firm conducts less R&D and has fewer projects over which to spread the fixed costs of hierarchy. Hence, with fewer firms, firms will invest in more layers of hierarchy, and engage in more accurate project evaluations.

²² We ignore integer constraints, and the analysis provides a heuristic examination of the forces at work.

The analysis also suggests that in a model of Bertrand competition with differentiated products and asymmetric market shares, the firms with larger shares—and thus larger potential gains from successful cost-reducing innovation—will invest in larger hierarchies. The possibility that smaller firms will choose smaller hierarchies raises a number of interesting issues about the types of research conducted by small and large firms within an industry.

VI. CLOSING REMARKS

Drawing inspiration from some of Joe's pathbreaking work, we have explored how both external and internal factors might affect firms' R&D portfolios. Much work remains to be done. We hope that this essay—written by a modified hierarchy that at times threatened to become a Sah and Stiglitz *polyarchy*²⁴—will be instrumental in stimulating that work by developing some promising paths along which others may proceed.

²³ Again, this is a heuristic argument because Equation (8) need hold only at the respective optima.

²⁴ Projects within this essay were sequentially evaluated and were rejected if they received two negative evaluations. In a polyarchy, a project is accepted as soon as it receives a favorable evaluation.

REFERENCES

- Anton, James J., and Dennis A. Yao (1994) "Expropriation and Inventions: Appropriable Rights in the Absence of Property Rights," *American Economic Review*, **84**(1), 190-209.
- Arrow, Kenneth J. (1962) "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, Richard R. Nelson (ed.), Princeton: Princeton University Press.
- Christensen, Clayton M. (1997), *The Innovator's Dilemma*, Boston: Harvard Business School Press.
- Diamond, Peter A., and Joseph E. Stiglitz (1974) "Increases in Risk and in Risk Aversion," *Journal of Economic Theory*, **8**(3), 337-360
- Dixit, Avinash and Joseph E. Stiglitz (1977), "Monopolistic Competition and Optimal Product Diversity," *American Economic Review*, **67**, 297-308.
- Henderson, Rebecca (1993) "Underinvestment and Incompetence as Responses to Radical Innovation: Evidence from the Photolithographic Alignment Equipment Industry," *The RAND Journal of Economics*, **24**(2), 248-270.
- Gilbert, Richard J. and David Newbery (1982), "Preemptive Patenting and the Persistence of Monopoly," *American Economic Review*, **72**(2), 514-526.
- Hoven, John and Daniel L. Rubinfeld (2001), "Innovation and Antitrust Enforcement", in J. Ellig (ed.), *Dynamic Competition and Public Policy: Technology, Innovation, and Antitrust Issues*, New York: Cambridge University Press.
- Katz, Michael L. (1980), "Multiplant Monopoly in a Spatial Market," *Bell Journal of Economics*, **11**(2), 519-535.
- Katz, Michael L. and Carl Shapiro (1987), "R&D Rivalry with Licensing and Imitation," *American Economic Review*, **77**(3), 402-20.
- Kreps, David M. (1990) *A Course in Microeconomic Theory*, Princeton: Princeton University Press.
- Reinganum, Jennifer (1989), "The Timing of Innovation: Research, Development, and Diffusion" in *The Handbook of Industrial Organization*, R. Schmalensee and R.D. Willig (eds.), Amsterdam: North Holland Publishing.
- Robinson, Constance K. (1999), "Leap-frog and Other Forms of Innovation," Address before the American Bar Association, Chicago, IL, June 10.

- Ross, Stephen A. (1981) "Some Stronger Measures of Risk Aversion in the Small and the Large with Applications," *Econometrica*, **49**(3), 621-638.
- Rothschild, Michael, and Joseph E. Stiglitz (1970) "Increasing Risk: I. A Definition," *Journal of Economic Theory*, **2**(3), 225-243.
- _____ and _____ (1971) "Increasing Risk II: Its Economic Consequences," *Journal of Economic Theory*, **3**(1), 66-84.
- Sah, Raaj Kumar, and Joseph E. Stiglitz (1985) "Human Fallibility and Economic Organization," *The American Economic Review*, **75**(2) Papers and Proceedings of the Ninety-Seventh Annual Meeting of the American Economics Association (May), 292-297.
- _____ and _____ (1986) "The Architecture of Economic Systems: Hierarchies and Polyarchies," *The American Economic Review* **76**(4), 716-727.
- _____ and _____ (1987) "The Invariance of Market Innovation to the Number of Firms," *The Rand Journal of Economics*, **18**(1), 98-108.
- _____ and _____ (1988) "Committees, Hierarchies and Polyarchies," *The Economic Journal* **98**(391), 451-470.
- Vives, Xavier (1999), *Oligopoly Pricing*, MIT Press, Cambridge, Mass.
- Yi, Sang-Seung (1999), "Market Structure and Incentive to Innovate: the Case of Cournot-Oligopoly", *Economics Letters*, **65**, 379-388.




Figure 1. Total number of R&D projects v. number of firms for Nash-Cournot case




Figure 2. Dependence of total number of R&D projects on the number of hierarchy levels