2003

Subjective Rationality, Self-Confirming Equilibrium and Corporate Strategy

Michael D Ryall, *Melbourne Business School*

Subjective Rationality, Self-Confirming Equilibrium, and Corporate Strategy

Michael D. Ryall
W. E. Simon School, University of Rochester, Rochester, New York 14627
ryall@simon.rochester.edu

This paper presents a formal theory of subjective rationality and demonstrates its application to corporate strategy. An agent is said to be subjectively rational when decisions are consistent with the available facts and, where these are lacking, with the agent's own subjective assessments. A self-confirming equilibrium arises when agents' subjectively rational actions generate events that are consistent with their own expectations. Equilibrium strategies may be suboptimal because certain counterfactual beliefs may be erroneous and yet fail to be contradicted by events observed in equilibrium. This weakening of the stronger rationality assumptions inherent in many of the more familiar equilibrium ideas appears well suited to applications in strategy. In particular, performance advantage may be sustained by a firm when its subjectively rational competitors persistently employ suboptimal self-confirming strategies.
(*Subjective Rationality; Self-Confirming; Equilibrium; Strategy; Persistent Advantage*)

1. Introduction

This paper presents a formal theory of *subjective rationality* and demonstrates its application to certain foundational issues in strategy. An agent pursuing a well-defined set of objectives is said to be subjectively rational when she selects actions that appear optimal given the available facts and, where these are lacking, given her own subjective assessments (beliefs). When the situation is dynamic, it is natural to assume that an agent's beliefs remain in force so long as they are not contradicted by unfolding events. Of course, agents do not learn what *would have* occurred had some alternative course of action been pursued. Since counterfactual outcomes are not observed, counterfactual beliefs are not tested—and herein lie both the opportunity for gain and the potential for trouble.

Subjective rationality leads immediately to the notion of *self-confirming equilibrium*. A self-confirming equilibrium (hereafter, SCE) arises when subjectively rational strategies interact to generate events that are consistent with agent expectations. Persistently suboptimal strategies may be part of an equilibrium in which erroneous counterfactual beliefs fail to be refuted by equilibrium events. This weakening of the stronger rationality assumptions inherent in many of the more familiar equilibrium ideas appears well suited to applications in strategy.[1]

By way of example, consider the (here, admittedly oversimplified) pre-Columbian theory that the world was flat. Subjective optimization under this theory led sailors to keep their ships close to shore. However, by keeping their ships close to shore they never generated data refuting their erroneous counterfactual belief that, should they sail into the horizon, they would be swept off the edge of the earth. This equilibrium was stable for centuries, even though, presumably since the time of Marco Polo, European merchants understood the enormous payoff implications of finding a direct route to the Far East.

One issue of special concern to both practitioners and researchers in strategy is the source of sustainable

[1] By "more familiar," I mean, e.g., DeBreu (1959), Nash (1950), and Harsanyi (1967–1968).

0025-1909/03/4907/0936$05.00
1526-5501 electronic ISSN

competitive advantage under competition. One longstanding intuition is that competition is sufficient for the attainment of efficiency in equilibrium (e.g., Demsetz 1988). The idea is that even when managers misunderstand their environment or suffer from decision biases, competition—by making winners of the superior strategies regardless of the reasoning behind their implementation—propels an industry toward efficiency and the elimination of economic rent (Hirshleifer and Riley 1992). If this is so, the most managers can hope for is a brief period of competitive advantage. The implication is that strategy researchers should look to the economic logic of quasi rents for their theoretical foundations. Much that has been written in the field does, indeed, have this flavor (e.g., Lippman and Rumelt 1982).

This paper presents a new possibility: A sustained performance advantage for a firm whose subjectively rational competitors persistently employ self-confirming strategies that are not, contrary to their beliefs, optimal. Note the corollary: Under certain circumstances, the forces of competition may be insufficient to drive an industry to efficiency.

The notions of subjective rationality and SCE, while new to strategy, are not original to this paper. Indeed, several economists have proposed formal analyses whose interpretations are consistent with subjective rationality. The one presented here is most closely related to Kalai and Lehrer (1993, 1995), who motivate their approach by pointing out that the behavioral assumptions underlying traditional equilibrium ideas seem too demanding, "even for highly rational players engaged in moderate size problems."[2]

The novel idea in this paper is the application of subjective rationality to the theoretical foundations of strategy. Specifically, I demonstrate that subjective rationality admits the possibility of sustained value appropriation under competition. Not only may the positive economic profit of a given firm be a stable component of a dynamic equilibrium, but the source of this stability is a set of *beliefs* that—although they appear reasonable and remain uncontradicted by the available information—lead the firm's competitors down a path to inferior performance. The robustness of an SCE depends, critically, upon the relationship between beliefs and the outcomes generated along the equilibrium path.

In the following analysis, I illustrate some competitive situations in which a self-confirming advantage is fragile, and others in which it is less so. For example, self-confirming behavior may arise even in the presence of significant firm heterogeneity and exogenous "shocks" to key industry variables. In dynamic situations, rational learning (in the Bayesian sense) may actually reinforce beliefs that support self-confirming strategies.

A significant portion of the paper is devoted to analyzing competitive advantage in a multibusiness setting. This is an active area of research associated with a vast literature. One theory is that a positive relationship exists between multibusiness diversification and superior performance (for an early discussion, see Teece 1980). The reasoning is that when certain resources exhibit public-good-like qualities (i.e., can be applied to many markets without affecting productivity in any one market), the firm increases value by putting them to productive use in several, related markets. However, managerial faith in the notion that their resources exhibit scope economies can be self confirming and lead to overdiversification. As I demonstrate below, under a variety of conditions—even when firms compete using the same production technology—heterogeneous misperceptions regarding the public-good-like qualities of that technology can result in stable, nontrivial differences in firm performance. These results are consistent with several well known empirical findings (discussed below).

The remainder of the paper is organized as follows. The following section presents a motivating example along with a verbal introduction to the theory. Section 3 presents the mathematical formalism. In §4, I construct a simple model around the motivating example and use it to illustrate the basic concepts and mathematical objects. Several new analytical results

[2] Kalai and Lehrer actually use the term "subjective equilibrium" for their definition. The term "self-confirming" was first used by Fudenburg and Levine (1993) to label a related equilibrium idea. Here, "self-confirming" is used to emphasize the interrelationship between beliefs, actions, and outcomes inherent in such equilibria. For related discussions, see Abreu et al. (1990), Battigalli and Guatoli (1988), and Blume and Easley (1992).

are presented in §5, which considers competition between multimarket competitors. Beginning with the simple one-firm, one-period case and then extending it to both static and dynamic oligopolies with stochastic costs, I establish conditions under which misperceptions regarding the sensitivity of one's costs to the number of markets in which one operates can lead to equilibria in which heterogeneous performance is a stable feature. In §6, I discuss positive issues raised by the theory and speculate on extensions of possible interest to strategy researchers.

2. Motivation and Description of the Framework

Consider the following excerpts from the autobiography of Pepsi's former CEO, John Sculley:

> Coca-Cola ... owned one of the world's most distinctive trademarks in its 6.5-ounce, hourglass-shaped bottle. The bottle design nearly became the product itself. It made Coke easier to stack, more comfortable to grip, and more sturdy to withstand a vending machine's drop ... Convinced that the bottle was Coke's most important competitive advantage, Pepsi spent millions of dollars and many years studying new designs to no avail ... Pepsi executives thought of competition strictly in terms of a bottle.
>
> I initiated one of the company's first massive consumer-research studies, an extended, in-home product test ... To our astonishment, we discovered that no matter how much Pepsi they ordered, they would always consume it. It dawned on me that what we needed to do was design packages that made it easier for people to get more soft drinks into the home.
>
> It wasn't until we shifted ... to larger-sized packages that the marketing advantages to Coke of having such a unique bottle began to erode ... Coca-Cola couldn't successfully translate its valuable silhouette to the larger-sized plastic bottles. The result: A trademark familiar to more than three generations of Americans became virtually extinct![3]

Apparently, Coca-Cola enjoyed a remarkable run of market leadership because Pepsi management (three generations?) consistently failed to grasp certain fundamental realities of its competitive environment. Was this failure, as Russo and Schoemaker (1989) assert, the result of a "common pitfall" of bad decision making? The vast literature on the limits to rationality under uncertainty and complexity would certainly point us in this direction.[4] There is no question, if one accepts Sculley's description, that the belief that market success turned on the discovery of a superior bottle design was erroneous. Still, without being dismissive of an intellectually significant body of research, one wonders about the sustainability of performance advantages built upon the illogicality, inconsistency, and/or hubris of one's competitors.

Returning to Sculley's example, Coke's packaging *was* superior and *did* provide it with a significant advantage. Other things being equal, had Pepsi managed to create a superior bottle, it was indeed possible that the advantage would have shifted to them. The difficulty and cost of discovering such a bottle were, after 30 years of effort, certainly well understood. Moreover, until Sculley's pioneering market research, there was no evidence that convenience (in the form of container size) was more important to consumers than design. Under these conditions, any rational Bayesian decision maker with strong priors that market leadership was strictly associated with bottle shape would have done exactly as Pepsi management did.

Sculley, who arrived on the scene with a different set of beliefs, embarked on a rational program of (costly) experimentation, which generated data that refuted earlier assumptions. Faced with this new evidence, Pepsi immediately changed its strategy. Analyzed in this way, we are not forced to invoke analytical ineptitude on the part of Pepsi management to explain Coke's extended period of market leadership. Having erroneous priors may be unfortunate, but it is not irrational.

In the general case, a management team must assess the consequences of a large number of strategic options within an environment that is complex, dynamic, and only partially understood. To proceed, management must identify which uncertainties are of strategic importance and then assess them. This requires a theory, henceforth *environmental theory*, that provides the basis for predicting the consequences of

[3] Sculley (1987, pp. 21–22).

[4] This literature is too large to survey here. Machina (1987) does summarize early work relating to rationality under uncertainty.

one's strategic moves and the evolution of the competitive environment. More specifically, an environmental theory takes as inputs the relevant history of events up to the present plus a contemplated action and provides, as output, a probabilistic forecast of likely consequences. Armed with such a theory, the agent may assess his available options and choose the one that best meets his objectives.

Firms whose environmental theories are consistent with the available information are called *subjectively rational*. Subjective rationality is, clearly, a weaker condition than objective rationality, which requires that management know the true environmental process driving results. The maximization of firm value with respect to subjectively rational beliefs is termed *subjective optimization*.

The strategic actions taken by a firm influence the consequences it ultimately experiences. This implies a direct connection between firm behavior and subjective rationality. Subjective rationality requires consistency between the consequences predicted by one's theory and those actually experienced. However, the consequences one experiences are, at least in part, the result of one's behavior. The interrelationships form a closed feedback loop: Managers take actions that are optimal given their beliefs, but what they believe is itself dependent upon the actions they take. When an environmental theory implies optimal actions that lead to theory-consistent consequences, it is said to be *self confirming*.

Why might self-confirming theories be problematic? By definition, after all, the actual results they generate are completely consistent with the results they predict. The potential for trouble arises from the fact that the theory's counterfactual predictions—that is, the predictions of what *would have happened* had the firm taken some other course of action—are not observed. A SCE is one in which all of the participants in a competitive situation subjectively optimize on the basis of their own self-confirming theories.

3. Formalism

Assume that there are n agents, $n \geq 1$, engaged in a dynamic competitive interaction. Each agent's situation is characterized by a nonempty (possibly infinite) set of *actions* A_i, a set of *consequences* C_i, and a *discount factor* $\delta_i \in (0, 1)$. In period t, agent i chooses action $a_i^t \in A_i$, after which a consequence $c_i^t \in C_i$ is stochastically determined. The agent then receives the payoff $\pi_i(a_i^t, c_i^t)$. Note that a_i^t and c_i^t may be multivalued; that is, they may be vectors or profiles with multiple components. In particular, c_i^t summarizes everything observed by i at the conclusion of period t. Typically, this includes both private and public information. For example, agents may observe each other's product offerings, market shares, and prices (i.e., these would be common components of c_i^t for all i), while own costs, customer identities, and organizational structure may be private (i.e., is own costs would only appear in c_i^t).

The *period t action profile*, $a^t = (a_1^t, \ldots, a_n^t)$, is a list of the actions taken by each agent in that period. Similarly, the *period t consequence profile*, $c^t = (c_1^t, \ldots, c_n^t)$, is a list of the consequences actually experienced by each agent at the conclusion of the period. A key primitive of the model is the *consequence-generating process*, a function that maps action profiles to probability distributions over consequences. Specifically, $p(c^t \mid a^t)$ is the probability that the profile of consequences c^t occurs conditional upon agents taking actions a^t. Thus, p is an exogenously specified stochastic process that links actions to outcomes (deterministic outcomes are a special case).

Prior to choosing a_i^t, agent i considers his own *individual history* up to that point, defined as $h_i^t \equiv (a_i^1, c_i^1, \ldots, a_i^{t-1}, c_i^{t-1})$.[5] Assuming the competitive interaction has no predetermined stopping point, let H_i^∞ be the set of infinite individual histories for player i with typical element $h_i^\infty = (a_i^1, c_i^1, a_i^2, c_i^2, \ldots)$.[6] A strategy for an agent is a prespecified plan that indicates, for every possible history, a precise course of action. To keep things general, agents are allowed to randomize over their actions. Thus, the *strategy for agent i* is a function s_i, where $s_i(a_i^t \mid h_i^t)$ is the probability with which i chooses action a_i^t, having observed individual

[5] Again, h_i^t may contain information in common with the histories of other agents.

[6] Finite-length interactions are a special case.

history h_i^t. A *strategy profile*, denoted $s = (s_1, \ldots, s_n)$, is a list specifying a strategy for each agent.[7]

Now consider the situation from the perspective of an individual agent. At the start of period t, agent i recalls his experience up to that point, h_i^t, and assesses the personal consequences associated with each potential action. Notice that what actually happens to i given some action a_i^t depends both on the process by which is opponents pick their actions, s_{-i}, as well as the process that links everyone's actions to their joint consequences, p. Thus, agent is decision problem can be summarized by an *environment response function*, denoted $e_i(c_i^t \mid a_i^t, h_i^t)$, which gives the probability that agent i experiences consequence c_i^t given his choice of a_i^t after having observed the individual history h_i^t. Note that e_i is an implication of s_{-i} and p (see the example below).

Without getting into the technical details, I simply note that s_i and e_i induce a probability distribution, denoted μ_{s_i, e_i}, on the set of outcome paths H_i^∞. The net present value to i of a particular path h_i^∞ is given by $v(h_i^\infty) = \sum_{t=1}^{\infty} \delta_i^{t-1} \pi_i(a_i^t, c_i^t)$. Therefore, the expected net present value to player i of choosing s_i under e_i is given by $E_v(s_i, e_i) = \int v(h_i^\infty)\mu_{s_i, e_i}(dh_i^\infty)$.

To help fix ideas, consider the simple, one-shot interaction presented in Figure 1. Agent 1's feasible actions are up (u) or down (d). Agent 2 observes what 1 did and then chooses left (l) or right (r). Following the agents' action choices, one of four consequences occur (typically, in a simple game like this, the consequences are just the payoffs to each of the agents). Here, p is deterministic; e.g., if action profile $a = (u, l)$ is chosen, then consequence c_{I} occurs with certainty. Knowing p is not sufficient to calculate e_1 since, for example, the consequence to agent 1 of choosing u is indeterminate until one specifies a strategy for agent 2. Suppose agent 2's strategy is as depicted (play l with probability 0.5 if u is observed, otherwise play l with probability 0.7). Then, $e_1(c_{\mathrm{II}} \mid u) = 0.5$, $e_1(c_{\mathrm{IV}} \mid d) = 0.7$, etc. Finally, the distribution over consequences can be calculated only when everyone's strategies are known. Suppose s_1 is that 1 chooses u with probability α. Then, for example, $\mu_{s_1, e_1}(c_{\mathrm{I}}) = 0.5\alpha$.

[7] When convenient, I write $s = (s_i, s_{-i})$, where s_{-i} indicates the strategies chosen by players other than i.

Figure 1 A Simple One-Shot Interaction

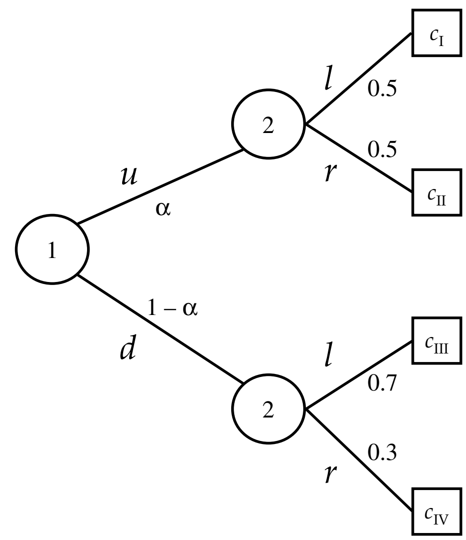

Agent i may or may not know e_i. Therefore, let $\tilde{e}_i$ denote agent is *subjective environmental theory*; so, $\tilde{e}_i(c_i^t \mid a_i^t, h_i^t)$ is the subjective probability i with which i believes consequence c_i^t will occur given the choice of a_i^t after observing personal history h_i^t.[8] Then, s_i and $\tilde{e}_i$ also induce a probability measure, in this case subjective, on H_i^∞. Let this distribution be represented by $\mu_{s_i, \tilde{e}_i}$. Then, is subjective expected net present value of strategy s_i is given by $E_v(s_i, \tilde{e}_i) = \int v(h_i^\infty)\mu_{s_i, \tilde{e}_i}(dh_i^\infty)$. The strategy s_i is said to be *subjectively optimal* under $\tilde{e}_i$ if, for all s_i', $E_v(s_i, \tilde{e}_i) \geq E_v(s_i', \tilde{e}_i)$. The strategy is *optimal* if, for all s_i', $E_v(s_i, e_i) \geq E_v(s_i', e_i)$.

Definition 1. Given a competitive interaction as described above, a SCE is a strategy profile $s = (s_1, \ldots, s_n)$ and a profile of subjective environmental theories $\tilde{e} = (\tilde{e}_1, \ldots, \tilde{e}_n)$ such that, for each agent, the following two conditions hold:

1. *Subjective optimization:* s_i is subjectively optimal with respect to $\tilde{e}_i$;
2. *Uncontradicted beliefs:* $\mu_{s_i, e_i} = \mu_{s_i, \tilde{e}_i}$.

The subjective optimization condition is self-explanatory. The uncontradicted beliefs requirement is that, in equilibrium, the observed events that unfold during play occur with the expected frequency.

[8] Most generally, $\tilde{e}_i$ may be constructed as the expectation implied by is subjective beliefs (probabilities) on a set of specific underlying theories, each considered by i to be a possible candidate for reality.

Equivalently, each agent's assessment of the probabilities with which histories are generated is correct. The key is that this degree of consistency does not imply correct counterfactuals; that is, there may be $s_i' \neq s_i$ for which $\mu_{s_i', e_i} \neq \mu_{s_i', \tilde{e}_i}$, thereby admitting the possibility that $E_v(s_i', e_i) > E_v(s_i, e_i)$.

It is important for the reader to take note that SCE is, inherently, a dynamic concept. The object in Condition 2 of the definition is a probability distribution on the space of *infinite play paths*. This condition requires correct conditional expectations given any finite history, which implies the possibility of equilibrium learning (e.g., certain equilibrium paths may cause agents to accurately refine their environmental theories). While some SCE will have the feature that agent behavior "settles down" to a state of static stability, others will not. In general, equilibrium strategies may be quite complex and imply subtle, dynamic interactions between the agents. Indeed, SCE is consistent with strategies that include dynamic experimentation.

4. The *P*-Challenge

To illustrate this formalism, I now present a stylized treatment of the Pepsi case discussed in §2. Assume that there are two firms, C and P, each of which supplies a cola product to a market that is differentiated on vertical and horizontal dimensions. The vertical dimension is packaging; that is, consumers agree on what constitutes a better package. The horizontal dimension is flavor; some people prefer C-cola while others prefer P-cola. Assume that the sequence of events is: (1) Soda packages are simultaneously chosen by each firm; (2) package choices are observed; and (3) product is sold and profit is determined.

Firm *i*s *product offering* is summarized by $a_i \equiv (a_i^d, a_i^s)$, where a_i^d is bottle design and a_i^s is bottle size, respectively. Assume $a_i^d, a_i^s \in \{0, 1\}$ with the following interpretation: a_i^d represents bottle design quality with 0 indicating basic and 1 indicating superior design; a_i^s represents bottle size, with 0 and 1 indicating, respectively, small and large sizes. An *action profile* is a pair $a = (a_P, a_C)$, which indicates the product offerings of both firms.

In keeping with Sculley's story, one package choice is not technologically feasible: High-quality designs must be produced in materials that preclude large bottle size (thereby ruling out $(1, 1)$ as a package choice). Assume low-quality/small packages $(0, 0)$ are weakly dominated in consumer preferences by $(1, 0)$ and $(0, 1)$. Let C begin with the ability to implement either basic or superior designs, $a_C^d \in \{0, 1\}$. P, on the other hand, may offer high-quality design only after expending a one-time product development investment of κ. Aside from this, assume the two firms' packaging technologies are symmetric. Let the relevant firm payoffs be given by Table 1.

Table 1

	C	
	(1, 0)	(0, 1)
P		
(1, 0)	$\pi-\kappa, \pi$	$\pi-\kappa+\Delta, \pi-\Delta$
(0, 1)	$\pi-\Delta, \pi+\Delta$	π, π

The interpretation is as follows. When both firms offer the same product, they split the market and each obtains an operating margin of $\pi > \kappa$ (which does not include product development costs). If firm i offers the high-quality/small package $(1, 0)$ when firm j offers the low-quality/large package, firm i has a profit adjustment of $+\Delta$ and firm j one of $-\Delta$. This represents the effect of changes in market share caused by the product asymmetry. Assume that consumers actually prefer convenient (larger) bottles; that is, $\Delta < 0$.

Let $\widetilde{\Delta}_i$ be the assessment by managers in firm i of Δ. The subjectively optimal package choice depends upon the sign of $\widetilde{\Delta}_i$. In particular, if $\widetilde{\Delta}_C > 0$, then offering $(1, 0)$ is a dominant strategy for C. Furthermore, if P shares this assessment ("shares" in the sense that $\widetilde{\Delta}_P > 0$) and if acquiring the high-quality design is cost effective ($\widetilde{\Delta}_P > \kappa$), then acquiring the high-quality bottle and offering it is also dominant for P. Thus, $a^* = ((1, 0), (1, 0))$ meets the subjective optimization condition of SCE, given beliefs $\widetilde{\Delta}_C > 0$ and $\widetilde{\Delta}_P > \kappa$. Moreover, a^* also meets the uncontradicted beliefs condition, since in this case Δ is not observed (firms offer identical product packages and split the market evenly).

A couple of additional points are worth mentioning. First, the subjective equilibrium described above is clearly not Nash: $(0, 1)$ is, objectively, a dominant strategy for both firms. Second, the subjective equilibrium in which high-quality designs are chosen results in a performance advantage for C (due to the design investment required of P). This advantage appears, to managers, to be a "quasi rent" resulting from the "investment barrier" associated with gaining a high-quality design. In reality, however, the source of the advantage is the pair of jointly self-confirming assessments regarding consumer preferences.

5. Self-Confirming Advantage in Multimarket Competition

An important open question in strategy is what can be said about the performance of multibusiness firms. Teece (1980) argues that specialized firms enjoy opportunities for diversification because, "there are *always* unused productive services which can be placed into employment" (emphasis added). He asserts that multiproduct diversification is chosen when transaction costs make it the economical choice over selling excess resource services on the open market.

Human capital, for example, is often identified as an important source of diversification economies. The argument is that know-how is often a common input to a variety of products and, therefore, exhibits a public-good-like quality; it can be used in multiple applications without depreciating its value. When such knowledge is tacit, it cannot be traded between firms and may, therefore, be a source of economic rent. This idea takes its normative form in Prahalad and Hamel (1990).

I say that this is an open question because empirical analyses designed to test the relationship between diversification and performance have produced mixed results. Montgomery and Wernerfelt (1988) find that Tobin's q is significantly negatively correlated with diversification. The authors interpret this finding as follows:

> …one would expect that the more widely a firm diversifies, the lower will be its average rents. Two points support this argument: first, wider diversification suggests the presence of less specific factors that normally yield less competitive advantage; second, a given factor will lose more value when transferred to markets that are less similar to that in which it originated.

The authors point out that an alternative explanation for their finding is that, "at some point firms *believed* that rents from diversification would be gained more easily than history has borne out" (emphasis added).

Montgomery and Hariharan (1991) find that a firm's propensity to enter new markets tends to be high given a high existing level of diversification. In other words, some firms can be classified as diversifying types. Moreover, these firms tend to maintain corporate focus; that is, they tend to enter industries that are close to their existing lines of business. The authors interpret this as evidence that firms with excess resource capacity typically seek ways to put them to productive use. However, using an extensive data set, McGahan (1999) finds that corporate focus is either unrelated or very loosely related to corporate performance.

In the analysis that follows, I demonstrate that these empirical findings are consistent with self-confirming equilibria in which some managers misjudge the sensitivity of their costs to market expansion and implement a suboptimal level of diversification. In such situations, the properly diversified firm may well sustain a performance advantage over its suboptimally diversified competitors—*provided the competitors' experiences jibe with their prior expectations*.

5.1. Dynamic Monopoly with Deterministic Costs

To lay the foundation for the multifirm case, we first consider a monopolist with the ability to produce two products. The problem facing the monopolist is its choice of product line: Should product 1, product 2, or both be produced? As in §4, this decision can be represented by a profile $a = (a_1, a_2)$ where $a_j = 1$ if the monopolist offers product j and 0 otherwise. The monopolist has discount rate δ.

Assume that there are two segments in the monopolist's consumer market. Let M_1^* denote the maximum aggregate margin (revenue less variable costs) available to the monopolist when only product 1 is offered. Similarly, let M_2^* and M_{12}^* be the maximum aggregate margins available when product 2 and both products 1 and 2 are offered, respectively. Assume: (i) $M_1^*, M_2^*, M_{12}^* > 0$ (offering either or both products is profitable), (ii) $M_1^*, M_2^* < M_{12}^*$ (consumers have heterogeneous preferences), and (iii) $M_{12}^* \leq M_1^* + M_2^*$ (consumers view the products as imperfect substitutes).

To offer a product, suppose the monopolist must develop an appropriate set of related competencies. Assume that the actual relationship between product offering and development costs is given by $e(a) = a_1 K_1 + a_2 K_2$. So, when only product 1 is offered, the development cost is K_1, whereas developing both products results in a cost of $K = K_1 + K_2$. In this example, then, e is the relevant consequence-generating process; the monopolist chooses a product mix a and generates product development expenses of $e(a)$.[9] Assume the parameters are such that $M_2^* - K_2 < 0$ but $M_{12}^* - K > 0$ (the objectively optimal strategy is to specialize in product 1).

Suppose the monopolist entertains two hypotheses regarding the relationship between cost and diversification. Hypothesis 1 is consistent with the true relationship, e. Hypothesis 2 is that competencies are like a public good—once acquired, they can be costlessly employed across all related products. Specifically, Hypothesis 2 is that costs equal zero if the firm does nothing, but are K otherwise; i.e., any feasible product mix can be developed for a fixed investment of K. In this case, it is critical that the monopolist's aggregate cost expectations are well calibrated; full diversification does imply development costs K. This assumption is relaxed later.

Let $\theta \in [0, 1]$ be the monopolist's initial prior that Hypothesis 1 is the true relationship. The monopolist updates beliefs in a Bayesian fashion. The following proposition says that in a world in which resources exhibit *no* scope economies, a strong belief that they do can result in suboptimal equilibrium behavior.

Proposition 1. *Given arbitrary initial prior $\theta \in [0, 1]$, play converges to one of two SCEs in at most two periods. If*

$$1 - \theta \geq \frac{K_2 - (M_{12}^* - M_1^*)}{K_2 - \delta(M_{12}^* - M_1^*)}, \qquad (5.1)$$

then the monopolist (*suboptimally*) *chooses full diversification in perpetuity. This is an SCE from period* 1 *on. If* (5.1) *fails, then the monopolist* (*optimally*) *offers product* 1 *in perpetuity. In the latter case, beliefs converge in period* 2 *and the SCE is attained.*

Thus, the erroneous belief that competencies behave like public goods can lead the subjectively rational monopolist to offer the full product line. By assumption, a cost level consistent with expectations is observed. The monopolist does not believe—and, after full diversification, observes no evidence to suggest—that development costs can be saved by not offering product 2. Unfortunately for our monopolist, there is such a savings, one that more than offsets the margin gained by offering the second product.

It is important to note (see the proof of Proposition 1) that the monopolist engages in a subjectively optimal level of experimentation; e.g., it is *not* the case that the monopolist suffers from a bias against seeking disconfirming evidence. The definition of SCE implies the adoption of dynamic strategies that include subjectively optimal experimentation schemes. The less patient the monopolist, the lower is the hurdle to making a bad decision (future potential benefits to experimentation are lower). This hurdle is also lower the greater the difference in margin between full and partial entry (current costs of experimentation are higher).

The suboptimal result is critically dependent upon the explicit assumption that the monopolist is well calibrated in terms of aggregate development cost as well as the implicit assumption that both markets can be entered simultaneously. Clearly, were entry to occur in some sequence, the firm would immediately realize that producing only product 1 is the optimal decision. I will have more to say about sequential entry in §5.3.

5.2. Static Oligopoly with Deterministic Costs

Assume a competitive interaction is structured as follows: (1) n firms simultaneously decide how many product variants to include in their product line, (2) firms observe the number of products in each others' product lines and then set prices, and (3) firms experience product development costs, gross margins, and profits. In this example, the action variable of

[9] In this (single-agent) situation, there is no distinction between the environment response function e and the consequence-generating process p.

interest is the number of products firm i opts to include in its product line, denoted a_i. Assume that a_i is chosen from the range $[0, k]$.[10] Let $a = (a_1, \ldots, a_n)$ be the profile detailing the number of differentiated products offered by each firm.

Assume that demand is such that (i) for all differentiated product profiles a, there exists a unique Nash equilibrium in the pricing stage, and (ii) the resulting gross margin for firm i as a function of a, denoted $m_i(a_i, a_{-i})$, is continuously differentiable in a, concave increasing in a_i and has $m_i(0, \cdot) = 0$. Note that this is consistent with the idea in Montgomery and Wernerfelt (1988) that increased diversification leads to positive but decreasing marginal returns.

Suppose the firms in the industry face a resource development technology that generates costs which are nondecreasing in the number of products offered. Specifically, assume firm *i*s *resource development technology*, denoted $r_i(a_i)$, is a continuous and quasi-convex function mapping the number of products offered by firm i to the cost level incurred for that development. Assume $r_i(0) = 0$. The *industry resource cost structure* is summarized by the profile $r = (r_1, \ldots, r_n)$, in which each firm may have a different technology.

Under these conditions, firm *i*s actual profit can be written as a function of the product offering profile: $v_i(a) = m_i(a_i, a_{-i}) - r_i(a_i)$.[11] Because v_i is differentiable in a and quasi-concave in a_i, for every industry resource cost structure r, there exists a Nash-equilibrium-differentiated product profile $a^r = (a_1^r, \ldots, a_n^r)$.[12]

[10] The assumption that firms can choose a fractional number of products may seem unrealistic. Keep in mind, however, that firms are permitted to choose probability distributions over their set of feasible actions. Therefore, a_i may be interpreted as the *expected* number of products scheduled for production by i.

[11] Note that the structure of the industry as elaborated above, combined with competitive strategies s_{-i}, implies an environment response function for firm i of $e^i(q_i^*, p_i^*, \bar{r}_i \mid s_i)$; i.e., e^i returns an actual vector of product demands and prices plus the development cost for any choice of product line and pricing. Thus, the expected return to firm i of a strategy k_i is given by

$$E_v(s_i, e^i) = \int \pi_i(s)\, d\mu_{s_i e^i}(s).$$

[12] This result depends upon the fact that $[0, k]$ is nonempty and compact.

As in the preceding section, assume that the only uncertainty facing firms is the behavior of product development costs. Firm i believes resource development costs are described by $\tilde{r}_i$ (also assumed to be continuously differentiable and quasi-concave in a_i). *Industry beliefs* are given by the profile $\tilde{r} = (\tilde{r}_1, \ldots, \tilde{r}_n)$. Firm *i*s *subjective* profit function is

$$\tilde{v}_i(a) = m_i(a) - \tilde{r}_i(a_i).$$

The following proposition states that any profile of beliefs regarding the behavior of resource development costs can support an SCE—provided that the subjectively optimal product line choices arising from these beliefs result in resource costs consistent with expectations.

Proposition 2. *Given an arbitrary profile of beliefs $\tilde{r}$, there exists a profile of product line offerings $\tilde{a}$ such that, for all i, $\tilde{a}_i$ optimizes $\tilde{v}_i(\cdot, \tilde{a}_{-i})$. If, for all i, $\tilde{r}_i(\tilde{a}_i) = r_i(\tilde{a}_i)$, then $(\tilde{a}, \tilde{r})$ constitutes an SCE.*

It should be noted that Proposition 2 does not imply suboptimal behavior. Even inaccurate beliefs may yet lead to behavior that is consistent with individual optimality; sometimes bad beliefs may lead to good behavior. The following proposition characterizes those situations in which bad beliefs lead to bad behavior. The potential for trouble arises when firms have bad assessments regarding the *shape* of their resource cost functions.

Proposition 3. *Let $(\tilde{a}, \tilde{r})$ be an SCE given the actual industry resource cost structure r. If $\partial r_i/\partial a_i > \partial \tilde{r}_i/\partial a_i$, then $\tilde{a}_i$ is suboptimal relative to $\tilde{a}_{-i}$. In particular, firm i offers an overdiversified product line relative to the objective optimum. Similarly, if $\partial r_i/\partial a_i < \partial \tilde{r}_i/\partial a_i$, then firm i underdiversifies.*

To understand the strategic implications of this proposition, suppose cost technologies are *identical* across firms (i.e., $r_1 = \cdots = r_n = r$). If firms have heterogeneous beliefs regarding their costs, the generic SCE exhibits heterogeneous diversification. The superior performers are the firms whose managers' beliefs are best aligned with reality.[13] Note well, superior performance in this situation does *not* arise from any real

[13] Relative advantage does not imply optimality. That is, firm i may outperform all competitors—even though its performance would have been even better had its beliefs been more accurate.

advantage in product development ability. Rather, it stems from erroneous yet self-confirming beliefs held by one's competitors. As we will see momentarily, such advantages may be sustainable in dynamic settings.

Thus, if subjective rationality is a reasonable approximation of reality, there is no reason to accept the economists' old saw that, "In a competitive environment, managers need not understand the true reasons for their success in order to drive an industry to an efficient equilibrium." This reasoning misses the mark. The pernicious problem for the underperformers in an SCE of this type is not that they choose optimal actions for the wrong reasons but, instead, that they choose suboptimal actions for the wrong reasons *and never have a clue that anything is amiss.*

Overdiversification arises when firm i underestimates the responsiveness of resource costs to product development activities. In particular, whenever firms believe their resource development technology exhibits the public goods property (i.e., when $\partial\tilde{r}_i/\partial a_i = 0$), full diversification is subjectively optimal. Figure 2 presents an example of this special case. Although development costs are linear in the level of product differentiation with slope γ, firm i believes they are constant. Here, $\partial r/\partial a_i = \gamma$ but $\partial\tilde{r}/\partial a_i = 0$, so the marginal cost of diversification is underestimated at every level. The result depicted is an SCE (since expected and actual costs equal R) in which firm i overdiversifies. The degree of suboptimality is worse the more responsive are development costs to the level of diversification (i.e., the larger is γ).

Figure 2 **Beliefs Induce Overdiversification in Equilibrium**

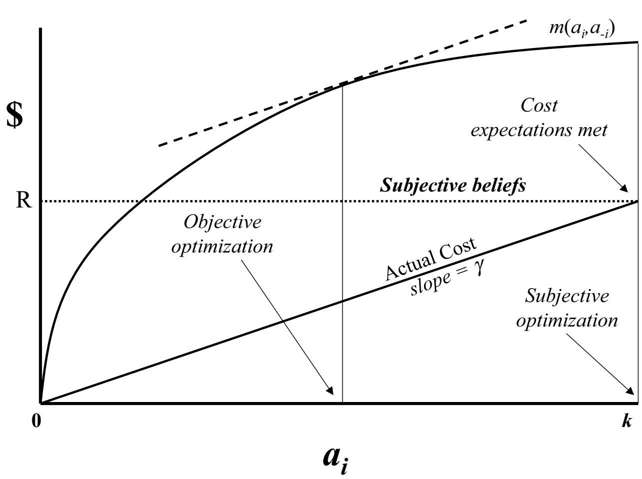

Considering Figure 2, it is tempting to question the likelihood of the proposed equilibrium since it relies on a "knife-edge" condition—the subjective cost curve intersects its objective counterpart exactly at the level of diversification that happens to be the subjective optimum. However, in the next section I extend this example to the dynamic case to show that problematic beliefs can be the end result of subjectively rational learning.

5.3. Dynamic Oligopoly with Stochastic Costs

Suppose individual resource development costs are linear (as in Figure 2) and, further, that this is common knowledge in the industry. Assume that the true cost relationship for firm i is stochastic and given by

$$r_i(a_i^t) = \gamma_i a_i^t + \kappa_i + \varepsilon^t,$$

where the ε^t are i.i.d. shocks with mean and variance $(0, \sigma_\varepsilon)$. Now, rather than an up-front resource investment, assume that r_i represents a per-period resource-maintenance cost.

In the beginning of period t, assume the managers of firm i are rational Bayesians who believe the slope and intercept of their own cost function are summarized by

$$\tilde{r}_i(a_i^t) = \tilde{\gamma}_i a_i^t + \tilde{\kappa}_i^t + \varepsilon^t,$$

where $\tilde{\kappa}_i^t$ is firm is belief regarding the intercept in period t. Assume the firm's initial priors are such that $R_0 < \gamma_i k$ (where $\gamma_i k$ is the true cost of developing a full product line). We require that $0 \leq \gamma_i, \tilde{\gamma}_i$ and $\sigma_\varepsilon, \tilde{\sigma}_0 < \infty$. For example, the situation in which firm i believes $\tilde{\gamma}_i = 0$ (managers hypothesize that resource costs are unresponsive to the degree of product differentiation) is depicted in Figure 3. Let $\tilde{a}_i^*$ solve $\partial m/\partial a_i = \tilde{\gamma}_i$ and assume that $m(\tilde{a}_i^*, a_{-i}) - (\tilde{\gamma}_i\tilde{a}_i^* + \tilde{\kappa}_i^0) > 0$ (so, entry into some market is subjectively optimal). Define $\kappa_i^* \equiv (\gamma_i - \tilde{\gamma}_i)\tilde{a}_i^* + \kappa_i$.

Condition 1. The beliefs of each firm satisfy

1. slope dogmatism: i believes $\Pr(\gamma_i = \tilde{\gamma}_i) = 1$, and
2. intercept grain-of-truth: κ_i^* is in the support of is beliefs.

Proposition 4. *Given Condition* 1, *the unique SCE occurs when each firm i chooses* $\tilde{a}_i^*$ *in each period. In the*

Figure 3 Actual Cost *Level* Is Learned

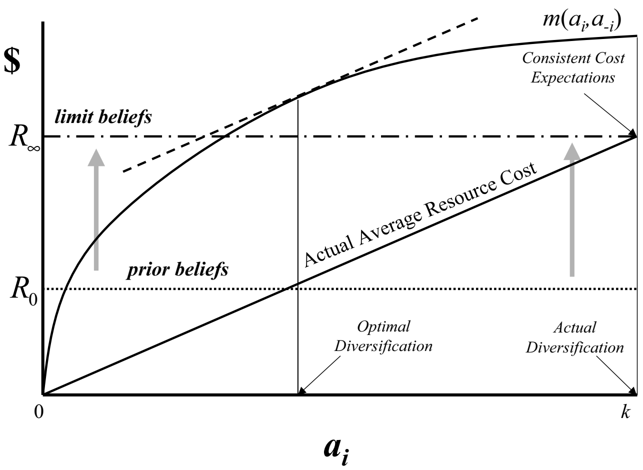

limit, each firm becomes well calibrated with respect to its own cost level. That is,

$$\lim_{t\to\infty} \tilde{\kappa}_i^t \to \kappa_i^*, \tag{5.2}$$

which immediately implies

$$\lim_{t\to\infty} \tilde{r}_i(\tilde{a}_i^*) = \gamma_i \tilde{a}_i^* + \kappa_i.$$

Diversification at $\tilde{a}_i^*$ is suboptimal when $\gamma_i \neq \tilde{\gamma}_i$. In such cases, perhaps in contradiction to intuition, repeated shocks to the firm's actual resource costs cause its initial self-confirming beliefs to be reinforced and refined—in the sense that the true cost of $\tilde{a}_i^*$, $r_i(\tilde{a}_i^*)$, is eventually known to a very high degree of precision. Although the observed data do correct certain errors in management's initial priors (regarding the cost level associated with the subjectively optimal level of diversification), it does not affect (indeed, by the assumption of dogmatic beliefs on this dimension, it *cannot* affect) managerial assessments regarding the shape of r. In this case, bad self-confirming beliefs on the part of one's competitors may be the *source* of sustained advantage.

Furthermore, when costs are stochastic, instantaneous entry is no longer a necessary condition for such an outcome (as it was in the previous cases). For example, suppose firm i sets a diversification goal of $\tilde{a}_i^*$ but is restricted to entering one market per period, perhaps due to technological or resource limitations. Then, it is still possible (though not guaranteed) that firm i will eventually enter a steady state diversified at $\tilde{a}_i^*$ with beliefs regarding the expected cost converging to $\gamma_i \tilde{a}_i^* + \kappa_i$. This is true because random cost outcomes during the early stages may be insufficient to overturn the erroneous beliefs. Stated more precisely, for every sequential implementation of $\tilde{a}_i^*$, there exists a cost history (typically many) such that $\tilde{a}_i^*$ is implemented as the subjectively optimal diversification strategy.

The assumption of purely dogmatic beliefs is clearly unrealistic. This assumption can be eliminated with somewhat weaker, but similar, results. Still, the preceding analysis does provide a rough approximation of how sufficiently "strong" beliefs may cause competitive trouble in certain situations. Furthermore, these results are consistent with the empirical findings mentioned above; that is, certain firms have a high propensity to diversify and that those that do tend to underperform their competitors.[14] The preceding analysis is also useful because it casts in sharp relief the conditions that lead to bad outcomes: Erroneous priors supported by a grain of truth (i.e., regarding costs at the subjectively optimal level of diversification).

6. Issues and Extensions

The preceding section demonstrates that performance heterogeneity can arise, under the right conditions, only as the result of diverse worldviews and not, in particular, as a result of differences in any "real" abilities (e.g., technological efficiencies or product qualities). Furthermore, this demonstration does not rely upon agent illogicality, biased decision processes, or any behavioral motivation other than profit maximization.

Up to this point, I have emphasized the ways in which subjective rationality may result in some firms earning sustained supranormal profits relative to their competitive peer group. However, such an outcome

[14] Of course, these findings are consistent with other hypotheses as well. It may be that diversification is prompted by slack resource capacity and that although such diversification is value maximizing, the marginal returns to diversification are decreasing. It may be (as Penrose 1959 and others have suggested) that managers like to build large organizations against the interests of their shareholders.

is not necessarily an implication of the theory. For example, Kalai and Lehrer (1995) show that in an oligopoly, if firm managers believe they are in a competitive market, the subjectively rational response is to produce the competitive level of output. This results in market prices equal to marginal cost (as predicted in a competitive model), even though the objectively rational equilibrium (Nash) predicts pricing above marginal costs. Thus, an industry may become highly competitive, in essence, because its participants believe it must be so.

Thus, like any analytical tool, SCE must be used carefully; assumptions about what agents do and do not know have a significant impact on the predictions of the theory. Indeed, any outcome can be supported as subjectively rational given complete ignorance of payoffs and sufficiently extreme beliefs; e.g., make any action profile a an equilibrium by endowing agents with the belief, "If I do not undertake action a_i, I will be abducted and eviscerated by space aliens." One would expect that in most applications of interest to strategy researchers, agents' environmental theories are sufficiently bounded by reality to exclude outcomes supported by wildly implausible beliefs.

Speculating on an extension *far* beyond the scope of the formal results contained in this paper, SCE may provide a nice foundation for a positive theory of performance advantage over the evolution of an industry. The thinking is as follows. At some point, a body of technological, scientific, and/or social knowledge becomes ripe for commercial exploitation. A number of firms attempt to do so, thereby creating a new industry. Initially, very little is known about how things work, competing environmental theories are diverse, and, in this pre-equilibrium stage, virtually anything goes. Over time, enough of the environment is fleshed out to reach equilibrium. In this stage, self-confirming beliefs constitute a stable "paradigm" by which managers conduct their business. Still, knowledge about the environment is coarse and performance heterogeneity is substantial.

Periodically, an innovator appears (à la Sculley) whose beliefs lead to subjectively rational experimentation that, in turn, refines managerial understanding of the world. Entrepreneurial profits accrue to the innovator for a time as the diffusion of knowledge leads to a new, more efficient equilibrium, and so on. Periods of relative stability are punctuated by innovation shake-ups (periods of disequilibrium induced by the failure of old paradigms to accurately predict outcomes). I conjecture that such innovations need not be based upon radical new technologies, but instead might employ existing technologies in previously unimagined ways. Open questions regarding the industry trajectory are: Is the history of firm knowledge always one of refinement? Is industry efficiency ever-increasing? Is performance heterogeneity stable, increasing, or decreasing over time? Are there certain industry primitives or specific firm strategies that increase the sustainability of self-confirming advantage?

Here is an example that seems consistent with my conjecture regarding innovation. According to Drucker (1985), by the 1950s, the ocean freight industry was in trouble—with air carriers making substantial inroads. Assessing this threat within the context of its own widely held beliefs regarding ocean transportation, the industry reacted by building faster ships that required less fuel and smaller crews. These efforts were unsuccessful. According to Drucker, *most outside observers* predicted the industry's imminent demise. However, by 1985 ocean freight traffic increased five-fold, costs dropped 60%, and port time, pilferage, and congestion were down 75%. This turnaround was due to the introduction of containerization—an innovation that depended upon no new technology and, presumably, could have been implemented decades earlier.[15]

Mulling this over, the careful reader may be wondering, "If SCE is, as advertised, designed to capture dynamic learning, why are 'periods of disequilibrium' a necessary ingredient in the preceding story? Can't we imagine a model in which the entire path of industry evolution is consistent with a single SCE?" The answer to the latter question is, of course, yes. In the revised story, agents begin with heterogeneous,

[15] As was pointed out by an anonymous referee, the switch to containerization required industrywide coordination—a difficult and expensive hurdle to make. So, although some shippers may have concluded long before that containerization was the way to go, it may have taken the harsh reality of air competition to force the doubters and recalcitrants into line.

diffuse prior beliefs and adopt subjectively optimal strategies; the equilibrium belief/strategy combination induces periods of substantial experimentation interspersed with periods of relative stability; along certain paths, some firms may fail while others thrive; sometimes a firm innovates in a way that fundamentally changes the way business is conducted.

The point here is that the uncontradicted beliefs requirement of SCE is a subtle one; e.g., it is consistent with agents accurately assessing the *chances* that a given firm will fail, that another will periodically innovate, and that said innovation occasionally induces shifts in industry dynamics. Entrants correctly assess entry to be a positive net present value action at the start—but this is not inconsistent with a correct conditional belief, *at a point along some path*, that later implies the subjective optimality of exit. Innovators need not be viewed as actors dropped into the environment from on high but, instead, as participants whose initial beliefs, updated under certain histories, cause them to experiment late in the game. Throughout, the evolution of performance advantage may be driven by beliefs that are fundamentally erroneous on key counterfactual assessments.

Empiricists may find this last point unsettling because the implication is: If a researcher is privy to no better information than that upon which managers make their own assessments, then he or she cannot distinguish performance advantages due to real differences in capabilities from those arising merely from differences in self-confirming worldviews. Under such circumstances, neither managers nor researchers have access to data disconfirming any subjectively rational theory.

This does not imply, however, that the theory is without empirical content. For example, a theory that assumes complete knowledge of the environmental response function is hard-pressed to explain the class of innovations just described (i.e., that are based upon the imaginative use of existing technologies).[16] At the same time, SCE does predict that managers undertake projects they sincerely believe are profit maximizing, that they implement strategies that are forward-looking and anticipatory, and that they update rationally in the face of unfolding events—all of which are inconsistent with many nonequilibrium theories that assume specific, logical flaws in agent decision making.

[16] Sculley and containerization are but two examples. Drucker's book is full of others.

Acknowledgments

Previously titled "When Competencies Are Not Core: Self-Confirming Theories and the Destruction of Firm Value." The author thanks N. Argyres, J. Baum, J. Brickley, R. Gibbons, H. Harmon, R. Henderson, J. Liebeskind, G. MacDonald, L. Marx, G. Saloner, C. Smith, O. Sorensen, J. Zimmerman, two anonymous referees, and participants at the University of Rochester theory workshop, the Stanford conference on strategy and the Strategic Research Forum for their insights and suggestions.

Appendix. Proofs of the Propositions

A.1. Proposition 1

The monopolist knows that $(0,0)$ and $(0,1)$ are suboptimal. If $a^1 = (1,1)$ is subjectively optimal, then $a^t = (1,1)$ is optimal for $t = 2, 3, \ldots$ because no information arises to change the monopolist's priors. Thus, the net present value when $a^1 = (1,1)$ is optimal is $\widetilde{V}_{(1,1)} = (M^*_{12} - K)/(1-\delta)$. If, on the other hand, $a^1 = (1,0)$ is chosen, the expected payoff in the current period is $\theta(M^*_1 - K_1) + (1-\theta)(M^*_1 - K)$. This choice reveals the true process with certainty (costs either equal K_1 or they equal K). Thus, with subjective probability θ, the monopolist discovers there are no scope economies and, optimizing thereafter, receives $M^*_1 - K_1$ per period and, similarly, with probability $(1-\theta)$ receives $M^*_{12} - K$. Thus,

$$\widetilde{V}_{(1,0)} = \theta\left[M^*_1 - K_1 + \delta\frac{M^*_1 - K_1}{1-\delta}\right] + (1-\theta)\left[M^*_1 - K + \delta\frac{M^*_{12} - K}{1-\delta}\right].$$

The monopolist chooses $(1,1)$ when $\widetilde{V}_{(1,1)} \geq \widetilde{V}_{(1,0)}$. Solving for the indifference point $\bar{\theta}$

$$\frac{M^*_{12} - K}{1-\delta} = \bar{\theta}\left[M^*_1 - K_1 + \delta\frac{M^*_1 - K_1}{1-\delta}\right] + (1-\bar{\theta})\left[M^*_1 - K + \delta\frac{M^*_{12} - K}{1-\delta}\right]$$

$$\bar{\theta} = \frac{(1-\delta)(M^*_{12} - M^*_1)}{K_2 - \delta(M^*_{12} - M^*_1)}.$$

If $\theta \leq \bar{\theta}$, then the monopolist chooses full diversification in perpetuity and, if not, only product 1 is produced in perpetuity.

A.2. Proposition 2

The conclusion follows directly from the existence of a fixed point in the subjective best-response correspondence and the definition of an SCE.

A.3. Proposition 3

The objective best response to $\tilde{a}_{-i}$ is the solution a^*_i to

$$\frac{\partial m(\cdot, \tilde{a}_{-i})}{\partial a_i} = \frac{\partial r_i}{\partial a_i}.$$

The subjective optimum $\tilde{a}_i^*$ solves

$$\frac{\partial m(\cdot, \tilde{a}_{-i})}{\partial a_i} = \frac{\partial \tilde{r}_i}{\partial a_i}.$$

Given the concavity of m, $\partial r_i/\partial a_i > \partial \tilde{r}_i/\partial a_i$ implies $\tilde{a}_i^* > a_i^*$.

A.4. Proposition 4

By construction, $\tilde{a}_i^*$ is the unique best response to the beliefs of firm i in each period. Condition (5.2) is an immediate consequence of Bayesian updating and the weak law of large numbers (and, of course, the assumption that κ_i^* is in the support of is beliefs).

References

Abreu, D., D. Pearce, et al. 1990. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica* **58**(5) 1041–1063.

Battigalli, P., D. Guatoli. 1988. Conjectural equilibria and rationalizability in a macroeconomic game with incomplete information. Mimeo, Instituto Economia Politica, Milan, Italy.

Blume, L. E., D. Easley. 1992. Rational expectations and rational learning. Cornell University, Ithaca, NY.

Debreu, G. 1959. *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*. Yale University Press, New Haven, CT.

Demsetz, H. 1988. *Efficiency, Competition and Policy*. Basil Blackwell, Inc, Cambridge, MA.

Drucker, P. F. 1985. *Innovation and Entrepreneurship; Practice and Principles*. Harper & Row, New York.

Fudenberg, D., D. K. Levine. 1993. Self-confirming equilibrium. *Econometrica* **61**(3) 523–545.

Gittens, J., D. Jones. 1974. A dynamic allocation index for the sequential allocation of experiments. J. Gani et al., eds. *Progress in Statistics*. North Holland, Amsterdam, The Netherlands, 241–266.

Harsanyi, J. 1967–1968. Games with incomplete information played by Bayesian players. *Management Sci.* **14** 159–182, 320–334, 486–502.

Hirshleifer, J., J. G. Riley. 1992. *The Analytics of Uncertainty and Information*. Cambridge University Press, Cambridge, MA.

Kalai, E., E. Lehrer. 1993. Subjective equilibrium in repeated games. *Econometrica* **61** 1231–1240.

——, ——. 1995. Subjective games and equilibria. *Games Econom. Behavior* **8** 123–163.

Lippman, S. A., R. P. Rumelt. 1982. Uncertain imitability: An analysis of interfirm differences in efficiency under competition. *Bell J. Econom.* **13** 418–438.

Machina, M. J. 1987. Choice under uncertainty: Problems solved and unsolved. *J. Econom. Perspectives* **1**(1) 121–154.

McGahan, A. M. 1999. The performance of US corporations: 1981–1994. *J. Indust. Econom.* **XLVII**(4) 373–398.

Montgomery, C. A., S. Hariharan. 1991. Diversified expansion by large established firms. *J. Econom. Behavior Organ.* **15** 71–89.

——, B. Wernerfelt. 1988. Diversification, Ricardian rents, and Tobin's q. *RAND J. Econom.* **19**(4) 623–632.

Nash, J. 1950. Equilibrium points in *N*-person games. *Proc. National Acad. Sci.* **36** 48–49.

Penrose, E. 1959. *The Theory of Growth of the Firm*, 2nd ed. Oxford University Press, New York.

Prahalad, C. K. G. Hamel. 1990. The core-competence of the corporation. *Harvard Bus. Rev.* 79–91.

Russo, J. E., P. J. H. Schoemaker. 1989. *Decision Traps*. Doubleday, New York.

Sculley, J. 1987. *Odyssey*. Harper and Row, Cambridge, MA.

Teece, D. J. 1980. Economies of scope and the scope of the enterprise. *J. Econom. Behavior Organ.* **1** 223–247.

Accepted by Rebecca Henderson; received May 2000. This paper was with the author $3\frac{1}{2}$ months for 3 revisions.