Econometrica, Vol. 61, No. 1 (January, 1993), 199–205

A "REPUTATION" REFINEMENT WITHOUT EQUILIBRIUM

BY JOEL WATSON [1]

THE ECONOMIC LITERATURE concerning agents' reputations has grown steadily since the seminal work of Kreps, Milgrom, Roberts, and Wilson.[2] Early work focused on how incomplete information leads to equilibria that are vastly different (but more intuitive) than those possible in the complete information game. Recently, however, game theorists have been studying how incomplete information might *refine* the set of equilibria.[3] One important class of games is that in which a single long-run agent plays a simultaneous move (stage) game with a sequence of opponents, each of whom plays only once, yet observes all previous play. Fudenberg and Levine (1989) study the reputation of the long-run player in this type of game. They argue that the "'most reasonable' equilibrium is the one which the long-run player most prefers." Their intuition is sustained when one perturbs the game with the "Stackelberg strategy." Fudenberg and Levine show that in the perturbed game the equilibrium payoffs of the long-run player are bounded below by a number that converges to the "Stackelberg payoff."

Fudenberg and Levine (and the others who have developed reputation models) take the notion of Nash equilibrium as fundamental in the analysis. However, it would seem as though our intuition about reputations relies little on equilibrium concepts. This leads to two questions. First, can meaningful reputations develop apart from equilibria? Second, if so, under what circumstances can reputations develop?

As I will demonstrate, equilibrium concepts are not required in order for players to establish significant reputations. I study (following Fudenberg and Levine (1989)) games in which a long-run player faces a sequence of short-run opponents. Like Fudenberg and Levine, I consider perturbations of the game involving the Stackelberg strategy. However, whereas they focus on equilibria, I will only require that players "best-respond" to their beliefs. Whenever the conjectures of the short-run players are "generally comparable" (e.g. contained in a compact set), I obtain the same refinement as do Fudenberg and Levine, but without an equilibrium assumption. What is important for reputations is that the beliefs of the short-run agents not be too dispersed in a sense to be made precise. Players can thus establish meaningful reputations from within the loose confines of individual rationality.

This paper borrows heavily from the work of Fudenberg and Levine (1989). In fact, their statistical result (their Lemma 1), which establishes the potential gain of building a reputation, requires no notion of equilibrium. It does require that the short-run agents hold the *same* belief, which is implied by equilibrium. I simply invoke their lemma in a more general setting (in which short-run players may hold different beliefs) and study the type of beliefs which allow it to refine the set of rational outcomes.

Note that a similar style of research has been followed on another front as well. Cho (1991) extends the Coase conjecture to a nonequilibrium setting. Cho's work is similar to mine in that we both take as fundamental a rationalizability notion. Our analyses require additional restrictions, however, and it is the nature of these restrictions in which our

[1] I am grateful to David Kreps, Marco LiCalzi, two referees, and the editor for comments. This is Chapter 2 of my Ph.D. dissertation for Stanford University Graduate School of Business.

[2] Kreps and Wilson (1982) and Milgrom and Roberts (1982) resolve Selten's (1977) chain-store paradox by studying the incumbent firm's potential reputation in an incomplete information game. Kreps, Milgrom, Roberts, and Wilson (1982) use the same technique to support cooperation in the finitely repeated prisoners' dilemma.

[3] Examples of such inquiries are Aumann and Sorin (1989), Fudenberg and Levine (1989, 1991), and Watson (1992a, b).

methodologies differ. My model includes an infinite number of short-run players and I must require that their beliefs not be too dispersed in order for a reputation by the long-run player to be viable. Cho (1991) studies learning schemes of the buyer in the standard one-sided offer bargaining model under one-sided incomplete information. He invokes rationalizability over the learning schemes of the buyer and the strategies of the seller. Cho restricts the buyer's forecasting schemes to be stationary and to satisfy a monotonicity condition, and argues that stationarity (rather than equilibrium) drives the Coase conjecture.

1. THE BASIC LONG-RUN / SHORT-RUN GAME

Suppose a long-run player (hereafter known as player 1) faces an infinite sequence of short-run players (players 2). Label each player 2 with the period in which this player faces player 1. That is, player 2_n faces player 1 in period n. Each period the long-run player and this period's short-run player play the finite stage game $G = \{A_1, A_2; u_1, u_2\}$. Each player 2 plays only once, but observes all previous play. In the stage game, player 1's action space is A_1 (from which he selects an action each period), while each player 2's action space is A_2. The corresponding sets of mixed actions are denoted ΔA_1 and ΔA_2. The stage game payoffs of player i ($i = 1, 2$) are given by u_i: $A_1 \times A_2 \to \boldsymbol{R}$, which is defined over action profiles and is extended to the space of independent, mixed action profiles by means of an expected payoff calculation.

Let BR: $\Delta A_1 \to \Delta A_2$ be the best-response correspondence of the players 2, and by a slight abuse of notation, let $a_i \in A_i$ denote the mixed strategy for player i that assigns all probability to action a_i. If some "best-responding" player 2_n holds conjecture $\alpha_1 \in \Delta A_1$ about player 1's action in period n, then this player will select a strategy from $BR(\alpha_1)$.

Player 1's *Stackelberg* (*stage game*) *payoff* is the greatest payoff player 1 can be guaranteed if he is able to commit to an action. That is, let

$$u_1^* \equiv \max_{a_1 \in A_1} \min_{\alpha_2 \in BR(a_1)} u_1(a_1, \alpha_2)$$

denote the (pure strategy) Stackelberg payoff, and let a_1^* satisfy

$$\min_{\alpha_2 \in BR(a_1^*)} u_1(a_1^*, \alpha_2) = u_1^*.$$

Action a_1^* is player 1's *Stackelberg action*. Also, let $\underline{u}_1 \equiv \min_{a \in A_1 \times A_2} u_1(a)$.

The repeated game described above shall be denoted G_δ^{LS}, where $\delta \in (0, 1)$ is player 1's discount factor. Let the long-run player's (stage game) payoff in period n be given by u_1^n, for all $n = 1, 2, \ldots$. The payoff of the long-run player in the supergame is simply the normalized, discounted sum of stage game payoffs: $(1-\delta)\sum_{n=1}^{\infty} \delta^{n-1} u_1^n$. The payoff of each short-run player is the stage game payoff in the period in which she faces the long-run player.

In each period, the long-run and short-run players can condition their actions on the entire past history of play. Let H_n denote the set of possible histories through period n. That is, $H_n \equiv (A_1 \times A_2)^n$. Also, let $H \equiv \bigcup_{n=0}^{\infty} H_n$ be the set of all possible histories, and let $H_\infty \equiv (A_1 \times A_2)^\infty$ be the set of infinite histories. Note that $H_0 \equiv \varnothing$ represents the "history" at the start of the game. Given some infinite history $h_\infty \in H_\infty$, let $T(h_\infty)$ be the set of finite histories that agree with h_∞. I will focus on histories in which player 1 always plays the Stackelberg action a_1^*. For simplicity, define H_∞^* as the set of all infinite histories in which player 1 always plays a_1^*, let $H^* \equiv T(H_\infty^*)$, and let $H_n^* \equiv H_n \cap H^*$.

A (supergame) strategy for player 1 is a mapping s_1: $H \to A_1$. A mixed (behavior) strategy for player 1 is a mapping σ_1: $H \to \Delta A_1$. Likewise, a strategy for player 2_n is a mapping s_2^n: $H_n \to A_2$, $n = 1, 2, \ldots,$ and mixed strategies are defined analogously. Shortly, I will consider the perturbed game in which player 1 may be forced to adopt the

Stackelberg strategy s_1^*, which takes the action a_1^* in each period, regardless of the history. That is, $s_1^*(h) \equiv a_1^*$, for all $h \in H$.

Fudenberg, Kreps, and Maskin (1990) show that a kind of Folk Theorem holds for the game just defined. More importantly for this paper, the set of *rationalizable strategies* (Bernheim (1984) and Pearce (1984)) leads to a wide range of outcomes as well. With common knowledge of individual rationality, there may be beliefs of the long-run player which lead to payoffs that are significantly below the Stackelberg payoff, even as δ approaches unity. I will show that best-response behavior (weaker than rationalizability) leads to a significant refinement in the expected payoffs of player 1 if one allows for a slight perturbation of the game. In fact, this payoff refinement is the same as that obtained by Fudenberg and Levine!

2. THE PERTURBED GAME AND RESULTS

In the perturbed version of G_δ^{LS}, player 1 may not be a "rational" player. With probability ε player 1 is the perturbation that adopts the Stackelberg strategy s_1^*. Other perturbations are allowed but do not change the results to follow, as long as player 1 is rational with some positive probability. So then, formally assume that before G_δ^{LS} is played, nature selects player 1's *type*. With probability ε player 1 is the Stackelberg strategy, and with probability γ player 1 is rational. (If $\gamma + \varepsilon < 1$ then other perturbations are selected with positive probability.) Denote this perturbed game $G_\delta^{LS}(\varepsilon, s_1^*)$.[4] I will focus on the expected payoff of the rational player 1 in this perturbed game.

At each information set in such a game (after all histories), the players entertain beliefs about the strategies of their opponents. Player 1 conjectures about the strategy that each player 2 employs, and each player 2 holds some belief about player 1's strategy and the strategies chosen by the other players 2. I will not need to characterize all of these beliefs. Rather, I will only need to formulate the conjecture of each player 2 concerning player 1's choice of action a_1^* (after each history).

Formally, each player 2 holds a *system of conjectures* $\pi \equiv \{\pi^h\}_{h \in H}$, which specifies this player's belief concerning the likelihood that player 1 will choose the Stackelberg action after each particular history $h \in H$. That is, given a system of conjectures π and a history $h \in H$, π^h is the probability that this player assigns to player 1 selecting a_1^* after h. Note that π does not represent a conjecture about player 1's *type*. It merely gives the probability that this agent believes player 1 (rational or a perturbation) will select the Stackelberg action. Let Π be the set of all systems of conjectures that are consistent with Bayes' Law and the perturbation of the game.

Each player 2 plays in only one period, and this period is the only one in which her conjecture is vital to her choice of action. However, players 2 will learn from all preceding play (and will update their conjectures accordingly). Furthermore, I will need to compare the conjectures of the players 2. For these reasons, I have defined systems of conjectures for all of the players 2. When it is necessary to identify a particular player 2_n, I will denote as π_n this player's system of conjectures in the game $G_\delta^{LS}(\varepsilon, s_1^*)$. Otherwise I will drop the subscript.

Assume that all the players know the form of the game being played. Since the game is perturbed, there are constraints on the system of conjectures $\pi \in \Pi$ of each player 2. First, it must be that $\pi^\varnothing \geq \varepsilon$. Furthermore, by Bayes' Law, $\pi^h \geq \varepsilon$ for each $h \in H^*$. Define the pseudo-metric d: $\Pi \times \Pi \to \boldsymbol{R}_+$ as follows.

$$\text{For } \pi, \mu \in \Pi, \qquad d(\pi, \mu) \equiv \sup\left\{|\pi^h - \mu^h| \;|h \in H^*\right\}.$$

[4] An alternative interpretation of the perturbed game, which Fudenberg and Levine (1989) adopt, is that "irrational" types are actually best-responding agents, but have different payoffs than their "rational" counterparts. For instance, the Stackelberg type may be thought of as having payoffs such that the Stackelberg strategy is dominant.

This pseudo-metric measures the difference in the conjectures about the Stackelberg action over all histories in which a_1^* is always played. (Π, d) is thus a pseudo-metric space.[5] Note, though, that Π is not compact.

I wish to assume a minimal amount of individual rationality. Essentially, I require that player 1 believe (or know) that each player 2 holds conjectures that are consistent with Bayes' Law and that each player 2 best-responds to her conjecture when called upon to play the stage game. I also require that all players know the form of the game, and that the rational player 1 selects a strategy that is a best-response to his belief concerning the strategies of his opponents. Note that these assumptions are much weaker than assuming that players reach an equilibrium in the game. In fact, these assumptions are weaker than *rationalizability*.[6] In particular, players may entertain beliefs that are not compatible with the common knowledge of individual rationality, much less compatible with each other.

In addition to the assumptions above, I require one further assumption concerning what the long-run player believes about the strategies of his opponents. A few more definitions will simplify matters. Given $\pi \in \Pi$ and $r > 0$, let $B_r(\pi) \equiv \{\pi' \in \Pi \mid d(\pi, \pi') < r\}$ be the ball of radius r centered at π. $\boldsymbol{P}$ shall denote the set of positive integers.

DEFINITION 1: Take a function $k: \boldsymbol{R}_+ \to \boldsymbol{P}$. A set $\Lambda \subset \Pi$ is said to be *of size* k if and only if for each $r > 0$, Λ can be covered by $k(r)$ balls of radius r.

Note that each compact set is of size k for some $k: \boldsymbol{R}_+ \to \boldsymbol{P}$, because compact sets are totally bounded (by the Bolzano-Weierstrass characterization of compact sets). Also, any subset of a compact set (any *conditionally compact* set) conforms to this definition.

I require player 1 to believe that the systems of conjectures of the short-run players are contained in *some* set of size k, for some $k: \boldsymbol{R}_+ \to \boldsymbol{P}$. This doesn't mandate that player 1 know in what set (of size k) the conjectures of the short-run players reside. I only require player 1 to believe that, whatever are the conjectures of the players 2, the conjectures are not too dispersed.

The nature of "dispersed conjectures" is not intuitively portrayed by Definition 1. For this reason, I think a few examples are in order. First, suppose $\Lambda \subset \Pi$ is a finite set. All finite sets are compact, so this certainly satisfies Definition 1. It also demonstrates how arbitrary the conjectures of the players 2 can be. Each player 2 can hold any system of conjectures, as long as the set of all the conjectures is finite.

However, much more can be accommodated by Definition 1. Suppose we start with a finite number of arbitrary systems of conjectures, $\mu_1, \mu_2, \ldots, \mu_L \in \Pi$. Then let $\Lambda \subset \Pi$ be defined as the convex hull of $\{\mu_1, \mu_2, \ldots, \mu_L\}$; that is, $\pi \in \Lambda$ if there is some x from the L-dimensional unit simplex such that $\pi = x \cdot (\mu_1, \ldots, \mu_L)$. By construction, Λ can be a very large (uncountably infinite) set and can include a wide variety of conjectures. In fact, Λ is easily seen to be compact and thus obeys Definition 1.

As a simple example, suppose we have some finite set of mixed strategies of player 1. Then if each player 2 holds a belief whose support is contained in this finite set, the corresponding set of systems of conjectures is of size k for some $k: \boldsymbol{R}_+ \to \boldsymbol{P}$. We can allow the support to be infinite, but then we must insist that the beliefs of the players 2 be convex combinations of a finite number of beliefs. An example which does not satisfy this restriction is the following. Suppose each player 2_n believes that the rational player 1

[5] Using equivalence classes defined by d, where π and μ are equivalent if and only if $d(\pi, \mu) = 0$, we have a metric space with metric d. See Royden (1988) for details.

[6] Rationalizability has not been formally defined for perturbed or Bayesian games, nor for infinitely repeated games. However, the notion of rationalizability (which would be present in any suitable definition) is based on the iterated deletion of those strategies which cannot be justified by best-response behavior. While rationalizability incorporates an infinite number of such iterations, I only require that behavior be consistent with two iterations.

will select the Stackelberg action through period $n-1$, but in period n will select another action. Notice, then, that $\pi_m^h - \pi_n^h \geq \gamma$ for all $m > n$ and all $h \in H_{n-1}^*$. Therefore $d(\pi_n, \pi_m) \geq \gamma$ for all $n \neq m$, which implies that the set of conjectures of the players 2 is not compact. My main result below will not hold in such a situation.

Let $\Gamma \equiv \{\pi_n | n \in \boldsymbol{P}\}$ be the set containing the systems of conjectures of every player 2. The assumptions above are summarized by the following definition.

Definition 2: Take a function k: $\boldsymbol{R}_+ \to \boldsymbol{P}$ and a given game $G_\delta^{LS}(\varepsilon, s_1^*)$. Players are said to hold *type* $\boldsymbol{R}$*-k beliefs* if and only if the following are satisfied:

(1) Player 1 believes (knows) that
 (a) Γ is of size k,
 (b) each player 2_n plays a best-response to her conjecture when she is called upon to play (in period n), and
 (c) the system of conjectures π_n of each player 2_n is consistent with the form of the game (specifically, the perturbation); and
(2) Player 1 knows the form of the game and selects a best-response to his conjectures about the strategies of the players 2.

Notice that player 1 need not know Γ. Also, it is not sufficient that player 1 believe that Γ is contained in some (possibly unknown) compact set. What is important is that player 1 believe that Γ is of size k for some given k. For example, it is insufficient that player 1 believe that Γ has finite cardinality; we need (in this case) that player 1 know that Γ is of cardinality M or less, for some given M.

Let $\underline{V}_1(\delta, \varepsilon, k)$ be the infimum expected payoff of the rational player 1 in the game $G_\delta^{LS}(\varepsilon, s_1^*)$ when players hold type $\boldsymbol{R}$-k beliefs. This infimum exists because the set of feasible payoffs for player 1 is bounded from below. The following result establishes the power of the Stackelberg perturbation and demonstrates that players can establish significant reputations independent of equilibria.

Theorem: *Take any function* k: $\boldsymbol{R}_+ \to \boldsymbol{P}$ *and let* $\varepsilon \in (0,1)$. *There exists a number* $l(k, \varepsilon)$ *such that*

$$\underline{V}_1(\delta, \varepsilon, k) \geq \delta^{l(k,\varepsilon)} u_1^* + (1 - \delta^{l(k,\varepsilon)}) \underline{u}_1.$$

That is, in the limit (*as* δ *approaches unity*) *the rational player* 1 *is guaranteed at least his Stackelberg payoff* (*in terms of his own expectations*).

Proof: Since BR is upper hemicontinuous and A_2 is finite, there is some $\rho \in (0,1)$ such that $[\alpha_1 \in \Delta A_1,\ \alpha_1(a_1^*) > \rho]$ implies that $BR(\alpha_1) \subset BR(a_1^*)$.[7] Therefore, given a system of conjectures π for some player 2_n, $\pi^h > \rho$ implies that player 2_n will play a best-response to a_1^* if she faces player 1 after history h.

For any set X, let $\#X$ denote the cardinality (number of elements) of X. Fudenberg and Levine (1989) prove the following result, which appears here in terms of the current notation.

Lemma: *Take any system of conjectures* $\pi \in \Pi$ *of a player* 2_n *in the game* $G_\delta^{LS}(\varepsilon, s_1^*)$, *and take any* $h_\infty \in H_\infty^*$. *Then*

$$\#\{\pi^h < z | h \in T(h_\infty)\} < \frac{\ln \varepsilon}{\ln z}$$

for every $z \in (0,1)$.

[7] See Fudenberg and Levine (1989).

This implies that, given a system of conjectures for some player 2_n, there is a bound on the number of times this player would *not* play a best-response to a_1^* (if called upon to play in all periods) if player 1 were to always play a_1^*.

Given the above facts, the result is not difficult to prove. Fix $\varepsilon \in (0,1)$ and take any function k: $\boldsymbol{R}_+ \to \boldsymbol{P}$. We presume, then, that player 1 believes Γ to be of size k. Let $\mu_1, \mu_2, \ldots, \mu_{k((1-\rho)/2)} \in \Pi$ be such that

$$\Gamma \subset B_{(1-\rho)/2}(\mu_1) \cup B_{(1-\rho)/2}(\mu_2) \cup \cdots \cup B_{(1-\rho)/2}(\mu_{k((1-\rho)/2)}).$$

Now take one such μ_t, $1 \le t \le k((1-\rho)/2)$. By the definition of the pseudo-metric d, $[\pi^h < \rho$ for some $\pi \in B_{(1-\rho)/2}(\mu_t)]$ implies that $\mu_t^h < (\rho+1)/2$. The lemma of Fudenberg and Levine (1989) establishes that, for any $h_\infty \in H_\infty^*$,

$$\#\left\{\mu_t^h < \frac{\rho+1}{2} | h \in T(h_\infty)\right\} < \frac{\ln \varepsilon}{\ln\left(\dfrac{\rho+1}{2}\right)}.$$

Therefore

$$\#\left\{\pi^h < \rho | h \in T(h_\infty),\, \pi \in B_{(1-\rho)/2}(\mu_t)\right\} < \frac{\ln \varepsilon}{\ln\left(\dfrac{\rho+1}{2}\right)}$$

for each $t = 1, 2, \ldots, k((1-\rho)/2)$.

Thus, for any $h_\infty \in H_\infty^*$, we have that

$$\#\left\{\pi_n^h < \rho | h \in T(h_\infty),\, n \in \boldsymbol{P}\right\} < k\left(\frac{1-\rho}{2}\right) \frac{\ln \varepsilon}{\ln(\rho+1) - \ln 2} \equiv l(k, \varepsilon).$$

That is, if player 1 always selects action a_1^*, there are at most $l(k,\varepsilon)$ periods (independent of δ) in which his opponent will not play a best-response to a_1^*. The worst case for player 1 is when these periods occur at the beginning of the game, where payoffs are least discounted. This gives the bound of the theorem. *Q.E.D.*

3. AN EXTENSION

The theorem is most powerful as a limiting result (when δ converges to one). As such, it can be embellished slightly. Here is a more robust interpretation, which is somewhat of a corollary. Suppose that we consider a convergent sequence of discount factors $\{\delta_i\} \to 1$ which defines a sequence of long-run/short-run games (all with the same stage game). Let $\pi_n^{(i)}$ be the system of conjectures of player 2_n in the game indexed by i, and let $\Gamma^i \equiv \{\pi_n^{(i)} | n \in \boldsymbol{P}\}$. For any set $\Lambda \subset \Pi$ let $w_i(\Lambda) \equiv (1-\delta_i)\Sigma\{\delta_i^{n-1} | \pi_n^{(i)} \in \Lambda\}$ be the discounted "weight" (in terms of player 1's payoffs) of the players whose systems of conjectures in game i are contained in Λ.

Suppose we have a sequence of sets $\{\Lambda^i\} \subset \Pi$ such that Λ^i is of size k_i: $\boldsymbol{R}_+ \to \boldsymbol{P}$, for each $i = 1, 2, \ldots$. Further suppose that $\{k_i\}$ is such that $\lim_{i\to\infty} \delta_i^{k_i(z)} = 1$ for all $z > 0$. This allows $k_i(z)$ to approach infinity, but not "too quickly." Assume that $w_i(\Lambda^i)$ converges to one as i approaches infinity. Finally, assume players hold type $\boldsymbol{R}$-k_i beliefs in game i. Then it is not difficult to show that there is a lower bound $\underline{V}_1^i$ on the expected payoff of the long-run player in game i, and $\lim_{i\to\infty} \underline{V}_1^i = u_1^*$.

Nuffield College, Oxford OX1 INF, United Kingdom

Manuscript received May, 1991; final revision received May, 1992.

REFERENCES

AUMANN, ROBERT J., AND SYLVAIN SORIN (1989): "Cooperation and Bounded Recall," *Games and Economic Behavior*, 1, 5–39.

BERNHEIM, B. DOUGLAS (1984): "Rationalizable Strategic Behavior," *Econometrica*, 52, 1007–1028.

CHO, IN-KOO (1991): "Rationalizability, Stationarity and Bargaining," mimeo, University of Chicago.

FUDENBERG, DREW, AND DAVID LEVINE (1989): "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57, 759–778.

——— (1991): "Maintaining a Reputation when Strategies are Imperfectly Observed," forthcoming, *Review of Economic Studies*.

FUDENBERG, DREW, DAVID M. KREPS, AND ERIC S. MASKIN (1990): "Repeated Games with Long-run and Short-run Players," *Review of Economic Studies*, 57, 555–573.

KREPS, DAVID M., PAUL MILGROM, JOHN ROBERTS, AND ROBERT WILSON (1982): "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma," *Journal of Economic Theory*, 27, 245–252.

KREPS, DAVID M., AND ROBERT WILSON (1982): "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253–279.

MILGROM, PAUL, AND JOHN ROBERTS (1982): "Predation, Reputation, and Entry Deterrence," *Econometrica*, 50, 443–460.

PEARCE, DAVID G. (1984): "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 52, 1029–1050.

ROYDEN, H. L. (1988): *Real Analysis*, 3rd ed. New York: Macmillan.

SELTEN, R. (1977): "The Chain-Store Paradox," *Theory and Decision*, 9, 127–159.

WATSON, JOEL (1992a): "Cooperation in the Infinitely Repeated Prisoners' Dilemma with Perturbations," Stanford Graduate School of Business.

——— (1992b): "Guided Beliefs: A Rationale for Strategy Perturbations," Stanford Graduate School of Business.