Quantitative Marketing and Economics, 3, 337–363, 2005.

How Verifiable Cheap-Talk Can Communicate Unverifiable Information

ROBERT BLOOMFIELD
VRINDA KADIYALI*

Johnson Graduate School of Management, Cornell University, Ithaca NY 14850
E-mail: rjb9@cornell.edu
E-mail: kadiyali@cornell.edu

Abstract. This study describes a "cheap-talk" model in which sellers can credibly convey unverifiable information by choosing whether or not to exaggerate verifiable information. We find that unexaggerated claims can communicate favorable unverifiable information if buyers are not too likely to verify claims, and sellers with better information care more about future prices than sellers with worse information. However, there is always another equilibrium in which sellers exaggerate all verifiable claims. Laboratory tests show that when buyers infrequently verify the sellers' claims, players converge to the equilibria close to the example provided in instructions. When buyers are very likely to verify claims, players fail to converge to any equilibrium. Both of these results are consistent with an evolutionary learning model, but inconsistent with the intuitive criteria of Cho and Kreps (1987). We discuss the implications of our results for both consumer and financial markets.

Key words. cheap talk, evolutionary game theory, signaling, quality, earnings management, disclosure

JEL Classification: C73, C92, G14, M3

Business settings provide many opportunities to make exaggerated claims about verifiable and unverifiable facts. Managers seeking to increase their firms' stock prices might exaggerate their verifiable reported earnings numbers by choosing aggressive accounting policies, while simultaneously making unverifiable optimistic claims about future market opportunities. A software producer can exaggerate verifiable claims about product speed and compatibility, while also exaggerating unverifiable claims about future reliability, customer support and plans for product upgrades.

We model why businesses do not always exaggerate verifiable claims, even though buyers often fail to verify them (Johnson and Russo, 1984; Moorthy et al., 1997). In particular, we examine the possibility that exaggerated verifiable claims signal low unverifiable quality. Past research supports this possibility. Firms that exaggerate past financial performance tend to have low future earnings growth (Lilien et al., 1988; Wahlen, 1994). Similarly, "ambulance chasers" appear to be more likely to exaggerate the benefits of their services than better-trained (and presumably superior) attorneys.[1] Using a game-theoretic model to

*Helpful comments were received by an anonymous reviewer, Mark Nelson, many doctoral students at Cornell University, and Accounting workshop participants at the University of Texas at Austin.

1 Rizzo and Zeckhauser (1990) show how established physicians may choose complete understatement over aggressive claims of newer entrants to show superior quality.

explore the link between verifiable and unverifiable claims, we identify two equilibria—one in which exaggeration of verifiable claims signals low quality, and one in which it does not. We use both traditional and evolutionary refinements to predict which equilibrium results in various settings. Our experimental test supports the evolutionary model.

In our model, one player sells an asset over two periods. The asset has a verifiable component of value and an unverifiable component. Before the first period, the seller makes a claim about the verifiable component. The claim is "cheap talk" (Crawford and Sobel, 1982; Farrell and Gibbons, 1989); exaggeration imposes no legal or regulatory costs on the seller. With some probability, the buyer verifies and adjusts this claim before the first period. By the second period, the buyer certainly knows the value. The price in each period is equal to the expected value of the asset to the buyer. We show that the accuracy of the verifiable claim can signal high unverifiable value if sellers of assets with high unverifiable value place greater weight on future prices than do sellers of assets with low verifiable value. Such a correlation seems likely in many settings. Managers who know future prospects are good are more likely to continue their association with their firm, and therefore care more about future stock prices. Similarly, sellers of products with high unverifiable value are more likely to see high sales in future periods, and therefore care more about prices in those periods. We show how the nature of the equilibria varies with the magnitude of this correlation, and the probability that the buyer verifies claims before purchasing in the first period.

Even when the signaling equilibrium exists, our model also has an equilibrium in which all sellers exaggerate and buyers believe (correctly) that the choice to exaggerate conveys no information. To predict which equilibrium is likely to arise under given market conditions, we consider two models of equilibrium selection. Cho and Kreps' (1987) intuitive criteria predict that the exaggeration equilibrium will never arise when verification is frequent enough to allow the signaling equilibrium to exist, but infrequent enough that the low-quality seller always prefers to exaggerate. In contrast, evolutionary analyses predict that, when verification is sufficiently frequent, which equilibrium arises is determined by the players' initial actions and beliefs. When there is a very high probability that buyers verify the verifiable claim, sellers do not select *any* equilibrium. Instead, they tend to behave erratically and unpredictably, never settling into any stable pattern of behavior, because, in each equilibrium, there is one strategy that is never chosen by either seller. As the sellers approach this equilibrium, slight differences in the (very small) probability with which they choose the disequilibrium strategy cause dramatic changes in the information content of this strategy, which can give sellers an incentive to choose it. As a result, sellers are driven away from the equilibria as they approach them. Such equilibria are said to be "locally unstable" with respect to the learning process (Hofbauer and Sigmund, 1988).

We provide an empirical test to discriminate between the intuitive criteria and evolutionary models by conducting a laboratory experiment involving 24 independent groups of subjects in which sellers make quality claims to buyers. One treatment manipulates the likelihood that the buyers in period 1 search for information to verify the seller's verifiable claim. The other treatment manipulates the starting point of seller behavior and buyer beliefs by providing two different numerical examples in the instructions to subjects. The results strongly support the predictions of the evolutionary learning model over the intuitive criteria model. When search is not too likely, the intuitive criteria rule out the equilibrium in which sellers always make exaggerated claims; however, that equilibrium arises almost every time with the evolutionary model, given the right initial conditions (i.e., a history of exaggeration). When search is very likely, neither the signaling nor exaggeration equilibria is selected, as predicted by the evolutionary model. The fact that different equilibria can arise as a function of initial conditions suggests that variations in quality claims across sellers and industries may be driven not only by their current economic environments, but also by historical accident.

In the next section, we present the model. Then, we describe the learning model to derive predictions on equilibrium selection. The fourth and fifth sections describe the experiment and its results. The final section summarizes and concludes.

2. The cheap talk model

2.1. The model

Consider a risk-neutral seller and two populations of risk-neutral buyers. The seller sells a product with a value of $V = V_0 + \Theta_I$, where V_0 is a verifiable component of value, Θ_I is an unverifiable component. E.g., in the value of a firm, current income under appropriate accounting methods can be verified by investors who can analyze financial statements, but any claims about future prospects are unverifiable. Similarly, the value of software can be decomposed into performance (speed and compatibility) and plans for service and future upgrades. Claims about performance are verifiable for savvy consumers, while claims about future service and upgrades are unverifiable. More generally, the value of a consumer good can be decomposed into "search quality," which is verifiable, and experience or credence quality, which can be verified only imperfectly even after purchase (Nelson, 1970; Darby and Karni, 1973).[2] Assume verifiable quality to be a continuous variable with an arbitrarily large variance, and uncorrelated with unverifiable quality. Hence, the buyer cannot infer whether the seller is exaggerating simply by observing the actual or claimed level of verifiable quality.[3] Unverifiable value is either Θ_H or Θ_L, with equal probability, with $\Theta_H - \Theta_L = S > 0$. A seller with an asset (or product) of unverifiable value Θ_H (Θ_L) is called a "high-quality". ("low-quality") seller, also denoted HQ (LQ).[4]

In period 0, the seller makes a claim regarding verifiable quality. We restrict the claim X to satisfy $X = V_0 + \alpha$, where $\alpha = \alpha_e$ indicates an exaggerated claim and $\alpha = \alpha_u$ indicates a less-exaggerated claim. Because the claim is a binary choice, only the (positive) difference between the two claims is relevant to our analysis. We denote this difference $\alpha_e - \alpha_u = A > 0$. We impose no restrictions on the signs of the separate two claims (e.g.,

2 Experience quality is defined as quality that is known perfectly after purchase. Credence quality is defined as quality that is not known even after purchase. We generalize these definitions by defining unverifiable quality as being unknown before the purchase and known imperfectly after the purchase.

3 If verifiable and unverifiable value were correlated, the report of verifiable value (and the true level of verifiable value, if the buyer searched) would also provide a noisy signal of unverifiable value in both periods. As a result, the expectation of value in both periods would also need to be conditioned on X or V. As with the event of failure or non-failure, this would have no qualitative effect on the analysis.

4 Allowing for different probabilities for each type of seller does not alter qualitative results of the model.

both could include some exaggeration, or both could include some understatement). For convenience, however, we refer to the claims as exaggerated and unexaggerated.

Unlike most signaling models (e.g., Simester, 1995 in product markets, and Titman and Trueman, 1988, in financial markets), there is no direct cost associated with either claim; thus, the quality claims choice is a form of "cheap talk" (Crawford and Sobel, 1982; Farrell and Gibbons, 1989). We do not allow the seller to make direct claims about unverifiable quality. As shown by Crawford and Sobel (1982) and Farrell and Gibbons (1989), such claims are not credible in our cheap talk model, because they cannot be verified and the incentives of the seller and buyer are diametrically opposed.

In period 1, the buyers pay $P_1 = \varepsilon E[V|X] + (1-\varepsilon)E[V|X, \alpha]$ for the asset. The variable ε may be interpreted as reflecting the extent to which buyers who fail to verify the verifiable claim influence the market price in period 1. Failure to verify the verifiable claim would be rational if doing so were costly, or if buyers lack the expertise or motivation to do so. Psychological biases such as overconfidence or gullibility might also inhibit verification. Competition among buyers who do verify the claim might lead prices to reflect the verified beliefs completely. However, arbitrage is rarely feasible in product markets, and even in financial markets, where arbitrage is more likely, market prices often appear to be influenced by less-informed or biased traders (Bloomfield, 2003). In period 2, the buyers pay $E[V|X, \alpha]$, as if the long-term presence of the product or financial report in the market reduces guarantees that the information is held by buyers.[5]

The seller earns a total payoff of $(1 - \pi_i)P_1 + \pi_i P_2$. We call π a "patience" factor that reflects not only the time value of money (a traditional discount rate probably similar for all types of sellers), but also other factors that could alter the sellers' interest in buyers' future valuations. We allow the two types of sellers to have different patience levels, to allow a correlation between seller type and patience. We make no ex-ante assumptions regarding the direction of that correlation. However, some examples suggest that high-quality sellers are likely to be more patient. Consider a manager of a firm who makes claims about the firm's value to buyers of the firm's stock. Assume that the managers' salary is directly proportional to the firm's stock price, and that the managers' quality determines the likelihood that he retains his job in period 2. A high quality manager faces a lower probability of being fired will get more benefit from a high future stock price. A low quality manager faces a higher probability of being fired and will not benefit from a high future stock price. As another example, consider an entrepreneur selling a new product. A high-quality product is more likely to enjoy increasing sales volume. An entrepreneur selling a product of high quality will expect product sales to rise more than one selling a product of low quality, and will therefore expect more benefit from a high future product price.

2.2. *Two perfect Bayesian equilibria (BNE)*

A seller of type i who exaggerates affects prices in two ways. First, exaggerating inflates the buyer's estimate of verifiable value if the buyer does not search. This benefit is quantified

5 Our qualitative results are unchanged if we use other price-setting rules that are proportional to expectations and reflect a weighted average of both the conditional and unconditional expectations.

as $(1-\pi_i)\varepsilon A$, (recall $A = \alpha_e - \alpha_u$) because it affects the period 1 price (which is weighted by $1-\pi_i$), by an expected amount of εA. Second, exaggerating potentially conveys some information about the seller's unverifiable information. The effect of known exaggeration on expected value is captured by the difference $\{E[V \mid X, \alpha_e] - E[V \mid X, \alpha_u]\}$. This difference is weighted by $[(1-\pi_i)(1-\varepsilon)+\pi_i]$ to reflect its impact on prices in both periods, weighted by the probability of verification $(1-\varepsilon)$ in period 1.

To determine whether a seller prefers to exaggerate, let p represent the buyer's estimate of the probability that a low quality (LQ) seller exaggerates, and let q represent the buyer's estimate of the probability that a high quality (HQ) seller exaggerates. Define $\delta(\pi_i, p, q)$ as the net benefit to making an exaggerated claim rather than a unexaggerated claim, given the buyer's expectations about p and q. Combining the two payoff effects in the preceding paragraph yields

$$\delta(\pi_{i,} p, q) = (1-\pi_i)\varepsilon A + [(1-\pi_i)(1-\varepsilon)+\pi_i]\{E[V|X, \alpha_e] - E[V|X, \alpha_u]\} \quad (1)$$

If $\delta(\pi_i, p, q)$ is positive, then a seller of type i prefers to exaggerate. If $\delta(\pi_i, p, q)$ is negative, then a seller of type i prefers to not exaggerate. If $\delta(\pi_i, p, q)$ is exactly 0, then a seller of type i is indifferent between the two different strategies.

Proposition 1 below shows that, regardless of the parameters of the game, there is always an "exaggeration" perfect BNE where both types of sellers always exaggerate (see the Appendix for proof of this and other propositions).[6] This equilibrium is supported by the buyers' (correct) belief that unexaggerated claims convey no information about the seller's type ($p = q$). This belief gives both types of sellers an incentive to exaggerate, without any countervailing cost. In equilibrium, both types of sellers exaggerate with certainty ($p^* = q^* = 1$), and the quality claim choice conveys no information.

Proposition 1. *For every parameterization of the game, there is an "Exaggeration" perfect Bayesian equilibrium with* $p^* = q^* = 1$.

The exaggeration BNE can exist even when all buyers verify the claim ($\varepsilon = 0$). In this case, the equilibrium can still be supported by the buyers' belief that a seller who does not exaggerate must have a low-quality unverifiable attribute. This belief is not "wrong" because in equilibrium sellers always exaggerate. If high-quality sellers are more patient than low-quality sellers, there can also exist a "signaling" equilibrium in which exaggerating always indicates a greater likelihood of low quality than does not exaggerating. This equilibrium results from the buyer's belief that the high quality seller never exaggerates ($q = 0$). A positive correlation of patience and unverifiable quality makes sense in many contexts. Using the examples above, a manager who knows that unverifiable attributes are of high

6 A perfect Bayesian equilibrium is an equilibrium concept for dynamic games with incomplete information (Kreps and Wilson, 1982). For any equilibrium to be Bayesian perfect, it has to be an equilibrium not just for that node of the game but for the entire game as well. Therefore, a perfect Bayesian Nash equilibrium refines Bayesian Nash equilibrium the same way that subgame perfection is applied to Nash equilibria in games of full information. This concept has been applied to signaling models, as well as cheap talk models, among other things.

quality is likely to assess a lower probability that she will be fired before the second period than one who knows that unverifiable attributes are of low quality (because those who make employment decisions typically have better information than investors). Similarly, firms whose products have unverifiable attributes of high quality are likely to expect greater sales volume in period 2 (due to word-of-mouth) than those whose products have unverifiable attributes of low quality. The following proposition sets out conditions for such a signaling equilibrium to exist.

Proposition 2. *Assuming that $\pi_H > \pi_L$, there exist three cutoffs c_1, c_2, and c_3, with $c_1 > c_2 > c_3$, with the following properties*:

(i) *If $\varepsilon > c_1$, then there exists no Perfect Bayesian Equilibrium (PBE) in which $q^* < 1$ or $p^* < 1$.*
(ii) *If $c_1 \geq \varepsilon \geq c_2$, then there exists a PBE in which $q^* = 0$ and $p^* = 1$*
(iii) *If $c_2 > \varepsilon > c_3$, then there exists a PBE in which $q^* = 0$ and $p^*(\varepsilon) \in (0, 1)$, with $p^*(\varepsilon)$ decreasing in ε*
(iv) *If $c_3 > \varepsilon$, then there exists a PBE in which $p^* = q^* = 0$. In this case, the PBE is supported by the off-equilibrium-path belief that a seller who exaggerates must be of low quality.*

Proposition 2 reveals that, given a positive correlation of patience and unverifiable quality (keeping other parameters fixed), the extent of verification determines the nature of the signaling equilibrium. Specifically, there exist four distinct regions of ε. In the highest region, no signaling equilibrium exists; the benefit of inflating the value estimates of a non-searching buyer outweighs any possibly benefit of communicating high quality by not exaggerating. In the next region, the likelihood of search is just high enough that the LQ seller prefers to exaggerate but the more-patient HQ seller does not (assuming that the buyer infers that exaggeration implies low quality). Thus, the signaling equilibrium in this region involves certain exaggeration by the LQ seller. In the third region the LQ seller prefers a mixed strategy involving some exaggeration and some mimicry of the HQ seller. Thus, the signaling equilibrium involves exaggeration by the LQ seller with probability $p(\varepsilon)$, where $p(\varepsilon)$ is decreasing in ε. In the lowest region, the benefit to exaggeration is outweighed by the cost of revealing low quality, so the signaling equilibrium involves no exaggeration by either type.[7]

Proposition 3 shows that there is never an equilibrium in which exaggeration would signal high quality, even if there were a negative correlation of patience and unverifiable quality.

Proposition 3. *There is no perfect Bayesian equilibrium in which $p^* < q^*$.*

7 More generally, when the incentives of the senders of cheap talk are not aligned with the incentives of the receivers, cheap-talk quality claim choices can be informative only if the sender reports to two distinct audiences that respond to the information differently (Farrell and Gibbons, 1989). For moderate values of ε, the sellers report to two potential audiences: buyers who verify and buyers who do not. For very small or very large values of ε, one audience overwhelms the other, and credible communication becomes impossible.

Such an equilibrium would imply that both seller types would recognize a benefit to claiming exaggerated quality, but would suffer no cost. As a result, both sellers exaggerate, so that exaggerated claims do not convey information. Because buyer's inferences must be validated in equilibrium, it is impossible to construct parameters that allow exaggerated quality claims to be a signal of high quality.

3. Predicting behavior in the game

3.1. Two prediction models

The preceding equilibrium analysis reveals that the signaling and exaggeration equilibria often exist simultaneously. In this section, we examine two competing models of equilibrium selection. The first is a traditional game-theoretic model that selects equilibria according to the reasonableness of player beliefs. The second is an evolutionary model with an arbitrary starting point for actions and a simple learning process for the players to follow. We then develop hypotheses by examining which equilibria the two models predict in two specific parameterizations of our game (see experiments in Section 4).

Turning now to more details, the traditional model of equilibrium used here is the Intuitive Criteria (IC) model of Cho and Kreps (1987), which places restrictions on what the buyer can believe when observing an action that should occur with probability 0 in equilibrium. In the context of the present game, the Intuitive Criteria apply either when sellers of both types always exaggerate (the exaggeration equilibrium) or never exaggerate (the signaling equilibrium when frequency of verification is very high). The IC eliminate these equilibria if they are supported by "unreasonable" beliefs when observing a claim that should be not be observed. A buyer's belief is unreasonable if it involves believing that the claim could have come from a seller who would *never* prefer to make that claim.

The evolutionary learning model assumes that sellers become more likely over time to choose strategies that would have generated high payoffs in previous rounds. This process is similar to Thorndyke's "law of effect," the basis for behaviorist psychology (Thorndyke, 1911; Herrnstein, 1997), except that the players get feedback from actions not chosen and chosen. These simple evolutionary models have been shown to be useful in predicting equilibrium selection (Blume, et al. 1994; Brandts and Holt, 1992; Camerer and Ho, 1998; and Erev and Roth, 1998).

The evolutionary model assumes that the rate at which each seller increases the probability of choosing a strategy is directly proportional to the payoff difference, $\delta(r_i, p_t, q_t)$ defined in the previous section, where the subscript indicates time, and the buyer correctly predicts the seller's strategy at time t. The true learning process is obviously much more complex (Camerer and Ho, 1998). However, this assumption allows us to construct the following two-equation dynamical system whose properties are well known (see Hofbauer and Sigmund, 1988; Weibull, 1995):

$$\frac{dp_t}{d_t} = \delta(\pi_L, p_t, q_t), \frac{dq_t}{d_t} = \delta(\pi_H, p_t, q_t). \quad (2)$$

If each seller's direction of motion always points toward the equilibrium from nearby outcomes, players are always driven closer to equilibrium once they are sufficiently close.

In the terminology of evolutionary game theory, the equilibria are *locally stable* (Hofbauer and Sigmund, 1988) with respect to the evolutionary dynamic (2). Such equilibria are reasonable predictions of the evolutionary model, which suggests that the learning process will be relatively robust to the introduction of noise and discreteness prevalent in real learning processes. Even though discreteness and noise can cause strategies to waver slightly in the neighborhood of the equilibrium, the learning process continually pushes the sellers back toward the equilibrium. If the equilibrium is not locally stable, it is *locally unstable*. In this case, players can be arbitrarily close to the equilibrium, but then be driven away. A small amount of noise or discreteness in the learning process typically makes it impossible for players to converge to an unstable equilibrium. Thus, the evolutionary model predicts that such equilibria will not occur. The Appendix provides a formal definition of local stability and instability.[8]

In many cases multiple equilibria may satisfy the intuitive criteria or local stability. The IC model does not provide any means of determining which equilibrium will be chosen when multiple equilibria can satisfy the criteria. However, because the evolutionary model is perfectly deterministic, it can provide a unique prediction for the convergence point of the process (if any exists) from any given starting point. Therefore, one can predict that equilibrium selection will depend on the starting point (p_0, q_0).

The IC and evolutionary models are very different in their assumptions about rationality. The IC model starts with the traditional assumptions regarding rationality, and then adds the additional requirements of beliefs being both "reasonable" and rational. The evolutionary model does not even involve a rational learning strategy—it could apply to the learning of rats and pigeons, as easily as it applies to humans.[9] We do not believe that either the IC or the evolutionary models represent perfect descriptions of human behavior in games. However, we do believe that the *qualitative* predictions of the models may be useful in predicting which equilibria will be more likely to arise.

To compare the models, we test the predictions of equilibrium selection in two different parameterizations of the game that allow clear discrimination between these two types of models. To test for the effect of starting points, we alter the instructions to create two different histories to alter the initial condition of the players' evolutionary path.

3.2. An infrequent-verification setting ($\varepsilon = 0.7$)

The infrequent-verification setting has ε that lies between the second and third cutoffs defined in Proposition 1. Hence, there is a signaling equilibrium with certain exaggeration by the LQ seller but none by the HQ seller ($p^* = 1, q^* = 0$), and an exaggeration equilibrium with certain exaggeration by both sellers ($p^* = q^* = 1$). The exaggeration equilibrium

8 Proposition 4 is valid for all dynamics in which seller i does not change strategies when $\delta(r_i, p_t, q_t) = 0$, but does increase (decrease) the probability of exaggerated claims whenever $\delta(r_i, p_t, q_t)$ is positive (negative). See Hofbauer and Sigmund (1988) and Weibull (1995) for alternative definitions of stability, and their relations to static equilibrium refinements.

9 The learning model is inconsistent with rational Bayesian behavior in a multiperiod setting, because the players are assumed to not use Bayes' theorem or to consider the effects of their current actions on other players' future actions.

is supported by the belief that an unexaggerated claim is at least as likely to come from an LQ seller as from an HQ seller. This equilibrium does not satisfy the intuitive criteria because an LQ seller always prefers to exaggerate, regardless of an HQ seller's strategy or the buyer's beliefs. Hence, the IC model requires the buyers to believe that a lack of exaggeration is a sure signal of high unverifiable quality. The signaling equilibrium satisfies the intuitive criteria because of the buyer's belief that exaggerated claims come from an LQ seller.

Panel A of Figure 1 illustrates the predictions of the evolutionary model in this setting. The horizontal axis depicts p_t, the probability that the LQ seller exaggerates at time t. The vertical axis depicts q_t, the probability that the HQ seller exaggerates at time t. The vectors show the direction of motion from various starting points, assuming that buyers expect strategies p_t and q_t. The dashed line represents the set of outcomes at which the HQ seller is indifferent between exaggeration and no exaggeration. For points above (below) this line, the HQ seller prefers to exaggerate (not exaggerate). There is no indifference curve for the LQ seller, because the high level of ε leads the LQ seller to prefer exaggerated claims for any set of buyer beliefs.

The vectors of motion shows that both equilibria are evolutionarily stable, because each seller's direction of motion always points toward the equilibrium from nearby outcomes. Thus, either equilibrium could occur according to the evolutionary model. Which equilibrium is selected depends on the starting point. If q_0 is sufficiently large relative to p_0 (as in the "exaggerated" starting point E shown in Panel A), exaggerated claims inflate the buyers' quality estimate at date 1, and the buyer infers that an exaggerated claim indicates *high* quality. Both types of sellers therefore increase their rates of exaggeration, till they reach the upper right corner, where both sellers always exaggerate. In contrast, if the sellers start at a point with p_0 much higher than q_0 (as in the "signaling" starting point S in Panel A), the HQ seller prefers not to exaggerate, because doing so provides a signal of high quality. This is a self-reinforcing process because exaggerating less increases the HQ seller's preference not to exaggerate. Ultimately, as shown in Panel A, the sellers converge to the Non-exaggeration Equilibrium in which HQ sellers signal their strength by not exaggerating.[10]

Proposition 4 of the Appendix shows that both equilibria are evolutionary stable in the infrequent-verification setting, as long as verification frequency is sufficiently low. Stability implies that the evolutionary process leads to the equilibrium from any starting point in a region close to the equilibrium. In principle, it is possible to identify the exact boundaries of this region (called the "basin of attraction"). However, the non-linearity of the dynamical system makes this task very difficult. Instead, we simply use numerical simulation to identify one point within each basin of attraction (points E and S), and predict that the players will be drawn to the exaggeration equilibrium from point E and the signaling equilibrium from

10 There is also a perfect Bayesian equilibrium in which the LQ seller always exaggerates and the HQ seller exaggerates with probability q(ε). This equilibrium, which lies at the intersection of the HQ seller's indifference curve and the boundary $p = 1$, is also unstable. A small deviation from equilibrium pushes the HQ seller even further from the equilibrium, toward one of the other two equilibria. As a result, the dynamic system never leads players to this equilibrium. Mixed strategy equilibria need not always be unstable or lack predictive power; see Amaldoss et al. (2000) for experimental support for a mixed strategy equilibrium in another marketing context.

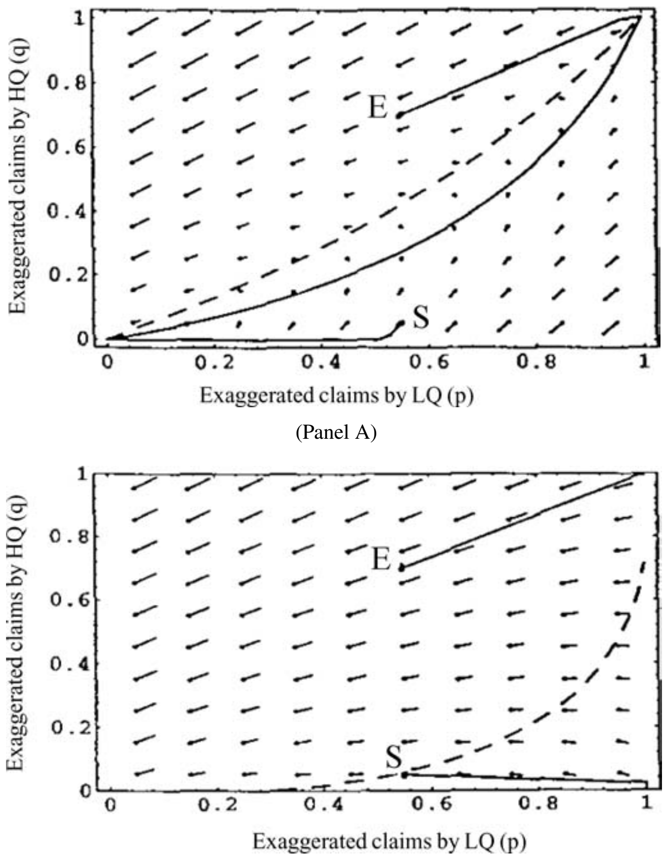

(Panel A)

(Panel B)

Figure 1. Dynamic forces and paths toward equilibrium. In each panel, the horizontal (vertical) axes represents the probability p (q) that the LQ (HQ) seller claims aggressively. The vector emanating from each point shows the direction of motion under the dynamics in equation (2) for representative outcomes, under the assumption that the buyers correctly anticipate the sellers' reporting choices. The length of the vector indicates the speed of motion. The dashed line represents the set of points for which the HQ seller is indifferent between exaggeration and understating quality claims. The solid line (not present in panel A) represents the set of points for which the LQ seller is indifferent between exaggeration and understatement. Above (below) the dashed line, the HQ seller increases (decreases) the probability of exaggerating. For all panels, $k = 6, r_H = 0.2/3, r_L = 1/4, A = 15, S = 10$. **Panel A**: Evolution of behavior in an Infrequent-verification setting ($\varepsilon = 0.7$), from an exaggeration starting point E and a signaling starting point S under the continuous dynamic (4). **Panel B**: Evolution of behavior in a Frequent-verification setting ($\varepsilon = 0.2$), from an exaggeration starting point E and a signaling starting point S under the continuous dynamic (4).

point S. The intuitive criteria model predicts that such starting points should have no effect, because only one equilibrium satisfies the criteria.

H1 (IC): Sellers in the Infrequent-Verification setting choose strategies more consistent with the signaling equilibrium than with the exaggeration equilibrium, regardless of history.

H1 (Evolution): Sellers in the Infrequent-Verification setting choose strategies more consistent with the signaling equilibrium given a signaling history, but choose strategies more consistent with the exaggeration equilibrium given an exaggeration history.

3.3. *A frequent-verification setting*

The frequent-verification setting involves a level of ε that is below the lowest cutoff identified in Proposition 1. As a result, there is a signaling equilibrium that involves no exaggeration by the either seller ($p^* = q^* = 0$), as well as an exaggeration equilibrium ($p^* = q^* = 1$). Both equilibria satisfy the intuitive criteria, because there is some set of beliefs that would lead each seller type to prefer exaggeration, and there is some set of beliefs that would lead each seller type to prefer no exaggeration. Thus, any belief about non-equilibrium behavior is reasonable.

On the other hand, neither equilibrium is evolutionarily stable. Panel B of Figure 1 shows the dynamic forces in a "Frequent-Verification" setting in which the buyer is likely to be aware of the seller's quality claims strategy in period 1 ($\varepsilon = 0.2$). As before, the dashed line indicates the set of beliefs that leaves the HQ seller indifferent between exaggeration and no exaggeration. The solid line indicates the set of indifference points for the LQ sellers. As shown in the figure, low values of ε lead to equilibria in which one strategy is not chosen by either seller. Near these equilibria, small differences in sellers' (very small) probabilities of choosing that strategy can dramatically alter the strategy's information content. For example, assume that the HQ seller exaggerates with probability 0.99, while the LQ seller exaggerates with probability 0.999. Even though the sellers are very close to the exaggeration equilibrium, an unexaggerated report is 10 times more likely to have come from an HQ seller. Thus, both sellers will sharply reduce their probability of exaggeration. If the sellers eventually come close to the signaling equilibrium, they can be driven away in a similar manner (when the LQ seller exaggerates slightly less than the HQ seller). Proposition 4 of the Appendix shows formally that both equilibria are locally unstable for sufficiently small values of ε. Because players are not expected to converge to locally unstable equilibria, a high value of ε is essential for long-term stability in quality claims, regardless of history.

This analysis leads to the following predictions:

H2 (IC): Sellers in the Frequent-Verification setting choose strategies consistent with equilibrium in the Infrequent-Verification setting, which may or may not depend on history.
H2 (Evolution): Sellers in the Frequent-Verification setting do not choose strategies consistent with either equilibrium, regardless of history.

4. The experiment

4.1. *Experimental design*

The experiment was constructed as a 2×2 factorial design. The treatment of primary interest manipulated the value of ε, yielding an Infrequent-Verification setting ($\varepsilon = 0.7$)

	Infrequent-Verification	Frequent-Verification
Signaling-History	1, 2, 3, 4, 5, 6 ---------- 13, 14, 15, 16, 17, 18	7, 8, 9, 10, 11, 12 ---------- 19, 20, 21, 22, 23, 24
Exaggeration-History	19, 20, 21, 22, 23, 24 ---------- 7, 8, 9, 10, 11, 12	13, 14, 15, 16, 17, 18 ---------- 1, 2, 3, 4, 5, 6

Figure 2. Experimental design. This figure indicates the order in which the 24 groups participated in different cells of the design. Numbers above (below) the dotted line in each cell indicate which groups participated in that cell in Part 1 (Part 2) of the experiment. For example, groups 1–6 participated first in the Infrequent-Verification/Signaling-History cell, and then in the Frequent-Verification/Exaggeration-History cell.

and a Frequent-Verification setting ($\varepsilon = 0.2$). The other treatment used two different numerical examples in the instructions to subjects (corresponding to points E and S respectively in Figure 1) to induce an Exaggeration-History and a Signaling-History starting point.

To collect data efficiently, each group of subjects participated in two settings, each differing with respect to both treatments (see Figure 2). One quarter of the groups played the Frequent-Verification/Exaggeration-History game first and the Infrequent-Verification/Signaling-History game second, while an equal number played the same two games in reverse order, to balance any order effects that might have arisen. Similarly, one-quarter of the groups played the Infrequent-Verification/Exaggeration-History first and the Frequent-Verification/Signaling-History game second, while an equal number played the same two games in reverse order. Requiring each group to play games that differ according to *both* treatments minimized any "holdover" effects that might have caused the outcome of the first game to influence the outcome of the second game. For each group, the first game played is referred to as "Part 1," while the second game is referred to as "Part 2."

4.2. The experimental task

In both the IC and Evolutionary models, equilibrium selection is driven by the beliefs of the buyer regarding seller strategies, and the payoffs received by the seller given those beliefs. In order to focus on these issues, the laboratory game differs in certain respects from the model presented above. First, each group of players includes 2 sellers (an "HQ" seller and an "LQ" seller), while the formal model includes only one seller with two possible types. The experiment included 24 such groups of subjects who each participated in a 150-minute session. Second, the buyer's incentive is simply to estimate verifiable and unverifiable quality accurately. This eliminates the complication of a price-setting phase (that might include negotiation or bidding). Neither of these changes alters the game's equilibria, or its basic nature. The sellers still choose whether or not to exaggerate, the buyer infers the sellers' type from observing whether reports are exaggerated or not, and the buyer's inferences influence the payoffs and preferences of the sellers in exactly the manner assumed in the model.

To make it easier for the subjects to understand the rules of the game, they are provided with a simple context. (We chose a financial-market context, but could just as easily have chosen a product-market context.) The sellers are managers of firms, and can choose whether or not to exaggerate the value of their firm to investors (see instructions in Appendix 2). The buyer is an equity analyst whose incentive is to estimate the firms' value as accurately as possible. It is possible that this context affects the subjects' behavior. However, these effects will not vary across cells of the design (because the context is the same across cells). Therefore, the context does not interfere with any inferences drawn from observing treatment effects.

In each round of play, each seller chooses how many times out of 100 it wishes to exaggerate its claim of verifiable quality. For example, in a given round, a seller could choose to exaggerate 65 claims and to not exaggerate 35 claims. This elicitation method provides more powerful data on players' true strategies than would be provided by eliciting a single choice in each round, while leaving the nature of the game essentially unchanged (Bloomfield, 1994).[11]

Each buyer sees a total of 200 claims per round (100 claims from each seller), without knowing which claim comes from which type of seller. The buyer is paid a salary that decreases linearly with the average error in estimated value of the products. The role of the buyer is to generate the three expectations that drive the seller payoffs: the expectations of unverifiable quality given an exaggerated or an unexaggerated claim, and the expectation of verifiable quality given that the buyer did not verify. These three expectations completely determine the payoffs of the sellers according to the equations used to construct Figure 1, with the parameter for ε depending on whether the players are in the Frequent- or Infrequent-Verification setting. Note that it is crucial to use a human participant as the buyer, because the buyer's expectations arise endogenously as a function of the buyer's beliefs about seller behavior. It is precisely these beliefs that must be tested empirically (to see, for instance, whether the beliefs are ever "unreasonable" as defined by the IC model).

Each game lasts 20 rounds. This ensures sufficient opportunity for quality claim strategies to evolve to a steady state (if a steady state would ever arise). After each round is completed, each seller learns the strategy chosen by the other seller, the buyer's three value assessments, the resulting payoff per exaggerated claim, and the resulting payoff per non-exaggerated claim. The buyer learns both sellers' strategies (i.e., how often they exaggerated), the optimal value assessments given these strategies, and the payoff to those assessments. Unlike many laboratory games, the players are paired with the same opponents for all 40 rounds (20 rounds in each of two parts). This fixed-pairing scheme allows players to learn other players' strategies and beliefs more effectively, providing a very strong test of our prediction that

11 If the sellers could choose any mixture, the similarity of the games would be exact. Because sellers can choose only whole percentages, the game becomes one in which both sellers can choose (or mix) among 101 strategies.

players will not converge to any equilibrium in the frequent-verification setting. It also allows for very powerful statistical tests, as each 3-player group is entirely independent of all other groups. Switching group assignment would impair this independence, and reduce our sample size. Fixed pairing might lead to behavior that differs from the predictions of both the evolutionary model and the equilibrium analysis of the one-shot game, because it allows players to employ multi-period strategies; we discuss the impact of such strategies in Section 5.4.

To manipulate history, the subjects review a detailed example before they begin play in the first setting (see Appendix 2). The example includes printouts of the feedback screens for the three players after the example outcome, and a one-page explanation of how the numbers in those screens are computed. The example strategies in the Exaggeration-History setting are exaggerated claim rates of 70 per round for the HQ seller and 55 per round for the LQ seller (corresponding to point "E" in Figure 1). The example strategies in the Signaling-History setting are exaggerated claim rates of 5 per round for the HQ seller and 55 per round for the LQ seller (corresponding to point "S" in Figure 1). The feedback screens and numerical examples show the relative payoffs to exaggerating and not exaggerating for each type of seller, given their choices and given that the buyer expected those choices.

All values in the experiment are denominated in "francs", a laboratory currency. Following the tenets of "induced value theory" that are standard in experimental economics (See Smith, 1976) we convert francs to cash so that winning more francs results in a greater cash payment, all else being equal.

5. Results

5.1. Test of hypotheses H1 (IC) and H1 (Evolution)

Hypotheses H1(IC) and H1(Evolution) suggest that the intuitive criteria and evolutionary models can be distinguished by examining which equilibria are most consistent with participants' behavior, and how this consistency varies with the history treatment. To test these predictions, we define a simple Euclidean metric to measure the distance of players from equilibrium. Letting LOW and HIGH denote the number of claims per 100 exaggerated by an LQ and HQ seller (analogous to p and q in Section 2), the Euclidean distance from the exaggeration equilibrium, denoted DISTE, is $[(100-\text{LOW})^2+(100-\text{HIGH})^2]^{1/2}$. The Euclidean distance from the signaling equilibrium, denoted DISTS, is $[(100-\text{LOW})^2+(0-\text{HIGH})^2]^{1/2}$ in the Low-Frequency Setting and $[(0-\text{LOW})^2+(0-\text{HIGH})^2]^{1/2}$ in the High-Frequency setting.

In the Low-Frequency setting, H1(IC) predicts that DISTE > DISTS, and that this difference does not depend on history. H1(Evolution) predicts that DISTE > DISTS when participants are given the signaling history, but that DISTS > DISTE when participants are given the exaggeration history. We therefore test whether DIFFDIST = DISTE − DISTS is greater than 0 in each Infrequent-Verification setting, and whether DIFFDIST varies with levels of history in that setting. For this and all other tests in Table 1, we compute the average of the dependent variable (DIST) over all rounds of play, and treat this average as a single

Table 1. Time-averaged strategy choices.

	Infrequent-verification (0)	Frequent-verification (1)	Effect of decreasing frequency
Signaling -history			
DISTS	65	69	−4
DISTE	48	84	−36 (0.0070)
DIFFDIST	17	−15	32
MINDIST	12	33	−11 (0.0419)
Exaggeration-history (1)			
DISTS	85	93	−8
DISTE	22	60	−38 (0.0024)
DIFFDIST	63 (0.001)	33	30
MINDIST	15	36	−21 (0.0003)
Effect of history			
DISTS	20 (0.0188)	24 (0.0820)	−4
DISTE	26 (0.0071)	24 (0.0942)	−2
DIFFDIST	46(0.010)	−48(0.0870)	−2
MINDIST	7	−3(0.0942)	−10

This table displays cell means for various dependent variables for all rounds of play in all four cells of the experiment. DISTS and DISTE are the Euclidian measure of distance from the exaggeration and non-exaggeration equilibrium. DIFFDIST is the difference between the two. MINDIST is the lowest distance to the equilibrium across all 20 rounds. One-tailed significance levels for differences (given in parentheses when significant at conventional levels) are computed by resampling, and determining the proportion of random samples with differences greater than the treatment effect.

observation. This technique, like a "repeated measures" analysis, avoids the concern that different choices by the same cohort do not represent independent observations.[12]

The results are in Table 1. In the Infrequent Verification/Signaling-History setting, DISTS is 65, DISTE is 48, i.e. DIFFDIST of 17. In the Infrequent Verification/Exaggeration-History setting, DISTS is 85, DISTE is 22, i.e., DIFFDIST of 63 (significant at $p < 0.005$). These two DIFFDIST measures are different from one another ($p < 0.01$). The high predictive power of the exaggeration equilibrium with the strong effect of history, allow us to reject the IC model.

An examination of the results by cohort clarifies these results. We focus now on the final five rounds of play, by which time players have presumably settled into preferred strategies (if they settle at all). Mean choices of HQ and LQ sellers (HIGH and LOW respectively) in these rounds are in Table 2. In the Infrequent-Verification/Signaling-History setting, HIGH is below 50 and LOW is above 50 in five of the 12 cohorts. In six of the remaining seven cohorts, HIGH and LOW are above 70. In contrast, in the

12 The crossover nature of the design makes it difficult to exploit all of the power normally accruing to a within-subjects analysis, because each cohort participates in two cells that differ in two ways from one another (we change both history and verification frequency). This makes it difficult for a cohort to serve as its own control, as can often be done in within-subject designs, sapping the analysis of some power. However, the design is fully balanced, and we use bootstrapping models to ensure that we get unbiased t-statistics. The slight loss of power only biases against finding significant differences across cells.

Table 2. Outcomes for each group, averaged over last five rounds of each part of the experiment. For each of the 48 outcomes (2 from each of the 24 groups), this table lists the percentage of claims exaggerated by the LQ seller (LOW) and the HQ seller (HIGH) over the last four rounds of each part of the experiment. The Nonexaggeration Equilibrium is ($LOW^* = 100$, $HIGH^* = 0$) in the Infrequent-Verification setting and ($LOW^* = HIGH^* = 0$) in the Frequent-Verification setting. The Exaggeration Equilibrium is ($LOW^* = HIGH^* = 100$) in both settings.

	Signaling-History Setting							
	Infrequent-verification setting				Frequent-verification setting			
	Part 1		Part 2		Part 1		Part 2	
	LOW	HIGH	LOW	HIGH	LOW	HIGH	LOW	HIGH
	70	24	100	5	100	100	60	60
	100	17	95	96	60	51	27	20
	100	100	100	100	87	100	20	0
	94	83	100	0	60	70	60	36
	94	100	82	1	44	37	79	59
	100	98	94	54	88	60	20	0
Part mean	93	70	96	43	73	70	44	29
Cell mean			94	57			59	50
	Exaggeration-history setting							
	97	76	100	100	14	60	20	4
	98	94	98	100	72	69	53	13
	99	99	80	100	76	56	45	4
	92	95	100	100	48	23	94	100
	100	70	80	80	100	79	81	83
	98	96	98	35	95	72	62	100
Part mean	97	88	93	86	68	60	59	51
Cell mean			95	87			63	55

Infrequent-Verification/Exaggeration-History setting, LOW is below 50 in only one of the 12 cohorts, while both HIGH and LOW are above 70 in 11 of the remaining cohorts. This indicates that the exaggeration equilibrium has strong drawing power. Even with a history consistent with signaling, only about half of the cohorts settle into a signaling equilibrium; the rest of the cohorts are drawn to the exaggeration equilibrium. In contrast, when given a history consistent with exaggeration, almost all cohorts are drawn to an exaggeration equilibrium.

5.2. *Test of hypothesis H2 (IC) and H2 (Evolution)*

The intuitive criteria model predicts that players will settle into an equilibrium if one that satisfies the criteria exists. The evolutionary model predicts that despite history, the locally unstable equilibria in the high-frequency setting will have lower predictive power than the locally stable equilibria in the low-frequency setting.

We discriminate between these hypotheses by testing whether the minimum distance from either equilibrium, denoted MINDIST = Min(DISTS, DISTE) is greater in the high-frequency setting than in the low-frequency setting, and whether MINDIST is influenced by the history treatment. As shown in Table 1, the high level of verification frequency results in a greater minimum distance of outcomes from equilibrium in both the Signaling-History setting (an increase of 11, $p < 0.05$) and in the Exaggeration-History setting (an increase of 21, $p < 0.01$). In contrast, history (which does not alter local stability) has no significant effect on MINDIST either directly or through an interaction with verification frequency. These results strongly support the evolutionary model over the IC model.

Behavior of individual cohorts confirms these results. Random allocation of average outcomes of the last five rounds would result in only 8% of the cohorts having both LOW and HIGH within 20 of an equilibrium. However, Table 2 shows that 19 of the 24 cohorts exhibit such behavior in the Low-Frequency setting, while only 5 of 24 cohorts do so in the High-Frequency setting.

Figure 3 displays the evolution of MINDIST. In both the Exaggeration and Signaling History settings, MINDIST declines more in the Infrequent-Verification setting than in the Frequent-Verification setting. Figure 4 shows similar behavior in the absolute strategy changes from round to round, with results collapsed across history settings. For HQ sellers, changes are similar until the last five rounds, by which time HQ sellers in the Infrequent-Verification setting start to settle into strategies, while HQ sellers in the Frequent-Verification setting do not. For LQ sellers, the results are even more dramatic: even in the early rounds, LQ sellers in the Infrequent-Verification setting change strategies much less from round to round than do LQ sellers in the Frequent-Verification setting.

5.3. A direct test of the learning model

As Subsections 5.1 and 5.2 show, the evolutionary model explains participant behavior in the game much better than a traditional model relying on the intuitive criteria. We now test more directly the evolutionary learning model described by equation (4). For each round t and seller type i, we regress the signed changes in HIGH and LOW with the computation of $\delta(r_i, p_{t-1}, q_{t-1})$, denoted PAYDIFF. Recall from equation (3) that $\delta(r_i, p_{t-1}, q_{t-1})$ depicts the payoff to exaggerating minus the payoff to not exaggerating, given the seller's type and the beliefs of the buyers. The learning model assumes that sellers increase (decrease) their rate of exaggeration when the last round's payoff to exaggeration was greater (less) than the payoff to not exaggerating. Thus, the coefficient on PAYDIFF should be positive.

As shown in Table 4, the coefficient is nominally positive in all four cells of the design, for both LQ and HQ sellers. It is statistically significant only in the Frequent-Verification setting. A possible explanation for the lack of statistical significance in the Infrequent-Verification setting is that the high speed with which sellers settle into an equilibrium makes the learning process hard to observe. Sellers never settle into an equilibrium in Frequent-Verification, providing a more powerful test of learning.

Despite the fact that the parameter on PAYDIFF has the expected sign in the Frequent-Verification setting, the explanatory power of the regressions is very small–the adjusted R^2 of the regressions exceeds 5% in only one cell. The participants' learning process is clearly far more complex than equation (4) suggests. The strong power of local stability to predict equilibrium selection and nonconvergence suggests that evolutionary models can still provide useful qualitative predictions, even when the detailed learning process is not well understood.

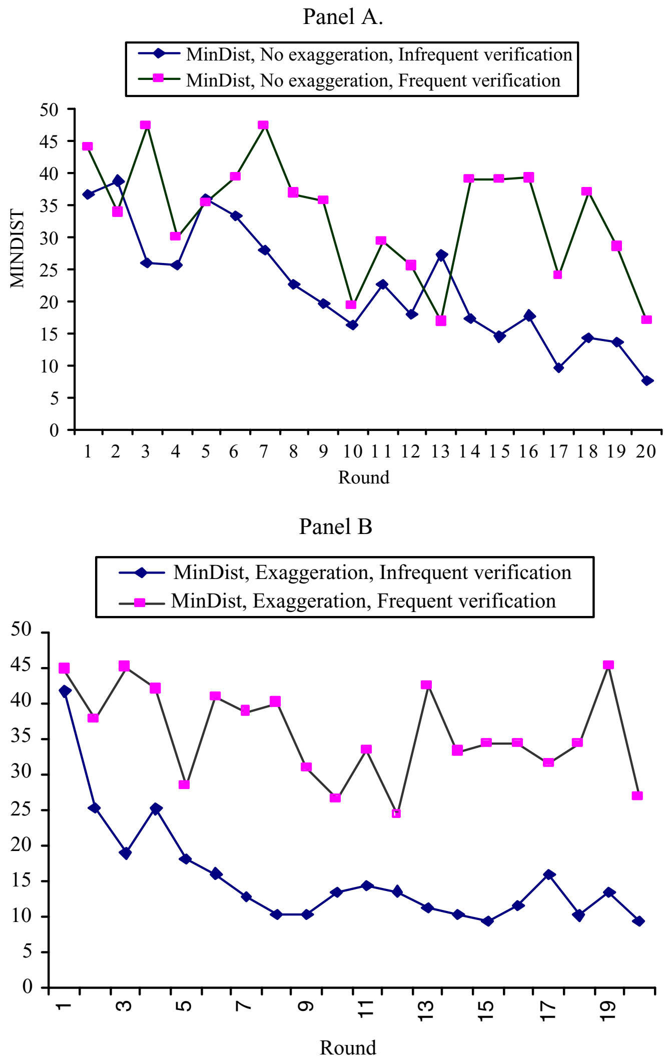

Figure 3. Evolution of distance from any equilibrium (MinDist) over time. This figure depicts the evolution over time of various dependent variables in the Low- and Infrequent-Verification settings. Panel A depicts the evolution of the MinDist to either the Exaggeration or the Signaling equilibrium with a no-exaggeration history, and Panel B with an exaggeration history.

5.4. *Implications of fixed-pairing.*

Our use of fixed pairings might allow players to use multi-period strategies. Such strategies weaken support for our predictions, because both evolutionary models and equilibrium analysis of the one-shot game assume that players do not consider such strategies. We find little evidence that multi-period strategies influence our data. In particular, we find

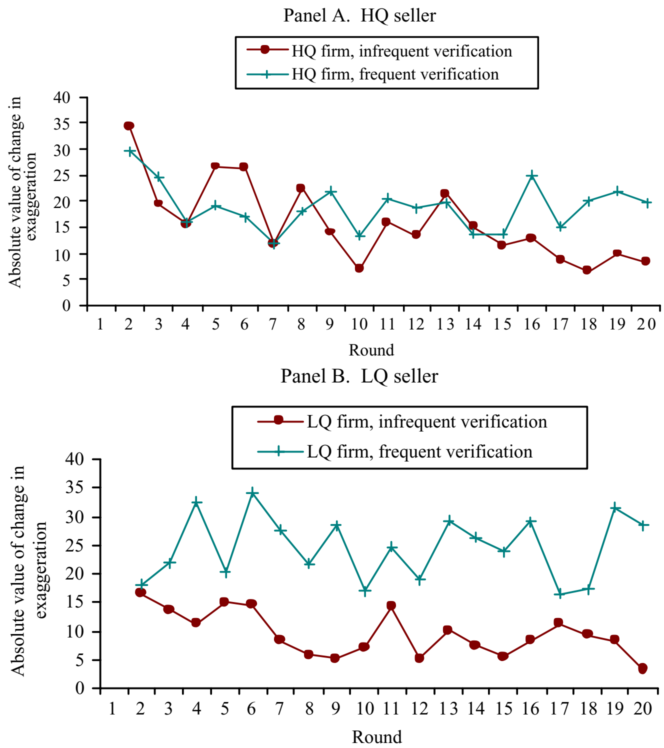

Figure 4. Evolution of absolute changes in exaggeration strategies of firms: Panel A depicts the evolution of average absolute strategy changes for the HQ seller in the frequent and infrequent verification settings. Panel B does the same for the LQ seller.

support for the equilibria of the one-shot game in the infrequent-search setting; fixed-pairing evidently does not lead to other (multi-period) equilibrium outcomes in this setting. The erratic behavior of players who are unable to converge to any stable outcome in the frequent-search setting is unlikely to reflect multi-period strategies because average payoffs to all player types are typically lower than they would be in the one-shot games, given the unpredictability of other players' actions.

We believe the evidence supporting our evolutionary model is more convincing because of a fixed-pairing scheme. A concern about evolutionary models is their assumption that players do not use multi-period strategies even though they are playing a multi-period game. Our experiment shows that evolutionary models can still have substantial predictive power even when the matching scheme allows such participants to implement multi-period strategies.

6. Summary and contributions

A choice to exaggerate verifiable attributes can convey credible information about unverifiable attributes, even if the choice poses no direct cost; the credibility of such a choice

Table 3. Regression results of the direct test of the learning model. The dependent variable here is the change in the exaggeration strategies of the firm. The independent variables are an intercept and Paydiff, the difference in payment from exaggerating versus not exaggerating.

Type of firm	Intercept (t-stat)	Paydiff (t-stat)	R-square
	Signaling history, frequent verification		
HQ	−5.83 (−1.04)	1.04 (1.34)	0.0035
LQ	−2.25 (−0.53)	0.25 (0.87)	−0.0011
	Signaling history, frequent verification		
HQ	−28.73 (−3.76)	13.78 (4.03)	0.0630
LQ	48.58 (−4.92)	11.40 (5.17)	0.1026
	Exaggeration history, infrequent verification		
HQ	−0.28 (−0.03)	0.17 (0.16)	−0.0043
LQ	−2.16 (−0.38)	0.23 (0.66)	−0.0025
	Exaggeration history, frequent verification		
HQ	−23.15 (−2.82)	9.54 (2.80)	0.0293
LQ	−30.12 (−3.38)	6.32 (3.42)	0.0449

depends on whether the buyers will verify the verifiable claim. As long as verification is sufficiently frequent, there is a signaling equilibrium where not exaggerating is a signal of strength. There is also an exaggeration equilibrium in which all sellers exaggerate, so exaggeration conveys no information. We show that an evolutionary learning model has much greater predictive power than a model based on the intuitive criteria (Cho and Kreps, 1987). For example, when verification is not too frequent, players often converge to the exaggeration equilibrium, even though it does not satisfy the intuitive criteria. When verification is not too frequent, both equilibria are evolutionarily stable, but only one satisfies the intuitive criteria; we find that players converge to whichever equilibrium is closest to their arbitrary starting point. When verification is more frequent, neither equilibrium is stable and both satisfy the intuitive criteria; we find that players don't converge to any outcome.

We make three contributions. First, we explore the important (for product marketing and financial reporting) link between verifiable and unverifiable claims, where there might be no incremental cost of exaggerating the verifiable claim. E.g., while the cost of buying advertising space (in an expensive medium) can be a signal of (good) quality, the advertising copy itself is a costless cheap-talk signal to the consumer. In our "cheap-talk" model it is only buyers' endogenous beliefs about the implications of verifiable claims that can make them credible signals. Our model is thus quite different from models with costly signals.

Second, the study shows the compelling power of an "exaggeration equilibrium" in which all sellers exaggerate their claims. With great foresight, Arthur Levitt, chairman of the Securities and Exchange Commission, asked managers to end "gamesmanship," perhaps concerned that might lead to an exaggeration equilibrium, making investors lose faith in financial reporting (Levitt, 1999). Pharmaceutical Executive (July 1998) reports "on the need for the pharmaceutical industry to "*take the high road*... The industry needs to set standards for direct-to-consumer advertising that will maintain a positive public image and *enhance credibility"* (italics added). The Electronic Retailing Association announced a new self-regulatory program "... to compete with outlandish claims" (Ives (2004)). Our model implies that regulators and high-quality sellers must immediately curtail any trend of widespread exaggeration to avoid a slippery slope leading to meaningless unbelievable claims, especially with infrequent verification by consumers.

Third, the superior predictive power of the evolutionary learning model has substantive implications in explaining marketing outcomes. The "hand of the past" matters, even if those circumstances have no direct relevance to agents' incentives. For example, if intense competition during a past business slump led to an exaggeration history, firms may find it hard to escape this equilibrium even when business conditions improve. Some markets may never reach equilibrium if all equilibria are evolutionarily unstable, leading to apparently unpredictable and erratic behavior, with sellers and buyers appearing to behave suboptimally (because they cannot predict their opponents' behavior or beliefs).

Future research might focus on how specific price-setting and patience-related mechanisms influence buyer expectation and seller behavior. It would also be promising to investigate whether a regulatory cost imposed on exaggeration might eliminate the exaggeration equilibria and make signaling equilibria stable even when verification is frequent. Finally, future research might examine and test how our learning model might affect equilibrium selection in other game theoretic models in marketing. Understanding the process by which players converge (or don't converge) to equilibrium represents an important step towards building richer and more predictive game-theoretic models.

Appendix 1: Propositions and proofs

Preliminaries

Let $F(p,q) = E[V|X,\alpha_u,p,q] - E[V|X,\alpha_e,p,q]$ represent the differences in expected values given two different claims. This allows equation (1) in the text to be rewritten as

$$\delta(\pi_{i,} p, q) = (1-\pi_i)\varepsilon A + [(1-\pi_i)(1-\varepsilon) + \pi_i]F(p,q) \tag{1'}$$

For values not satisfying $p = q = 0$ and $p = q = 1$, we have

$$F(p,q) = \left(\frac{1-q}{(1-q)+(1-p)} - \frac{q}{q+p}\right)S \tag{2'}$$

Note that both functions are continuous in p and q, increasing in p and decreasing in q, and take on maximum values at $p = 1, q = 0$. When $p = q = 0$ or $p = q = 1$, one of the terms in (2′) is undefined, and a value for that term must be assumed as part of the off-equilibrium-path beliefs.

Proposition 1. *For every parameterization of the game, there is an "Exaggeration" perfect Bayesian equilibrium with $p^* = q^* = 1$.*

Proof: Assume that buyers believe that $p = q = 1$ and that $\Pr(\Theta = \Theta_H|\alpha = \alpha_u) \leq 1/2$ (off-equilibrium-path beliefs). Then $F(p,q) \leq 0$ so that $\delta(\pi_i, p, q) > 0$, and both seller types prefer to exaggerate always. The buyers' beliefs are never contradicted (because they always see exaggeration), so $p^* = q^* = 1$ is a perfect Bayesian equilibrium. □

Proposition 2. *Assuming that $\pi_H > \pi_L$, there exist three cutoffs c_1, c_2, and c_3, with $c_1 > c_2 > c_3$, with the following properties*:

(i) *If $\varepsilon > c_1$, then there exists no Perfect Bayesian Equilibrium (PBE) in which $q^* < 1$ or $p^* < 1$.*
(ii) *If $c_1 \geq \varepsilon \geq c_2$, then there exists a PBE in which $q^* = 0$ and $p^* = 1$*
(iii) *If $c_2 > \varepsilon > c_3$, then there exists a PBE in which $q^* = 0$ and $p^*(\varepsilon) \in (0, 1)$, with $p^*(\varepsilon)$ decreasing in ε.*
(iv) *If $c_3 > \varepsilon$, then there exists a PBE in which $p^* = q^* = 0$. In this case, the PBE is supported by the off-equilibrium-path belief that a seller who exaggerates must be of low quality.*

Proof: Solving equation (1') for the value of ε that sets $\delta(\pi_{i,} p, q)$ to 0 yields

$$\varepsilon = F(p, q)/[(1 - \pi_i)(A + F(p, q))]. \tag{3'}$$

F takes on its maximum value of $F(p, q) = S$ at $(p = 1, q = 0)$, so the cost of exaggeration can never be larger. Setting $c_1 = S/[(1 - \pi_H)(A + S)]$ satisfies condition (i) because, for equal or higher values of ε the HQ seller will always prefer to exaggerate. Because $\delta(\pi_{i,} p, q)$ is increasing in π, the LQ seller also always prefers to exaggerate for any ε greater than c_1, so $p^* = q^* = 1$ in equilibrium.

If ε is greater than this value, then seller i will always prefer to exaggerate because

Setting $c_2 == S/[(1 - \pi_L)(A + S)] < c_1$ satisfies condition (ii) because, for higher values of ε the LQ seller will always prefer to exaggerate, while the HQ seller will wish not to exaggerate.

When the buyer believes that $p^* = q^* = 0$ and exaggeration indicates an LQ seller (an off-equilibrium-path belief, since exaggeration is never observed in equilibrium), $F(0, 0) = 1/2S$. The value of ε leaving seller of type i indifferent between exaggerating and not exaggerating is $\varepsilon = 1/2S/[(1 - \pi_i)(A + 1/2S)]$. Setting $c_3 = 1/2S/[(1 - \pi_L)(A + 1/2S)]$ satisfies condition (iv) because, for values of ε at or below this level, the LQ seller (and therefore also the HQ seller) will prefer not to exaggerate.

In between c_2 and c_3, HQ seller prefers not to exaggerate for any level of p. Any equilibrium must involve a mixed strategy for the LQ seller ($0 < p < 1$). Because $\delta(\pi_i, p, q)$ is monotonic, there is a single value $p^*(\varepsilon)$ that leaves the LQ seller indifferent between exaggerating and not exaggerating, which allows a mixture of the two strategies. Implicit differentiation shows that $p^*(\varepsilon)$ is declining in ε.

As a final remark, note that the cutoffs may be negative or greater than 1, given the values of the constants A, S, π_L, π_H and k. Because ε must lie within the interval [0, 1], some of the regions described in the proposition may be infeasible. Also, note that the proposition could be rewritten in terms of cutoffs for another parameter (such as A, S or S/A), fixing all remaining parameters including ε. □

Proposition 3. *There is no perfect Bayesian equilibrium in which $p^* < q^*$.*

Proof: If $p^* < q^*$ then $F(p^*, q^*)$ is negative. As a result, $\delta(\pi_i, p^*, q^*) > 0$ for all parameterizations. This implies that $p^* = q^* = 1$, which contradicts the premise. □

Definition 1. Let $H_\eta(p^*, q^*) = \{(p, q) : |p - p^*| < \eta \text{ and } |q - q^*| < \eta\}$ be an η-neighborhood of the equilibrium (p^*, q^*), with $\eta > 0$. An equilibrium (p^*, q^*) is locally stable with respect to dynamic (4) if and only if that dynamic $(p_t, q_t) \in H_\eta(p^*, q^*)$ implies that $(p_{t+\tau}, q_{t+\tau}) \in H_\eta(p^*, q^*)$ for all >0, and that $(p_t, q_t) \to (p^*, q^*)$ as $\tau \to \infty$. An equilibrium that is not locally stable is locally unstable.

Proposition 4. *(a) The Signaling Equilibrium* $(p^*, 0)$ *exists and is stable with respect to dynamic* (4) *if and only if* $c_3 \geq \varepsilon > c_1$. *(b) The Exaggeration Equilibrium* (1, 1) *is stable with respect to dynamic* (4) *if and only if* $\varepsilon > 1/2\ S/[(1 - \pi_H)(A + 1/2S)]$.

Proof: Under the definition above, local stability requires that the sign of $\delta(\pi_i, p, q)$ remains the same within some neighborhood of the equilibrium (p^*, q^*) if player i chooses a boundary strategy in equilibrium. Local stability also requires that the sign of $\delta(\pi_i, p, q)$ be negative above the equilibrium and positive below the equilibrium if player i chooses a mixed strategy in equilibrium.

Proof of (a): If $\varepsilon > c_1$ then the Signaling equilibrium does not exist. If $c_1 > \varepsilon > c_2$, then $\delta(\pi_H, p^*, q^*)$ is strictly negative, while $\delta(\pi_L, p^*, q^*)$ is strictly positive. The continuity of δ guarantees that these functions maintain these signs in some neighborhood of the equilibrium. If $\varepsilon < c_2$, then $\delta(\pi_H, p^*, q^*)$ is strictly negative, while $\delta(\pi_L, p^*, q^*)$ is 0 (the equilibrium involves mixed strategies). Because δ is decreasing in p, the sign of $\delta\ (\pi_L, p, q)$ is negative immediately above p^* and positive immediately below p^*. The equilibrium is therefore locally stable. If $c_3 \geq \varepsilon$, then $\delta(\pi_i, p^*, q^*)$ is negative for both i at the equilibrium (0, 0). However, $\delta(\pi_i, p^*, \eta) > 0$ for all $\eta 0$, so the equilibrium is unstable.

Proof of (b): Local stability for the corner equilibrium (p^*, q^*) requires that $\delta(\pi_H, 1, q)$ is positive for all values of q more than $1 - \eta$, while $\delta(\pi_L, \text{p}, 1)$ is positive for all values of p more than $1 - \eta$. The latter is always true given that $p < q$ implies that $\delta(\pi_i, p, 1)$ is positive (see proposition 2). The former is true if and only if $\delta(\pi_H, 1, 1) > 0$, which occurs when ε is greater than the expression given in the proposition. □

Appendix 2: Excerpts from instructions to subjects

In this experiment, participants form groups of two "Managers" and one "Analyst."

The value of a firm

The total value of each firm is equal to the sum of two numbers:

(i) A "base value," denoted "B," and
(ii) A "Hidden-Value," which depends on the firm's prospects. The Hidden-Value is always +10 for one of the manager's firms and always −10 for the others. As a result, all of the

firms of one manager (the "High-Value" Manager) are worth B + 10, while all of the firms of the other manager (the "Low-Value" Manager) are worth B − 10.

The managers' task

Each manager gives a report about the firm's base value to a news service. However, the manager is not able to report this base value exactly–instead, the manager must report a number 15 francs higher than base value *or* 15 francs lower than base value. Reporting a number 15 francs too high is called INFLATION while reporting a number 15 francs too low is called DEFLATION.

The analyst's task

The analyst gets information from the News Service and estimates the value of each firm. Half of the time, the analyst will learn only the base value that was reported by the manager. The other half of the time, the news service corrects the report, telling the analyst the true base value, and whether the manager inflated or deflated the report. However, the analyst never learns whether the firm is managed by the High-Value or Low-Value Manager.

When there is no Correction provided by the news service, the analyst simply enters a "Bias Adjustment" to adjust the reported value for the amount of inflation/deflation expected. The bias adjustment could be as high as +15 (if the analyst believes both managers deflate always) or as low as −15 (if the analyst believes both managers inflate always). Because the firm is as likely to have a high Hidden-Value as a low Hidden-Value, it isn't possible to make any adjustment for Hidden-Values. Therefore, the resulting price, called the "UNCORRECTED price," is

UNCORRECTED price = Reported Base Value + Bias adjustment

When there is a Correction provided by the news service, the analyst knows whether the manager inflated or deflated. The analyst can then determine a "Hidden-Value Adjustment" which reflects the probability that the firm was provided by the High-Value or Low-Value Manager. For example, the Hidden-Value Adjustment could be as high as +10 (if the analyst believes that only the High-Value Manager would have reported that way) or as low as −10 (if the analyst believes that only the Low-Value Manager would have reported that way). The resulting price is called the "CORRECTED Price," and is

CORRECTED price = True Base Value + Hidden-Value Adjustment

Analyst's payoffs

The analyst is paid based on the accuracy of his Bias and Hidden-Value adjustments, as follows:

UNCORRECTED: 5 FrancsError in Bias Adjustment
CORRECTED: 5 FrancsError in Hidden-Value Adjustment

Managers' payoffs

Each Manager receives a base salary of 1 franc for each firm, plus or minus a proportion of the difference between the price and the base value, as shown in the following table. If the price is greater than base value, the Manager will earn more than 1 franc; if the price is less than the base value, the Manager will earn less than 1 franc. The exact proportion depends on the Hidden-Value of the manager's firm, and on whether the price is Corrected or Uncorrected. Halfway through the experiment, the commissions will switch from Schedule 1 to Schedule 2; you will be informed when this happens.

	Commission schedule for high-value manager	Commission schedule for low-value manager
Commission Schedule 1	28% if uncorrected 72% if corrected	56% if uncorrected 44% if corrected
Commission Schedule 2	8% if uncorrected 92% if Corrected	16% if uncorrected 84% if corrected

100 firms at a time

The experiment will be just as described above, except that each manager will be asked how many of 100 reports to inflate, and how many to deflate. Half of the inflated reports and half of the deflated reports will be corrected. The analyst must estimate the value of all 200 firms, by deciding on one Bias Adjustment, one Hidden-Value Adjustment for inflated reports, and one Hidden-Value Adjustment for deflated reports.

Each participant's payoff for each round is simply equal to the sum of the payoff for each of the firms.

Sequence of events

Before the experiment begins, the proctor will review the instructions aloud. You will then begin the task under commission schedule 1. After you make your decisions for many rounds (learning the outcome after each round) decisions many times, you will switch to commission schedule 2. Before starting each commission schedule, you will have a chance to review an example of a representative outcome under that commission schedule.

Converting points to money

Your winnings will be equal to the greater of \$5 and

\$20 × (Francs you won)/Average francs won *by subjects in your role*)

Your winnings (in francs) for the first three rounds are multiplied by 5. Thus, you should think very carefully about what you want to do, because they are 5 times more important than later rounds. **Your winnings are not affected by whether the others in your group do well or poorly.** It is entirely possible that all three members of your group will do well (or poorly) relative to other Managers and Analysts. No matter what happens, you will receive at least $5 for completing the session.

Infrequent-Verification/Exaggeration-History Example

(Other examples were included in the instructions, but suppressed here and are available from the authors on request)

Inflation rates High-value manager 70% Low-value manager 55%		
Bias adjustment	Hidden-value adjustment (inflation)	Hidden-value adjustment (deflation)
	Understanding the analyst's screen	
• On average, 125 of the 200 reports are INFLATED. • Because this is more than half (62.5%), the reports tend to be biased UPWARD, and the analyst should use a NEGATIVE Bias Adjustment (−3.6)	• Of the 125 inflated reports, 70 are from the High-Value Manager • Because this is more than half (56%), inflated reports tend to have a HIGH Hidden-Value, and the analyst should use a POSITIVE Hidden-Value Adjustment (+1.2)	• Of the 75 deflated reports, only 30 are from the High-Value Manager • Because this is less than half (40%), deflated reports tend to have a LOW Hidden-Value, and the analyst should use a NEGATIVE Hidden-Value Adjustment (−2)

Note that the Analyst's adjustments are quite close to the best adjustments. As a result, the Analyst earns close to the maximum of 1000 francs.

Given the Analyst's Adjustments, the payoffs to inflating and deflating are computed as follows:

	High-value manager (price × commission × #firms)	Low-value manager (price × commission × #firms)
	Understanding the managers' screens	
Assuming inflation for all 100 firms	$11 \times 28\% \times 50 = 154$ $1 \times 72\% \times 50 = 36$ + Salary 100 Total 290	$11 \times 56\% \times 50 = 308$ $1 \times 44\% \times 50 = 22$ + Salary 100 Total 430
Assuming deflation for all 100 firms	$-19 \times 28\% \times 50 = -266$ $-2 \times 72\% \times 50 = -72$ + Salary 100 Total −238	$-19 \times 56\% \times 50 = -532$ $-2 \times 44\% \times 50 = -44$ + Salary 100 Total −476

A manager who inflates some reports and deflates others will receive a weighted average of the two totals above, as shown at the bottom of the managers' computer screens.

References

Amaldoss, W., R. Meyer, J. Raju, and A. Rapoport. (2000). "Collaborating to Compete," Marketing Science 19(2), 105–126.

Bloomfield, Robert. (1994). "Learning a Mixed Strategy Equilibrium in the Laboratory." Journal of Economic Behavior and Organization 25, 411–436.

Bloomfield, Robert. (2003). "The Incomplete Revelation Hypothesis: Implications for Financial Reporting." Accounting Horizons 16(3), 233–244.

Blume, A., D. DeJong, Y. Kim, and G. Sprinkle. (1994). "Evolution of the Meaning of Messages in Sender-Receiver Games: An experiment." Working paper, University of Iowa.

Brandts, J. and C. Holt. (1992). "An Experimental Test of Equilibrium Dominance in Signaling Games." American Economic Review 82, 1350–1365.

Camerer, C. and T. Ho. (1998). "Experience-Weighted Attraction Learning in Games: A Unifying Approach." Econometrica.

Cho, I. and D. Kreps. (1987). "Signaling Games and Sequential Equilibria." Quarterly Journal of Economics 102, 179–221.

Crawford, V.P. and J. Sobel. (1982). "Strategic Information Transmission." Econometrica 50, 1431–1451.

Darbi, M. and E. Karni. (1973). "Free Competition and Optimal Amount of Fraud." Journal of Law and Economics 16, 67–88.

Erev, I. and A. Roth. (1998). "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategies Equilibria." American Economic Review 88 (4), 848–882

Farrell, J. and R. Gibbons. (1989). "Cheap-Talk With Two Audiences". American Economic Review 79, 1214–23.

Herrnstein, R. (1997). "The Matching Law." In H. Rachlin and D. Laibson (eds.), Papers in Psychology and Economics. Cambridge: Harvard University Press.

Hofbauer, J. and K. Sigmund. (1988). The Theory of Evolution and Dynamical Systems. London: London Mathematical Society.

Ives, N. (2004). Infomercials Clean Up Their Pitch (But Wait, There's More). New York Times, April 12.

Johnson, E.J. and J. Edward Russo. (1984). "Product Familiarity and Learning New Information." Journal of Consumer Research 11, 542–550.

Kreps, D. and R. Wilson. (1982). "Sequential Equilibria." Econometrica 50, 863–894.

Levitt, A. (1999). Quality Information: The Lifeblood of Our Markets. Address delivered at the Economic Club of New York on October 18.

Lilien, S., L. Mellman. and V. Pastena. (1988). "Accounting Changes: Successful versus Unsuccessful Sellers." The Accounting Review 63, 642–656.

Moorthy, S., Ratchford, B.T. and D. Talukdar. (1997). "Consumer Information Search Revisited: Theory and Empirical Analysis." Journal of Consumer Research 23, 263–277.

Nelson, P. (1970). "Information and Consumer Behavior." Journal of Political Economy 94, 796–821.

Rizzo, J. and R. Zeckhauser. (1990). "Advertising and Entry: The Case of Physician Services." Journal of Political Economy 98(3), 476–500.

Simester, D. 1995. "Signaling price image using advertised prices." Marketing Science 14(2), 166–188.

Smith, V. (1976). "Experimental Economics: Induced Value Theory." American Economic Review 66, 274–279.

Thorndyke, E. (1911). Animal Intelligence. New York: Macmillan.

Titman, Sheridan and Brett Trueman. (1988). "An Explanation for Accounting Income Smoothing." Journal of Accounting Research (Supplement) 26, 127–144.

Wahlen, J. (1994). "The Nature of Information in Commercial Bank Loan Loss Disclosures." The Accounting Review 3, 455–478.

Weibull, J. (1995). Evolutionary Game Theory, Cambridge: The MIT Press.