Feints

KENNETH HENDRICKS
Department of Economics
University of Texas
Austin, TX 78712
hendrick@eco.utexas.edu

R. PRESTON MCAFEE
Humanities & Social Sciences
California Institute of Technology
Pasadena, CA 91125
mcafee@hss.caltech.edu

In both economic and military situations, agents may try to mislead rivals about their true types or plans, whatever they may be. We consider a simple model in which one player attacks and the other player defends. We show that such environments have two types of possible equilibrium behavior, depending upon the signaling technology. If the signal is not very revealing about the attacker's plans, then the attacker always invests more resources in attack than in misdirection. If the technology is revealing, then the attacker does not always feint, but when he feints, he invests more than half of his resources into misdirection. Comparative statics also depend on whether the technology is revealing.

> *"Always mystify, mislead and surprise the enemy, if possible."*
> *—General Thomas J. "Stonewall" Jackson*

> *"Create havoc in the east and strike in the west."*
> *—Sun Tze*

1. Introduction

In the past two decades, economists have used signaling models to rationalize and understand a number of frequently practised business strategies. For example, Milgrom and Roberts (1982) show that limit pricing can be explained as an attempt by an incumbent firm to signal

We thank audiences at Northwestern, Arhus, Heidelberg, Rochester, Wisconsin, ITAM, UC-San Diego and, in particular, Hugh Neary, Eric Maskin, Larry Samuelson, Guofu Tan, anonymous referees, and an associate editor for helpful discussions and comments. Ken Hendricks thanks SSHRCC for financial support.

Journal of Economics & Management Strategy, Volume 15, Number 2, Summer 2006, 431–456

that its costs are low and hence, that entry is likely to be unprofitable. Their model can, with small changes, be reinterpreted to explain predatory pricing. In a model, where entrants use their current profits to infer future profits, Fudenberg and Tirole (1986) show that the incumbent firm will try to manipulate the entrant's inference by lowering its price, thereby encouraging the entrant to exit the market. Saloner (1987) uses a variant of the Milgrom-Roberts model to explain why predatory pricing can lower a firm's costs of acquiring a rival. Following the seminal work of Milgrom and Roberts (1986), a large literature has developed showing how "burning money" in the form of image advertizing can signal consumers that the featured product is high quality.

This paper develops a new kind of signaling model to study feints, a commonly used strategy in conflicts. Feints are offensive actions conducted with the purpose of deceiving rivals as to the location (or time) of the main offensive action. They are ubiquitous in military campaigns and sports, but also occur frequently in political and legal conflicts. In business, feints are likely to occur when a firm is trying to enter a new market or market a new product but wishes to conceal the true nature of its investment plans to prevent a rival from counterattacking effectively while it is still vulnerable. For example, by the end of the 1870s, the Standard Oil Trust had locked up most of the transportation and refining of oil in Pennsylvania, at that time the world's largest producing region. A group of producers attempted to circumnavigate Standard Oil by building a 110-mile pipeline to connect the oil regions to the Pennsylvania and Reading Railroad. This plan was audacious, because no one had ever attempted such a large pipeline, but also because the group attempted to keep Standard Oil in the dark about the scale and location of the pipeline. Fake surveys were taken in order to confuse Standard Oil of the route of the pipeline so that it could not block it. In May 1879, the Tidewater pipeline opened. It "not only caught Standard by surprise, but also meant that its control of the industry was suddenly again in jeopardy" (Yergin, 1991, p. 43). Standard acquired Tidewater in 1883. Other examples include pharmaceutical firms patenting "dead-end" products (Langinier, 2001) to mislead competitors' R&D expenditures and oil firms surveying tracts that they have no intention of bidding in order to confuse potential rivals as to which areas they should pursue.

A famous military example was the Allied invasion of Nazi-occupied France. The most natural locations for the invasion were Pas de Calais and Normandy. The office of the Chief of Staff of the Supreme Allied Commander (COSSAC) decided on Normandy. However, in order to mislead the Nazi military into continuing to fortify Calais at the expense of Normandy, the Allies developed an ambitious plan named

Operation Fortitude. Operation Fortitude was a feint. As part of this operation, a mythical "U.S. 1st Army Group," led by George Patton, was created. This army consisted of plywood airplanes and inflatable tanks located near Dover, and a vast armada of rubber landing craft located in the Thames River estuary. Patton's mythical force produced large volumes of coded radio traffic for German interception, and there were extensive real military maneuvers near the location of the ersatz army. False information was fed to known enemy agents. The operation worked. The German high command, led by Field Marshall Wilhelm Keitel, believed that the invasion of Normandy was a feint, a diversion from the real invasion to come at Calais. Nineteen German divisions remained at Pas de Calais after D-Day, the initiation of the Normandy invasion. The Allies continued bombing Calais even after the invasion began, with two bombs dropped at Calais for every bomb dropped at Normandy, successfully maintaining the ruse for several critical days.[1]

We model feints as a sender-receiver game with noisy signals. The sender has to allocate resources between at least two alternatives. The sender's type determines which alternative is preferred. The receiver cannot observe the sender's type but he does observe a noisy signal of the sender's allocation and can draw inferences about his type from the observed signal. The receiver's inference matters to the sender because the receiver's gain is his loss. Thus, each sender type has an incentive to mislead the receiver. The distinguishing feature of our model is that the way in which each type tries to mislead the receiver is to pretend to behave like the *opposite* type. For example, the Allies wanted the Germans to think that they were attacking Pas de Calais but, if they had attacked Pas de Calais, they would have wanted to be perceived as attacking Normandy. Similarly, Tidewater wanted to be perceived as building the pipeline on the route covered by the fake surveys, but had it intended to build the pipeline along this route, it would have wanted to be perceived as building the pipeline on the route that it actually chose.

As in most signaling models, we are interested in knowing when and how much firms are likely to invest in signal distortion, and whether rivals are misled into taking actions that favor the sender. Our major result distinguishes two cases. If the signal generated by the sender's action is not very "revealing" (in the sense to be made precise), then the sender pursues a pure strategy and allocates most of his resources to the preferred alternative. Feints are small and deterministic. In contrast, when the signal is very revealing, the sender cannot easily conceal his actions and small feints are not effective. Moreover, large, deterministic feints are also not effective. The receiver, knowing the sender's strategy,

1. For more details, see Hatfield (1997).

understands that the signal is likely to be generated by the feint and infers that the preferred alternative is the opposite of the one indicated by the signal. He responds accordingly, which makes large feints a poor strategy. The equilibrium strategy for senders is to randomize, sometimes investing all of his resources and other times investing less than one-half of his resources, in the preferred alternative. The randomization ensures that his adversary is uncertain about the sender's type, and about the scale of the sender's investment conditional on his type. Thus, in our model, all sender types feint, and the feints are sometimes quite successful. Our model can explain both frequent, small feints and infrequent, large feints.

The comparative static results for the two cases are strikingly different. The risk aversion of the sender's rival works to the sender's advantage when the signal is noisy but is irrelevant when the signal is revealing. It might seem that a more revealing signal would always hurt the sender because the sender is best off when he can act secretly without generating any signal. This intuition is basically correct for noisy signaling technologies. But, when the technology is revealing, the informativeness of the signal works to the sender's advantage. The reason is that a more revealing signal makes it easier for the sender to feint, and hence more profitable.

Crawford and Sobel (1982) and Sobel (1985) show that no communication occurs in sender-receiver games where messages are costless ("cheap talk") and the receiver's gain is the sender's loss. Farrell (1993) used the term "babbling" to describe equilibria of cheap talk games in which the sender's message is uninformative. Crawford (2003), in a paper motivated by our analysis, introduces boundedly rational types into a zero-sum, cheap talk game and shows that their presence yields equilibria in which the sophisticated (fully rational) sender lies and fools a sophisticated (fully rational) receiver some of the time. "Babbling" equilibria do not arise in our model because signaling is costly. If the receiver ignores the signal, the sender's best reply is to apply more resources to the preferred alternative, because resources devoted to fooling the receiver are mostly wasted. But, in that case, the receiver should not ignore the signal but allocate more resources at the alternative signaled. Thus, in the equilibrium of our model, the signal is informative, in that the receiver's Bayes update differs from his prior and he allocates more resources to the alternative indicated by the signal.

Our model and results contribute to signaling theory. The standard assumption in signaling models in industrial organization is that all sender types wish to be perceived as the same type. For example, in Milgrom and Roberts' (1982) model of entry and in Fudenberg and Tirole's (1986) model of signal jamming, all incumbent firm types want to be perceived as having the lowest cost because this inference makes it more likely that the entrant will stay out.[2] Bernheim (1994) and Banks (1987) study environments in which all sender types wish to be perceived as the middle type. In Crawford and Sobel (1982), each sender type wants to be perceived as a higher type, but not necessarily the highest type. Bernheim and Severinov (2003) study a model of bequests in which the most desired type may be higher or lower than the sender's type, but the desired type is still an increasing function of the sender's true type. Models in which the desired type is monotone increasing in actual type typically possess a separating equilibrium in pure strategies in which essentially all sender types invest in the signal. By contrast, in our signaling model, the desired type is a *decreasing* function of the sender's true type: "high" types want to be perceived as "low" types and "low" types want to be perceived as "high" types.[3] We are not aware of anyone who has studied this class of signaling games, with or without noisy signals.

The paper is organized as follows. In Section 2, we develop a simple model of feints. In Section 3, we characterize the equilibrium. In Section 4 we provide some comparative static results. We discuss applications in Section 5 and conclude in Section 6.

2. The Signaling Model

The generic signaling model consists of two players, one sender and one receiver, in which the sender's type is private information. The sender acts first, and the receiver acts second after observing a noisy signal about the action taken by the sender. The payoffs to each player depend upon the sender's type and the actions taken by the two players. In our model, each player's action consists of a division of resources between two investment alternatives, a and b. Let x denote the resources that the sender allocates to alternative a. We normalize the sender's total available resources to one, so that the balance $1 - x$ is allocated to alternative b. The receiver then receives a private, binary signal $S \in \{\alpha, \beta\}$ about the allocation of the sender's resources. The signal yields information about the relative allocation of resources, that is, which alternative is getting more resources. Define

$$p(x) = \Pr\{S = \alpha \mid x\}$$

2. Matthews and Mirman (1983) introduced noisy signals into the standard signaling game.

3. Von Neumann and Morgenstern (1944) introduced the term "inverted signalling" to describe environments in which each sender type wants to signal to the receiver that it is the *opposite* type. Their main example is poker, where a weak player wants to give a (false) impression of strength and a strong player wants to give a (false) impression of weakness. Of course, they consider only zero-sum, complete information games.

as the probability that the receiver observes signal α conditional on the sender allocating x units to alternative a. We will impose the following regularity conditions on the signaling technology p.

Assumption 1: (i) p is strictly increasing and differentiable, (ii) $p(z) + p(1 - z) = 1$, (iii) $p(0) = 0$.

Condition (i) states that the probability of generating signal α increases with the amount of resources allocated to alternative a. Condition (ii) imposes symmetry. It implies that the probability of the receiver getting signal α when the sender allocates z units to alternative a is the same as the probability of getting signal β when the sender allocates z units to alternative b. Note that this condition implies that $p(\frac{1}{2})$ is equal to $\frac{1}{2}$. Condition (iii) states that if the sender allocates no resources to alternative a, then the receiver is certain to get signal β. By symmetry, if the sender allocates all of his resources to alternative a, the receiver is certain to get signal α. In other words, the sender has to allocate resources to both alternatives if he wants the receiver to be uncertain about which alternative is getting more resources.[4]

In response to the signal, the receiver divides his resources between the two alternatives. We assume that the receiver has the same amount of resources available as the sender. Let y_s denote the resources allocated to alternative a conditional on obtaining signal s, with the balance $1 - y_s$ allocated to b.

The sender's payoff is a weighted average of his investments in the two alternatives:

$$\begin{aligned}\pi(x; q, y_\alpha, y_\beta) &= q[x - (p(x)y_\alpha + (1 - p(x))y_\beta)] \\ &\quad + (1 - q)[1 - x - (p(x)(1 - y_\alpha) + (1 - p(x))(1 - y_\beta))] \\ &= (2q - 1)[x - (p(x)y_\alpha + (1 - p(x)y_\beta)],\end{aligned}$$

where $q \in [0, 1]$. Here q indexes the sender's type and measures his preferences for alternative a over alternative b. In the absence of a receiver, high types obtain a higher payoff from investing in alternative a and low types obtain a higher payoff from investing in alternative b. The cumulative distribution of q is denoted by F. It is assumed to be atomless, with support [0, 1], and a density function f which is symmetric around $\frac{1}{2}$. In the presence of a receiver, the sender's profits from each alternative increases with the size of his own investment and decreases with the size

4. Condition (iii) is not essential. If the signal is random even without an expenditure of resources, the sender receives some free feinting. This randomness may be enough for the sender, in which case no feinting occurs, but otherwise the analysis is very similar to the present results. This issue was explored in an earlier version of the paper.

of the receiver's investment. We allow the sender to randomize over the allocation of resources given his type q and denote such an allocation by $X(q)$.

The receiver weights the alternatives in the same way as the sender. When the receiver believes that $q = \hat{q}$, then his payoff from allocating y to alternative a is given by

$$v(y) = \hat{q}U(y) + (1 - \hat{q})U(1 - y),$$

where U is strictly increasing, a strictly concave function. Note that the receiver's reward from each alternative does not directly depend upon the sender's investment. Hence, the receiver always wants to allocate more resources to the alternative that he believes has the higher weight, regardless of the sender's allocation of resources. In fact, it is straightforward to show that the receiver's optimal response $y(\hat{q})$ is strictly increasing in $\hat{q}$.

The assumption that the receiver's preferences across alternatives mirror those of the sender creates a conflict between the two players. The sender wants a mismatch, in which he invests all of his resources and the receiver invests none of his resources in the preferred alternative. But doing so invites a response from the receiver, who also wants to invest heavily in the preferred alternative, which could eliminate any reward to the sender. The following assumption implies that the receiver's best reply to the sender's strategy of investing only in the preferred alternative is to match the sender's action and invest all of his resources into the same alternative.

Assumption 2: $E[q \mid q > \frac{1}{2}]U'(1) + (1 - E[q \mid q < \frac{1}{2}])U'(0) > 0.$[5]

Assumption 2 will prove useful in ruling out uninteresting boundary cases.

In general, the receiver does not know q but must infer it from the signal generated by the sender's allocation. He updates his beliefs using Bayes' rule. The posterior density of q given signal α is given by

$$f(q \mid S = \alpha; X) = \frac{Ep(X(q))f(q)}{\int_0^1 Ep(X(t))f(t)\,dq}. \tag{1}$$

The posterior density of q given signal β is given by

$$f(q \mid S = \beta; X) = \frac{(1 - Ep(X(q)))f(q)}{1 - \int_0^1 Ep(X(t))f(t)\,dt}. \tag{2}$$

5. Note that symmetry of f implies that $E[q \mid q < \frac{1}{2}] = 1 - E[q \mid q > \frac{1}{2}]$.

Let $\hat{q}_s(X)$ denote the expectation of q conditional on signal s and strategy X. Then, given signal s and strategy X, the receiver's expected payoff from allocating y to alternative a is

$$v_s(y; X) = \hat{q}_s(X)U(y) + (1 - \hat{q}_s(X))U(1 - y).$$

A Bayesian Nash equilibrium is a profile $\{X^*, y^*_\alpha, y^*_\beta\}$ that satisfies the usual conditions. The receiver's response must be optimal given his beliefs and his beliefs must be Bayes consistent with the signaling technology and the sender's equilibrium strategy. Thus, y^*_s maximizes $v_s(y; X^*)$ for $s = \alpha, \beta$. The sender's allocation must be an optimal response to (y^*_α, y^*_β), which implies that every strategy x^* in the support of $X^*(q)$ maximizes $\pi(x; y^*_\alpha, y^*_\beta)$.

In the above model, the sender's type is exogenous. In an earlier version of this paper, we endogenized the sender's type by allowing him to first choose the weight q, restricting it to be 0 or 1, and then to choose his allocation of resources. The endogenous type model is well suited to military examples, as we discuss in a later section. The analysis is quite similar, primarily because, in equilibrium, the sender randomizes between the two alternatives, which in turn leads to a subform or subgame equivalent to the model with exogenous types. Consequently, the results are also similar.

3. The Equilibrium

We begin our analysis by distinguishing between two kinds of signaling technologies. Figure 1 illustrates a signaling technology that is not very responsive to the differential in investment. The thin dashed line is the diagonal and the bold line is the graph of $p(x)$. For most allocations, both signals are nearly equally likely. Hence, as long as the sender allocates some resources to both alternatives, the signal is not very informative about x. Figure 2 illustrates a signaling technology that is very sensitive to any differences in the amounts invested. The receiver is almost certain to get signal α if x exceeds $\frac{1}{2}$ and signal β if x is less than $\frac{1}{2}$. Thus, for most allocations, the receiver can identify the alternative in which the sender invests relatively more resources.

The above description suggests the following criterion for comparing different signaling technologies. Let p_1 and p_2 denote two signaling technologies where the probability of generating signal α under p_1 is higher than that of p_2 for any x that exceeds $\frac{1}{2}$. By symmetry, the probability of generating signal β when x is less than $\frac{1}{2}$ is also higher under p_1 than under p_2. In this case, we will say that p_1 is a more revealing signaling technology than p_2.

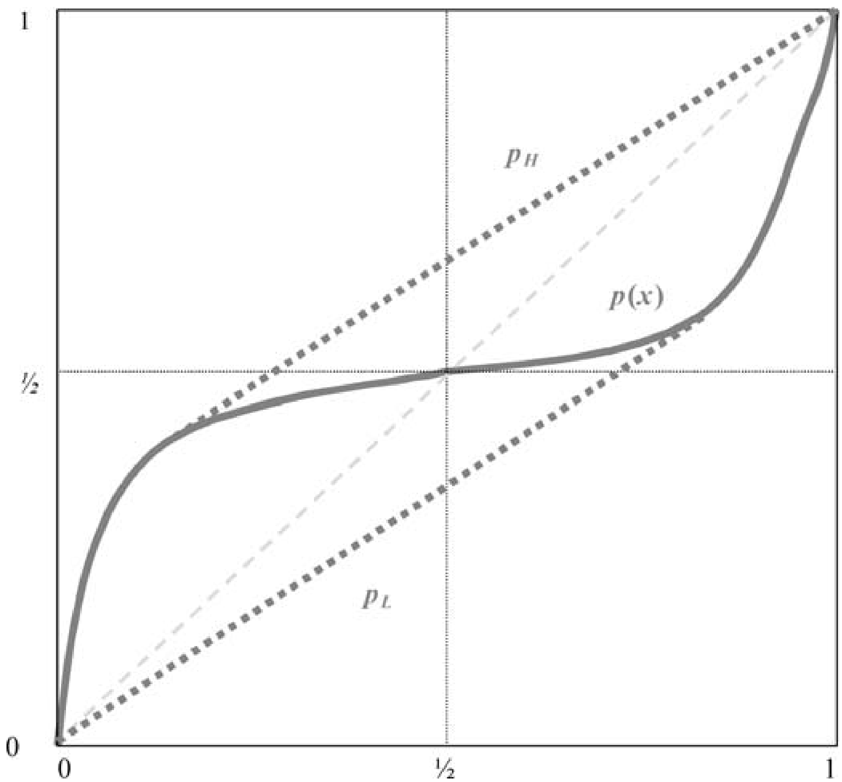

FIGURE 1. THE NOISY TECHNOLOGY CASE

Definition 1: *p_1 is more revealing than p_2 if $\forall x \in [0, 1], |p_1(x) - \frac{1}{2}| > |p_2(x) - \frac{1}{2}|$.*

In what follows, we will distinguish between two types of signaling technologies: noisy and revealing. We define them relative to the identity signaling technology, $p(x) = x$.

Definition 2: *Suppose $p_2(x) = x$. Then p_1 is a noisy signaling technology if it is less revealing than p_2 and it is a revealing signaling technology if it is more revealing than p_2.*

In a noisy technology, $p(x)$ exceeds x for any x less than $\frac{1}{2}$ and, in a revealing technology, $p(x)$ is less than x for any x less than $\frac{1}{2}$. Symmetry implies that the opposite inequalities hold for any x between $\frac{1}{2}$ and 1. A noisy signaling technology is illustrated in Figure 1 and a revealing technology is illustrated in Figure 2.

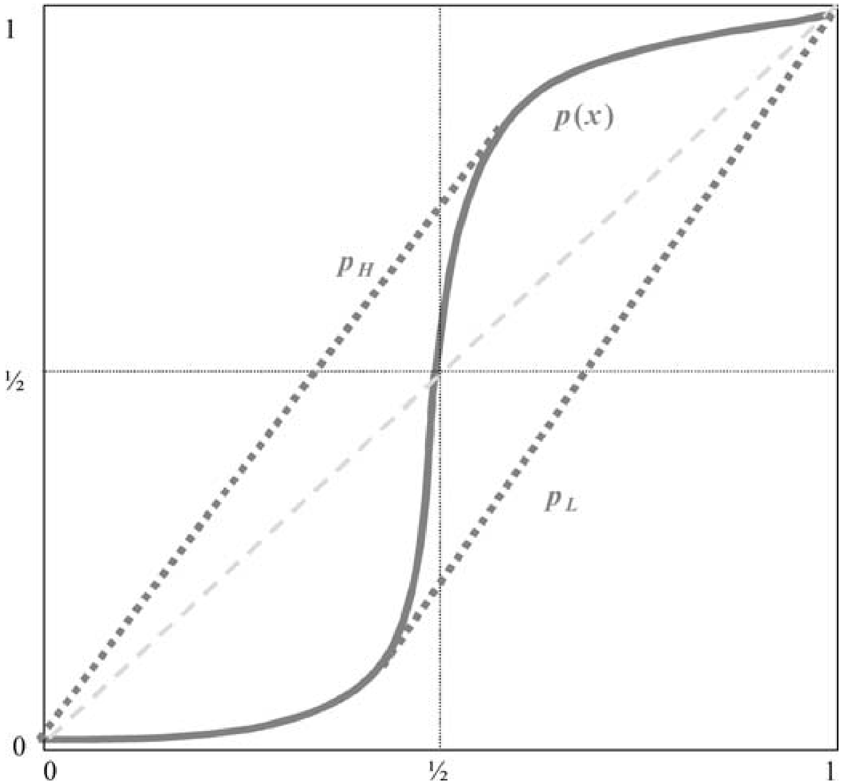

FIGURE 2. THE REVEALING TECHNOLOGY CASE

By randomizing the allocation of resources, the sender can achieve signaling probabilities that differ from p. The set of probabilities that can be generated using mixed strategies is the convex hull of the graph of p. Define

$$p_L(z) = \min_Z E(p(Z)) \quad s.t. \quad E(Z) = z$$

and

$$p_H(z) = \max_Z E(p(Z)) \quad s.t. \quad E(Z) = z$$

as, respectively, the lower and upper bounds of the convex hull of the graph of p. These bounds are depicted as the thick dashed lines in Figures 1 and 2. Note that, by Assumptions 1(ii) and (iii), the convex hull includes the diagonal. The following lemma, which follows from Assumption 1(ii), establishes that the boundaries are symmetric.

Lemma 1: $p_L(z) + p_H(1 - z) = 1.$

To solve for the equilibrium of the signaling game, we proceed by backward induction. Given the sender's strategy, the receiver's best reply is easily characterized. Differentiating v_s with respect to y_s yields first-order conditions,

$$\frac{U'(y_s)}{U'(1-y_s)} = \frac{1-\hat{q}_s}{\hat{q}_s}, \quad s = \alpha, \beta. \tag{3}$$

Manipulating these first-order conditions yields the following result.

Lemma 2: *$y_\alpha > y_\beta$ if and only if $\hat{q}_\alpha > \hat{q}_\beta$.*

Consider next the sender's best reply to an arbitrary receiver strategy. If $y_\alpha < y_\beta$, then the sender's payoff is strictly decreasing in x for $q < \frac{1}{2}$, and strictly increasing in x for $q > \frac{1}{2}$. Thus, the sender's best reply in this case is to allocate all of his resources to the alternative that has the higher weight in his utility function. But this implies that $\hat{q}_\alpha > \hat{q}_\beta$, which contradicts the previous lemma. Hence, a necessary condition for existence of a pure strategy equilibrium is that $y_\alpha > y_\beta$.

Given any such strategy, and assuming an interior solution, the sender's best reply must satisfy

$$p'(x) = \frac{1}{(y_\alpha - y_\beta)}, \quad (2q-1)p''(x) > 0.$$

There are (generically) two possible solutions to the first-order condition: one lies in the half interval $[0, \frac{1}{2}]$ and the other lies in the half interval $[\frac{1}{2}, 1]$. Let $\hat{x}(y)$ denote the solution that exceeds $\frac{1}{2}$. Symmetry implies that the other solution is $1 - \hat{x}(y)$. For q greater than $\frac{1}{2}$, the second-order condition holds at $\hat{x}(y)$ if p is noisy and at $1 - \hat{x}(y)$ if p is revealing; the converse is true for q less than $\frac{1}{2}$.

In the case of noisy technologies, the best reply for a sender is to allocate more than half of his resources to alternative a if q exceeds $\frac{1}{2}$ and to alternative b if q is less than $\frac{1}{2}$. Because high types are more likely to generate signal α, and the prior expectation is $\frac{1}{2}$, the receiver believes that he is more likely to be playing against high types conditional on signal α. Similarly, because low types are more likely to generate signal β, the receiver's posterior beliefs conditional on signal β is that he is more likely to be playing low types. Given these beliefs, it follows from the receiver's first-order conditions that he allocates more than half of his resources to alternative a if he obtains signal α and to alternative b if he obtains signal β, satisfying the necessary condition for existence.

Proposition 1: *There is an equilibrium in pure strategies if and only if the signaling technology is noisy. The equilibrium is unique and consists of the sender allocating more resources to his preferred alternative, and the receiver allocating more resources to alternative a or b depending upon whether he obtains signal α or β.*

To establish necessity, suppose the signaling technology is revealing. Assumption 2 implies that the best reply of the sender to a strategy in which y_α exceeds y_β is to invest less than half of his resources in his preferred alternative. The sender does so in hopes that his action will induce the receiver to invest most of his resources in the other alternative. But the receiver, knowing the sender's strategy, understands that the signal he has received is more likely to have been generated by the feint, indicating that the preferred alternative is the opposite of the one implied by the signal. Hence, he is not fooled, and allocates more of his resources to alterative a conditional on obtaining signal β, and to alternative b conditional on obtaining signal α. But, given this response, the sender's optimality conditions are not satisfied.

The failure of a pure strategy equilibrium to exist in the case of revealing technologies means that we need to search for an equilibrium in mixed strategies. We will look for an equilibrium in which the sender randomizes between investing a small and a large amount in his preferred alternative. This randomization will cause the receiver to be uncertain as to whether the sender is a high or low type, and hence about which of the two alternatives is his preferred one. The main difficulty in constructing a mixed strategy equilibrium with a continuum of types is that each sender type has to be indifferent between investing a small or a large amount of resources in his preferred alternative.

The next lemma establishes that we can restrict our search for mixed strategies to points on the lower and upper bound of the convex hull of the graph of p

Lemma 3: *If $q > \frac{1}{2}$ then $(EX^*(q), Ep(X^*(q))) \in p_L$; if $q < \frac{1}{2}$ then $(EX^*(q), Ep(X^*(q))) \in p_H$.*

A formal proof is given in the Appendix. The idea is simple. For q greater than $\frac{1}{2}$, consider any allocation strategy X that generates a point that is not on p_L. By definition of p_L, there is an alternative strategy $\hat{X}$ on p_L such that $Ep(\hat{X}) = Ep(X)$ and $E\hat{X} > EX$. It then follows from the linearity of the sender's profit function that $\hat{X}$ is a more profitable strategy than X. Hence there exists at least one pure strategy in the support of $\hat{X}$ that yields higher profits than X. The argument for q less than $\frac{1}{2}$ is similar.

Given the above lemma, the support of $X^*(q)$ for q larger than $\frac{1}{2}$ is $\{x_1, 1\}$ where x_1 solves $p_L(x) = p(x)$.[6] A necessary condition for x_1 to be a best reply against a strategy (y_α, y_β) is that it satisfies

$$p'(x_1)(y_\alpha - y_\beta) = 1.$$

Note that the second-order conditions for a (local) optimum are satisfied at x_1 because it is less than $\frac{1}{2}$. Furthermore, because $p'(x_1) = p_L'(z)$ for $z > x_1$, the sender's payoff from 1 is the same as his payoff from x_1 (assuming the latter is a best reply). By symmetry of p, the support of $X^*(q)$ for q less then $\frac{1}{2}$ is $\{0, 1 - x_1\}$ where $1 - x_1$ solves $p_H(x) = p(x)$ for x greater than $\frac{1}{2}$.

By definition of a mixed strategy, the high types must be indifferent between choosing x_1 and 1. Similarly, low types must be indifferent between choosing 0 and $1 - x_1$. These two indifference relationships yield two equations, which uniquely determine the receiver's strategy. Solving the equations for y_α and y_β yields

$$y_\alpha^* = \frac{1}{2} + \frac{1 - x_1}{2(1 - p(x_1))}, \quad y_\beta^* = 1 - y_\alpha^*. \tag{4}$$

The final step of the construction is for the sender to randomize in such a way that the receiver's beliefs induces him to choose y_α^* conditional on obtaining signal α and y_β^* conditional on obtaining signal β. Let θ denote the probability that high types choose x_1 and assume that this is also the probability that low types choose $1 - x_1$. The first-order condition for y_α^* determines the value of θ^*. It solves

$$\frac{U'\left(y_\alpha^*\right)}{U'\left(1 - y_\alpha^*\right)} = \frac{1 - \hat{q}_\alpha(\theta^*)}{\hat{q}_\alpha(\theta^*)}.$$

where

$$\hat{q}_\alpha(\theta^*) = \left(1 - \theta^* + \theta^* p(x_1)\right) E\left[q \mid q > \frac{1}{2}\right] + \left(\theta^*(1 - p(x_1)\right) E\left[q \mid q < \frac{1}{2}\right]. \tag{5}$$

Given Assumption 2, the solution for θ is unique and lies between 0 and 1. Furthermore, $\hat{q}_\beta = 1 - \hat{q}_\alpha$, which implies that the sender's other first-order condition is also satisfied.

We have established the following proposition.

Proposition 2: *Suppose the signaling technology is revealing. There is a unique equilibrium in which the sender randomizes between investing all of*

6. We assume that x_1 is unique.

his resources and investing an amount that is less than half of his resources in his preferred alternative, and the receiver responds to the signal generated by investing more than half of his resources in alternative a given signal α and alternative b given signal β.

We interpret the investment by the sender in the less preferred alternative as a feint. In the absence of signals, the equilibrium would consist of the sender investing all of his resources into the preferred alternative, and the risk-averse receiver dividing his resources equally between the two alternatives. With signals, the sender cannot pursue his desired objective without revealing his plans. As a result, he tries to disguise his true objective by investing some of his resources into the less preferred alternative. The sole purpose of this investment is to mislead the receiver into allocating his resources away from the more preferred alternative.

Propostions 1 and 2 provide a striking characterization of equilibrium behavior. When signals are not very informative, the sender always feints, the feint is relatively small, and is successful only some of the time. The signal is informative because it changes the receiver's beliefs and he responds by allocating more resources to the alternative signaled. When signals are informative, the sender only feints some of the time, but if he feints, it is always a large feint, using more than half of his resources. One can show that the expected size of the investment exceeds $\frac{1}{2}$, so it is rational for the receiver to respond to signal α by investing more than half of his resources in alternative a and to signal β by investing more than half of his resources in alternative b. As a result, whether the sender invests a large or a small amount in his preferred alternative, it tends to be matched by the receiver.

An interesting special case arises when the receiver is risk neutral. In this case, the receiver always fully invests in alterative a or b depending upon whether $\hat{q}_\alpha$ is greater than or less than $\frac{1}{2}$. Hence, to support the particular division needed to make the sender indifferent, the receiver has to be indifferent over all allocations. He is indifferent if and only if the randomization on the part of the sender has the effect of making the signals uninformative (i.e., $\hat{q}_\alpha = \hat{q}_\beta = \frac{1}{2}$). Regardless of which signal is received, the receiver's posterior beliefs are the same as his prior beliefs. Despite this indifference, the receiver has to respond differently to the signal received.

4. Comparative Statics

The key primitives of our model are the signaling technology and the preferences of the receiver. In this section we examine how equilibrium behavior and payoffs vary with these factors. Proofs of the propositions can be found in the Appendix.

4.1 The Value of Risk Aversion

Is a risk-averse receiver a weaker opponent than a risk-neutral receiver? The answer differs depending upon whether the signaling technology is revealing or noisy. In the revealing case, the sender's payoff is

$$\pi^*(q) = (2q - 1)\left[\frac{1}{2} - \frac{1 - x_1}{2(1 - p(x_1))}\right].$$

Thus, his payoff depends upon p but does not depend on the receiver's preferences. The reason is that the receiver's strategy is determined by the condition that the sender must be indifferent between full investment and the partial investment of x_1.

The receiver's strategy does depend upon his preferences when the technology is noisy. Consider two receivers, 1 and 2, with utility functions U_1 and U_2 and assume that receiver 1 is more risk averse than receiver 2. Applying Theorem 1 of Pratt (1964), for any $y_\alpha > \frac{1}{2}$,

$$\frac{U_1'(y_\alpha)}{U_1'(1 - y_\alpha)} < \frac{U_2'(y_\alpha)}{U_2'(1 - y_\alpha)}.$$

Thus, given any pure strategy by the sender, the more risk-averse receiver always divides his resources more evenly between the two alternatives, which yields a higher payoff to the sender. In equilibrium, the sender takes advantage of the receiver's risk aversion by investing more in his preferred alternative and earning a higher payoff.

Proposition 3: *If the technology is noisy, the sender prefers to play against a more risk-averse receiver. If the technology is revealing, the sender's payoff does not depend upon the receiver's degree of risk aversion.*

4.2 The Value of Stealth

In our model, the sender takes the signaling technology as exogenous. But, he may have some scope for choosing among signaling technologies. For example, in modern warfare, the signaling technology depends upon the method of attack. Adversaries are likely to have a harder time detecting the target of an air attack than an armed forces attack, and an air attack that uses conventional bombers is more easily detected than an air attack that uses stealth bombers. The question that arises is, in what circumstances, if any, does the sender prefer a more revealing technology?

One approach to this question is to study the limit cases. As the signaling technology gets very noisy, $p(x)$ converges to

$$\underline{p}(x) = \begin{cases} 0 & \text{if } x = 0 \\ \dfrac{1}{2} & \text{if } 0 < x < 1 \\ 1 & \text{if } x = 1. \end{cases}$$

In the limit, a negligible feint makes the likelihood of generating signals α and β equally likely. The signal is essentially uninformative. The receiver splits his resources equally between the two alternatives and the sender fully invests in his preferred alternative. Hence, in the limit, the expected payoff to the sender converges to

$$\pi^*(q) = \begin{cases} q - \dfrac{1}{2} & \text{if } q > \dfrac{1}{2} \\ \dfrac{1}{2} - q & \text{if } q < \dfrac{1}{2}. \end{cases}$$

By contrast, as the technology becomes more revealing, $p(x)$ converges to

$$\bar{p}(x) = \begin{cases} 0 & \text{if } \dfrac{1}{2} > x \geq 0 \\ \dfrac{1}{2} & \text{if } x = \dfrac{1}{2} \\ 1 & \text{if } 1 \geq x > \dfrac{1}{2}. \end{cases}$$

In the limit, the amount that the sender needs to invest to generate the false signal for certain approaches (but is never equal to) $\frac{1}{2}$. High types randomize between investing fully and an amount slightly less than $\frac{1}{2}$ in alternative a; low types randomize between investing fully and an amount that is slightly less than $\frac{1}{2}$ in alternative b. As a result, if the receiver observes signal α, he does not know whether he is facing a high type who is investing fully in alternative a or a low type who is feinting. Similarly, if he observes signal β, he does not know whether he is facing a high type who is feinting or a low type who is investing fully in alternative b. Recall that, to ensure that the sender is willing to randomize in this way, the receiver's strategy must satisfy (4). This yields the limit values $y^*_\alpha = 3/4$ and $y^*_\beta = 1/4$.[7] Substituting these values into the sender's profit function gives equilibrium payoffs

$$\pi^*(q) = \begin{cases} \frac{1}{2}\left(q - \frac{1}{2}\right) & \text{if } q > \frac{1}{2} \\ \frac{1}{2}\left(\frac{1}{2} - q\right) & \text{if } q < \frac{1}{2}. \end{cases}$$

Thus, the sender is better off with a very noisy technology than a very revealing technology. The intuition is that the sender has to expend at least half of his resources generating the false signal when the signaling technology is very revealing whereas it is free when the signaling technology is very noisy. Note that the sender's payoff is continuous in q and equal to zero for $q = \frac{1}{2}$.

The dividing line between the two classes of signaling technologies is the identity technology, $p(x) = x$. In this case, there is a continuum of equilibrium allocations, all of which yield an expected payoff of zero to the sender.

A second approach to the question of whether the sender prefers a noisier technology is to fix the signaling technology and consider an increase in the noise. The following proposition establishes that stealth does not pay if the technology is revealing.

Proposition 4: *Suppose p_1 and p_2 are revealing technologies. Then the sender prefers p_1 to p_2 if p_1 is more revealing than p_2.*

The intuition is that a more revealing technology makes it is easier to generate the false signal so the sender can reduce the amount invested in the feint. The proof is easily seen using Figure 2. Let p_{Li} denote the lower bound of the convex hull of the graph of p_i and let x_{1i} denote the lower endpoint of the linear segments on p_{Li}, $i = 1, 2$. The slope of the chord connecting $(x_{1i}, p(x_{1i}))$ to $(1, 1)$ is

$$\frac{1 - p(x_{1i})}{1 - x_{1i}}.$$

Because p_1 is more informative than p_2, the slope of the chord on p_{1L} is larger than the slope of the linear segment on p_{2L}. From (4), this implies that the amount invested in alternative a conditional on signal α, denoted $y^*_{\alpha i}$, is smaller when the technology is p_1 than when the technology is p_2.

7. The value of θ^* must satisfy (5). For example, if q is uniformly distributed, this equation becomes

$$\frac{U'(3/4)}{U'(1/2)} = \frac{1 + 2\theta}{3 - 2\theta}.$$

Because the equilibrium payoff to the sender for revealing technology p_i is $1 - y^*_{\alpha i}$, the hypothesis is established. A more formal proof is given in the Appendix.

Although the sender's payoff for a noisy technology is clearly larger than his payoff for the identity technology, we have not been able to show that the value of stealth is always positive without restricting the preferences of the receiver. If the receiver is risk neutral, then the receiver's equilibrium strategy is always to fully invest in the alternative signaled (i.e., $y_\alpha = 1 - y_\beta = 1$). As a result, the sender's equilibrium payoff is independent of the receiver's preferences. This allows one to compare the sender's profits for different noisy technologies.

Proposition 5: *Suppose p_1 and p_2 are noisy technologies and the receiver is risk neutral. Then the sender prefers p_1 to p_2 if p_1 is less revealing than p_2.*

5. Applications

We have presented the model in a generic form because many of the specific details on payoffs that vary with applications are not essential to the analysis. Feinting is a useful strategy in games where the second mover wants to direct its resources at the first mover's choice of location in product/geographical space. The first mover has an incentive to deceive the second mover about its location choice so that the second mover's response will be more muted (or delayed) than it would have been otherwise. In this section, we discuss two classes of applications.

5.1 Military Games

The most obvious and direct application of the model is to military conflicts. Here the sender is the attacker, the receiver is the defender, and the alternatives are locations. The attacker's payoff is a weighted average of the outcomes of the battles at each location, as measured by the differential in the size of the forces. The attacker's type, q, is his belief that the battle at location a is likely to be decisive to winning the war. The attacker would like to commit all of his forces to the more decisive battle but, in preparing and carrying out his plan of attack, the attacker cannot avoid generating information about the relative size of the forces being sent to each location. A larger force requires more support, more extensive supply lines, and cannot move as quietly or as quickly as a small force. If the attacker is certain to engage in an all-out attack, it will be met by an all-out defense. The attacker must pursue a plan of attack so that the defender cannot be certain as to which of the two battles is the decisive one. As a result, he will split his forces between the two battles, sending relatively more forces to the battle that he thinks is more likely to be decisive based on the information generated by the attacker's division of forces. The attacker, knowing how the defender will respond, has an incentive to manipulate the defender's beliefs by allocating some of his forces to the less important battle, and hopefully send a signal that will draw most of the defender's forces away from the more important battle.

An interesting special case studied in a previous version of this paper assumes that q is either 0 or 1 (i.e., only one battle matters). In this model, the attacker allocates forces to each location so as to make the defender uncertain as to which location is going to be attacked. The diversionary force threatens but does not actually fight so the cost of diverting forces away from the battle is an opportunity cost. The payoffs to the attacker and the defender are measured by the difference in the size of the forces at the location of the attack. The attacker's payoffs are also allowed to vary across locations to reflect differences in the locations. For example, Calais was an easier target to attack than Normandy. None of these modeling details had any qualitative effect on equilibrium behavior. The attacker randomizes over which location to attack and then, depending upon the signaling technology, divides his forces using the strategies described in Propositions 1 and 2.

Our results distinguish between two kinds of situations. When the signaling technology is noisy, the defender finds it difficult to identify which location is being attacked by the larger force even though the diversionary force is small. Hence, the attacker can afford to attack his preferred location with most of his forces. The signal is not very informative so the defender's strategy is basically to divide his forces equally between the two locations and marginally increase the forces allocated to the location signaled. The Allied invasion of Normandy serves as an example of this kind of situation. When the defender can easily detect whether the attack force targeted at a location is small or large, the attacker sometimes attacks his preferred location with all of his forces and sometimes he uses the main force as a diversion and attacks with a small force. The uncertainty created by this strategy makes the signal less informative. The defender may know which location is targeted by the larger force, but he cannot be certain as to which battle is going to be the decisive one. Despite this uncertainty, he responds by allocating substantially more forces to the location threatened by the perceived larger force. The response means that the expected payoff to the attacker from a sneak attack is similar to his payoff from an all-out attack.[8]

5.2 Entry Games

In the standard signaling model of entry, the sender is an incumbent firm and the receiver is a potential entrant who has to decide whether or not to enter the market and compete against the incumbent firm. It is willing to enter against high-cost incumbents but not against low-cost incumbents. Thus, the high-cost incumbent has an incentive to imitate the low-cost incumbent's behavior and try to mislead the entrant into staying out of the market. However, in many cases, the roles of the sender and receiver are reversed: it is the entrant who wants to mislead the incumbent firm. For example, Capital One, an innovative credit card issuer, wanted to offer new products targeted at specific customer segments. In order to do so, they collected detailed information on its customers and then tested thousands of offerings to identify the new markets. The success of this strategy depended upon the company's ability to keep the industry giants such as Citibank from matching its product offerings. Wary of drawing Citibank's attention, Capital One pursued a more costly marketing strategy that relied upon telephone solicitation and direct mail rather than high-profile advertising campaigns. According to a former executive of Capital One, competitors may come across its mailings but "Citibank would never know what customers we are targeting, unless we told them. As a result, we have largely been able to stay under the radar screen." Palm pursued a somewhat similar strategy in marketing the Pilot. The company introduced its product in 1996, carefully positioning it as a complement to the PC, not a substitute, and playing down its potential in order not to draw unwelcome attention from competitors like Microsoft. The strategy apparently worked; Microsoft underestimated the profitability of the handheld market and delayed its attack on Palm until it was too late.[9]

The Capital One and Palm stories are examples of innovative firms undermarketing new products to disguise their entry into an incumbent's markets. But they are not examples of feints because, if the firms had not been innovative, they would have had no reason to pretend that they were innovative. But, if the decision by the entrant is

8. A classic example in sports of such a strategy is the running play in American football known as the "naked quarterback bootleg." While the entire offensive line and running backs move in one direction, the quarterback fakes a handoff to a running back, tucks the ball behind his leg, and jogs off in the opposite direction. He has no blockers to help him gain any yards but, if the feint is successful, none is required.

9. The two cases are studied in detail in Yoffie and Kwak (2001). We thank Lucy White for drawing our attention to this book.

not whether to enter the market but which market to enter, then the entry game becomes a feinting game. Entry involves a substantial investment in marketing, and the payoffs to the potential entrant are a weighted average of the difference between its marketing expenditures and those of the incumbent in each market. The market weight represents the entrant's type, and is a measure of its comparative advantage in entering a particular market. If the entrant ignores the response of the incumbent firm, then the entrant's best strategy is to use all of its scarce resources to enter the market in which it has a comparative advantage. However, it knows that pursuing such a strategy would generate a quick and overwhelming response from the incumbent. A better strategy would be to exploit the incumbent firm's uncertainty about which market is the real target by investing in both markets. This will force the incumbent to split its resources between the two markets, and weaken its defense of the target market. Hopefully, the incumbent will allocate relatively more resources to defending the other market. Whether the entrant has to expend a small or a large amount on diversionary marketing will depend upon the observability of its expenditures. The important issue is that, even if the incumbent firm can easily identify the market in which the entrant is launching a bigger campaign, the equilibrium does not involve head-to-head competition. The bigger marketing campaign is sometimes a diversion, and knowledge that this might be the case forces the incumbent firm to expend resources in both markets.

A variant of the above model that we studied in a previous version of this paper involves new product entry over time. Suppose an innovative firm has perfected a new product and its rival, knowing that the entrant has been working on a new product, is positioned to produce a close substitute as soon as the entrant's new product hits the market. Instead of engaging in a head-to-head competition with its rival, the innovator could pursue a more subtle strategy. It could market an inferior version of its new product first, in the hope that its rival will expend all of its scarce resources on developing and marketing a close substitute for the inferior product. Once the rival does so, then the entrant comes out with the original product, calling it the new, improved version, and captures the market. Of course, in equilibrium, the imitator firm anticipates the innovative firm's strategy, and does not expend all of its resources reacting to the first product marketed by the entrant. It keeps some resources in reserve on the chance that the first product was a decoy and it has to react to the new, improved version. Nevertheless, by pursuing such a strategy, the innovator forces its rival to delay or reduce its response to the initial product offering, which increases profits even when it is not a decoy product. The fake surveys conducted by Tidewater to decoy Standard Oil into blocking the wrong pipeline route are an example of this strategy.

Many companies are accused of product announcements whose purpose is to mislead. Such products are known as vaporware. In most cases, the purpose is to mislead consumers. For example, Microsoft has been accused of announcing products to deter consumers from switching to superior products (see McAfee (2002), pp. 348–349). These examples are not well described by the present theory because the interests of the consumers align with the firms' interests in the state of the world where the product exists. However, some instances may involve misdirection of the research of rival companies. Microsoft has recently dropped the file system known as WinFS, now almost 10 years overdue, from the new operating system code-named Longhorn.[10] Because there are no significant competing products to the Windows XP filesystem NTFS, and the delay so long as to suggest intent, it is plausible to think the WinFS vaporware is aimed at competitors, especially Linux, rather than consumers. Even without the new file system, Longhorn itself has been pushed back two years, so that some are suggesting it be called "longwait." At trade shows, firms have allegedly showcased modified versions of their new products lest rivals copycat their designs before the market opens. Langinier (2000) cites anecdotal evidence of firms patenting "deadends" in order to send competitors in the wrong research direction and provides an interesting example from the pharmaceutical industry in which the patenting firm could have tried to pursue a feinting strategy.

6. Conclusion

We present a signaling theory of feints based on rational players, which emphasizes costly actions taken to mislead. When misdirection is costly, the sender will not attempt to mislead unless there is a payoff to fooling the receiver; thus feints must be at least partially successful to occur. Our theory shows that there are two types of feints, and the type that arises depends on the accuracy or noisiness of the signaling technology. In a noisy world, the receiver does not respond strongly to the signal (because the signal is not very informative), so typically large efforts at misdirection are not worthwhile, and investment in a feint is modest. In contrast, if the world is not noisy, which we call revealing, signals are meaningful and the receiver would respond very strongly to the

10. See the forthcoming article by Leander Kahney, "Vaporware Phantom Haunts Us All," in *Wired Magazine*; http://www.wired.com/news/culture/0,1284,66195,00.html, January 7, 2005.

signal were feints not possible. Such strong responses by the receiver invite large attempts to fool the receiver, and equilibrium dictates a mixed strategy in attempts to fool the receiver with a large diversion of resources when an attempt occurs. The "sneak attack" is a part of a natural equilibrium that arises with positive probability when the receiver can react strongly to the signal. The theory also presents striking comparative static results. More noise benefits the sender in a noisy world, but not in a revealing world. In a revealing world, making the technology less noisy makes the impact of the feint greater and less costly.

We have tried to keep the model as simple as possible in order to uncover the essential logic of feints. The model can be extended in a number of directions. For example, the signal does not have to be binary. In fact, we have been able to show that, under certain conditions, the model with a continuum of signals is isomorphic to a model with binary signals. The essential properties are those of the payoff functions that imply that the sender's optimal perceived type is a decreasing function of its type. As a result, the sender's best reply to a nondecreasing strategy by the receiver may not be nondecreasing, and a pure strategy equilibrium may not exist. The primary assumption that accounts for the simplicity of the analysis is the assumption that the sender's payoff is linear in his action. We are currently exploring the case in which the sender's payoff is strictly concave.

Appendix

This appendix contains the proofs that are not in the text.

Proof of Lemma 1. By definition

$$\begin{aligned} p_L(z) &= \min_Z E(p(Z)) \quad s.t. \quad E(Z) = z \\ &= 1 - \max_Z (1 - E(p(Z))) \quad s.t. \quad E(1-Z) = 1 - z \\ &= 1 - \max_Z E(p(1-Z)) \quad s.t. \quad E(1-Z) = 1 - z \\ &= 1 - p_H(1-z). \end{aligned}$$

The third equality follows from symmetry of p. □

Proof of Proposition 1. Substituting the sender's best reply into (1) and (2), it follows from the symmetry of p and f that $\hat{q}_\alpha = 1 - \hat{q}_\beta$. If the technology is noisy, then the receiver's expectation of q given signal α is given by

$$\hat{q}_\alpha = \frac{p(1-\hat{x})\int_0^{1/2} qf(q)\,dq + p(\hat{x})\int_{1/2}^{1} qf(q)\,dq}{p(1-\hat{x})\frac{1}{2} + p(\hat{x})\frac{1}{2}}.$$

Applying condition (ii) of Assumption 1, we need to show that

$$(1-p(\hat{x}))\int_0^{1/2} qf(q)\,dq + p(\hat{x}))\int_{1/2}^{1} qf(q)\,dq > \frac{1}{4}.$$

Exploiting the symmetry of f about $\frac{1}{2}$, the above inequality is equivalent to showing that

$$\left[p(\hat{x}) - \frac{1}{2}\right]\left[\int_{1/2}^{1} qf(q)\,dq - \int_0^{1/2} qf(q)\,dq\right] > 0,$$

which holds since $\hat{x}$ exceeds $\frac{1}{2}$. On the other hand, if the technology is revealing, then

$$\hat{q}_\alpha = \frac{p(\hat{x})\int_0^{1/2} qf(q)\,dq + p(1-\hat{x}))\int_{1/2}^{1} qf(q)\,dq}{p(1-\hat{x})\frac{1}{2} + p(\hat{x})\frac{1}{2}}.$$

In this case, an argument similar to the one given above leads to a contradiction. □

Proof of Lemma 3. Consider any strategy X_a that generates a point that is not on p_L. By definition of p_L, there is an alternative stratgey $\hat{X}_a$ on p_L such that $Ep(\hat{X}_a) = Ep(X_a)$ and $E\hat{X}_a > EX_a$.

$$\begin{aligned} E\pi_a(\hat{X}_a, y_\alpha, y_\beta) &= E\hat{X}_a - [Ep(\hat{X}_a)y_\alpha + (1 - Ep(\hat{X}_a))y_\beta] \\ &> EX_a - [Ep(X_a)y_\alpha - (1 - Ep(X_a))y_\beta] \\ &= E\pi_a(X_a, y_\alpha, y_\beta). \end{aligned}$$

But this implies that there is at least one pure strategy in the support of $\hat{X}_a$ which yields higher profits than X_a^*, which contradicts the hypothesis that X_a^* is a best reply to (y_α, y_β). The argument for (ii) is similar.

Proof of Proposition 4. When technology p_i is revealing, the equilibrium payoff to the sender is $1 - y_{\alpha i}^*$. Therefore, we need to show that the equilibrium response by the receiver is smaller against a more revealing technology. Let p_{Li} denote the lower bound of the convex hull of the graph of p_i and let x_{1i} denote the solution to $p_{Li}(x) = p_i(x)$. Then, for any $z \in (1/2, 1]$,

$$\begin{aligned} p_{L2}(z) &= \min_{x,\theta}\{\theta p_2(x) + 1 - \theta \mid \theta x + 1 - \theta = z\} \\ &> \min_{x,\theta}\{\theta p_1(x) + 1 - \theta \mid \theta x + 1 - \theta = z\} \\ &= p_{L1}(z), \end{aligned}$$

where the inequality follows from Definition 1. This implies that

$$\frac{1 - p_{L1}(z)}{1 - z} > \frac{1 - p_{L2}(z)}{1 - z}.$$

The definition of x_{1i} implies,

$$p'_{Li}(x_{1i}) = \frac{1 - p_{Li}(z)}{1 - z}.$$

Hence, $p'_{L1}(x_{11}) > p'_{L2}(x_{12})$. It then follows from equation (4) that $y^*_{\alpha 1} < y^*_{\alpha 2}$. Q.E.D. □

Proof of Proposition 5. Without loss of generality suppose the sender's type is $q \geq 1/2$ and let $\hat{x}_1$ and $\hat{x}_2$ denote the sender's best replies to a strategy in which the receiver allocates all of resources to the alternative signaled when the signaling technology is p_1 and p_2 respectively. When both technologies are noisy, $\hat{x}_1$ and $\hat{x}_2$ exceed $\frac{1}{2}$ and are equilibrium allocations. Therefore, letting π_i denote the sender's profits from technology i,

$$\begin{aligned} \pi_1(\hat{x}_1) &= (2q - 1)[\hat{x}_1 - p_1(\hat{x}_1)] \\ &\geq (2q - 1)[\hat{x}_2 - p_1(\hat{x}_2)] \\ &> (2q - 1)[\hat{x}_2 - p_2(\hat{x}_2)] = \pi_2(\hat{x}_1). \end{aligned}$$

The first inequality follows from the definition of $\hat{x}_1$ as a best response and the second from the definition of less revealing. Q.E.D. □

References

Banks, J. S., 1987, "A Model of Electoral Competition with Incomplete Information," *Journal of Economic Theory*, 55, 647–661.

Bernheim, B. D., 1994, "A Theory of Conformity," *Journal of Political Economy*, 102, 841–927.

——, and S. Severinov, 2003, "Bequests as Signals: An Explanation for the Equal Division Puzzle," *Journal of Political Economy*, 111(4), 733–764.

Crawford, V., 2003, "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions," *American Economic Review*, 93, 133–149.

——, and J. Sobel, 1982, "Strategic Information Transmission," *Econometrica*, 50, 1431–1451.

Farrell, J., 1993, "Meaning and Credibility in Cheap-Talk Games," *Games and Economic Behavior*, 5, 514–531.

Fudenberg, D. and J. Tirole, 1986, "A 'Signal-Jamming' Theory of Predation," *Rand Journal of Economics*, 17(2), 366–376.

Hatfield, T., 1997, "The Crucial Deception," *University of Texas Discovery Magazine*, 14(2), 17–25 .

Langinier, C., 2000, "Innovation, Improvement and Strategic Patenting Decision," Mimeo.

McAfee, R. P., 2002, *Competitive Solutions: The Strategist's Toolkit*, Princeton University Press.

Matthews, S.A. and L. Mirman, 1983, "Equilibrium Limit Pricing: The Effects of Private Information and Stochastic Demand," *Econometrica*, 51(4), 981–996.

Milgrom, P. and J. Roberts, 1982, "Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis," *Econometrica*, 50(2), 443–460.

——, and ——, 1986, "Price and Advertising Signals of Product Quality," *Journal of Political Economy*, 94, 796–821.

Pratt, J.W., 1964, "Risk Aversion in the Small and in the Large," *Econometrica*, 32, 122–136.

Saloner, G., 1987, "Predation, Mergers, and Incomplete Information," *The RAND Journal of Economics*, 18(2) (Summer), 165–186.

Sobel, J., 1985, "A Theory of Credibility," *Review of Economic Studies*, LII(4), 557–573.

Von Neumann, J. and O. Morgenstern, 1944, *Theory of Games and Economic Behavior*, Princeton University Press.

Yergin, D., 1991, *The Prize*, Simon & Shuster, New York.

Yoffie, D.B. and M. Kwak, 2001, *Judo Strategy*, Havard Business School Press.