Dynamic Deception[*]

Axel Anderson — Georgetown

Lones Smith[†] — Michigan

September 15, 2009

Abstract

We model dynamic deception exercises in Ponzi schemes, war, terrorism, industrial/political espionage, and insider trading. We introduce a new class of competitive continuous time games with an informed player having an unbounded action intensity, whose actions are partially observed by a rival, or sequence of rivals. We find extremely strong square-root logarithm decreasing returns to deception.

We provide economic formulations and analyses of espionage, disinformation, and obfuscation. Notably, we establish a globally concave value of information (*espionage*). *Disinformation* is always substantial, and pushes confidence in the wrong state above 86%. Finally, *obfuscation* rises as the public grows more uncertain.

[*]This paper was begun when Axel was visiting the University of Michigan. We thank seminar participants at Texas, and especially Max Stinchcombe. Earlier less focused versions of this paper were presented at GAMES 2008, Decentralization (2008), Stoneybrook (2008), George Washington, Georgetown, Michigan, Toronto, and Western Ontario.

[†]Lones thanks the NSF for research funding.

1 Introduction

The most valuable commodity I know of is 'information'.

— Gordon Gekko ("Wall Street", 1987)

Deception is in the air: Our inboxes are swamped with spam Nigerian email schemes or other phishing emails that generate \$20 billion a year. Political subterfuge and cover-ups persist around the world — like Iran's falsified election outcome. On Wall Street, multi-billion dollar long-running Ponzi schemes are revealed. In the corporate world, several deep-rooted deceptions by companies like Enron are discovered.

A resilient feature of many such deceptions is their durability. How long do they last? How do their profits relate to their longevity? How should the deceivers and deceived adapt over time as the deception wears thin? What are the incentives of the deceived to investigate them? What other more active forms can deception assume?

This paper develops and explores a new economic theory of dynamic deception. We step back from the specific details of the many flavors of dynamic deception in favor of a unified analysis. To do this, we introduce and fully solve a model of the dynamic exploitation of misinformation. We then slightly twist this framework to define and flesh out the economics of espionage and sleuthing, disinformation, and obfuscation.

We begin with a class of potentially infinite-horizon purely competitive games in continuous time. An informed party knows critical payoff-relevant information, and can only earn profits if his actions reflect this information. But he cannot blatantly do so, since his actions are partially observed by his long-lived rival, or sequence of rivals. The informed player's action space and payoffs are unbounded at any moment in time, and not subject to any explicit decreasing returns. Resource bounds are not a salient feature of our economic landscape. For instance, before the D-Day invasion, cost was essentially no object for the Allies. This affords a laser focus on the information revelation properties of deception. By assuming two states, Gaussian noise, and symmetric and purely competitive games, we can fully solve for the unique Markov equilibrium.

We treat the informed player as a firm that balances greater dividends from exploiting his edge against the cost of greater capital losses from losing it. In equilibrium, he always exploits his advantage precisely so fast that his "dividends" from using his information edge are twice his "capital losses" from eroding it. We find that this forces a *square root law* for all substitution effects: To wit, if the informed player is four times as impatient, or if the variance of noise concealing his actions doubles, he exploits his edge only twice as intensely; further, his value halves in the first case, and doubles in the second. The one exception is when public beliefs change, where a *square root logarithm law* describes the response. If the public is twice as deceived, its squared payoff rises by a constant. This speaks to the strong decreasing returns to deception that should be a comfort to those who feel that the Ponzi schemers of the world are overly powerful. We precisely characterize how fast the informed player's intensity vanishes or explodes with the strength of his informational advantage.

The equilibrium belief path turns out to be independent of how noisily actions are observed, and this permits some strong conclusions. For instance, the passage of time initially roughly corresponds to the accumulation of money, so that "time is money". But there are diminishing returns, as the informed player burns his information rents faster, the smaller they are. Namely, the expected time until the monopoly secures 99% of his profits is a concave function of his value. We find that this duration is inverse to the interest rate — inclusive of the informed player's exogenous discovery rate. We derive deception durations, consistent with casual empiricism from Ponzi schemes and business scandals like Enron. If the interest rate and annual discovery rates sum to just 10%, then a deceiver takes fifteen years before securing 99% of his profits. But higher discovery rates, as may arise in wartime or crime, abbreviate this duration.

Having understood this baseline model, we next use it to explore the economics of three deception activities. First, we touch on espionage, or more generally, any investigative sleuthing done by the uninformed players. In our world, how does their value of information behave? For if it is non-concave, we should never see small amounts of commercial or military espionage. The value of Gaussian information in a static decision-theory context is poorly behaved, for instance — indeed, its marginal value initially vanishes, rises and then falls. This is a standard source of demand discontinuities for information.[1] But in our dynamic competitive world, when the public learns more about the state of the world, the informed party exploits his advantage less intensively. This induces a diminishing returns to informational sleuthing absent from decision theory contexts. We prove that except at extreme beliefs, the value of

[1] See Radner and Stiglitz (1984), Chade and Schlee (2002) and Keppo, Moscarini, and Smith (2008).

information is globally concave, and behaves like any other good.

For our third contribution, we introduce and flesh out two forms of more active deception: To this point in the paper, the model focuses on the optimal exploitation of exogenously noisy actions signals. This is the touchstone of *misinformation*, namely, inaccurate information that is spread unintentionally. We now turn to *disinformation* — which is inaccurate information that is instead deliberately spread. We offer an economic formulation of disinformation and from it deduce some simple insights.

Disinformation relies on seemingly informative signals that push beliefs away from the truth. For example, Operation Mincemeat in WWII sought to convince Germany that the Allied 1943 invasion would occur in Greece and Sardinia rather than Sicily, as it did. British intelligence suitably dressed and equipped a corpse with fake "secret" invasion plans, and released it by submarine to wash up on a beach in Spain (Montagu 1953). Such a ruse only succeeded as an out of equilibrium surprise — for if the Germans were aware of this possibility, they would have ignored it. So proceeding, we assume that the informed player has access to costly but misleading signals. With these units of disinformation, we derive a non-concavity: To wit, there is no small deception. On the contrary, we show that disinformation stratagems should bring the uninformed party past the point of 86% confidence in the wrong state. Disinformation is non-monotonic in the public beliefs: It is not optimal when the uninformed party is near certainty — either very deluded or quite in-the-know.

Having explored the economics of espionage and deception, we conclude with a common and orthogonal form of deception. *Obfuscation* is the costly concealment of one's action, by raising the signal noise. For instance, military invasions often rely on diversions to distract the enemy's attention. Enron always refused to release balance sheets with its earnings statements. Meanwhile, a benefit of private versus public equity is weaker reporting requirements; this facilitates obfuscation. Since the informed party may obfuscate more in one state than another, his rival learns about the state also from the volatility. Loosely, a loud explosion at one location may either signal an invasion there or an increase in diversionary activity that precedes an invasion elsewhere. Overcoming the technical hurdles of this learning, we find that unlike disinformation, the more uncertain is the uninformed party, the greater the incentive to obfuscate.

While dynamic deception has not been studied so far, economists have explored strategic deception before. The first work we know are the related clever papers by Hendricks and McAfee (2006) and Crawford (2003). They model it as a zero sum sender-receiver game in which the attacker's actions are imperfectly signaled. The former analysis is fully rational while Crawford adds a behavioral sender-type to better capture deception. In their one-shot game, they find a mixed strategy equilibrium with positive weight on "feinting" — reminiscent of what we call disinformation. By contrast, our informed player always chooses a pure strategy in equilibrium.

We employ tools from the nascent literature on continuous time dynamic games with unobserved actions and incomplete information. We are unaware of other closed form solutions or even partial characterizations in this vein.[2] Our paper differs from reputation models as there are no behavioral types committed to an action. Both informed player's types are rationally motivated. As Hendricks and McAfee point out, this kills the monotonicity that signaling games rely on. And unlike here, one may exploit a reputation and pool with the crazy type without revealing private information.

An extreme parametric case in our model captures insider trading, since it is always optimal to buy in the good state and sell in the bad. In this way, our paper is a bridge between modern finance and the game theory literature on reputation, and an invitation for the two fields to learn from each other. For the market price that appears in their setting becomes a mixed strategy for the uninformed player in this strategic setting. We show that an insider trading model by Back and Baruch (2004) has the largest adverse selection edge for which an equilibrium exists.[3] As a bi-product, we offer an alternative uniqueness proof for their setting, using methods from zero-sum games.

That the equilibrium value of information may offer surprises to economists echoes the well-known general equilibrium finding of Hirshleifer (1971), also fleshed out in Schlee (2001). They show how an information release may inhibit risk-sharing, thereby harming everyone. Their framework was static and did not involve the strategic application of market power like ours does. Finally, Bassan, Gossner, Scarsini, and Zamir (2003) specifically explore the class of normal form games with a positive value of information for all players. This class excludes the competitive games we study.

We now turn to the model and equilibrium analysis, and then our three applications.

[2]See the seminal work by Faingold and Sannikov (2007) for a good summary of this literature.

[3]The only other paper we have seen seeking strategic twists on the insider trading models is Caldentey and Stacchetti (2007).

2 The Model

We study a dynamic game of incomplete information and partially observed actions. We introduce the model in three illustrative examples — with one *long-run informed player* ("he") and either a continuous time stream of *uninformed players* ("she"), or a single long-run uninformed player. In all cases, the informed player's action intensities are unbounded, and he can only profit from his information if the intensities vary in it.

(a) STRATEGIC DECEPTION IN CRIMINAL ENDEAVORS. Much of the computer world $[0, 1]$ uses Internet Explorer (IE) in Windows. This web browser is subject to ongoing attacks by programmers. Assume that a computer hacker has discovered a massive vulnerability of IE; he can siphon financial information from computers around the world. The security hole is found with chance $\phi\ dt > 0$ in any interval $[t, t + dt]$, and so will be fixed after time t with chance $e^{-\phi t}$. The hacker discounts payoffs at the implied interest rate $r = \phi + i$, namely, the sum of ϕ and the bank interest rate i.

Computer users deploy the current version or the newer beta version of IE. A fraction $p(t)$ of the population uses the beta version at time t. The optimal attack depends on the target computer's version of IE: The hacker has written two programs, each designed to best attack one version. His computer bank uses these hacks at respective *attack intensities* α and β, i.e. the rates at which current and beta version are infected. The hacker loses $x > 0$ if his hacking efforts meet incorrect browsers, and gains expected flow payoff x so long as he hacks the correct web browser. This equality reflects opportunity costs and expected criminal penalties of hacking, and the low yield rate on successful hacks. So faced with fraction p, this part of the net flow payoff is

$$\alpha[(1 - p)x - px] + \beta[px - (1 - p)x] = (\alpha - \beta)(1 - 2p)x \tag{1}$$

The hacker's payoffs also reflect his private information about the efficacy of each program, captured by two *states* $\theta = 0, 1$. The current version's hack is better in state $\theta = 0$, while the beta version hack is better in state $\theta = 1$. We assume an additive flow benefit of $z > 0$ of the more effective program in each state. The larger is z the stronger is the hacker's adverse selection edge. On top of (1), he earns extra flow payoff:

$$-z(\alpha - \beta) \text{ in state } \theta = 0 \qquad \text{and} \qquad z(\alpha - \beta) \text{ in state } \theta = 1 \tag{2}$$

Computer users do not know the state, but have access to public information about computer attacks. The more the hacker exploits the vulnerability, the more people post to electronic bulletin boards about being hacked. Since there are many smaller hacks ongoing that confuse the reports, there emerges a noisy public signal $Y(t)$ about the *intensity differential* $\Delta(t) = \alpha(t) - \beta(t)$, obscured by a Weiner noise process $W(t)$:[4]

$$dY = \Delta\ dt + \sigma\ dW. \tag{3}$$

So $\sigma W(t)$ is a driftless Brownian motion with *volatility* σ (or variance σ^2). This process constitutes *misinformation* about the informed player's activity, veiling his intensity. Computer users see the signal process (3) unfold, and make inferences about the state θ.

(b) Insider Trading (Back and Baruch 2004). An informed insider knows if an asset is worth $\theta = 0$ or $\theta = 1$ per share. He may buy or sell continuously over $[0, \infty)$ until payoffs are revealed at the random market closure time, arriving at a constant rate $r > 0$. The insider chooses a *net order flow* (i.e. buys minus sells) volume $\Delta(t) \in \mathbb{R}$ at each time t. So the absolute net order flow $|\Delta(t)|$ is the net purchase rate if $\Delta(t) > 0$, and the net sales rate if $\Delta(t) < 0$. *In this finance setting, the choice variable p is no longer a mixed strategy but the current market price.* The flow payoff from insider trading is $\Delta(\theta - p)$, or the sum of (1) and (2) in the special case $z = x = 1/2$.

Rather than assume competitive market makers, we venture that prices are set by a sequence of players (the "market"), constrained to choose a single price p at each moment, and take all trade volume at this price. The market knows neither the state nor the payoffs of past trades, but sees the market net order flow dY, i.e. the sum of the insider's net order flow and noise trader volume. As is standard, we model noise trade volume as a scaled Wiener process $\sigma W(t)$, yielding the public signal process (3).

The hacker model nontrivially generalizes the insider trading model. For the insider prefers to sell at any price when $\theta = 0$, and buy at any price when $\theta = 1$. Namely, $\Delta < 0$ is myopically dominated for the long run player if $\theta = 1$, and $\Delta > 0$ is dominated if $\theta = 0$. In contrast, the hacker model when $z < x$ has no dominated strategies. In state $\theta = 0$, he prefers $\Delta > 0$ for any low enough mixture p satisfying $(1-2p)x - z > 0$, and prefers $\Delta < 0$ for high enough mixtures p with $(1-2p)x - z < 0$.

[4]Unlike Hendricks and McAfee (2006) and Crawford (2003), this Gaussian signal process has full-support, thereby obscuring all actions.

(c) Commercial Espionage. Inspired by the 2009 movie "Duplicity", imagine that two rival high technology firms compete in a market. Firm 1 is quite large — it has an essentially unlimited research budget — while Firm 2 has a fixed research budget. They each choose how much to invest in two alternative R&D paths: a or b. Firm 1 invests α and β dollars into these paths, respectively, while Firm 2 splits its research budget $(1-p, p)$ on (a, b). Firm 1 profits most if Firm 2 copies it, while Firm 2 prefers a different strategy. Firm 1 also knows about a patent discovery. It either has not or has made a major breakthrough that renders strategy a better or not (states $\theta = 0, 1$). Payoffs and information are as in (1), (2), and (3), with one key exception: In this case, *a single uninformed player discounts the stream of payoffs at rate* $r > 0$.

Given the surprise payoff gain x in (1) and private information factor z in (2), the *information ratio* is $\xi \equiv z/x$. This is a quotient of the private information gains and the fixed gain from action surprise. The general model posits a zero-sum game, with an informed player's payoff $\Delta u_\theta(p)$, given by (1) plus (2), for the *unit flow payoffs*:

$$u_0(p) = 1 - 2p - \xi \qquad \text{and} \qquad u_1(p) = 1 - 2p + \xi \tag{4}$$

His flow payoff $v_\theta(q)$ is the *product* of his unit flow payoff $u_0(q)$ and his intensity differential $\Delta_\theta(q)$. If $u_\theta(p) < 0$, then he earns a positive flow payoff by choosing $\Delta < 0$.

By restricting to a state-symmetric world, we are able to give very sharp and simple characterizations throughout the paper. We often need only give results for one state.

3 The Unique Markov Equilibrium

We explore Markov equilibria where strategies depend on the public belief $q(t)$ that the state is $\theta = 1$. Since the observation process is common knowledge, so too are the induced public beliefs. Their evolution over time depends on: (i) the *actual* mixed strategies of the informed player, (ii) the state contingent mixed strategies *expected by the uninformed player*, and (iii) random observational noise.

Let $\delta_\theta(q)$ be the informed player's *expectation of the intensity differential* Δ *in state* θ *at public belief* q. Unconditional on θ, she expects the intensity differential:

$$\delta(q) \equiv (1-q)\delta_0(q) + q\delta_1(q). \tag{5}$$

We can now express the public belief process, derived via Bayes rule in the Appendix.[5]

Lemma 1 (Public Beliefs) *For the realized intensity differential* Δ*, beliefs obey:*[6]

$$dq(t) = \left[q(1-q)(\delta_1 - \delta_0)(\Delta - \delta)/\sigma^2\right] dt + \left[q(1-q)(\delta_1 - \delta_0)/\sigma\right] dZ \tag{6}$$

This process obtains in and out of equilibrium, and has many intuitive properties central to our theory. First, belief volatility is greater for smaller signal noise σ^2. This makes sense, since a stronger inference can be drawn from any given intensity Δ. Next, if the uninformed player expects the same intensity in the two states ($\delta_0 = \delta_1$), then she ignores signal process innovations dY and public beliefs never change; essentially, the uninformed player attributes all fluctuations in the signal process to noise. Third, if the informed player chooses the unconditionally expected intensity differential $\Delta = \delta(q)$, then beliefs do not drift; however, they still are volatile, as the uninformed player does not know what the intensity differential Δ is, and so nonetheless seeks to learn from the signal innovations. Finally, when the realized intensity Δ differs from its expectation $\delta(q)$, the belief drift responds linearly, while belief volatility is unaffected. Unexpected changes in the observation process cannot raise volatility, since they are unobserved.

We will assume for now the case of a stream of uninformed players, each around for a moment in time. Each does not know the state, and so cannot impact the belief evolution process by her strategy, except insofar as it depends on the belief q. So if the informed player only conditions his intensity on the current public belief, his rivals' choice at any time t cannot impact the future. At each time t, the informed player chooses an intensity differential $\delta_\theta(q(t))$ in state θ. Simultaneously,[7] the uninformed player sets the mixed strategy $p(q(t))$, having seen the choices of the uninformed players and the signal process $Y(s)$ for times $s < t$. Each therefore minimizes her myopic expected loss:

$$q\delta_1(q)u_1(p(q)) + (1-q)\delta_0(q)u_0(p(q)) \tag{7}$$

In light of (4), the uninformed player is indifferent across mixtures p in this linear

[5] Henceforth, we will often suppress the q arguments below, whenever it is clear.

[6] This Wiener process Z is formed from the uninformed player's perspective. It is an expectation of the two possible realizations of the Wiener process W, each consistent with one of the states $\theta = 0, 1$.

[7] The sequential versus simultaneous distinction choices in this continuous setting is immaterial, as the realized observation process sample paths are continuous: Knowing $Y(s)$ on $[0, s)$ tells us $Y(t)$.

minimization exactly when she expects the zero intensity differential in (5):

$$\delta(q) = 0$$

The informed player engages in a far-sighted optimization. His *informed value* $V_\theta(q)$ is the present value of his payoffs in state θ given the belief q. It is positive precisely due to the prevailing informational asymmetry and the ongoing misinformation. The Bellman equations as usual can be intuitively captured by the asset value equation that "the return equals the dividend plus the expected capital gain" — where the "dividend" is the flow payoff, while the capital gains are expected value losses. So:

$$rV_\theta = \max_{\Delta} \Delta u_\theta(p) + q(1-q)(\delta_1 - \delta_0)(\Delta - \delta)V_\theta'/\sigma^2 + \frac{1}{2}q^2(1-q)^2(\delta_1-\delta_0)^2 V_\theta''/\sigma^2 \quad (8)$$

given the drift and variance in (6). The informed player best responds to the mixture p. Since his payoff is linear in Δ, he must be indifferent across all intensity differentials Δ. While we cannot deduce the sign of Δ from a second order condition, we can show that in equilibrium, he always chooses $\Delta > 0$ for state $\theta = 1$ and $\Delta < 0$ for $\theta = 0$. So he secures a myopic gain and flow capital loss from eroding his information edge. Thus:

$$u_\theta(p) = -q(1-q)(\delta_1(q) - \delta_0(q))V_\theta'(q)/\sigma^2 \quad (9)$$

in states $\theta = 0, 1$. Inserting (5) and (9) into equation (8), and putting $\Delta = \delta_\theta$, we get:

$$r\sigma^2 V_\theta(q) = -\delta(q)q(1-q)(\delta_1(q) - \delta_0(q))V_\theta'(q) + \frac{1}{2}q^2(1-q)^2(\delta_1(q)-\delta_0(q))^2 V_\theta''(q) \quad (10)$$

A *Markov equilibrium* is a 5-tuple $(V_0, V_1, \delta_0, \delta_1, p)$, such that: (a) strategies and values depend solely on the public belief q, which obeys the law (6), and (b) the uninformed player's mixture p minimizes (7), and (c) the informed player's intensity differential δ_θ maximizes his present value of payoffs V_θ, namely, δ_θ and V_θ obey (9) and (10). The present value of the uninformed player's losses is then $V(q) \equiv qV_1(q) + (1-q)V_0(q)$.

We first observe that the insider trading model is the extreme parametric case in our class of unbounded intensity competitive models for which an equilibrium exists.

Proposition 0 *There is no equilibrium for any private information ratio $\xi > 1$.*

Understanding this proves instructive, insofar as it exploits the unbounded action space: The logic is that the informed player can earn infinite profits with any stronger private information ratio. When $\xi > 1$, the informed player's flow payoff equals

$$\begin{cases} \Delta x(1-2p-\xi) \geq \Delta x(1-\xi) > 0 & \text{in state } \theta = 0, \text{ if } \Delta < 0 \\ \Delta x(1-2p+\xi) \geq \Delta x(\xi - 1) > 0 & \text{in state } \theta = 1, \text{ if } \Delta > 0 \end{cases}$$

He could earn unbounded flow payoffs by exploding his intensity differential $\Delta \to \mp\infty$, respectively — for any uninformed strategy. The uninformed player can only defend herself against action surprise provided the private information ratio is not excessive.

From now on, we assume an information ratio $0 < \xi \leq 1$, and deduce a unique Markov equilibrium, expressed in terms of the *Probit function* $F(q)$ — the inverse of the standard normal cumulative distribution function. As Figures 1 and 2 illustrate,

Proposition 1 (Equilibrium) *Fix an information ratio* $0<\xi\leq 1$ *and belief* $0<q<1$.
(a) There is a unique Markov equilibrium $(V_0, V_1, \delta_0, \delta_1, p)$.
(b) In this equilibrium, the uninformed player chooses the mixture $p(q) = 1/2+\xi(q-1/2)$ *while the informed player opts for the intensity differentials:*

$$\sigma\sqrt{r/\pi}e^{-\frac{1}{2}F(q)^2}/q = \delta_1(q) > 0 > \delta_0(q) = -\sigma\sqrt{r/\pi}e^{-\frac{1}{2}F(q)^2}/(1-q) \tag{11}$$

(c) The uninformed player's flow losses have present value $V(q)$*, where:*

$$V(q) = \left(\xi\sigma/\sqrt{r\pi}\right) e^{-\frac{1}{2}F(q)^2}$$

(d) The informed player's unit flow payoffs $u_1(q) = -u_0(1-q) = 2\xi(1-q)$ *are positive, while his value functions obey* $V_1(q) = V_0(1-q)$*, where* $V_0(q) = V(q)+(\xi\sigma/\sqrt{r\pi})\, qF(q)$.

Our existence argument in the Appendix is constructive. We prove uniqueness by exploiting the zero sum structure of the game, since we have found a saddle point.

By Proposition 1, the informed player incurs capital losses $E(dV) < 0$. Neither a growth stock nor a blue chip, it is instead a dying firm, slowly burning its information capital stock. In a key identity, its expected dividend (7) is twice the return $2rV(q)$.

Corollary 1 *Expected dividends are identically twice the return, or twice capital losses.*

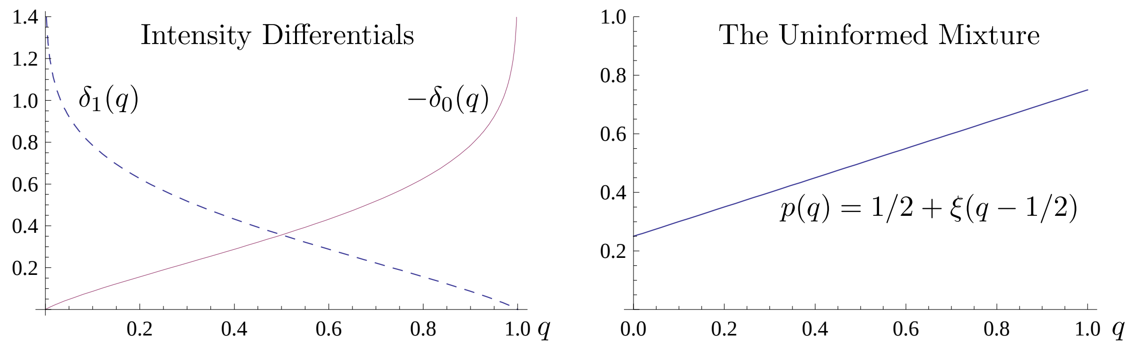

Figure 1: **Equilibrium Strategies:** The informed player's intensity differentials δ_1 (dashed) and $-\delta_0$ (solid) are on the left, and uninformed player's mixture p on the right, for an informational ratio $0 < \xi < 1$. When $\xi = 1$, this coincides with the diagonal, or "price equals expected value" for insider trade. Parameters are $(r, \sigma, \xi) = (0.1, 1, 1/2)$.

4 Behavior and Payoffs

A. Changing Information Ratio

This section fleshes out the economics of Proposition 1. With a unit private information ratio $\xi = 1$, the uninformed mixture chance p equals the expected state q — the analogue of price equals expected value ($p(q) = q$) in finance. So this classic equality is a best response by a market seeking to minimize losses from insider trading. The uninformed player's strategy ensures herself zero expected losses, unconditional on the state. To wit, Back and Baruch (2004), and thus Kyle (1985), have formally found a saddle point in a zero-sum game. This is an orthogonal route to their equilibrium.

As the ratio ξ diminishes, the uninformed player tempers her response to the belief q, as the mixture function $p(q)$ flattens. So unlike a "price", $p(q)$ generally does not converge to 0 or 1 when public beliefs q do. Instead, the uninformed players choose an interior mixture in order to guard themselves from the fixed portion of the action surprise. The computer users will employ each browser with a boundedly positive chance, even as it becomes a near sure thing that it is more vulnerable to the hack.

In equilibrium, the informed player's strategy is not affected by the ratio ξ. The insider's trading behavior in Back and Baruch (2004) captures the informed player's *behavior* in the entire class of games for any $\xi > 0$. The terrorist hacker and the industrial espionage firm alike act like insiders trying to secretly unload a long position in a bad stock. In all cases, the uninformed player must be indifferent between his actions. This does not depend on the private information ratio ξ. By contrast, the informed player's payoffs are scaled by the ratio ξ, inflating the adverse selection portion of his payoffs. Our inquiry does not ask questions posed in finance of the insider.

B. Changing Noise and Impatience — Square Root Substitution

Given his private information, the informed hacker maximizes his immediate flow profits with an unbounded intensity differential $\delta_0 = -\delta_1 = \infty$. But in equilibrium, he trades off current and future payoffs, and instead chooses a finite intensity. By holding $\delta_1 - \delta_0$ in check, he optimally exploits the misinformation from the signal process: The veil over his activity lifts the greater is his intensity differential.

A more impatient informed player relatively values current payoffs more, and so more vehemently exploits his private information. As the information depreciation rate ϕ explodes, so too does the interest rate $r = \phi + i$, and therefore the intensity differentials. For the informed player's required return rises in the interest rate, and this can only be achieved with a rising dividend and so a rising intensity differential.

The same dynamic substitution effect also arises with noisier action observations. When the signal volatility rises, the informed trader accommodates by accelerating his trades. In equilibrium, the capital losses from doubling the signal variance σ^2 scale down by the same factor as the return rises from doubling the interest rate r.

Corollary 2 (Square Root Substitution) *Intensity differentials are scaled by* $\sigma\sqrt{r}$.

While greater volatility σ affects the opportunity costs of intensity differentials, while more impatience r changes preferences, both have the same effect on actions because they are equivalent by a time rescaling. If variance σ^2 doubles, then observations are as informative as if the model ran in "double time" — i.e., as if the interest rate doubled.

Values also transparently reflect all model parameters via the constant $\xi\sigma/\sqrt{r}$. For if the hacker did not adjust his strategies, his value would move like a bond price $1/r$. But with the square root substitution effect, values in Proposition 1 behave like $1/\sqrt{r}$. Eg., the informed player's payoffs roughly halves if the discovery rate ϕ quadruples.

The equality of dividends and twice the return forces the square root dynamic substitution effect. Indeed, if a greater discovery rate ϕ quadruples his interest rate, he must double his dividend and absorbs a capital loss twice as large. His value scales by half, and his return doubles. The same logic directly applies when the volatility σ halves. When noise quadruples, the intensity differential, value, and return all double.

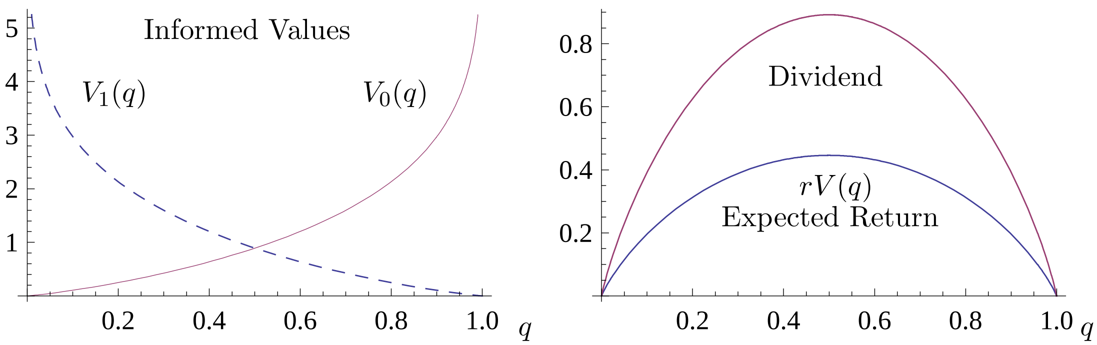

Figure 2: **Value Functions:** On the left are informed values V_0 and V_1. At right is the dividend and expected return rV, when $(r,\sigma,\xi) = (0.1, 1, 1)$. We see how $V_0(q)/V_1(q) > q/(1-q)$ for beliefs $q > 1/2$, and that dividends are twice the return, as in Corollary 1.

C. Changing Public Beliefs — Square Root Log Substitution

Let us now consider how behavior and payoffs respond as the public beliefs q evolve over time. As seen in Figure 2, the informed player's values rise in q as the uninformed player grows more convinced in a lie, so that $V_1'(q) < 0 < V_0'(q)$, with both vanishing near correct beliefs. To sustain the equality of the dividend and twice the return, the intensity differentials are monotone and vanish as in Figure 1.

The distinction between our dynamic setting and a purely static one is seen in the ratio of values. In a static bet, if the likelihood ratio favors the high state, the state-contingent payoffs inversely adjust: one can make a 300% return if a 3:1 bet pays off. But here, the hacker betrays his information as he uses it, and his equilibrium value ratio cannot so fully reflect public beliefs. Precisely, in the Appendix we prove the inequality:

$$\frac{V_0(q)}{V_1(q)} \lesseqgtr \frac{q}{1-q} \qquad \text{for } q \lesseqgtr 1/2 \tag{12}$$

We can further sharpen our insights by focusing on behavior and values near the extreme public beliefs 0 and 1, when the uninformed player is either most deceived or most informed. The Probit function formulas in Proposition 1 then greatly simplify:

Corollary 3 (Square Root Logarithm Law of Deception) *Fix state* $\theta = 0$.
(a) As the uninformed player grows convinced in a lie, or $q \uparrow 1$, *the informed player's payoff and intensity differential both explode like* $\sqrt{-\log(1-q)}$.
(b) As the uninformed player grows convinced in the truth, or $q \downarrow 0$, *the informed player's intensity differential vanishes like* $q\sqrt{-\log q}$, *and his payoff like* $q/\sqrt{-\log q}$.

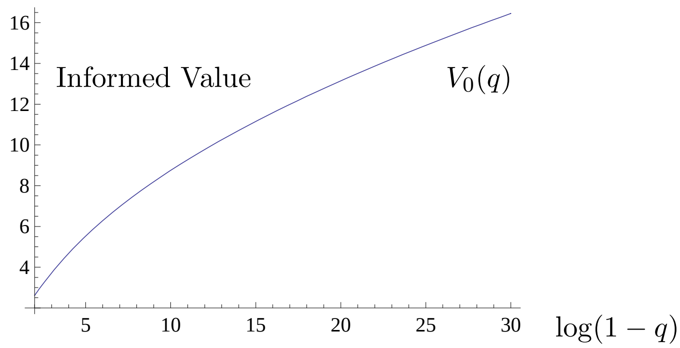

Figure 3: **The Square Root Logarithm Law of Deception.** The informed player has tremendous diminishing returns to greater deception. His value has a square root shape in the logarithmic confidence in the correct state, like $\log(1-q)$ in state $\theta = 0$.

This *Square Root Logarithm Law* speaks to the extreme decreasing returns to misinformation, underscoring the ephemeral quality of deception. At high deception levels, when the odds in favor of the wrong state double, the informed player's payoff does not even rise by a constant — rather, his squared payoff roughly does. His intensity differential behaves identically, rising less than linearly in the public odds $q/(1-q)$ in state $\theta = 0$. The flip side arises with vanishing deception, where the odds in favor of the wrong state double. In this case, the informed player's payoff more than doubles — *initially increasing returns to deception*, as seen at the left in Figure 2.

An informative signal is unambiguously beneficial for the uninformed player: Since this signal effects a *mean preserving* change in beliefs, she likes fair belief gambles, and is therefore "risk-loving". So by the converse to Jensen's inequality, her value function $-V(q)$ is convex, and her present value of losses $V(q)$ is then concave, as in Figure 2.

While the informed and unformed players have diametrically opposed preferences in the game, they share the same risk preferences. For not only are uninformed values $V(q)$ in Figure 1 strictly convex in public beliefs, so too are informed values $V_\theta(q)$. In other words, both rivals strictly benefit from mean preserving spreads in the public belief. How is this? Surely an informative signal is unambiguously *bad* for the informed player. Notice, however, that it does not effect a fair belief gamble for him — since *he knows the state*, information pushes public beliefs toward the truth.[8]

[8]Since $V_0' > 0 > E_0[dq]$ and $V_1' < 0 < E_1[dq]$ by (6) and $\delta_0 < 0 < \delta_1$, the expected change in the value V_θ from any dq in beliefs $V_\theta' E_\theta[dq] + V_\theta'' E_\theta[dq^2]$ can be negative despite the value convexity.

5 Time and Money

"The art of using deceit and cunning grow continually weaker and less effective to the user." — John Tillotson, Archbishop of Canterbury (1691–94)

The misinformation veil on the informed player's actions ultimately allows him to convert time into money, as his edge erodes with the passage of time. We now ask how rapidly the informed player monetizes his private information edge. How long before the jig is up? We find that the Archbishop's instinct above well-captures how long would a maximizing Ponzi schemer or other deceiver sustain his deceptive activities.

For our first insight, recall from Corollary 1 that capital losses at all times coincide with the return. We can prove that this sustained capital loss forces the firm to shrink in expectation at precisely the rate of interest.

Corollary 4 *Uninformed players expect the future value after delay t is $e^{-rt}V(q)$.*

For an alternative interpretation, we can examine how values depend on time. Let $\mu_\theta(q)$ and $\varsigma(q)$ denote the equilibrium drifts and volatility of public beliefs (6) — where volatility is state-independent. After substituting the trading rates (11), we find:

$$\mu_1(q) = \frac{2r}{qF'(q)^2} = -\mu_0(1-q) > 0 \qquad \text{and} \qquad \varsigma^2(q) = \frac{2r}{F'(q)^2}. \tag{13}$$

We see at once that *in equilibrium, the public belief drift and volatility is independent of the observation noise σ.* By the square root substitution effect, *as the volatility σ rises, the informed player exactly compensates by proportionately scaling up his intensity differential* (11). In other words, the ease of misinformation exactly cancels out. This allows us to draw some sharp conclusions valid for any volatility σ.

The relationship between time and money depends on the equilibrium public belief process (13). As seen in the Figure 4, the public belief volatility $\varsigma^2(q)$ is symmetric, peaking when the uninformed player is most unsure of the state at $q = 1/2$. While Bayes rule damps both belief drift and volatility near extremes, the state contingent belief drifts in (13) are asymmetrically tear-drop shaped — for intensity differentials in (11) are monotone. Beliefs drift fastest toward the truth when the public is mostly deceived (namely 73%, i.e. peaking at $p = 0.27$ when $\theta = 1$), and halt at focused beliefs.

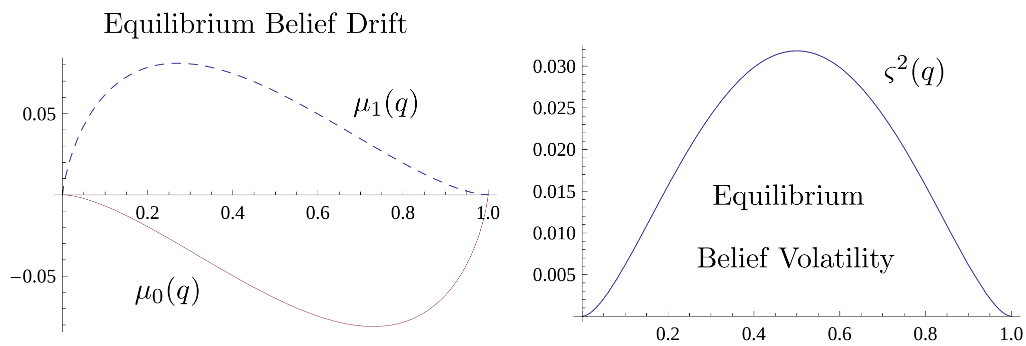

Figure 4: **Belief Evolution in Equilibrium:** The graph on the left depicts the skewed state contingent belief drift, namely $\mu_0 < 0$ (solid) and $\mu_1 > 0$ (dashed), given the interest rate $r = 0.1$. The graph on the right is the belief variance $\varsigma^2(q)$. The equilibrium process is noticeably fastest when the public belief is mostly wrong.

Back and Baruch (2004) show that the insider's private information is eventually fully embedded in the price, and the truth is revealed. But since the public belief drift vanishes as the belief converges to the truth, learning slows down. We argue in the appendix that the information is only fully revealed after infinite expected time. Because of this, we settle for "nearly complete learning". We also focus on the unconditional belief process, for this is what an outsider observing everything could infer. Let $T(q)$ be the expected time until the belief process (6) starting at $q_0 = q$ leaves any strictly interior symmetric interval $(1 - \bar{q}, \bar{q})$, where $1/2 < \bar{q} < 1$. If we interpret $T(q)$ as the value of an asset paying a dividend of 1 per unit time, with zero interest rate (zero return), zero drift, and volatility $\varsigma^2(q)$, then we find that it obeys the differential equation:

$$1 + \frac{1}{2}\varsigma^2(q)T''(q) = 0 \tag{14}$$

and $T(\bar{q}) = T(1-\bar{q}) = 0$. Conveniently, neither $T'(q)$ nor $T''(q)$ depends on $\bar{q}$, since the function $T(q)$ is only vertically shifted by $\bar{q}$.[9] We next show that T is not only a concave function of beliefs by (14), but also of "money", i.e. the fraction $\nu = V(q)/V(1/2)$ of peak expected value $V(1/2)$. Indeed, let $\top_\epsilon(\nu)$ be the expected time (from the market's perspective) starting at value $\nu V(1/2)$ until earning a fraction $1 - \epsilon$ of the value. Or, this is the market's expected time until it loses a fraction $1 - \epsilon$ of its maximal loss.

[9]For by the intuitive logic of stopping times, if we chose the smaller barrier $\epsilon' < \epsilon$, then $T_\epsilon(q)$ rises by a fixed amount. Indeed, the expected time to first hit ϵ' or $1 - \epsilon'$ equals the expected time to first hit ϵ or $1 - \epsilon$ plus the expected time from that moment to first hit the next barrier ϵ' or $1 - \epsilon'$.

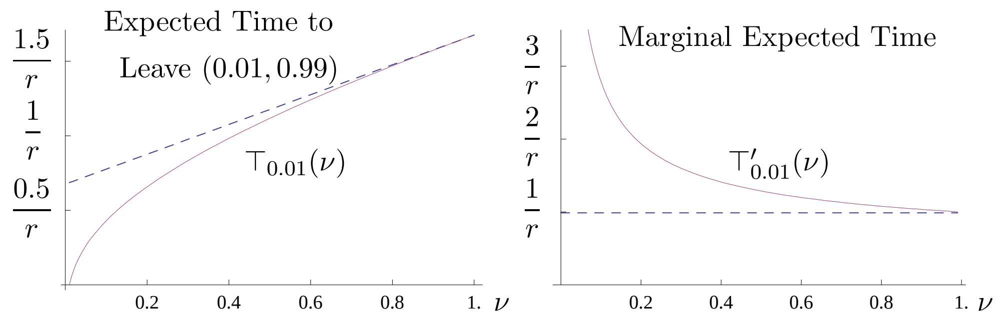

Figure 5: **Deception Duration.** At left is the expected time until beliefs first leave an interval $(0.01, 0.99)$, drawn as a function of the fraction of the value (solid). Its best linear approximation at $\nu = 1$ (dashed) has slope $1/r$ at $\eta = 1$ ("time is money"). At right is its derivative — a fundamental function, independent of all parameters σ, r, ξ.

Proposition 2 (The Diminishing Returns to Time) *For any* $0 < \epsilon < 1/2$*, the expected time* $\top_\epsilon(\nu)$ *until the value* $V(q)$ *hits* $\epsilon V(1/2)$ *is increasing and concave in the initial fraction* $\nu \leq 1$ *of peak value. Further, its initial slope in* ν *is* $r\top'_\epsilon(1) = 1$.

So the informed player finds that his deception exercise grows less profitable over time, consistent with the Archbishop's insight. This diminishing returns makes sense, as he confronts increasingly worse "terms of trade". For the mixture p that he faces drifts up in state $\theta = 1$, and down in state $\theta = 0$, namely $E_1[dp] > 0 > E_0[dp]$ by (13).

Second, in a different manifestation of the constant firm shrink rate in Corollary 4, we initially have $r\top'_\epsilon(1) = 1$, or "time is money". Starting at belief $q = 1/2$, he earns $r\%$ of this peak value the first year. His payoff falls over time, by concavity, but he still earns at least $\frac{1}{2}r\%$ of his peak value through year $0.7/r$. All told, the market expects him to earn 99% of this value within a little more than $1.5/r$ years (Figure 5). *With no discovery risk, and a 10% bank interest rate, the informed player earns 99% of his peak value within 15 years.* With a positive exogenous discovery risk, the deception ends sooner. Let's briefly compare this with classic examples. Wikipedia lists a host of Ponzi schemes, with durations ranging from a few months to 14 years. Near the shorter end, Charles Ponzi's infamous scheme lasted about two years from 1918–20. The accounting deception by Enron may have begun in the early 1990s, and persisted until 2001. Madoff's deception may have lasted the longest, perhaps fifteen years — he admitted that he stopped trading and started fabricating earnings in the mid 1990s.

Since time is money, one might be tempted to think that a longer-lasting monopoly is a stronger one. We now show that this conjecture depends on the source of the informed player's strength. First, when the signal volatility σ rises, so does the informed player's value, by Proposition 1. But this does not affect equilibrium public beliefs (13), and thus not the expected duration. So for parametric changes in the signal process, longer is neither stronger nor weaker. We illustrate this with a likely apocryphal story: Everyone knew that Nathan Rothschild would learn the outcome of the Battle of Waterloo first.[10] After victory, he allowed his known agents to observably sell consols on the London Exchange to deceive the market. Then his more clandestine shills secretly bought them cheaply. Within twenty-four hours, he earned a fortune. Such quick profits are consistent with our theory if his strength owed to a large volatility σ.

Second, the same stronger-is-not-longer conclusion holds with a greater information edge ξ, but for different reasons. The informed player's actions do not depend on ξ, and so neither does the equilibrium belief evolution, nor thus duration; however, the informed player earns a higher payoff per "trade" as ξ rises, and so his value rises.

Finally, a more patient informed player is simultaneously a more profitable one and a longer one. Here, the strength and length of an adverse selection edge coincide. For instance, Madoff's Ponzi scheme was both long-lasting and lucrative, and it appears that a major source of this power was his patience: He perceived an extremely low exogenous discovery rate. In the end he was not found, but instead turned himself in.

6 Information or Espionage

We now turn to information acquisition by the public that undermines the informed player's edge. This subsumes a host of economic situations — from trying to learn about the hacker in our example, to military or industrial espionage, as one firm seeks to catch wind of his rival's critical advantage. In this information acquisition exercise, we confront an old problem — that while information is valuable, it critically differs from other goods: Under standard conditions, its marginal value in static contexts is initially zero, and so the value of information cannot possibly be concave.[11] This

[10] See www.thetruthseeker.co.uk/article.asp?ID=840. It surely inspired the climactic trading scene in the Academy Award winning movie "Wall Street".

[11] See Radner and Stiglitz (1984) and Chade and Schlee (2002). As Keppo, Moscarini, and Smith (2008) show, this non-concavity even holds for two state models with Gaussian information, like here.

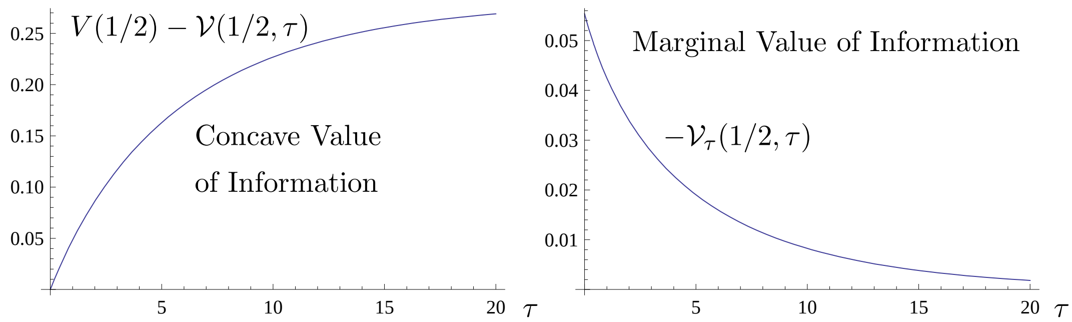

Figure 6: **The Concave Value of Information to the Uninformed.** Except near extreme beliefs, the static value of information is a globally concave function of the precision of signals for the uninformed player, and the marginal value of information is initially finite, and then globally falling. So the law of demand holds.

immediately implies that no one ever buys just a little information. So is it true in our dynamic setting that no one ever has an incentive to engage in just a little espionage?

For a fair comparison, assume that the uninformed player can access an information process similar to that afforded in equilibrium by the passage of time: At any moment in time, she can instantaneously see the realization of a signal process just like (3), except with fixed unit drift and volatility. The belief process evolves according to $dQ = Q(1-Q)dZ$ in "metaphorical time" on $[0, \tau]$, for a Wiener process Z. If the uninformed player ends with the random posterior belief $Q(\tau)$, then she earns the "terminal reward" $V(Q(\tau))$. If we define $\mathcal{V}(q, \tau) \equiv E[V(Q(\tau))|Q(0)=q]$, then the value of information is the reduction $V(q) - \mathcal{V}(q, \tau)$ in the uninformed player's expected loss.

Proposition 3 (A Mostly Concave Value of Information)
(a) The initial marginal value of information is positive and finite. It peaks at the public belief $q = 1/2$, and vanishes as the belief q tends to 0 or 1.
(b) The marginal value of information is forever falling at interior public beliefs in $(q^*, 1-q^*)$, but rises and then falls for beliefs outside $(q^*, 1-q^*)$, for some $0 < q^* < 1/2$.

To be precise, we find the threshold $q^* \approx 0.077$, which is independent of all parameters. So the value of information is globally concave so long as beliefs lie in $(0.077, 0.923)$.

In our dynamic competitive environment, with standard informational units, the classsic informational non-concavity disappears except at extreme beliefs. Even a little spying or sleuthing can be worthwhile in wartime, industrial competition, or fighting crime. Madoff's fund investors would profit more from the first bit of news about his fund than from later bits, and therefore could have an incentive to procure even a small amount of information. Why? The very act of public information acquisition changes the economic environment. Intuitively, if investors grow more informed, then Madoff would exploit his edge less intensely. This provides an extra source of diminishing returns to information purchases that is sufficient to secure a concave value of information except at extreme beliefs. This restores the standard law of demand, as seen in Figure 6. As with standard economic goods, demand is zero above a choke-off price for information, and rises continuously from zero as the price falls below this threshold.[12]

Finally, Proposition 3(b) asserts that near focused beliefs — such as when word of Madoff was breaking — information experienced a non-concavity. At this point, individuals had an incentive to engage in large time-intensive sleuthing, or none at all.

7 Disinformation

So far, we have explored *misinformation* — i.e., false or inaccurate information that is spread *unintentionally.* The exogenous Wiener noise process that has partially veiled actions is the source of misinformation: Beliefs sometimes drift away from the truth due to chance Wiener process increments. But Rothschild's stunningly profitable purported deception exercise after Waterloo was instead an example of *disinformation* — namely, the *intentional* dissemination of false information to sway the opinions of its recipients.

This suggests an economic formulation of disinformation in our dynamic world: taking current dividend losses for capital gains, i.e. a more favorable public belief. Yet this cannot occur in equilibrium because if the informed player knows his rival's strategy, then at the worst, he weakly learns nothing about the state. In other words, the public belief cannot drift away from the truth in equilibrium. While out-of-equilibrium exercises were used in the strategic refinements literature, they are often frowned upon. But there is no more compelling venue for such a thought experiment than in the study of deception.[13] In fact, we find precise limits on the use and level of disinformation.

So inspired, imagine that the informed player has a *wholly unexpected opportunity*

[12]That information may have counterintuitive properties in a general equilibrium context is the message of Hirshleifer (1971) and Schlee (2001). By contrast, we argue that an equilibrium setting (now in game theory) *rescues an intuitive economic feature to information.*

[13]Or, through the lens of Occam's Razor, we are approximating a richer and less tractable model that also tracks the chance that the informed player has access to another unlikely deception technology.

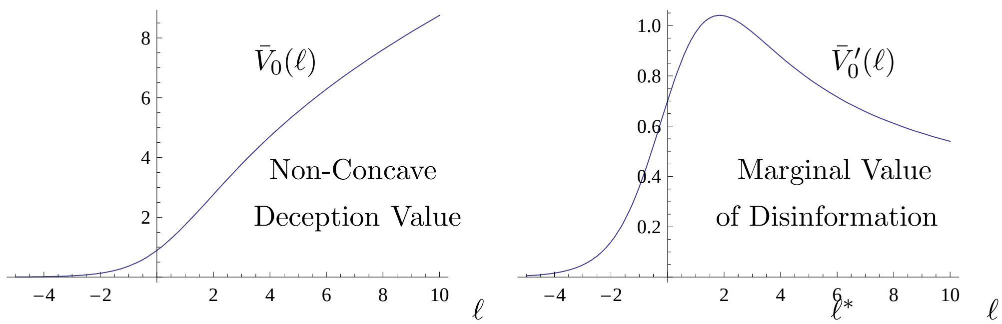

Figure 7: **The Deception Value and the Marginal Value of Deception.** This graph is drawn for state $\theta = 0$ as a function of the log likelihood ratio ℓ. Any disinformation exercise logically must push public beliefs further from the truth.

to engage in disinformation. In the hacker example, he could use the less advantaged hack, "buying" false signals for the uninformed player. Disinformative signals are still evaluated with noise. For instance, in Operation Mincemeat discussed earlier, the body might not have washed up on the beaches of Spain. The British also worried that the Germans might see through the ruse, and this would have been a strong contrary signal.

The informed player's disinformation exercise differs from the information purchase by uninformed players in §6, since it is unexpected, and is achieved by way of misleading signals. If we imagine these signals as the basic cost units, since Bayes rule is additive in log-likelihood ratio space, *the difference of* $\ell(q) - \ell(q_0)$ *of log-likelihood ratios* $\ell(q) \equiv \log(q/(1-q))$ *quantifies the level of disinformation* (Appendix A.9). Written as a function of ℓ, we suggestively call $\bar{V}_\theta(\ell) \equiv V_\theta(e^\ell/(1+e^\ell))$ the *deception value.*

Proposition 4 (Convex-Concave Deception Value) *Fix state* $\theta = 0$*. There is a threshold log-likelihood ratio* ℓ_0^* *such that the deception value* $\bar{V}_0$ *is convex for* $\ell < \ell_0^*$*, and concave for* $\ell > \ell_0^*$*. This threshold* ℓ_0^* *is invariant to the model parameters* σ, r, ξ*.*

This result asserts that the deception value behaves like the standard static value of Gaussian information in decision theory — first convex and then concave (Figure 7). The threshold $\ell_0^* \approx 1.845$ between the concave and convex regions of the function $\bar{V}_0(\ell)$ is unaffected by all parameters since the informed values V_θ are linear in $\xi\sigma/\sqrt{r}$.

For analytic insight into Proposition 4, the approximation $\bar{V}_0(\ell(q)) \sim q/\sqrt{-\log q} \sim e^{\ell(q)}/\sqrt{-\ell(q)}$ from Corollary 2 suggests increasing returns for public beliefs q near 0, or very negative log-likelihood ratios ℓ. Likewise, diminishing returns for large positive $\ell(q)$ are hinted by $\bar{V}_0(\ell(q)) \sim \sqrt{-\log(1-q)} \sim \sqrt{\ell(q)}$ for high public beliefs q near 1.

This shape of the deception value function yields an immediate tangible economic implication — *there is no small deception:* If its marginal cost is constant, the optimal level of disinformation is either zero or boundedly positive. For instance, Nathan Rothschild engaged in a large deception. British disinformation exercises in advance of D-Day seeking to convince the Germans of a Calais invasion — Operation Fortitude — were quite substantial in scope. These included inflatable rubber tanks, plywood artillery, mock wireless traffic, and phantom groups like the First U.S. Army Group, led by George Patton. The accounting fraud that Enron committed was clearly a large scale disinformation exercise involving offshore entities to absorb losses.

But not only must disinformation be large, we can predict where it should occur. For the threshold $\ell_0^* \approx 1.845$ maps to a belief cut-off $q_0^* \approx 0.864$ — independent of all model parameters. So any disinformation exercise must aim to move public beliefs past the point where concavity sets in — and so beyond 86% confidence in the wrong state. In the end, eg., Operation Fortitude led the Germans to hold most of their reserves waiting for a Calais attack that never came. And when stock market confidence in Enron was high, it helped sustain its record-setting deception enterprise by creating offshore Special Purpose Entities to hide losses. It also enlisted the complicity of the accounting firm Arthur Anderson to so disguise their finances. These activities were not without cost — Arthur Anderson's cooperation was bought through excessive payments and explicitly steering business their way without competitive bids.

8 Obfuscation

An important cross between misinformation and disinformation is *obfuscation* — the deliberate and costly concealment of actions. Like disinformation, obfuscation is intentional. But it controls the volatility rather than drift of public signals. In the business world where information is legally verifiable ex post, disinformation is usually illegal, while information concealment can be defended as good business practice. In war and competition, obfuscation is often achieved by elaborate diversion activities that distract the opponent's attention. An easy modification of our model captures this facility.

We now assume that the informed player can adjust the signal volatility. By paying the flow cost $c(\sigma)$, he may maintain volatility σ. Assume a free baseline level of noise $\bar{\sigma} \geq 0$; above this threshold, he pays the smooth, strictly increasing, and convex flow cost $c(\sigma)$ for $\sigma \geq \bar{\sigma}$, with initial zero marginal cost.

The intensity choices reflect the same marginal tradeoffs as earlier in (9), except that the state contingent Bellman equations (10) reflect the flow cost $c(\sigma)$. We next assume WLOG that the uninformed player perfectly infers the chosen volatility (see the appendix). His obfuscation level cannot reflect his information, for otherwise he would immediately reveal the state, and thereby earn nothing. We thus assume that prior to learning the state, the informed player chooses a common volatility schedule $\sigma(q)$, and two state contingent intensity differentials δ_θ. Let $v(q)$ denote the flow dividend in (7), after choosing the optimal intensity differential. Then just focusing on the obfuscation choice,

$$rV(q) = v(q) + \max_{\sigma \geq \bar{\sigma}} \tfrac{1}{2} q^2(1-q)^2(\delta_1(q) - \delta_0(q))^2 V''(q)/\sigma^2 - c(\sigma) \qquad (15)$$

Zero volatility is not optimal, for then actions are observed and the value vanishes, whereas any constant and positive volatility secures a positive value. From this, we infer from (15) that the value is strictly concave always — for if $V''(q) \geq 0$, then $\sigma = 0$ maximizes (15). More subtly, volatility cannot be explode provided the average obfuscation costs $c(\sigma)/\sigma$ explode in σ.[14] The volatility choice is therefore interior, where the marginal cost balances the expected capital gains from more obfuscation:

$$c'(\sigma) = -q^2(1-q)^2(\delta_1(q) - \delta_0(q))^2 V''(q)/\sigma^3 \qquad (16)$$

Differentiating the right side of (15) a second time, we see that it is globally concave:

$$3q^2(1-q)^2(\delta_1(q) - \delta_0(q))^2 V''(q)/\sigma^4 - c''(\sigma) < 0,$$

So the obfuscation in each state exceeds the baseline $\bar{s}$, is unique, and finite, and obeys the first order condition. Now substitute the FOC into the Bellman equation (15) to

[14] The only possibility is that the intensity differentials and obfuscation jointly explode. But then flow payoffs must explode and therefore the differential to obfuscation ratio explodes. This entails dominating flow costs under our cost function assumption.

discover:

$$v(q) = rV(q) + \tfrac{1}{2}\sigma c'(\sigma) + c(\sigma) \tag{17}$$

We prove in the Appendix that the dividend less the return, $v(q) - rV(q)$, has the same shape here as in Figure 2 — specifically, concave, and peaking at $q = 1/2$. So:

Proposition 5 (Obfuscation) *When the public grows more uncertain, the informed player obfuscates more: obfuscation $\sigma(q)$ is hill-shaped, peaking at belief $q = 1/2$.*

In contrast to disinformation, obfuscation falls as the public beliefs grow more focused on either state of the world. When the informed player is most uncertain about the world, he has the greatest ex ante incentive to guard his dividend gains behind a cloud of smoke. All in all, obfuscation and disinformation are neither complements nor substitutes.

9 Conclusion

This paper has fully solved and explored a competitive model of dynamic deception. In all cases, to use information is to lose it: The informed player only earns payoffs from his edge by differentially exploiting his information in the two states. This model offers bridging insights for game theory and finance, subsuming insider trading as an extreme case. Using this as a springboard, we have economically formalized and fleshed out three different forms of deception — espionage, disinformation, and obfuscation.

A Appendix: Omitted Proofs

A.1 Belief Processes: Proof of Lemma 1

Moscarini and Smith (2001) and Back and Baruch (2004) prove a result related to Lemma 1 via filtering theory. We pursue a more accessible derivation exploiting the simple structure of the binomial world, adapted from Bolton and Harris (1999). This allows for off equilibrium path behavior.

Consider the random change in beliefs from $q(t) = q$ over $[t, t + dt]$. Let dY be the signal change observed by the uninformed player over this time interval. By Bayes' rule:

$$dq(t) \equiv q(t+dt) - q = \frac{q(1-q)(f_1(dY) - f_0(dY))}{qf_1(dY) + (1-q)f_0(dY)} \tag{18}$$

If the actual intensity differential is Δ, while the uninformed expects state contingent intensities δ_θ, then the probability "density" $f_\theta(dY)$ of increment dY is a normal r.v. with mean $\delta_\theta(q)dt$ and variance $\sigma^2 dt$. We exploit the first order approximation:

$$f_\theta(dY) \propto e^{-\frac{(dY-\delta_\theta dt)^2}{2\sigma^2 dt}} \approx e^{-\frac{1}{2}+\delta_\theta dY/\sigma^2} \propto e^{\delta_\theta dY/\sigma^2} \approx 1 + (\delta_\theta/\sigma^2)dY$$

Where the first approximation follows from $(dY)^2 = \sigma^2 dt$. Substituting back into (18):

$$\begin{aligned} dq(t) &\approx \frac{q(1-q)(\delta_1(q) - \delta_0(q))dY/\sigma^2}{1 + \delta(q)dY/\sigma^2} \\ &\approx q(1-q)(\delta_1(q) - \delta_0(q))dY/\sigma^2 \left[1 - \delta(q)dY/\sigma^2\right] \end{aligned}$$

Again using $dY = \Delta dt + \sigma dW$ and $(dY)^2 = \sigma^2 dt$, we discover:

$$dq(t) \approx q(1-q)(\delta_1(q) - \delta_0(q))\left(\Delta - \delta(q)\right)/\sigma^2 dt + q(1-q)(\delta_1(q) - \delta_0(q))/\sigma dW$$

A.2 Existence and Uniqueness: Proof of Proposition 1

We first assume short-run uninformed players, and then show that this is also the unique equilibrium when the short-run player is long-lived.

STEP 0: PROPERTIES OF THE PROBIT FUNCTION. Given the error function $\Phi(x) = (2/\sqrt{\pi})\int_0^x e^{-t^2}dt$, the Probit function is defined by $F(q) \equiv \sqrt{2}\Phi^{-1}(2q-1)$. Next, differentiating the identity $\Phi(F(q)/\sqrt{2}) \equiv 2q-1$ yields $F'(q) \equiv 2/\Phi'(F(q)/\sqrt{2})$, i.e.

$$F'(q) \equiv \sqrt{2\pi}e^{F(q)^2/2} \tag{19}$$

$$F''(q) \equiv F(q)F'(q)\sqrt{2\pi}e^{F(q)^2} = F(q)F'(q)^2 \tag{20}$$

STEP 1: SHORT-RUN PLAYERS BEST RESPOND. Given intensity differentials $\delta_0(q)$ and $\delta_1(q)$, the uninformed player's expected loss from mixture p is:

$$q\delta_1(q)(1-p) + (1-q)\delta_0(q)(-p) = q\delta_1(q) - p\delta(q).$$

Since we have derived $\delta(q) = 0$, indifference prevails.

Step 2: Differential Equations for Values and strategies. Define the value difference $\Lambda(q) \equiv V_0(q) - V_1(q)$. Since $\delta(q) = 0$, the Bellman (10) becomes:

$$r\sigma^2 V_\theta(q) = \frac{1}{2} q^2 (1-q)^2 (\delta_1(q) - \delta_0(q))^2 V_\theta''(q) \tag{21}$$

Combining the two state contingent first order conditions (9), we get:

$$\delta_1(q) - \delta_0(q) = \frac{2\xi\sigma^2}{q(1-q)\Lambda'(q)} \tag{22}$$

Define $\rho = r/(\xi^2\sigma^2)$, and substitute (9) into (21) to get:

$$\rho V_\theta(q) = 2\frac{V_\theta''(q)}{\Lambda'(q)^2} \tag{23}$$

Subtracting this at $\theta = 0, 1$ yields a second order ordinary differential equation for Λ:

$$\rho \Lambda(q) = 2\frac{\Lambda''(q)}{\Lambda'(q)^2} \tag{24}$$

Substitute (22) into the indifference equation (5) and the FOCs (9), to get:

$$(1-q)\delta_0(q) = -\frac{2\xi\sigma^2}{\Lambda'(q)} = -q\delta_1(q) \quad \text{and} \quad p(q) = \frac{1}{2} + \xi\left[\theta - \frac{1}{2} + \frac{V_\theta'(q)}{\Lambda'(q)}\right] \tag{25}$$

We finish the construction by verifying the solution to the state $\theta = 0$ differential equation (23) (the state $\theta = 1$ analysis is symmetric), the solution to (24); and the proposed $p(q)$ equation. Since we have already imposed the belief process, and the indifference conditions for each player, we are done.

Step 3: $V_0(q) = q\Lambda(q) + 2[\rho\Lambda'(q)]^{-1}$ Solves (23). Differentiating the suggested V_0 function, $V_0'(q) = q\Lambda'(q) + \Lambda(q) - 2\Lambda''(q)/(\rho\Lambda'(q)^2)$. The last two terms cancel by (24). So:

$$V_0'(q) = q\Lambda'(q). \tag{26}$$

If we differentiate (26) and divide both sides by $\Lambda'(q)^2$, we find:

$$\frac{V_0''(q)}{\Lambda'(q)^2} = \frac{q\Lambda''(q)}{\Lambda'(q)^2} + \frac{1}{\Lambda'(q)}.$$

Substitute for $\Lambda''(q)/\Lambda'(q)^2$ from (24) to get:

$$\frac{V_0''(q)}{\Lambda'(q)^2} = \frac{\rho}{2}\left[q\Lambda(q) + \frac{2}{\rho\Lambda'(q)}\right],$$

This verifies that $V_0(q) = q\Lambda(q) + 2[\rho\Lambda'(q)]^{-1}$ solves (23).

STEP 4: THE SOLUTION TO THE ODE (24). One can verify the general solution:

$$\Lambda(q) = \frac{2}{\sqrt{\rho}}\,\Phi^{-1}(c_1(q-1/2) + c_1c_2),$$

By symmetry, $V_0(1/2) = V_1(1/2)$, where $\Lambda(1/2) = 0$ and so $c_2 = 0$. To see that our proposed solution lies in this class, set $c_1 = 2$ and use the identity $F(z) = \sqrt{2}\,\Phi^{-1}(2z-1)$.

STEP 5: THE IMPLIED INFORMED PLAYER STRATEGY AND BELIEF PROCESS. As $q = V_0'(q)/\Lambda'(q)$ from (26), the mixture $p(q)$ in (25) becomes $p(q) = 1/2 + \xi(q-1/2)$. Substituting the equilibrium strategies for δ_θ into the belief diffusion (6) produces the variance coefficient $(2r)/F'(q)^2$ and the state contingent drifts:

$$\mu_0(q) = -\frac{2r}{(1-q)F'(q)^2} \quad \text{and} \quad \mu_1(q) = \frac{2r}{qF'(q)^2}.$$

The equilibrium flow payoffs in part (d) follow from (4).

STEP 6: UNIQUENESS. Let $V \equiv qV_1 + (1-q)V_0$. Let (δ, p) be an equilibrium of the game with payoffs $(V, -V)$. Assume a different equilibrium payoff $(V^*, -V^*)$ generated by some strategy (δ^*, p^*). The uninformed player does weakly better at (δ^*, p^*) than at (δ^*, p). Because this is a zero-sum game, a gain for one player is a loss for another — the informed player does weakly better at (δ^*, p) than at (δ^*, p^*). Allowing him to re-optimize, he does even better at (δ, p) than at (δ^*, p). Since $V^* \neq V$ by assumption, we must have $V > V^*$. But our starting point was arbitrary, and the same logic establishes that $V^* > V$. So $V^* = V$ for all equilibria. Because this holds for all initial beliefs q, across all sequential equilibria, there is a unique equilibrium value $V(q)$, and it is Markovian, and therefore described by our formulas. So our strategies uniquely follow from the given value functions and their derivatives.

STEP 7: A LONGRUN UNINFORMED PLAYER. Finally, if the uninformed player is long-lived, then she cannot beat a short-run uninformed player if the informed player holds to his strategy. But since she can do no worse if she mimics the short-run uninformed players, this is an equilibrium, and it is unique by the same logic.

A.3 The Ratio of Informed Values: Proof of Inequality (12)

STEP 1: AN EQUIVALENT PROBIT INEQUALITY. Assume the case $q < 1/2$. Since $V_0(q)/V_1(q) = V_0''(q)/V_1''(q)$ by (23), the desired inequality (12) holds exactly when

$$\frac{V_0''(q)}{V_1''(q)} \leq \frac{q}{1-q}$$

This is equivalent to $H(q) = (1-q)V_0''(q) - qV_1''(q) \leq 0$, and thus

$$H(q) = (1-q)[F'(q) + qF''(q)] - q[F'(1-q) + (1-q)F''(1-q)] \leq 0$$

because $V_0''(q) \equiv V_1''(1-q) \equiv F'(q) + qF''(q)$. Since $F(1-q) \equiv -F(q)$, we have $F'(1-q) \equiv F'(q)$ and $F''(1-q) \equiv -F''(q)$, and therefore we may rewrite

$$H(q) \equiv 2q(1-q)F''(q) - (2q-1)F'(q)$$

STEP 2: PROOF OF PROBIT INEQUALITY. Substituting $z \equiv F(q)/\sqrt{2}$, so that $q(z) = (1+\Phi(z))/2$, and exploiting (19) and (20):

$$H(q(z)) = \sqrt{2}e^{z^2}\left[e^{z^2}\pi z(1-\Phi(z)^2) - \sqrt{\pi}\Phi(z)\right] \equiv e^{2z^2}\sqrt{2\pi}h(z)$$

Then $H(q) \geq 0$ for all $q > 1/2$ holds precisely when $h(z) \geq 0$ for all $z > 0$, where

$$h(z) = 1 - \Phi(z)^2 - \frac{e^{-z^2}\Phi(z)}{\sqrt{\pi}z}$$

Now, $h(0+) = 1-2/\pi > 0$ (by l'Hopital's rule) and $\lim_{z\to\infty} h(z) = 0$. Since h is smooth, if by contradiction $h(z) < 0$ for some $z > 0$, then there exists a local extremum, where

$$0 = h'(z) \equiv -\frac{4}{\sqrt{\pi}} e^{-z^2}\Phi(z) - \frac{2}{\pi z} e^{-2z^2} + \frac{e^{-z^2}\Phi(z)}{\sqrt{\pi} z^2} + 2\frac{e^{-z^2}\Phi(z)}{\sqrt{\pi}} = 0$$

for some $z > 0$. This equality requires that $\Phi(z) = \chi(z)$, where

$$\chi(z) \equiv \frac{2z/\sqrt{\pi}}{1-2z^2} e^{-z^2}$$

But $\Phi(z) = \chi(z)$ is impossible at any $z > 0$, since $\Phi(0) = 0 = \chi(0)$, and for all $z > 0$,

$$\sqrt{\pi}\chi'(z)/2 = \frac{e^{-z^2}}{1-2z^2} - \frac{2z^2 e^{-z^2}}{1-2z^2} + \frac{4z^2 e^{-z^2}}{(1-2z^2)^2} = e^{-z^2}\frac{1+4z^2}{(1-2z^2)^2} > e^{-z^2} = \sqrt{\pi}\Phi'(z)/2$$

A.4 Focused Public Beliefs: Proof of Corollary 3

STEP 1: STRATEGIES. By Dominici (2003), the asymptotic behavior of $F(q)$ is:[15]

$$F(q) \sim \begin{cases} -\sqrt{-\log(2\pi q^2) - \log(-\log(2\pi q^2))} & q \text{ near } 0 \\ \sqrt{-\log(2\pi(1-q)^2) - \log(-\log(2\pi(1-q)^2))} & q \text{ near } 1 \end{cases}$$

These expressions imply the asymptotic behavior of $F'(q)^{-1} = e^{-\frac{1}{2}F(q)^2}/\sqrt{2\pi}$:

$$F'(q)^{-1} \sim \begin{cases} 2q\sqrt{-\pi \log q} & q \text{ near } 0 \\ 2(1-q)\sqrt{-\pi\log(1-q)} & q \text{ near } 1 \end{cases}$$

This immediately yields the resulting approximations in (a) and (b) for strategies.

STEP 2: VALUES. The asymptotic behavior of V follows directly from that of $F'(q)^{-1} \to 0$ as $q \to 1$. Finally, we consider $V_0(q)$ as $q \to 0$. By (26), we have $V_0'(q) \propto \left(\sqrt{-\log q}\right)^{-1}$ as $q \to 0$. Integrating this, we find:

$$V_0(q) \propto \int_0^q (\sqrt{-\log s})^{-1} ds = \sqrt{\pi}\left(1 - \text{erf}\left(\sqrt{-\log q}\right)\right)$$

[15] We say that "$A(q)$ behaves like $B(q)$" for $q \approx 0$ or $q \approx 1$ if their ratio tends to 1 near 0 or 1. We write this as $A \sim B$. When we say $A \propto B$, then we mean that their ratio tends to a positive constant.

This behaves like $q/\sqrt{-\log q}$, for by l'Hopital's rule, $\Phi'(t)=(2/\sqrt{\pi})e^{-t^2}$, and $\Phi(\infty)=1$:

$$\sqrt{\pi}\lim_{q\to 0}\frac{\sqrt{-\log q}\left(1-\Phi\left(\sqrt{-\log q}\right)\right)}{q}=\sqrt{\pi}\lim_{t\to\infty}\frac{(1-\Phi(t))t}{e^{-t^2}}=1$$

A.5 Learning the Truth in Infinite Expected Time

Lemma 2 (Beliefs) *Beliefs converge to the truth, but in infinite expected time.*

STEP 1: DEFINING THE "SCALE FUNCTIONS". We just consider state $\theta = 0$. As in §15.3 in KT81, define the "scale functions" s_0 and S_0 as follows:

$$s_0(x)\equiv e^{-2\int_0^x \frac{\mu_0(q)dq}{\varsigma^2(q)}}=e^{2\int_0^x \frac{dq}{1-q}}=(1-x)^{-2}\quad\text{and}\quad S_0(x)\equiv\int_0^x\frac{dy}{(1-y)^2}=\frac{1}{1-x}$$

STEP 2: BELIEFS CONVERGE TO THE TRUTH. We satisfy Lemma 6.1 of KT81, that if $S_0(b)-\lim_{\epsilon\to 0}S_0(\epsilon)<\infty$ for some $0<b<1$ then the truth is revealed in the limit.

STEP 3: THE TRUTH IS NOT REVEALED IN FINITE EXPECTED TIME. By Lemma 6.2 in §15.6 of KT81, the true state is not revealed in finite time if for some $0<c<1$:

$$\Sigma(c)\equiv\lim_{b\to 0}\int_b^c\left[\int_y^c[q^2(1-q)^2(\delta_1(q)-\delta_0(q))^2 s_0(q)]^{-1}dq\right]s_0(y)dy=\infty.$$

Using the expressions from Proposition 1 for the equilibrium intensity differentials δ_θ:

$$\Sigma(c)=\lim_{b\to 0}\int_b^c\left[\int_y^c(1-q)^2\Lambda'(q)^2dq\right](1-y)^{-2}dy\geq(1-c)^2\lim_{b\to 0}\int_b^c\left[\int_y^c\Lambda'(q)^2dq\right]dy$$

Next, by the asymptotic expressions in our Appendix A.4, for small q we have:

$$\Lambda'(q)^2\propto F'(q)^2\sim-\left[q^2\log q\right]^{-1}$$

and so $\Lambda'(q)^2>q^{-2}$ for all small enough $q>0$, say $q<c$. For such q, we have

$$\Sigma(c)\geq\lim_{b\to 0}\int_b^c\left[\int_y^c q^{-2}dq\right]dy=\lim_{b\to 0}(-\log b)=\infty$$

A.6 Diminishing Time Returns: Proof of Proposition 2

Step 1: A Critical Ratio of Derivatives. Transform the ODE defining V, namely, $rV(q) = v(p(q)) + \frac{1}{2}\varsigma^2(q)V''(q)$, using $rV(q) = v(p(q))/2$ (Corollary 1), to get:

$$-rV(q) = \tfrac{1}{2}\varsigma^2(q)V''(q)$$

Dividing this by (14) we find:

$$\frac{1}{rV(q)} = \frac{T''(q)}{V''(q)} \tag{27}$$

Step 2: Time is Money at $q = 1/2$. Define $\nu(q) = V(q)/V(1/2)$ and $\top(\nu(q)) = T(q)$. Totally differentiating the second identity and rearranging, we find:

$$r\top'(\nu) = r\frac{T'(q)}{V'(q)}V(1/2) \tag{28}$$

Since $T'(1/2) = V'(1/2) = 0$, we evaluate this limit using (27):

$$r\top'(1) = r\lim_{q\to 1/2}\frac{T'(q)}{V'(q)}V(1/2) = r\lim_{q\to 1/2}\frac{T''(q)}{V''(q)}V(1/2) = 1$$

Step 3: $\top$ is concave. Assume WLOG that $q > 1/2$. Differentiate (28) to get:

$$\top''(\nu) = V(1/2)\frac{T''(q) - T'(q)V''(q)/V'(q)}{V'(q)^2} \tag{29}$$

and so, $\top'' \le 0$ iff:

$$\frac{T''(q)}{V''(q)} > \frac{T'(q)}{V'(q)}$$

Using $T'(1/2) = 0$ and then (27) we have:[16]

$$T'(q) = \int_{\frac{1}{2}}^{q} T''(s)ds = \int_{\frac{1}{2}}^{q}\frac{V''(s)}{rV(s)}ds > \int_{\frac{1}{2}}^{q}\frac{V''(s)}{rV(q)}ds = \frac{V'(q)}{rV(q)}$$

[16] Using our solution and (28), convert (29) to an ODE in $\top'$ and $\top''$, and then a precise formula:

$$r\top(\nu) = \log(\nu)\sum_{n=0}^{\infty}\frac{b_n\left[\log(\nu)\right]^n}{n!} + \text{constant} \qquad \text{and} \qquad r\top'(\nu) = \frac{\sqrt{\pi}\,\text{erfi}(\sqrt{-\log(\nu)})}{2\sqrt{-\log(\nu)}}$$

where $b_{n+1}/b_n = (n+1)^2/[(n+3/2)(n+2)]$, and $b_0 = 1$. (Details upon request.)

where the strict inequality owes to $V'' < 0$ and $0 < V(q) < V(s)$ for $1/2 \le s \le q$, and the final equality follows from $V'(1/2) = 0$. Since, $V'(q) < 0$ for $q > 1/2$ we conclude using by (27):

$$\frac{T'(q)}{V'(q)} < \frac{1}{rV(q)} = \frac{T''(q)}{V''(q)}$$

A.7 The Shrinking Firm: Proof of Corollary 4

Define $\mathcal{V}(q,t) \equiv E[V(Q(t)) \mid Q(0) = q]$. By the Feynman-Kac equation ((5.4) on p. 214 of Karlin and Taylor (1981), or KT81), the time return when forecasting is just the expected capital gain, since there are no dividends:

$$\mathcal{V}_t(q,t) = \frac{1}{2}q^2(1-q)^2\mathcal{V}_{qq}(q,t) = \frac{1}{2}q^2(1-q)^2\frac{\partial^2}{\partial q^2}E[V(Q(t))|Q(0) = q] \tag{30}$$

Focus on the inner expectation. Since expected capital losses are minus the return,

$$E[V(Q(t+dt)) \mid q] = E[V(Q) + dV(Q) \mid q] = E[V(Q) - rV(Q)dt \mid q] \tag{31}$$

We can now subtract the expression (30) at $t + dt$ and t, getting $\mathcal{V}_{t+dt} - \mathcal{V}_t = -r\mathcal{V}_t dt$, or $\mathcal{V}_{tt} = -r\mathcal{V}_t$. In other words, $\mathcal{V}_t(q,t) = e^{-rt}\mathcal{V}_t(q,0)$, as desired.

A.8 Value of Information: Proof of Proposition 3

Step 1: Initial Marginal Value of Information. As we proved in A.7, the expected terminal reward obeys the Feynman-Kac equation in "time" τ:

$$\mathcal{V}_\tau(q,\tau) \equiv \frac{1}{2}q^2(1-q)^2\mathcal{V}_{qq}(q,\tau) \tag{32}$$

with boundary condition $\mathcal{V}(q,0) = -V(q)$. So $\mathcal{V}_\tau(q,0+) = -\frac{1}{2}q^2(1-q)^2V''(q) > 0$. Twice differentiate the boundary condition $\mathcal{V}_{qq}(q,0) = V''(q)$, and put $\tau = 0$ into (32).

Step 2: Re-Expressing the Marginal Value of Information. Proceeding as in (31), let the time increment $d\tau$ lead to a stochastic increment dQ in beliefs. Then:

$$\mathcal{V}_\tau d\tau = E_q[V(Q(\tau + d\tau))] - E_q[V(Q)] = E_q[V'(Q)(dQ) + V''(Q)(dQ)^2/2] \tag{33}$$

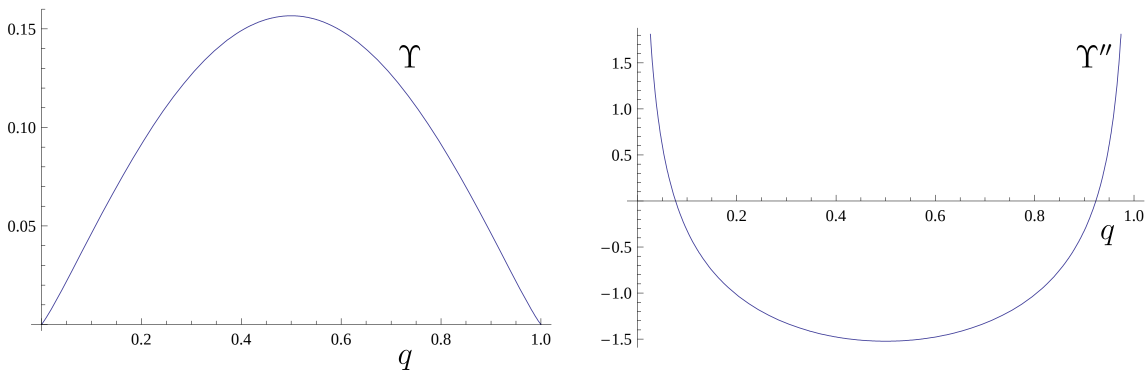

Figure 8: **Quasi-Concavity of $\Upsilon(q)$ and Graph of $\Upsilon''(q)$.** Here we use the known probit function, and plot the derived function $\Upsilon(q)$ and its second derivative $\Upsilon''(q)$.

Now, $E_q[dQ] = 0$ and $(dQ)^2 = Q^2(1-Q)^2 d\tau$. And by (20), the second derivative of $V(Q) = k/F'(Q)$ is $V''(Q) = kF'(Q)$. If we put $\Upsilon(q) \equiv q^2(1-q)^2 F'(q)$, then (33) gives:

$$\mathcal{V}_\tau(q) = kE_q[\Upsilon(Q)] \tag{34}$$

Step 3: The Marginal Value is Quasi-concave and Eventually Vanishes. By (34), the marginal value of information equals $\mathcal{V}_\tau(q,\tau) = \int_0^1 \Upsilon(x) g(x,\tau,q)dx$, where $g(x,\tau,q)$ is the density over public beliefs $q(\tau) = x$ at time τ given prior $q(0) = q$. As computed in Keppo, Moscarini, and Smith (2008):

$$g(x,\tau,q) = \sqrt{\frac{q(1-q)}{x^3(1-x)^3 2\pi\tau}} \exp\left(-\frac{1}{8}\tau - \frac{(\log[x/(1-x)] - \log[q/(1-q)])^2}{2\tau}\right)$$

We now exploit a theory unique to Karlin (1968). By his Theorem 1.2.1,[17] the density g is *totally positive* (TP) of any order jointly in (q,τ). By the "variation diminishing property" (see Theorem 5.3.1), $\mathcal{V}_\tau$ is quasi-concave in τ as $\Upsilon(q)$ is quasi-concave in q, as seen in Figure 8. Since beliefs converge to the truth in the limit, and $\Upsilon(0) = \Upsilon(1) = 0$, we have:

$$\lim_{\tau\to\infty} \mathcal{V}_\tau(q,\tau) = q\Upsilon(1) + (1-q)\Upsilon(0) = 0$$

Step 4: The Initial Slope of the Marginal Value of Information. By Steps 1 and 2, $\mathcal{V}_\tau(q, 0+) = -q^2(1-q)^2 V''(q)/2 = \Upsilon(q)$. As this holds for all q, we may

[17] This says that any function like ours of the form $L(x,y) = h_1(x)h_2(y)\exp[h_3(x)h_4(y)]$ is TP.

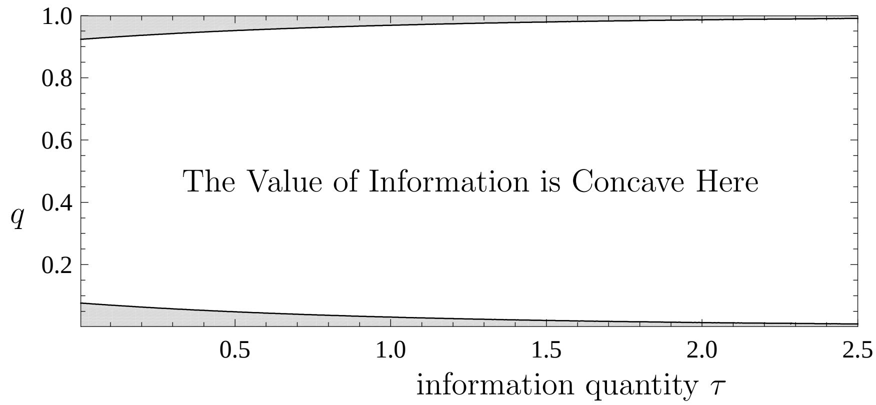

Figure 9: **Concavity of Value of Information.** This depicts the small regions with extreme public beliefs in q-τ space where the marginal value of information is rising.

twice differentiate both sides to get: $\mathcal{V}_{\tau qq}(q,0+) = \Upsilon''(q)$. Differentiating (32), we find:

$$\mathcal{V}_{\tau\tau}(q,\tau) = \frac{1}{2}q^2(1-q)^2\mathcal{V}_{qq\tau}(q,\tau) \quad \Rightarrow \quad \mathcal{V}_{\tau\tau}(q,0+) = \frac{1}{2}q^2(1-q)^2\Upsilon''(q)$$

Figure 8 reveals that $\Upsilon'' < 0$ inside an interval $(q^*, 1-q^*)$ and $\Upsilon'' > 0$ otherwise. Numerically evaluating the crossing point yields $q^* \approx 0.077$.

Step 5: Putting it Together. The marginal value of information is quasi-concave, and so cannot rise after falling. When $q \in (q^*, 1-q^*)$, the marginal value falls for all τ. Since the marginal value eventually tends to zero, when $q^* \notin (q^*, 1-q^*)$, the marginal value rises initially, but eventually crosses a threshold in τ and falls ever after.[18]

A.9 The Log Likelihood as a Cardinal Information Measure

In light of equation (6), the informed player can induce the public belief evolution $dq = q^2(1-q)^2(dt + dZ)$ in state $\theta = 0$, for Wiener process Z. This belief process has a nonlinear drift and volatility. But, by Ito's Lemma, we have a simple law of motion $d\ell = [\mu_0(q)\ell'(q) + \frac{1}{2}\varsigma^2(q)\ell''(q)]dt + \varsigma(q)\ell'(q)dZ = -\frac{1}{2}dt + dZ$.

[18] For the purposes of numerically evaluating $\mathcal{V}_{\tau\tau}$, we proceed as in Step 2:

$$\mathcal{V}_{\tau\tau} = E_q[\Upsilon(Q+dQ) - \Upsilon(Q)] = E_q[(dQ)^2\Upsilon''(Q)] = \int_0^1 x^2(1-x)^2\Upsilon''(x)g(x,\tau,q)dx$$

This function was used to produce the contour plot in Figure 9.

A.10 The Value of Disinformation: Proof of Proposition 4

Twice differentiate $\bar{V}_0$ to discover:

$$\bar{V}_0''(\ell) = \frac{\sqrt{\pi}e^{2\ell+\xi(\ell)^2}}{\sqrt{\rho}(1+e^\ell)^5}\left[2+e^\ell-e^{2\ell}+2\sqrt{\pi}\xi(\ell)e^{\ell+\xi(\ell)^2}\right] \equiv \frac{\sqrt{\pi}e^{2\ell+\xi(\ell)^2}}{\sqrt{\rho}(1+e^\ell)^5}h(\ell).$$

where $\xi(\ell) = \Phi^{-1}(\tanh(\ell/2))$. So $\bar{V}_0''(\ell)$ has the same sign as h. Replacing $\ell \mapsto \ell(q)$,

$$h(\ell(q)) = (1-q)^{-2}\left(2-3q+q(1-q)\sqrt{2\pi}F(q)e^{\frac{1}{2}F(q)^2}\right) \equiv (1-q)^{-2}g(q).$$

The asymptotic properties of $F(q)$ yield $g(0) = 1$ and $g(1-\eta) < 0$ for small η.

We claim that g is concave on $q < 1/2$ and convex on $q > 1/2$. Indeed, use (19) to express $g''(q)$ as a function of F and q alone, and then appeal to $F(q) \gtrless 0$ for $q \gtrless 1/2$. The function g is continuous, and must cross 0 at a unique point. We can verify numerically that this occurs at $q \approx 0.864 \Leftrightarrow \ell \approx 1.845$.

A.11 Uncertain Obfuscation is Inferred at Once

Let the uninformed player observe a signal process like (3) with an unknown volatility σ over any finite interval $[a, b]$ of time dt. She may partition this interval into n equal segments of width $(b-a)/n$, and let $X_i^{(n)}$ denote n times the ith squared Wiener increment. As a squared normal $N(0, \sigma^2)$ random variable, each $X_i^{(n)}$ is an independent draw from a χ-squared distribution with mean σ^2 and variance $2\sigma^2$. By a version of the strong law of large numbers for a discounted finite variance, the average of the n terms $X_i^{(n)}$ almost surely converges to σ^2 as letting $n \to \infty$.

A.12 Unfocused Obfuscation: Proof of Proposition 5

Differentiating (17), the derivative $\sigma'(q)$ shares the sign of $v'(q) - rV'(q)$:

$$\sigma'(q) = \frac{2\left[v'(q) - rV'(q)\right]}{\sigma c''(\sigma) + 3c'(\sigma)}$$

Easily, $v(q) - rV(q) \geq 0$ is symmetric about $q = 1/2$, and satisfies the boundary conditions $v(0) - rV(0) = v(1) - rV(1) = 0$. So if $v(q) - rV(q)$ is concave, then

$\sigma'(q) \gtrless 0$ for $q \lessgtr 1/2$. To establish concavity, we modify (23) for the flow volatility cost:

$$\rho V_\theta(q) + c(\sigma) = 2\frac{V_\theta''(q)}{\Lambda'(q)^2} \tag{35}$$

The difference $rV(q) - v(q)$ is convex in q if the objective function of the Bellman equation (15) is. Using (22), we can simplify the objective function as follows:

$$q^2(1-q)^2(\delta_1(q) - \delta_0(q))^2 V''(q)/\sigma^2 = 2\xi^2\sigma^2 \frac{V''(q)}{\Lambda'(q)^2}.$$

To see that this is convex in q, twice differentiate $V(q) = qV_1(q) + (1-q)V_0(q)$:

$$V''(q) = qV_1''(q)+(1-q)V_0''(q)+2(V_1'(q)-V_0'(q)) = qV_1''(q)+(1-q)V_0''(q)-2\Lambda'(q) \tag{36}$$

since $\Lambda(q) \equiv V_0(q) - V_1(q)$, by definition. Thus, equation (35) yields:

$$\frac{2V''(q)}{\Lambda'(q)^2} \quad = \quad 2q\frac{V_1''(q)}{\Lambda'(q)^2} + 2(1-q)\frac{V_0''(q)}{\Lambda'(q)^2} - 4/\Lambda'(q) = \rho V(q) - 4/\Lambda'(q) + c(\sigma)$$

Twice differentiate the RHS/ρ in q, and conclude (24) and (35) and (36), to get:

$$\left(V'(q) + 4\frac{\Lambda''(q)}{\rho\Lambda'(q)^2}\right)' = (V'(q) + 2\Lambda(q))' = V''(q) + 2\Lambda'(q) = qV_1''(q) + (1-q)V_0''(q) > 0$$

References

Back, K., and S. Baruch (2004): "Information in Securities Markets: Kyle Meets Glosten and Milgrom," *Econometrica*, 72(2), 433–465.

Bassan, B., O. Gossner, M. Scarsini, and S. Zamir (2003): "Positive Value of Information in Games," *International Journal of Game Theory*, 32, 17–31.

Bolton, P., and C. Harris (1999): "Strategic Experimentation," *Econometrica*, 67(2), 349–74.

Caldentey, R., and E. Stacchetti (2007): "Insider Trading With Stochastic Valuation," SSRN working paper.

Chade, H., and E. Schlee (2002): "Another Look at the Radner-Stiglitz Nonconcavity in the Value of Information," *Journal of Economic Theory*, 107, 421–452.

Crawford, V. P. (2003): "Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions," *American Economic Review*, 93(1), 133–149.

Dominici, D. (2003): "The Inverse of the Cumulative Standard Normal Probability Function," arXiv Working Paper Cornell University.

Faingold, E., and Y. Sannikov (2007): "Reputation Effects and Degenerate Equilibria in Continuous-Time Games," August mimeo.

Hendricks, K., and R. P. McAfee (2006): "Feints," *Journal of Economics and Management Strategy*, 15(2), 431–456.

Hirshleifer, J. (1971): "The Private and Social Value of Information and the Reward to Inventive Activity," *The American Economic Review*, 61, 561–574.

Karlin, S. (1968): *Total Positivity*. Stanford University Press.

Karlin, S., and H. Taylor (1981): *A Second Course in Stochastic Processes*. Academic Press, San Diego, CA.

Keppo, J., G. Moscarini, and L. Smith (2008): "The Demand for Information: More Heat than Light," *Journal of Economic Theory*, 138(1), 21–50.

Kyle, A. S. (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53(6), 1315–1335.

Montagu, E. (1953): *The Man Who Never Was: World War II's Boldest Counter-Intelligence Operation*. Naval Institute Press.

Moscarini, G., and L. Smith (2001): "The Optimal Level of Experimentation," *Econometrica*, 69(6), 1629–1644.

Radner, R., and J. Stiglitz (1984): "A Nonconcavity in the Value of Information," in *Bayesian Models in Economic Theory*, ed. by M. Boyer, and R. Kihlstrom, pp. 33–52. North-Holland.

Schlee, E. E. (2001): "The Value of Information in Efficient Risk Sharing Arrangements," *The American Economic Review*, 91(3), 509–524.