Dynamic Deception

Axel Anderson   Lones Smith

Georgetown   Wisconsin

Summer 2013

A Framework for Deception

The most valuable commodity I know of is information.
— Gordon Gekko ("Wall Street", 1987)

- Two sides play a competitive game over time.
- One player knows the "state of the world". The other player / sequence of players — the "public" — does not.
- To profit from his informational advantage, the informed player must condition his actions on it.
- But acting in accordance with his information reveals it to the other side: "'use it and lose it"
- Fundamental tradeoff: extracting value today vs. eroding your informational edge tomorrow.

Related Work

- Repeated, Constant Sum Games: Aumann and Maschler (1966):
  - Infinite horizon repeated game with no discounting.
  - One informed and one uninformed player.
- Finance: Informed trade by insiders: Kyle (1985), Glosten and Milgrom (1985), Back and Baruch (2004).
- Dynamic Models of Reputation: Following Selten. Kreps and Wilson (1982), Milgrom and Roberts (1982), Fudenberg and Levine (1992), Cripps Mailath Samuelson (2004), Faingold and Sannikov (2007).
- We build a simple bridge between the finance and game theoretic models of reputation, and draw new conclusions for dynamic behavior and values.

Stage Game

	state $\theta = 0$	
	a	b
A	$-1-\xi$	$1-\xi$
B	$1+\xi$	$\xi-1$

	state $\theta = 1$	
	a	b
A	$\xi-1$	$1+\xi$
B	$1-\xi$	$-1-\xi$

- Underlying Competitive Structure: Matching pennies.
- The state θ is known only to row.
- The *information edge* $\xi > 0$.
- If $\xi > 1$, then row has a dominant strategy in each state.

- Sports Example: Penalty Kicks (kicker vs. goalie)
- War Example: D-Day invasion in Normandy or Calais

One-Shot Insider Trading Interpretation: $\xi = 1$

state $\theta = 0$

	$\pi = 1$	$\pi = 0$
buy	-1	0
sell	1	0

state $\theta = 1$

	$\pi = 1$	$\pi = 0$
buy	0	1
sell	0	-1

- An asset has values 0 and 1 in states 0 and 1
- The insider chooses to buy or sell a unit.
- The uninformed mixes between $\pi = 0$ or $\pi = 1$ — the mixture p is a relative *price* on the informed actions.
- This is the knife-edged case of our model in which row has a weakly dominant strategies in each state ($\xi = 1$).

Infinitely Repeated, Undiscounted Game with Observed Actions

	a	b
A	$-1+\xi(2q-1)$	$1+\xi(2q-1)$
B	$1+\xi(1-2q)$	$-1+\xi(1-2q)$

- Aumann-Maschler-Stearns (1960s) introduce the one-shot game above, i.e. assuming symmetric common knowledge belief q that $\theta = 1$.
- When the value of this one-shot game is concave in q, then the infinitely repeated game has the same value.
- Our simple symmetric game has a constant value $\Rightarrow$ no benefit to private information in the infinitely repeated undiscounted game. In every period:
  - Uninformed Mixture: $p(q) = 1/2 + \xi(q - 1/2)$
  - Informed Mixture: $1/2$.

Timing and States

- Time: Continuous on $[0, \infty)$.
- Discount rate $r = i + \phi$, where
  - $i > 0$ is the bank interest rate / impatience
  - $\phi \geq 0$ is the constant rate of exogenous stochastic ending (eg. market closure)
- State: $\theta \in \{0, 1\}$ fixed for all time.
  - The informed player knows θ.
  - An uninformed player (or sequence of players) has at any time a *public belief* $q = Pr\{\theta = 1\}$.

Flow Payoffs with Intensities

state $\theta = 0$

	a	b
A	$-1-\xi$	$1-\xi$
B	$1+\xi$	$\xi-1$

state $\theta = 1$

	a	b
A	$\xi-1$	$1+\xi$
B	$1-\xi$	$-1-\xi$

- The uninformed (column) player chooses $p(t) \in [0, 1]$.
- The informed player chooses activities A and B with intensities $\alpha(t) \in [0, M]$ and $\beta(t) \in [0, M]$.
- State contingent flow payoffs for the informed:

$$u_0(p) \equiv (\alpha - \beta)(1 - 2p - \xi) \quad \text{State:} \quad \theta = 0$$
$$u_1(p) \equiv (\alpha - \beta)(1 - 2p + \xi) \quad \text{State:} \quad \theta = 1$$

- Payoffs only depend on the *intensity difference* $\Delta \equiv \alpha - \beta \in [-M, M]$.

Imperfectly Observed Actions (Gaussian Noise)

- If there is a sequence of uninformed players, then each cannot observe prior payoffs — eg. finance, herding.
- If there is a single uninformed player, then she cannot observe her payoff until the game ends.
- Commonly Observed Signal: $dY = \Delta\, dt + \sigma\, dW$ (where W is Weiner noise)
- Insider trading assumption: only net orders (buys minus sells) observed.
- Limit of finite signal garbling when α actions misinterpreted as β at the same rate as the opposite.
- An impatient informed player is tempted to "chisel away" his information advantage for myopic short term gains.

Game Theory Finance Bridge

Our model compared to AMS:

- We move to continuous time and add discounting and noise to a game in their concealing class.
- We have generalized strategies to allow for an intensity interpretation (subsuming mixed strategies, $M = 1$).

Our model compared to insider trading (Back and Baruch):

- We constrain the intensity of the informed “trader”
- We allow for any information edge $\xi \in (0, 1]$ (vs. $\xi = 1$).
- Surface difference?
  - Us: observational noise. Them: noise traders.
  - Us: uninformed player forced to play. Them: profit maximizing market maker with free entry.

Public Belief Evolution via Bayes Rule

- Beliefs will be our state variable $\Rightarrow$ we track them.
- Over a small *dt* interval of time:

$$q(t+dt) = \frac{q(t)P(dY|1)}{q(t)P(dY|1) + (1-q(t))P(dY|0)}$$

- $P(dY|\theta)$ depends on public's expectation of intensity δ_θ.
- dY depends on the actual intensity differential Δ.
- Drift: $\mu(\Delta, q) = q(1-q)\left(\delta_1(q) - \delta_0(q)\right)\left(\Delta - E[\delta]\right)/\sigma^2$
  - valid for Δ in or out of equilibrium
  - drift is linear in Δ.
  - Public thinks $E[dq] = 0$ in equilibrium
- Variance: $\varsigma^2(\Delta, q) = q^2(1-q)^2\left(\delta_1(q) - \delta_0(q)\right)^2/\sigma^2$

Optimality for a Myopic Uninformed Player

- The uninformed has no impact on the belief evolution.
- Thus, she myopically best responds to q and δ_θ.
- Her expected flow loss at each instant is:

$$v(q) = q\delta_1(q)u_1(p(q)) + (1 - q)\delta_0(q)u_0(p(q))$$

- Since this is linear in p, we have:

$$\text{Indifference:} \quad \delta(q) = q\delta_1(q) + (1 - q)\delta_0(q) = 0.$$

- Valid for both short and long run players. Thus, our equilibrium captures both cases.

Optimality for the Patient Informed Player

- Return on information = dividend + E[capital gain].

$$rV_\theta(q) = \max_{\Delta} \Delta u_\theta(p) + \mu(\Delta, q) V'_\theta(q) + \frac{1}{2}\varsigma^2(\Delta, q) V''_\theta(q)$$

- Informed player balances flow rewards and costs.
  - Flow Benefit: Δu_θ (value extracted today)
  - Flow Cost: $-\mu_\Delta V'_\theta$ (edge eroded tomorrow)

- When unconstrained, these must exactly balance:

$$u_\theta(p) = -q(1-q)(\delta_1(q) - \delta_0(q)) V'_\theta(q)/\sigma^2$$

Unique Markov Equilibrium

- We construct the unique Markov (in beliefs q) Equilibrium.
- Qualitative behavior is determined largely by the deception parameter $\psi \equiv r\sigma^2/M^2$.

High Deception

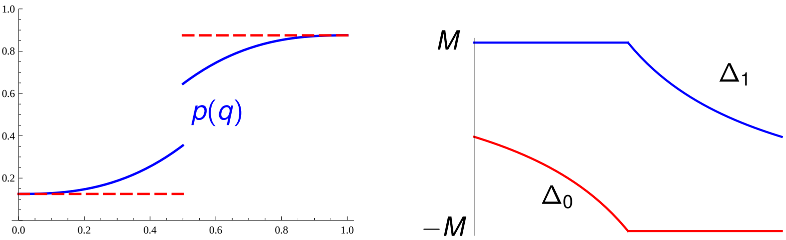

- When $\psi \equiv r\sigma^2/M^2 > 1$ the uninformed "price" $p(q)$:
  - Is convex for $q < 1/2$ and concave for $q > 1/2$.
  - Jumps up at $q = 1/2$
  - As $\psi \to 0$, $p(q)$ converges to the one-shot asymmetric solution (dashed red).
- For all deception parameters $\psi > 1$ the informed acts exactly as in the one-shot asymmetric info game.

Low Deception

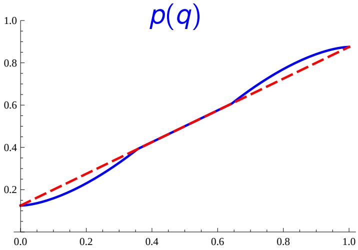

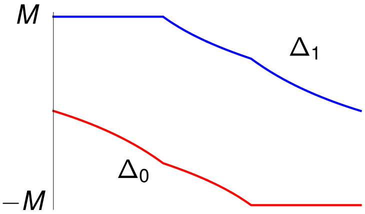

- When $\psi < 1$, there is a confounding region, on which:
  - The uninformed behaves as in AMS. (dashed red)
  - The intensity constraint does not bind on either type.
  - The informed behaves as in insider training models.
- Outside of this region:
  - The uninformed shades toward the one-shot price.
  - The informed behaves as in the one-shot game.

Convergence to AMS Solution

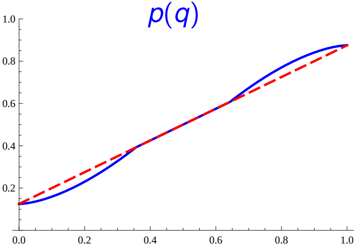

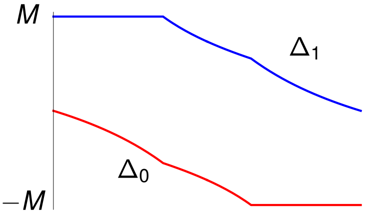

- The confounding region is symmetric $[q^*(\psi), 1 - q^*(\psi)$.
- The cutoff q^* is monotonic in ψ with $\lim_{\psi \to 0} q^*(\psi) = 0$.
- Thus, p converges to the AMS solution as $\psi \to 0$.
- The informed limit differs for $r\sigma^2 \to 0$ and $M^2 \to \infty$.
- The AMS limit ($\Delta_\theta = 0$) corresponds to $r\sigma^2 \to 0$.

Convergence to Insider Tradiing

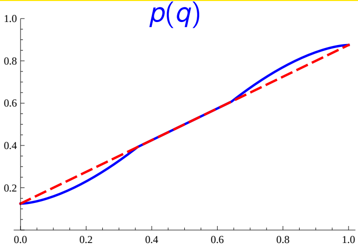

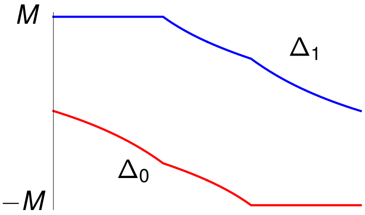

- The market price in Back and Baruch (2004) corresponds to the AMS price with $\xi = 1$.
- Thus, when $\xi = 1$ the informed p converges to the B& B price as $\psi \to 0$.
- As $M \to \infty$, the informed intensities converges to the insider trading strategies in B& B as well.
- Altogether, insider trading is a special case of our model when $\xi = 1$ and $M = \infty$.

Uninformed Bias and Mean Reversion

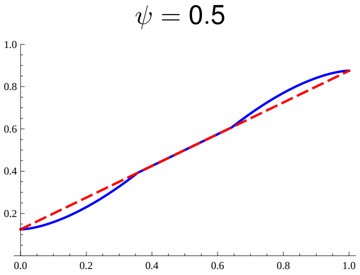

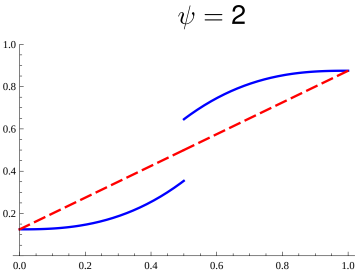

- Generally, p (blue) is biased toward the likely state relative to the AMS solution (dashed red).
- This cross sectional bias has a time series implication:
  - The uninformed "price" $p(q)$ mean reverts
  - Realized actions display negative serial correlation

Uninformed Bias Intuition

- Why does the constraint cause a bias in $p(q)$?
- For low q, the constraints binds in state 1 but not 0:
  - The uninformed expects a lower net intensity $(\delta_1 - \delta_0)$.
  - Thus beliefs are less sensitive to actions.
  - Decreasing Δ in state 0 has a smaller impact on q.
  - To maintain the FOC in State 0: the flow benefit to $\Delta < 0$ (sales) must fall $\Rightarrow$ the "price" p must fall.
- Symmetric reasoning holds for $q > q^*$, save $\uparrow p$ to lower the flow benefit to $\Delta > 0$ (buys).

- As the time $t \to \infty$:
  - Public beliefs converge to the truth.
  - But the truth is not revealed in finite expected time.
- How quickly does the informed monetize his informational advantage?
- Set $\nu(q) \equiv V(q)/V(1/2)$, i.e. the fraction of peak expected value.
- Let $\top_\varepsilon(v)$ be the expected time until $\nu(q)$ falls to $\varepsilon < v$ starting at v.

Time and Money

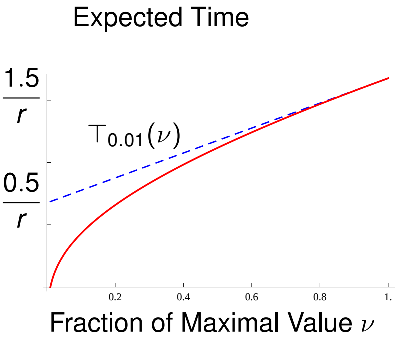

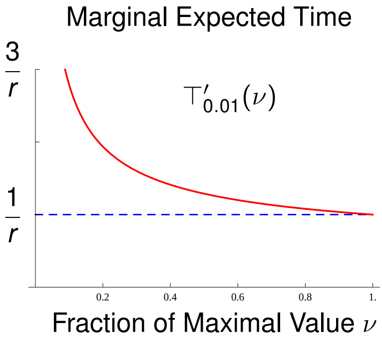

- Diminishing returns to time: Informed burns his information rent faster the higher it is.
- Intuition: He faces increasingly worse terms of trade, as public catches on.
- "Time is money:" $r\top'_{\varepsilon}(\nu) = 1$ when $\psi < 1$.

Application: The Market Value of Information

- Allow the public at any moment access to an alternative Gaussian information source.
- Measure information by "metaphorical time" τ, the length of "time" he sees this Gaussian process.
- Assume this observation process is public information.
- If the uninformed player ends with the random posterior belief $Q(\tau)$, then she earns "terminal reward" $V(Q(\tau))$.
- $\mathcal{V}(q,\tau) \equiv E[V(Q(\tau))|Q(0)=q]$
- The value of market information is the reduction $V(q) - \mathcal{V}(q,\tau)$ in the uninformed player's expected loss.
- Standard decision theory result: Marginal value of information initially rises and then falls.

Marginal Value of Market Information

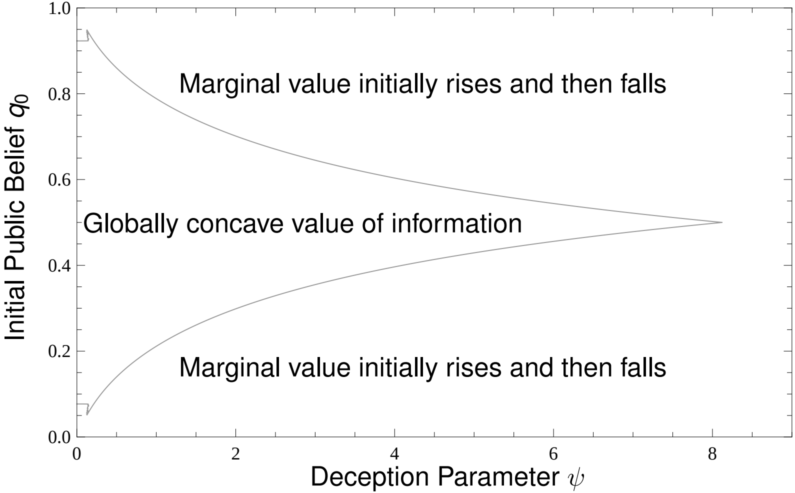

- When $\psi \equiv r\sigma^2/M^2$ is low and the initial belief q_0 interior, the market value of information is globally concave.

Obfuscation by the Informed

- What if the informed player can obscure his actions?
- Curiously, the informed player does not always benefit from increased noise, σ:

  The conditional value V_θ falls in observational noise σ for sufficiently accurate public beliefs.
- Since state contingent noise instantaneously reveals the state, we explore pooling equilibria.
- We allow the informed player to maintain unconditional noise σ, at flow cost $c(\sigma)$.
- Noise below $\underline{\sigma} \geq 0$ is free, $c'(\underline{\sigma}) = 0$, while above $\underline{\sigma}$, c is smooth, strictly increasing, and convex, with $c(\sigma)/\sigma$ unbounded.

The Optimal Level of Obfuscation

- New Bellman equation:

$$rW(q) = w(q) + \max_{\sigma \geq \bar{\sigma}} \tfrac{1}{2} q^2 (1-q)^2 (\delta_1(q) - \delta_0(q))^2 W''(q)/\sigma^2 - c(\sigma)$$

- Under our assumptions this problem is globally concave, with an interior solution satisfying the FOC.
- Combining the FOC and the Bellman Equation we find.

$$w(q) - rW(q) = \tfrac{1}{2} \sigma c'(\sigma) + c(\sigma)$$

For large intensity bounds M:

- The informed obfuscates less as the public grows more certain: $\sigma(q)$ is quasiconcave, peaking at $q = 1/2$
- If $4c''(\sigma) + \sigma c'''(\sigma) > 0$, then $\sigma(q)$ is concave $\Rightarrow$ obfuscation drifts down.