Repeated 0-Sum Incomplete Information Games OR Bayesian Persuasion as Verified Cheap Talk

Lones Smith

Spring, 2025

Repeated Zero-Sum Games with Incomplete Information

- Aumann and Maschler
  - 1966: "Game Theoretic Aspects of Gradual Disarmament"
  - 1995 (book): "Repeated Games with Incomplete Information"
- One long-lived player knows the zero sum stage game repeatedly played, and his opponents do not
- Financial market: Seller knows an asset value; buyers do not
- Patient Player 1 earns G^a or G^b in states a (prior p) and b
- Player 1 picks Row and (sequence of) Players 2 picks columns
- Example 1:

$$G^a = \begin{pmatrix} 0 & 0 \\ 0 & -1 \end{pmatrix} \text{ and } G^b = \begin{pmatrix} -1 & 0 \\ 0 & 0 \end{pmatrix}.$$

- Unique Bayes-Nash equilibrium is completely revealing:
  - Player 1 chooses *Top* row in state a and *Bottom* row in state b

$$G(p) = \begin{pmatrix} -p & 0 \\ 0 & p-1 \end{pmatrix}$$

  - value function for infinitely repeated game is $u(p) = -p(1-p)$
  - Revealing the state creates a game with value $v(p) = 0$

Repeated Zero-Sum Games with Incomplete Information

- Example 2:

$$G^a = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \text{ and } G^b = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

- Unique Bayes-Nash equilibrium:
  - Player 1 plays a concealing strategy, independent of the state
  - Assume Player 2 cannot see past payoffs
  - Then Player 2 never learns about the state.
  - Assume $q = 1/2$. Then Player 2 acts as if game is

$$\tfrac{1}{2}G^a + \tfrac{1}{2}G^b = \begin{pmatrix} 1/2 & 0 \\ 0 & 1/2 \end{pmatrix}$$

  - Its value is $1/4$ – where players randomize equally on actions
  - Payoffs are hidden $\Rightarrow$ Player 1 can act as if one-shot game is

$$G(p) = \begin{pmatrix} p & 0 \\ 0 & 1-p \end{pmatrix}$$

  - value function for infinitely repeated game is $u(p) = p(1-p)$
  - Concealing the state creates a game with value $v(p) = u(p)$

Repeated Zero-Sum Games with Incomplete Information

- Example 3:

$$G^a = \begin{pmatrix} 4 & 0 & 2 \\ 4 & 0 & -2 \end{pmatrix} \quad \text{and} \quad G^b = \begin{pmatrix} 0 & 4 & -2 \\ 0 & 4 & 2 \end{pmatrix}$$

- Player 1 plays a revealing strategy:
  - Player 2 plays *Middle* in G^a, *Left* in G^b. Player 1 gets payoff 0
- Player 1 plays a concealing strategy:

$$\frac{1}{2}G^a + \frac{1}{2}G^b = \begin{pmatrix} 2 & 2 & 0 \\ 2 & 2 & 0 \end{pmatrix}$$

  - Player 2 always plays *Right* column. So Player 1 gets payoff 0
- Player 1 plays a partially revealing strategy
  - State a, picks T with chance 3/4, and B with chance 1/4
  - State b, picks T with chance 1/4, and B with chance 3/4.
  - 50-50 expected game: $(3/4)G^a + (1/4)G^b$ or $(1/4)G^a + (3/4)G^b$

$$\frac{3}{4}G^a + \frac{1}{4}G^b = \begin{pmatrix} 3 & 1 & 1 \\ 3 & 1 & -1 \end{pmatrix} \quad \text{or} \quad \frac{1}{4}G^a + \frac{3}{4}G^b = \begin{pmatrix} 1 & 3 & -1 \\ 1 & 3 & 1 \end{pmatrix}$$

  - Player 1 plays *Top* at left, *Bottom* at right, and gets payoff 1!

Repeated Zero-Sum Games with Incomplete Information

- Example 3: value function for infinitely repeated game $v(p)$ is dashed line above solid line payoff $u(p)$

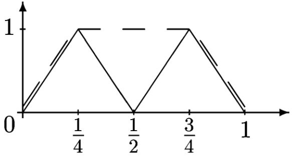

Theorem (Aumann Maschler)

The value of the infinitely repeated game with concealed actions is $v(p) = cav(u(p))$ *— namely, the least concave function* $\geq u(p)$

- Player 2 cannot see past payoffs? No discounting?
- Kohlberg (1975), "The information revealed in infinitely-repeated games of incomplete information" models it as a Sender-Receiver Game with commitment
- Almost no one reads this paper
  - Recall: Leif Erikson, discovers North America, 1000AD. Leaves
- Kamenica and Gentzkow (2011), "Bayesian Persuasion"

Bayesian Persuasion: Verifiable Free Communication

- Sender is a Prosecutor (P) and Receiver is a Judge (J)
- Two states: a defendant is either guilty or innocent.
  - Prosecutor and Judge think that 30% of defendants are guilty
- Judge must convict or acquit a defendant.
- Payoffs:
  - J earns 1 for a just action (convicting if guilty, and acquitting if innocent) and 0 for an unjust action
  - $\Rightarrow$ J convicts if guilt probability ≥ 0.5 (prove graphically!)
  - P earns 0 if defendant is found innocent, 1 if guilty.
  - $\Rightarrow$ P's and J's preferences align if guilty, and oppose if innocent
- Babbling equilibrium: Judge convicts no one if the Prosecutor sends useless signal (guilty prior is $0.3 < 0.5$ for all accused)
- Judge convicts 30% if Prosecutor sends perfect information

	guilty g	innocent i
acquit a	P(a|g) =0	P(a|i) =0
convict c	P(c|g) =1	P(c|i) =1

Bayesian Persuasion: Verifiable Free Communication

- Assume the Prosecutor conducts an investigation
- This is vetted "verifiable cheap talk" (rules of evidence)
- This yields a clever *type-dependent signal*

	guilty g	innocent i
acquit a	P(a|g) =0	P(a|i) =4/7
convict c	P(c|g) =1	P(c|i) =3/7

$\Rightarrow$ Prosecutor reveals the guilt of everyone guilty *and also smears* $3/7$ *of the innocent*

- Judge is barely willing to convict if Prosecutor says "convict":

$$P(g|c) = \frac{P(c|g)P(g)}{P(c|g)P(g) + P(c|i)P(i)} = \frac{1 \times 0.3}{1 \times 0.3 + \frac{3}{7} \times 0.7} = 0.5$$

$\Rightarrow$ Prosecutor gets conviction rate of $1 \times 0.3 + \frac{3}{7} \times 0.7 = 60\%$

- Designer signals belongs to the "information design" literature

Concavification Picture of Bayesian Persuasion

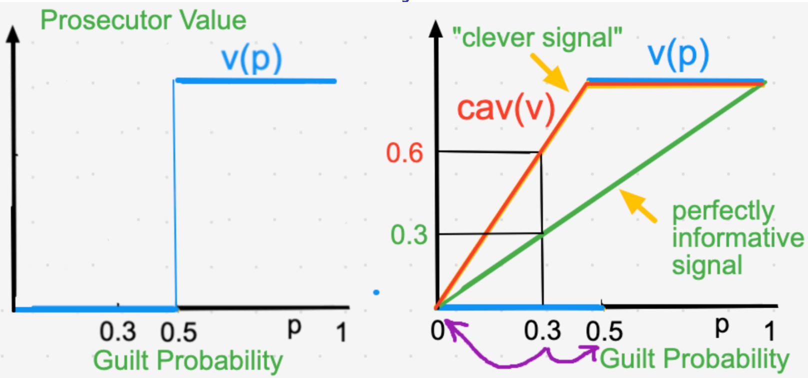

- A perfectly informative signal sends 30% of accused to 0, and 70% to 1, and thus has value 0.3: it splits the prior 0.3 to posteriors 0 and 1, namely, subject to the *martingale equality*
- Clever signal splits 0.3 to 0 and 0.5. Sender's optimal value is

$$cav(v) = \inf\{\text{functions } f \text{ on } [0,1] | f(p) \geq v(p), f \text{ is concave}\}$$

- Note: minimum of a set of concave functions is concave

How to Design Clever Signals that "Split" Beliefs

- Assume $1/2 < p < 3/4$
- To split a prior belief p into a higher posterior belief $3/4$ and a lower posterior belief $1/4$, we use a binary signal:

	L	H
s	$1-a$	$1-b$
t	a	b

- A posterior of $3/4$ corresponds to H:L odds of $3:1$
- By the easier odds formulation of Bayes rule, the L-favorable s signal and the H-favorable t signal obey:

$$\frac{p}{1-p}\frac{1-b}{1-a} = \frac{1}{3} \quad \text{(see } s\text{)} \qquad \text{and} \qquad \frac{p}{1-p}\frac{b}{a} = \frac{3}{1} \quad \text{(see } t\text{)}$$

- In other words:

$$a = \frac{4p-1}{8(1-p)} \qquad \text{and} \qquad b = \frac{12p-3}{8p}$$

- Finally, $p > 1/4$ implies $a, b > 0$, and $p < 3/4$ implies $a, b < 1$
- Exercise: what is the chance that the posterior is $3/4$?

Repeated Competitive Games of Incomplete Information

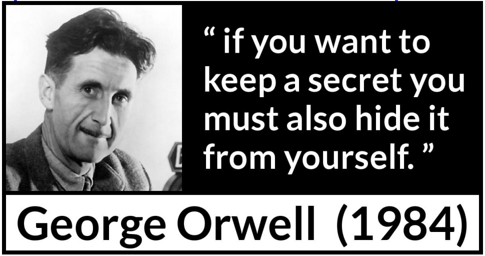

- This work dates to Aumann and Maschler (1966) on "Repeated Competitive Games of Incomplete Information"
- We don't need rules of evidence: this arises in dynamic equilibria when players have a desire to maintain a reputation
- Aumann (1995): "negotiating strategy used by the Americans in a series of arms control conferences might implicitly send signals to the Russians about the nature of the US arsenal"
- In "The Imitation Game", Turing knew: do not openly use the Enigma machine to save a ship about to be sunk
- Omitted Important Topic: **Reputation**

Bonus Question: An Incomplete Information Trading Game

- Example: A Broker who knows if shares are worth \$6 or \$10
- A Trader thinks the chance of the high state (\$10) is $p \in [0, 1]$
- He buys or sells shares. He may also pick an \$8 price; only in this case, the Broker simultaneously chooses buy or sell (blue)
- The Broker's payoffs (and negative Trader payoffs) are:

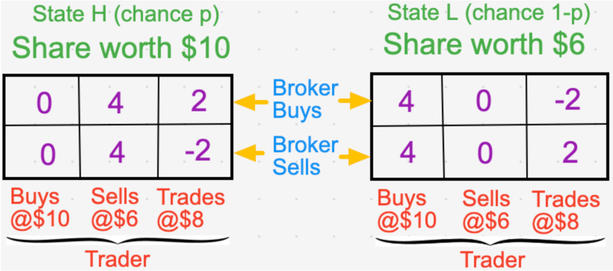

Equilibrium of the Trading Game

- If the Broker plays as if he does not know the game:

G(p)

	Buys	Sells	Trades
Buys	4(1-p)	4p	4p-2
Sells	4(1-p)	4p	2-4p

- Buy dominates Sell if $p > 1/2$ & Sell dominates Buy if $p < 1/2$
- *Nash equilibrium*: Broker buys if $p > 1/2$ and sells if $p < 1/2$
- Trader Sells, Trades, and Buys on $[0, \frac{1}{4}]$, $[\frac{1}{4}, \frac{3}{4}]$, $[\frac{3}{4}, 1]$, resp.
- Values of special cases

G(0)⇢v(0)=0

	B	S	T
B	4	0	-2
S	4	0	2

G(1/4)⇢v(1/4)=1

	B	S	T
B	3	1	-1
S	3	1	1

G(1/2)⇢v(1/2)=0

	B	S	T
B	2	2	0
S	2	2	0

Value of the Broker Trader Game

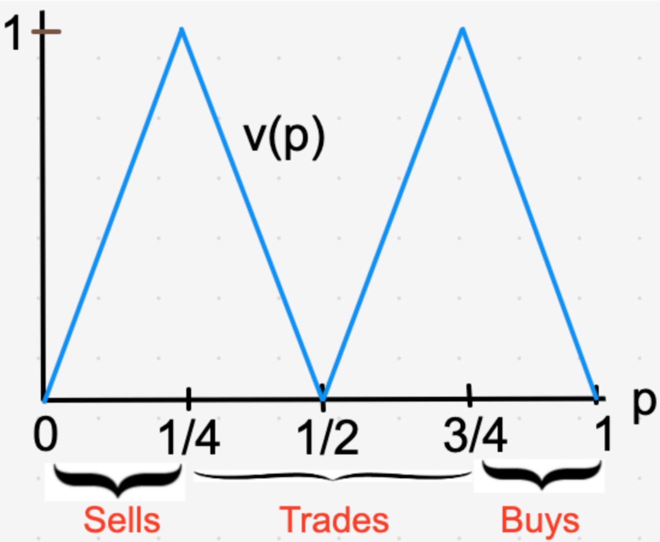

Concavification Picture of the Communication Game

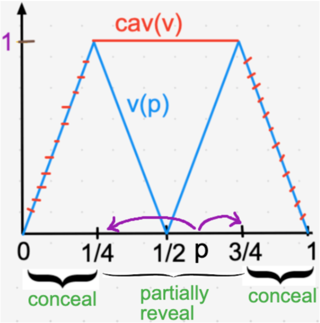

- As noted, splitting beliefs is profitable if the Sender value is locally strictly convex, not if it is locally concave
- As in the Bayesian persuasion model, $cav(v)$ is the best value attainable by the informed party (the Broker)
  - For $p \leq 1/4$ or $p \geq 3/4$ the Broker reveals nothing
  - For $1/4 \leq p \leq 3/4$, the Broker reveals credible information that splits prior p to $1/4$ and $3/4$ (partially reveals the state)