

Representation and value of information

Lecture 3

One player's belief

Represent information.

Compare and choose experiments.

Several players

Compare information structures.

Understand strategic value.

Four lectures

- 1 Knowledge and agreement**
Partitions; common knowledge; Aumann's theorem
- 2 Solution concepts**
Minimax; rationalizability; Bayesian games; correlated equilibrium
- 3 Representation and value of information**
Experiments; information choice; strategic value
- 4 Higher-order beliefs and universal representation**
Email game; universal type space; interim rationalizability

I. Information representation for one player

Represent and compare

I.1 Representations

Partitions; experiments; posteriors.

Bayes-plausibility.

I.2 Comparison

Blackwell's Theorem:
information $\iff$ usefulness.

Choose information

I.3 Bayesian persuasion

Choose a test to influence a decision.

I.4 Rational inattention

Choose information net of its cost.

I.5 Repeated revelation

Choose what the opponent learns.

Weather forecasts: Alice or Bob?

Rain R or sun S , each with probability $\frac{1}{2}$.

The decision

Take an umbrella (U), or leave it (N).

$g(a, k)$	R	S
U	0	0
N	-1	1

The information

Alice predicts rain, sun, or says she is unsure.

Bob always predicts rain or sun.

Whose forecast gives the higher expected payoff?

Forecasts are statistical experiments

An experiment gives a signal distribution $\alpha_k \in \Delta(X)$ in each state $k \in K$.

Alice: experiment α

	x_R	x_I	x_S
R	.6	.2	.2
S	.2	.1	.7

x_I : unsure

Bob: experiment β

	y_R	y_S
R	.7	.3
S	.25	.75

Prior $p \in \Delta(K)$; retain states with $p(k) > 0$.

Partitions and experiments: two representations

From a partition to signals

World $\omega \in \Omega$, law P , payoff state $\kappa(\omega) \in K$.

Signal $x(\omega)$ labels the cell of $\mathcal{P}$ containing ω .

$$\alpha_k(x) = P(x(\omega) = x \mid \kappa(\omega) = k).$$

From signals to a partition

Take $\Omega = K \times X$, with

$$P(k, x) = p(k)\alpha_k(x).$$

The payoff state is k ; the observed cell is


$$K \times \{x\}.$$

The signal identifies a cell, rather than the payoff state.

From signals to posterior beliefs

Signal x has probability m_x and induces posterior q_x :

$$m_x = \sum_{k \in K} p(k) \alpha_k(x), \quad q_x(k) = \frac{p(k) \alpha_k(x)}{m_x} \quad (m_x > 0).$$




The induced law μ assigns mass $\mu\{q\} = \sum_{x: q_x=q} m_x$ to posterior q .

Decision values are convex in beliefs

For finite actions A and payoffs $g(a, k)$,

$$v(q) = \max_{a \in A} \sum_k q(k) g(a, k), \quad V(\alpha, g) = \mathbb{E}_\mu v(q).$$



Convexity

Maximum of affine functions.


Umbrella decision

$q = \Pr(S)$:

$$v(q) = \max\{0, 2q - 1\}.$$

Alice and Bob: distributions of posteriors

$q = \Pr(S \mid \text{forecast})$; the same decision value $v(q)$ for both forecasts.



Bayes-plausibility: which posterior distributions are feasible?

Let μ be a finite distribution of posteriors $q \in \Delta(K)$, with prior p .

Bayes-plausibility

μ is induced by an experiment at prior p

$$\iff \mathbb{E}_\mu q = p.$$

This condition is called **Bayes-plausibility**.

In coordinates

Posteriors q_j with probabilities m_j satisfy

$$\sum_j m_j q_j(k) = p(k) \quad \text{for every state } k.$$

Bayes-plausibility: proof in two directions

Experiment $\rightarrow$ posterior split

Bayes's rule gives

$$m_x q_x(k) = p(k) \alpha_k(x).$$

$$\sum_x m_x q_x(k) = p(k) \sum_x \alpha_k(x) = p(k).$$

Posterior split $\rightarrow$ experiment

Given $\sum_j m_j q_j = p$, set

$$\alpha_k(j) = \frac{m_j q_j(k)}{p(k)} \quad (p(k) > 0).$$

Signal j has probability m_j and posterior q_j .

Free information cannot hurt

$$\mathbb{E}_\mu v(q) \geq v(\mathbb{E}_\mu q) = v(p).$$

For the umbrella decision, Alice wins

Take the umbrella when $q \leq \frac{1}{2}$; leave it when $q \geq \frac{1}{2}$.

Alice

$$\begin{aligned} V(\alpha, g) &= .45 \times \left(2 \times \frac{7}{9} - 1 \right) \\ &= \boxed{.25}. \end{aligned}$$

Bob

$$\begin{aligned} V(\beta, g) &= .525 \times \left(2 \times \frac{5}{7} - 1 \right) \\ &= \boxed{.225}. \end{aligned}$$

Is Alice better for every decision problem?

Garbling: Alice can reproduce Bob's forecast

Quentin replaces Alice's "unsure" by a fair coin; other forecasts pass unchanged.



$$\underbrace{\begin{pmatrix} .6 & .2 & .2 \\ .2 & .1 & .7 \end{pmatrix}}_{\alpha} \underbrace{\begin{pmatrix} 1 & 0 \\ \frac{1}{2} & \frac{1}{2} \\ 0 & 1 \end{pmatrix}}_Q = \underbrace{\begin{pmatrix} .7 & .3 \\ .25 & .75 \end{pmatrix}}_{\beta}.$$

Blackwell's Theorem (1953)

Finite states and signals; common prior p , with $p(k) > 0$.

Blackwell's Theorem

More informative $\iff$ more useful.

More informative

There is a garbling $Q_x \in \Delta(Y)$ with

$$\beta_k(y) = \sum_x \alpha_k(x) Q_x(y).$$

Q depends only on the signal x .


More useful

For every finite action set A and payoff function g ,

$$V(\alpha, g) \geq V(\beta, g).$$

Proof: simulate any decision rule

Suppose $\beta = \alpha Q$. A decision rule after β is $\tau_y \in \Delta(A)$.



Use the rule $\sigma_x(a) = \sum_y Q_x(y) \tau_y(a)$ after α .

$$\sum_x \alpha_k(x) \sigma_x(a) = \sum_y \beta_k(y) \tau_y(a).$$

The same action distribution in every state; the same expected payoff.

Converse: a decision problem tests each candidate garbling

Let $\mathcal{Q}$ be the garblings from X to Y and $\mathcal{D} = [0, 1]^{Y \times K}$ the payoff tables with action set Y .

$$\gamma(Q, g) = \sum_{k, y} p(k) \left[\sum_x \alpha_k(x) Q_x(y) - \beta_k(y) \right] g(y, k).$$

Under α

Use Q as a decision rule: choose action y after signal x .

Under β

Use the identity rule: choose action y after signal y .

More usefulness implies $\forall g \in \mathcal{D}, \max_{Q \in \mathcal{Q}} \gamma(Q, g) \geq 0$.

Converse: minimax finds one garbling for all tests

The sets $\mathcal{Q}$ and $\mathcal{D}$ are convex and compact; γ is bilinear.

$$\max_{Q \in \mathcal{Q}} \min_{g \in \mathcal{D}} \gamma(Q, g) = \min_{g \in \mathcal{D}} \max_{Q \in \mathcal{Q}} \gamma(Q, g) \geq 0.$$


One Q^* passes every test. Test a single pair (k, y) :

$$g(y', k') = \mathbf{1}_{\{(y', k') = (y, k)\}} \implies \sum_x \alpha_k(x) Q_x^*(y) \geq \beta_k(y).$$

Both sides sum to 1 over y , so every inequality is an equality.

$$\beta = \alpha Q^*. \quad \text{More useful} \implies \text{more informative.}$$

I.3–I.5 Choosing information: three applications




Bayesian persuasion: which test should the company choose?

Regulator

Approve a drug only if it is sufficiently likely to be effective.

Pharmaceutical company

Maximize the probability that the drug is approved.



Common prior: effective and ineffective are equally likely.

The company chooses the test before learning the state.

The regulator's decision

Regulator's payoffs		
	Effective	Ineffective
Approve	1	-9
Reject	0	0

Let $q = \Pr(\text{effective} \mid \text{test result})$.

Expected payoff from approval = $q - 9(1 - q) = 10q - 9$.

Approve if and only if $q \geq .9$.

At $q = .9$, the regulator approves: indifference favors the company.

Two benchmark tests

No information

Posterior $q = \frac{1}{2}$.

The regulator rejects.

Approval probability: 0.

Full revelation

Posterior $q \in \{0, 1\}$.

Only effective drugs are approved.

Approval probability: $\frac{1}{2}$.

Can a less informative test produce more approvals?

A test that produces more approvals

	Pass	Fail
Effective	1	0
Ineffective	$\frac{1}{9}$	$\frac{8}{9}$

Pass

$$\Pr(\text{Pass}) = \frac{5}{9}$$

$$q_{\text{Pass}} = \frac{1/2}{5/9} = \frac{9}{10}: \text{approve.}$$

Fail

$$\Pr(\text{Fail}) = \frac{4}{9}$$

$$q_{\text{Fail}} = 0: \text{reject.}$$

Approval probability: $\frac{5}{9} > \frac{1}{2}$.

Why can approval not exceed 5/9?

Bayes-plausibility

$$\mathbb{E}_{\mu} q = \frac{1}{2}$$

Approval requires


$$q \geq \frac{9}{10}$$

$$\frac{1}{2} = \mathbb{E}_{\mu} q \geq \frac{9}{10} \Pr(\text{approval}).$$

$\Pr(\text{approval}) \leq \frac{5}{9}$; the pass-fail test attains the bound.

Choosing posteriors: the geometry of the test

Company payoff at posterior q : $u_S(q) = \mathbf{1}_{\{q \geq .9\}}$.



Split $p = \frac{1}{2}$ into 0 and $\frac{9}{10}$, with weights $\frac{4}{9}$ and $\frac{5}{9}$.


Concavification: the optimal Bayes-plausible split

Payoff $u(q)$ at posterior q ; $\text{cav } u$ is its smallest concave majorant.

Concavification theorem

$$(\text{cav } u)(p) = \max_{\mu: \mathbb{E}_{\mu} q = p} \mathbb{E}_{\mu} u(q)$$

The maximization is over **Bayes-plausible** distributions of posteriors.



Finite K ; bounded upper semicontinuous u . In the binary case, two posteriors suffice.

Concavification: the upper bound

$$W(p) = \sup_{\mu: \mathbb{E}_{\mu} q = p} \mathbb{E}_{\mu} u(q), \quad h = \text{cav } u.$$


For every Bayes-plausible distribution μ ,

$$\begin{aligned} \mathbb{E}_{\mu} u(q) &\leq \mathbb{E}_{\mu} h(q) && u \leq h, \\ &\leq h(\mathbb{E}_{\mu} q) && \text{Jensen's inequality,} \\ &= h(p) && \text{Bayes-plausibility.} \end{aligned}$$

$$W(p) \leq (\text{cav } u)(p).$$

Concavification: successive splits

$$p = \lambda p_1 + (1 - \lambda)p_2, \quad \lambda \in [0, 1].$$



The terminal posteriors average to p : both stages form a single feasible split.

$$W(\lambda p_1 + (1 - \lambda)p_2) \geq \lambda W(p_1) + (1 - \lambda)W(p_2).$$

No information gives $W \geq u$.

W is a concave majorant of u : $W \geq \text{cav } u$.

Bayesian persuasion: the general solution

Receiver actions A ; payoff states K ; payoffs $g_R(a, k)$ and $g_S(a, k)$.

Receiver's best responses

$$A^*(q) = \arg \max_{a \in A} \mathbb{E}_q g_R(a, k)$$

Indifference favors the sender.

Sender's payoff at q

$$u_S(q) = \max_{a \in A^*(q)} \mathbb{E}_q g_S(a, k)$$

$$\max_{\mu: \mathbb{E}_\mu q = p} \mathbb{E}_\mu u_S(q) = (\text{cav } u_S)(p)$$

Kamenica–Gentzkow (2011); commitment to the experiment.

I.4 Rational inattention: choosing information at a cost

Umbrella problem: $q = \Pr(S)$, $v(q) = \max\{0, 2q - 1\}$, $p = \frac{1}{2}$.

Benefit

Choose an experiment, then act.

Expected payoff: $\mathbb{E}_\mu v(q)$.

Cost


More precise information requires more attention.

Charge for the reduction in uncertainty.

Which posterior split gives the highest payoff net of information costs?

Entropy measures uncertainty

$$H(q) = - \sum_{k \in K} q(k) \log_2 q(k), \quad 0 \log_2 0 = 0.$$



The information cost

Price per unit of uncertainty reduction: $C > 0$.

$$C[H(p) - \mathbb{E}_\mu H(q)].$$

Rational inattention: change the function, keep the method

$$\max_{\mu: \mathbb{E}_{\mu} q = p} \{ \mathbb{E}_{\mu} v(q) - C[H(p) - \mathbb{E}_{\mu} H(q)] \}$$

Decision value
 $v(q)$

+

Posterior entropy
 $CH(q)$


=

Concavify
 $u_C = v + CH$

$$\text{Net value} = (\text{cav } u_C)(p) - CH(p).$$

The umbrella problem: optimal partial attention

$$p = \frac{1}{2}, C = \frac{1}{2}, u_C(q) = \max\{0, 2q - 1\} + \frac{1}{2}H(q).$$



Split $\frac{1}{2}$ equally into $\frac{1}{5}$ and $\frac{4}{5}$; act differently after the two signals.

Partial attention beats both extremes

The umbrella problem with $p = \frac{1}{2}$ and $C = \frac{1}{2}$.

Experiment	Gross payoff	Cost	Net payoff
No information	0	0	0
Full revelation	$\frac{1}{2}$	$\frac{1}{2}$	0
$\frac{1}{5}/\frac{4}{5}$ split	.300	.139	.161

I.5 Repeated games: choosing what to reveal

Zero-sum game; states k_1, k_2 ; prior $p = \Pr(k_1)$.

Player 1

Knows the state.

Maximizes the average payoff.

Player 2

Does not know the state.

Minimizes Player 1's average payoff.

Choose actions
simultaneously



Observe the
public actions



Player 2 updates
the belief p

Player 2 observes actions, not stage payoffs. The state remains fixed.

The benchmark: non-revealing play

Player 1 uses the same mixed action in both states.

Write g^k for Player 1's stage payoff in state k .

$$g^p(a_1, a_2) = p g^{k_1}(a_1, a_2) + (1 - p) g^{k_2}(a_1, a_2).$$

$$u(p) = \text{Val}(g^p)$$

The value of the average matrix.

Example 1: concealment can preserve a gain

Player 1's stage payoff; $p = \frac{1}{2}$.

State k_1

	L	R
T	1	0
B	0	0

State k_2

	L	R
T	0	0
B	0	1

Example 1: a first-stage gain can destroy future gains

Full revelation

Play T in k_1 , B in k_2 .

Stage 1: guarantee $\frac{1}{2}$.

Later: Player 2 knows the state and avoids the positive cell.

Long-run payoff: 0.

Concealment

Mix T and B equally in both states.

No learning from the action.


Player 2 cannot avoid both positive cells.

Payoff every stage: $\frac{1}{4}$.

A better one-shot payoff need not be a better repeated-game payoff.

Example 1: concavity favors concealment

Non-revealing matrix: $\begin{pmatrix} p & 0 \\ 0 & 1-p \end{pmatrix}$.



Example 2: revelation can eliminate a loss

Player 1's stage payoff; $p = \frac{1}{2}$.

State k_1

	L	R
T	-1	0
B	0	0

State k_2


	L	R
T	0	0
B	0	-1

Play B in k_1 and T in k_2 : reveal the state and obtain 0 every stage.

Under concealment, the best payoff is $u(\frac{1}{2}) = -\frac{1}{4}$.

Example 2: the best split fully reveals the state

Non-revealing matrix: $\begin{pmatrix} -p & 0 \\ 0 & -(1-p) \end{pmatrix}$.



Example 3: can partial revelation beat both extremes?

Player 1's stage payoff; $p = \frac{1}{2}$.

State k_1

	L	M	R
T	4	0	2
B	4	0	-2


State k_2

	L	M	R
T	0	4	-2
B	0	4	2

Concealment gives 0; full revelation gives 0 in the long run.

Example 3: reveal a noisy clue, then repeat the action

Choose the first action using these probabilities; play it forever.



Both posteriors occur with probability $\frac{1}{2}$; later actions reveal nothing further.

Example 3: the noisy clue guarantees 1

After T

Player 2's belief is $p_T = \frac{3}{4}$.

Payoffs against (L, M, R) :

$(3, 1, 1)$.

After B

Player 2's belief is $p_B = \frac{1}{4}$.

Payoffs against (L, M, R) :

$(1, 3, 1)$.

	Concealment	Full revelation	Partial revelation
Long-run payoff	0	0	1

Example 3: computing the non-revealing value

At belief p , the average payoff matrix is

$$g^p = \begin{array}{c|ccc} & L & M & R \\ \hline T & 4p & 4(1-p) & 4p-2 \\ B & 4p & 4(1-p) & 2-4p \end{array}.$$


The first two columns do not depend on Player 1's row.

Choose the row that maximizes the R payoff: $|4p - 2|$.

$$u(p) = \min\{4p, 4(1-p), |4p-2|\}.$$

Example 3: concavification selects partial revelation

$$u(p) = \min\{4p, 4(1-p), |4p-2|\}.$$



$(\text{cav } u)(\frac{1}{2}) = 1$: split into $\frac{1}{4}$ and $\frac{3}{4}$, with equal weights.

Aumann–Maschler (1995): the long-run value

Let $V_N(p)$ be the value of the N -stage game, with average payoffs.

$$\lim_{N \rightarrow \infty} V_N(p) = (\text{cav } u)(p)$$

Achieve the envelope

Induce a posterior split.

Then play optimally without further revelation.

Bound the value

Beliefs form a martingale.

Continuing revelation cannot sustain an extra gain in the long-run average.

II. Value of information for n players

II.1 The difficulty: information changes deviations and equilibrium.

II.2 Two comparisons

Refinement

More contingent strategies.

Faithful reproduction

More equilibrium outcomes.

II.3–II.5 Payoff implications

Zero-sum Better guarantees.

No-trade Opting out.

Positive value Common maximum.

II.6 Synthesis: more information $\iff$ more strategies.

Investment: does information help the first mover?

Two projects, A and B ; one is good, each with probability $\frac{1}{2}$.



Investing in the good project

Alone: payoff 3.

With the other firm: payoff 1.

Investing in the bad project

Payoff 0.

Information receipt is common knowledge.

Neither firm is informed: Firm 2 chooses the other project

Firm 1's choice carries no news about quality.

Firm 2 follows

The chosen project is good with probability $\frac{1}{2}$.

$$\mathbb{E}[g_2] = \frac{1}{2} \times 1 = \frac{1}{2}.$$

Firm 2 chooses the other project

The other project is good with probability $\frac{1}{2}$.

$$\mathbb{E}[g_2] = \frac{1}{2} \times 3 = \frac{3}{2}.$$

Firm 1 may choose either project or any mixture.

Backward induction gives expected payoffs $(\frac{3}{2}, \frac{3}{2})$.

Firm 1 is informed: Firm 2 follows

Firm 1's choice

Good project: payoff at least 1.

Bad project: payoff 0.

Firm 1 always chooses the good project.

Firm 2's response

The observed choice reveals quality.

Follow: payoff 1.

Choose the other project: payoff 0.

Equilibrium payoffs: $(1, 1)$.

Information lowers Firm 1's payoff from $\frac{3}{2}$ to 1.

Ignoring information does not restore the old equilibrium

Keep the old profile

Firm 1 chooses independently of quality.

Firm 2 chooses the other project.

Expected payoffs remain $(\frac{3}{2}, \frac{3}{2})$.

Use the new information

Against Firm 2's old response, Firm 1 chooses the good project.

Firm 2 chooses the bad project.

Firm 1 obtains 3.

The old profile is still feasible; it has a new profitable deviation.

An information structure describes what each player knows

$$\mathcal{I} = (\Omega, P, \kappa, (\mathcal{P}_i)_i), \quad k = \kappa(\omega) \in K.$$

Information

Common prior: $\omega \sim P$.

Player i observes $\mathcal{P}_i(\omega)$.

$\Sigma_i^{\mathcal{I}}$: plans $\sigma_i : \Omega \rightarrow \Delta(A_i)$ constant on each cell.

Payoffs

Payoff structure $G = ((A_i)_i, (g_i)_i)$.

Payoff $g_i(a, k)$.

Expected payoff $\gamma_i^{\mathcal{I}}(\sigma)$ averages over P and mixing.

Refinement gives a player more contingent strategies

Keep (Ω, P, κ) and the other players' partitions fixed.

$\mathcal{I}$ is more informative for player i than $\mathcal{J}$
if $\mathcal{P}_i^{\mathcal{I}}$ refines $\mathcal{P}_i^{\mathcal{J}}$.

$$\Sigma_i^{\mathcal{J}} \subseteq \Sigma_i^{\mathcal{I}}, \quad \Sigma_j^{\mathcal{J}} = \Sigma_j^{\mathcal{I}} \quad (j \neq i).$$

$$\gamma^{\mathcal{I}}(\sigma) = \gamma^{\mathcal{J}}(\sigma) \quad \text{for every old strategy profile } \sigma.$$

A second comparison: which equilibrium outcomes can be reproduced?

Private signals: $x = (x_i)_i \in X = \prod_i X_i$, $y = (y_i)_i \in Y = \prod_i Y_i$.

$$\mathcal{I} : (k, x) \sim \lambda, \quad \mathcal{J} : (k, y) \sim \eta.$$

Let $D(\mathcal{I}, G)$ be the set of equilibrium distributions of actions in G with information structure $\mathcal{I}$.

$\mathcal{I}$ is *richer* than $\mathcal{J}$ if

$$D(\mathcal{I}, G) \supseteq D(\mathcal{J}, G) \quad \text{for every compact payoff game } G.$$

Faithful reproduction preserves the beliefs behind incentives

Player i interprets a signal using $\phi_i : X_i \rightarrow \Delta(Y_i)$.



Compatible

The interpreted joint law of (k, y) is η .

Faithful

For each player, x_i adds no news about (k, y_{-i}) beyond y_i .

The interpretations use independent private randomization.

Coin versus die: a public die reproduces a public coin

Two players; no payoff uncertainty. Both observe the same fair die d .

Die: $d \in \{1, 3, 5\}$



Both report heads
probability $\frac{1}{2}$

Die: $d \in \{2, 4, 6\}$



Both report tails
probability $\frac{1}{2}$

Faithful: parity already determines the other player's interpreted signal.

Every equilibrium with the public coin can be played using the public die.

A public coin cannot reproduce a public die

Players observe one public coin; any additional randomization is private and independent.

Target: the same fair die for both

Reports $y_1, y_2 \in \{1, \dots, 6\}$ must satisfy

$$\Pr(y_1 = y_2) = 1,$$

$$\Pr(y_i = j) = \frac{1}{6} \quad (j = 1, \dots, 6).$$

Conditional on heads or tails

The two reports are independent.

To agree with probability 1, both must choose the same deterministic face.

At most two common faces; a fair public die requires six.

A hidden die implements the traffic-light device

A fair die selects the lights. Each driver observes only their own light.

Hidden die d	1, 2	3, 4	5, 6
Driver 1			
Driver 2			

$$\Pr(R, R) = \Pr(R, G) = \Pr(G, R) = \frac{1}{3}, \quad \Pr(G, G) = 0.$$

Making the die public: compatible, but not faithful

Both drivers now observe d , then compute their own light from the same table.

Only Driver 1's red light

The other light is red or green, equally likely.

$$\Pr(y_2 = R \mid y_1 = R) = \frac{1}{2}.$$

Driver 1 also observes the die

At $d = 1$, both lights are red.

$$\Pr(y_2 = R \mid d = 1) = 1.$$

The die reveals the other driver's light.

The joint law of the lights is preserved; the conditional beliefs are not.

Why faithfulness matters: the same lights, a new deviation

Week 2's game: C = concede; T = take priority.

	C	T
C	6, 6	2, 7
T	7, 2	0, 0

Red: C Green: T

On green, T gives 7; C gives 6.

Only red observed

$$\mathbb{E}[C \mid R] = 4 > \frac{7}{2} = \mathbb{E}[T \mid R].$$

Public die $d = 1$

Against C , deviate to T : $7 > 6$.

Private lights: obedience is an equilibrium. Public die: it is not.

Gossner's comparison theorem

Gossner (2000); finite signals, compact actions, continuous payoffs.

$\mathcal{I}$ is richer than $\mathcal{J} \iff \mathcal{J}$ is faithfully reproducible from $\mathcal{I}$.

From reproduction to equilibrium

Interpret the signals, then follow a $\mathcal{J}$ -equilibrium.

The outcome law is preserved.

Faithfulness preserves best responses.

From equilibrium to reproduction

Use games that reward truthful reports of beliefs.

Universal outcome inclusion forces a compatible interpretation that preserves those beliefs.

11.3–11.5 When do incentives yield payoff conclusions?

Three different routes

11.3 Zero-sum

Opposed interests.

Minimax turns a better guarantee into a better equilibrium value.

11.4 No-trade

No aggregate gains.

Opting out forces every equilibrium payoff to zero.

11.5 Positive value

A common maximum.

One equilibrium attains every player's best feasible payoff.

Zero-sum games: information improves the informed player's guarantee

Payoffs $g(a, k)$ to Player 1 and $-g(a, k)$ to Player 2.

Write $\gamma^{\mathcal{I}}(\sigma_1, \sigma_2)$ for Player 1's expected payoff.

$$\text{Val}(\mathcal{I}, G) = \max_{\sigma_1 \in \Sigma_1^{\mathcal{I}}} \min_{\sigma_2 \in \Sigma_2^{\mathcal{I}}} \gamma^{\mathcal{I}}(\sigma_1, \sigma_2).$$

Refine Player 1's information, or coarsen Player 2's:

Player 1's value weakly increases.

By the Minimax Theorem, this value is also the equilibrium payoff.

Proof: enlarge the maximizer's set or shrink the minimizer's set

Refine Player 1's information, holding Player 2's fixed.

Write $\gamma = \gamma^{\mathcal{I}}$; on old profiles it also equals $\gamma^{\mathcal{J}}$.

$$\Sigma_1^{\mathcal{J}} \subseteq \Sigma_1^{\mathcal{I}}, \quad \Sigma_2^{\mathcal{J}} = \Sigma_2^{\mathcal{I}}.$$

$$\underbrace{\max_{\sigma_1 \in \Sigma_1^{\mathcal{I}}} \min_{\sigma_2 \in \Sigma_2^{\mathcal{I}}} \gamma(\sigma_1, \sigma_2)}_{\text{Val}(\mathcal{I}, G)} \geq \underbrace{\max_{\sigma_1 \in \Sigma_1^{\mathcal{J}}} \min_{\sigma_2 \in \Sigma_2^{\mathcal{J}}} \gamma(\sigma_1, \sigma_2)}_{\text{Val}(\mathcal{J}, G)}.$$


The same inequality holds when Player 2's strategy set shrinks.

Minimax turns a better guarantee into a better equilibrium payoff.

Betting: two players both think the bet is attractive

Three equally likely worlds. Both must accept for a transfer to occur.

Player 1: $\{\omega_1, \omega_2\}$ or $\{\omega_3\}$



Player 2: $\{\omega_1\}$ or $\{\omega_2, \omega_3\}$

$g_2 = -g_1$ if both accept; both get 0 otherwise.

The opponent's willingness to bet changes the calculation

If the opponent always accepts

Player 1 at $\{\omega_1, \omega_2\}$:

$$\frac{1}{2}(-1 + 2) = \frac{1}{2} > 0.$$

Player 2 at $\{\omega_2, \omega_3\}$:

$$\frac{1}{2}(-2 + 5) = \frac{3}{2} > 0.$$

Account for selective acceptance

Player 1 at ω_3 rejects a loss of 5 whenever trade is possible.

Player 2 cannot count on that gain of 5 at $\{\omega_2, \omega_3\}$.

Optimistic conditional beliefs alone do not support an equilibrium bet.

No-trade equivalence theorem

Every player can opt out

An action $n_i \in A_i$ gives

$$g_i(n_i, a_{-i}, k) = 0$$

for every a_{-i} and k .

No gains from trade

For every action profile and state,

$$\sum_i g_i(a, k) \leq 0.$$

Every Bayesian Nash equilibrium is equivalent to no trade in expected payoffs:
every player's expected payoff is 0.

No-trade equivalence: proof

Let w_i be player i 's expected equilibrium payoff.

Individual rationality

Deviation to n_i gives 0.

$$w_i \geq 0 \quad \text{for every } i.$$

Aggregate feasibility

Average the no-gains condition.

$$\sum_i w_i \leq 0.$$

$$w_i = 0 \text{ for every } i.$$

Zero expected gains do not imply zero transfers.

A betting equilibrium with trade

A fair coin determines the transfer; neither player observes it before accepting.

If both accept

	Player 1	Player 2
Heads	1	-1
Tails	-1	1

Both accepting is an equilibrium

For each player,


$$\mathbb{E}[g_i \mid \text{accept}] = 0$$

and refusing also gives 0.

A nonzero transfer occurs in every state; expected payoffs equal those of no trade.

When can more information benefit every player?

Let F_I be the expected payoff vectors attainable by strategy profiles; $D_I = \text{co } F_I$.



One Pareto-maximal point w^* .

Every feasible payoff satisfies

$$w_i \leq w_i^* \quad \text{for every } i.$$

All players can attain their maximal payoff together.

A sufficient condition for positive value of information

Bassan–Gossner–Scarsini–Zamir (2003)

Assume $F_{\mathcal{I}}$ is compact and the Pareto frontier of $D_{\mathcal{I}}$ is the singleton $\{w^*\}$.

Coarsen every player's partition, keeping (Ω, P, κ) fixed.

w^* is an equilibrium payoff under $\mathcal{I}$.

It dominates every equilibrium payoff under every coarser structure $\mathcal{J}$.

The theorem selects a dominating equilibrium; other equilibria under $\mathcal{I}$ may be worse.

Proof: the best feasible point is itself an equilibrium

1. w^* is an extreme point of $D_{\mathcal{I}}$.

Compactness gives $w^* \in F_{\mathcal{I}}$: some profile σ^* attains it.

2. Every unilateral deviation produces a feasible payoff.

No player can exceed w_i^* , so σ^* is an equilibrium.

3. Coarser plans remain available under $\mathcal{I}$.

$$F_{\mathcal{J}} \subseteq F_{\mathcal{I}}.$$

Every payoff under $\mathcal{J}$, including every equilibrium payoff, is dominated by w^* .

What the singleton condition characterizes

For a given game

A singleton Pareto frontier is sufficient.

A game may still have positive value of information with several Pareto points.

Uniformly over a payoff polytope

Fix a payoff polytope D .

All games with compact feasible set and convex hull D have positive value of information

$\iff$ the Pareto frontier of D is a singleton.

The converse characterizes feasible-payoff geometry across games.

Adding a strategy can also hurt its owner

Player 1 has only C

	C	D
C	3, 3	-1, 2

Equilibrium: (C, C) .

Payoffs: $(3, 3)$.

Player 1 also has D

	C	D
C	3, 3	-1, 2
D	4, -1	0, 0

Equilibrium: (D, D) .

Payoffs: $(0, 0)$.

More options can improve deviations and worsen equilibrium outcomes.

Gossner: information and strategies have the same representation

From information to strategies

Refine one player's information.

Old contingent plans remain available,
with unchanged payoffs.

From strategies to information

Enlarge one player's strategy set in a
finite game.

A suitable Bayesian environment
unlocks those strategies by revealing a
private key.

Strategic equivalence: identify contingent plans with their effective mixed strategies.

The converse: an atomless key unlocks the new action

Let g be the preceding two-row payoff matrix. Draw θ uniformly on $[0, 1]$.

Player 1 chooses $h \in \{C\} \sqcup [0, 1]$; Player 2 chooses C or D .

$$e_{\theta}(h) = \begin{cases} D, & h = \theta, \\ C, & \text{otherwise.} \end{cases} \quad \text{Payoffs: } g(e_{\theta}(h), a_2).$$

Report h
or choose baseline C



$h = \theta$: effective D
Otherwise: effective C

Player 2 never sees the key. This construction uses continuous states and actions.

Knowing the key reproduces exactly the extra strategy

Player 1 does not see θ

Any mixed report is independent of θ .

$$\Pr(h = \theta) = 0.$$

Every strategy is effectively C .

Player 1 sees θ

Choose C , or report $h = \theta$.

Every contingent plan induces an effective mixture over C, D .

Every mixture is attainable.

The two information structures reproduce the two original normal-form games.

Use a key-labelled report for each extra action to represent any finite expansion.

Main references

Comparing information

Blackwell (1951, 1953).
Gossner (2000), *Comparison of Information Structures*.

Choosing information

Sims (2003), rational inattention.
Kamenica and Gentzkow (2011),
Bayesian persuasion.

Repeated revelation

Aumann and Maschler (1995).
Sorin (2002), *A First Course on Zero-Sum Repeated Games*.

Strategic value

Bassan, Gossner, Scarsini and Zamir
(2003), *Positive Value of Information in Games*.

Olivier Gossner, course notes: *Value of information*.

Conclusion

One player: three key results

Bayes-plausibility

Feasible posteriors: $\mathbb{E}_\mu q = p$.

Blackwell's Theorem

Information $\iff$ usefulness.

Concavification

Optimal posterior split: $(\text{cav } u)(p)$.

Several players: incentives

More information expands contingent strategies.

Equilibrium payoff effects depend on incentives.

The representation theorem unifies information and strategies.

Lecture 4: which beliefs about others matter for strategic predictions?

The three-world bet: equilibrium calculation

Three equally likely worlds. If both accept, Player 1's payoffs are $(-1, 2, -5)$; Player 2 receives the opposite transfer. Otherwise both get 0.

Player	Information cell	Acceptance probability	Gain from accepting
1	$\{\omega_1, \omega_2\}$	x	$(-y + 2t)/2$
1	$\{\omega_3\}$	z	$-5t$
2	$\{\omega_1\}$	y	x
2	$\{\omega_2, \omega_3\}$	t	$(-2x + 5z)/2$

$x > 0 \implies y = 1, t > 0 \implies z = 0 \implies t = 0$: contradiction.

With $x = 0, z > 0 \implies t = 1 \implies z = 0$: contradiction.

$x = z = 0, 0 \leq 2t \leq y \leq 1$. No trade in this example.