Information Design

Lecture Notes

Daniele Condorelli

Acknowledgements

I thank Michael Challis, Özlem Toplar, Carlos Viché, Yanjun Gao, Desmond Fairall, and Shantanu Chadha for helping with the creation of these notes.

Comments and corrections: d.condorelli@gmail.com.

September 2026

Contents

Lecture 1

Bayesian Persuasion

1 Introduction

In most economic models, information is part of the environment. Information design asks what outcomes can be obtained by changing it, holding preferences and the prior fixed.

A decision maker can ignore a free signal, so information cannot lower her maximum expected payoff when the decision problem is unchanged. This may not be true when someone else is taking a decision simultaneously or after. In general, not giving all information to a decision maker may benefit someone with a different objective than the decision maker.

There are two interpretations of information design:

- *Design.* A designer commits to an information structure and chooses it to maximise an objective.
- *Bounds.* An analyst knows the game and some of the players' information, and asks which outcomes are consistent with what they might additionally know.

The first interpretation is central to Kamenica and Gentzkow's persuasion model. The second is central to Bergemann and Morris's analysis of games. Both impose two restrictions: beliefs must be consistent with Bayes' rule, and players must want to take the proposed actions.

Lecture 1 expresses persuasion as a choice of posterior beliefs. Lecture 2 compares experiments, characterises feasible posterior means, and introduces information costs. Lecture 3 writes the incentive constraints in games. Lecture 4 applies these methods to market segmentation. Lecture 5 allows the buyer to choose the distribution of values or her learning, and then allows a platform to choose matches. The examples follow the same steps: describe feasible outcomes, bound the objective, and construct a choice attaining the bound.

2 Information Design: an Example

Let's start with an example. A defendant is guilty, G, or innocent, I. The judge chooses whether to acquit, a, or convict, c:

$$u_J(a, I) = u_J(c, G) = 1, \qquad u_J(a, G) = u_J(c, I) = 0.$$

The judge convicts when indifferent. The prosecutor wants conviction in either state:

$$u_P(c, \omega) = 1, \qquad u_P(a, \omega) = 0.$$

The common prior is $\mathbb{P}(G) = 0.3$, $\mathbb{P}(I) = 0.7$. Without information, the judge acquits. The prosecutor obtains zero.

The prosecutor commits to an experiment before the state is realised. A signal is drawn according to a publicly known rule and observed by the judge. For two signals, the rule can be written as

	s_1	s_2
G	q	$1-q$
I	r	$1-r$

where the entries are conditional probabilities given the state. Full revelation gives conviction with probability 0.3.

Consider instead

	i	g
G	0	1
I	4/7	3/7

Bayes' rule gives

$$\mathbb{P}(g) = 0.3 + 0.7(3/7) = 0.6, \qquad \mathbb{P}(G \mid g) = \frac{0.3}{0.6} = \frac{1}{2}, \qquad \mathbb{P}(G \mid i) = 0.$$

The judge convicts after g and acquits after i. The defendant is guilty with probability 0.3, but is convicted with probability 0.6.

The judge's expected payoff is

$$0.3 + 0.7(4/7) = 0.7.$$

The judge is indifferent between using and ignoring this experiment. The prosecutor benefits from pooling some innocent defendants with the guilty.

3 Kamenica and Gentzkow's Persuasion

For now, states and actions are finite. The environment consists of:

- a state space Ω and common prior $\mu_0 \in \Delta(\Omega)$;
- a receiver with actions A and payoff $u(a, \omega)$;
- a designer with payoff $v(a, \omega)$.

Here $\Delta(X)$ denotes the probability distributions on X, and δ_x denotes the distribution assigning probability one to x. We write $\mathbf{1}_E$ for the indicator of an event E. We restrict Ω to the support of the prior, so $\mu_0(\omega) > 0$.

An information structure is a finite signal set S and a kernel

$$\Pi_\omega(s) = \mathbb{P}(s \mid \omega), \qquad \Pi_\omega(s) \geq 0, \quad \sum_s \Pi_\omega(s) = 1.$$

Given the prior, it induces the joint distribution

$$P(\omega, s) = \mu_0(\omega)\Pi_\omega(s).$$

Conversely, a joint distribution with state marginal μ_0 defines the kernel by division by $\mu_0(\omega)$.

The timing is:

1. The designer publicly commits to (S, Π).
2. The state and signal are drawn.

3. The receiver observes the signal and chooses an action.
4. Payoffs are realised.

The commitment is to an experiment, not to a report chosen after seeing the state. The designer takes no further action.

4 Bayesian Updating and Posteriors

For every signal of positive probability,

$$\mu_s(\omega) = \mathbb{P}(\omega \mid s) = \frac{\mu_0(\omega)\Pi_\omega(s)}{\sum_{\omega'} \mu_0(\omega')\Pi_{\omega'}(s)}.$$

Thus μ_s is a vector of nonnegative numbers summing to one. Write $q_s = \mathbb{P}(s)$. The experiment induces a distribution $\tau \in \Delta(\Delta(\Omega))$ over posterior beliefs. If different signals induce the same belief, their probabilities are added.

Example. Let $\mu_0 = (1/3, 1/3, 1/3)$. Consider the following *joint* probabilities:

	ω_1	ω_2	ω_3	q_s
s_1	1/3	1/9	0	4/9
s_2	0	2/9	0	2/9
s_3	0	0	1/3	1/3

The posteriors are

$$\mu_{s_1} = (3/4, 1/4, 0), \quad \mu_{s_2} = (0, 1, 0), \quad \mu_{s_3} = (0, 0, 1).$$

They occur with probabilities $4/9, 2/9, 1/3$, respectively.

Work backwards. At posterior μ, the receiver chooses

$$a^*(\mu) \in \arg\max_{a \in A} \sum_\omega u(a, \omega)\mu(\omega).$$

We use designer-favourable tie-breaking: among best replies, select one that maximises the designer's expected payoff. This convention matters at a threshold, as in the prosecutor example.

The designer's payoff at posterior μ is

$$V(\mu) = \sum_\omega v(a^*(\mu), \omega)\mu(\omega).$$

The ex-ante payoff is

$$\sum_s q_s V(\mu_s) = \mathbb{E}_\tau[V(\mu)].$$

The signal labels have no further role. The problem is to choose a distribution of posteriors.

5 Splitting Lemma

Bayes' rule implies that posteriors average to the prior:

$$\sum_s q_s \mu_s(\omega) = \sum_s P(\omega, s) = \mu_0(\omega) \quad \text{for every } \omega.$$

This restriction is called *Bayes plausibility*.

Proposition 1.1 (Splitting lemma)**.** *A finitely supported distribution τ over posterior beliefs is induced by an experiment if and only if*

$$\mathbb{E}_\tau[\mu] = \mu_0.$$

Proof. Necessity was shown above. For sufficiency, label the beliefs in the support by s, with probabilities q_s, and set

$$\Pi_\omega(s) = \frac{q_s \mu_s(\omega)}{\mu_0(\omega)}.$$

The entries are nonnegative and each row sums to one by Bayes plausibility. The probability of signal s is $\sum_\omega \mu_0(\omega)\Pi_\omega(s) = q_s$. Bayes' rule then gives posterior μ_s. □

The same construction allows an arbitrary distribution τ of posteriors if signals need not be finite. For a measurable set of beliefs B, set

$$\mathbb{P}(\mu \in B \mid \omega) = \frac{1}{\mu_0(\omega)} \int_B \mu(\omega)\, d\tau(\mu).$$

Bayes plausibility makes this a probability distribution in every state; the induced joint distribution is $\mu(\omega)\, d\tau(\mu)$, so the posterior after signal μ is μ.

For example,

$$\frac{4}{9}(3/4, 1/4, 0) + \frac{2}{9}(0, 1, 0) + \frac{1}{3}(0, 0, 1) = (1/3, 1/3, 1/3).$$

The corresponding kernel is obtained by multiplying each entry of the joint table by three.

The splitting lemma is both a feasibility test and a construction. Once a useful split is found, the experiment follows from the formula in the proof.

6 Concavification

The designer's problem is

$$\max_\tau \mathbb{E}_\tau[V(\mu)] \quad \text{subject to} \quad \mathbb{E}_\tau[\mu] = \mu_0.$$

Define $\operatorname{cav} V$ as the smallest concave function weakly above V. For binary states, beliefs are points on $[0, 1]$. A split between two posteriors gives a point on the chord joining their payoffs.

Proposition 1.2 (Concavification)**.** *The value of persuasion is* $(\operatorname{cav} V)(\mu_0)$.

Geometric argument. Let $C(\mu)$ be the supremum of payoffs from finite splits with mean μ. Combining two splits shows that C is concave. No information gives $C \geq V$. Conversely, for any concave function $h \geq V$, Jensen gives

$$\mathbb{E}_\tau[V(\mu)] \leq \mathbb{E}_\tau[h(\mu)] \leq h(\mu_0).$$

Thus C is the smallest concave majorant of V: $C = \operatorname{cav} V$. Geometrically, finite splits give convex combinations of points $(\mu, V(\mu))$, whose upper boundary is $\operatorname{cav} V$. The argument below shows that the supremum is attained. □

Existence. Combine signals leading to the same action. Write $x(\omega, a)$ for the resulting joint probabilities. They satisfy

$$x(\omega, a) \geq 0, \quad \sum_a x(\omega, a) = \mu_0(\omega), \quad \sum_\omega x(\omega, a)[u(a, \omega) - u(b, \omega)] \geq 0 \quad (a, b \in A).$$

The last inequality says that action a is optimal after recommendation a. These constraints define a nonempty compact set: no information and a prior-optimal action are feasible. The linear objective $\sum_{\omega,a} x(\omega, a)v(a, \omega)$ therefore attains a maximum. Conversely, any such x defines an experiment with obedient action labels. Designer-favourable ties can only increase its payoff. Since every experiment also gives a feasible x, the two maxima coincide.

Number of posteriors. Suppose an optimal split uses $k > |\Omega|$ beliefs μ_j with positive weights q_j. The vectors are linearly dependent: some nonzero d satisfies $\sum_j d_j \mu_j = 0$, and hence $\sum_j d_j = 0$. Small changes from q to $q \pm \varepsilon d$ preserve feasibility. Optimality implies $\sum_j d_j V(\mu_j) = 0$; otherwise one direction improves the payoff. Increase ε until a weight is zero and repeat. At most $|\Omega|$ beliefs remain. In particular, two posteriors suffice for binary states.

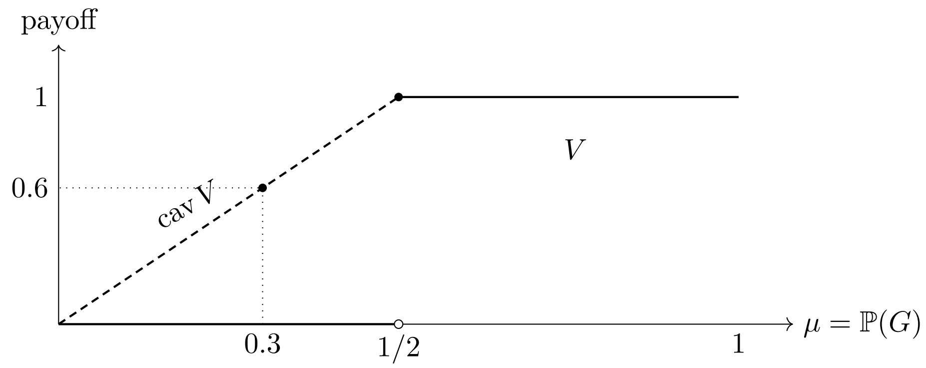

In the prosecutor example,

$$V(\mu) = \mathbf{1}_{\{\mu \geq 1/2\}}, \qquad (\mathrm{cav}\, V)(\mu) = \min\{2\mu, 1\}.$$

At prior 0.3, use posteriors 0 and 1/2, with weights 0.4 and 0.6. The value is 0.6.

Some useful cases:

- If V is concave, no information is optimal.
- If V is convex, full revelation is optimal. Indeed, $V(\mu) \leq \sum_\omega \mu(\omega)V(\delta_\omega)$.
- Strict convexity gives a strict improvement over no information when the prior is not degenerate.

6.1 Supporting bounds and duality

Let a be affine in the posterior and satisfy $a(\mu) \geq V(\mu)$. Every feasible experiment then satisfies

$$\mathbb{E}_\tau[V(\mu)] \leq \mathbb{E}_\tau[a(\mu)] = a(\mu_0).$$

If the experiment uses only beliefs where $a = V$, the bound is attained. This proves optimality without comparing the candidate with every experiment.

In the prosecutor example, $a(\mu) = 2\mu$ is such a bound. It touches V at 0 and 1/2, precisely the posteriors used by the optimal experiment.

More generally, write $a(\mu) = \sum_\omega \lambda_\omega \mu(\omega)$. The finite-state dual problem is

$$\min_{\lambda \in \mathbb{R}^{|\Omega|}} \sum_\omega \lambda_\omega \mu_0(\omega) \quad \text{subject to} \quad \sum_\omega \lambda_\omega \mu(\omega) \geq V(\mu) \text{ for all } \mu.$$

An affine constant can be absorbed into λ, since beliefs sum to one. The prior is in the relative interior of the simplex. The supporting-hyperplane theorem for a finite concave function therefore gives an affine a with $a \geq \operatorname{cav} V$ and $a(\mu_0) = (\operatorname{cav} V)(\mu_0)$. This a is feasible in the displayed problem and attains the persuasion value; the preceding upper bound shows that no smaller value is possible. Thus the dual and concavification values agree. We will use a related bound for posterior means in Lecture 2.

Reference. Kamenica and Gentzkow (2011), Proposition 1 and Corollary 2, for the posterior formulation and concavification. Dworczak and Kolotilin (2024), Section 3, for the general persuasion duality and its regularity conditions.

Lecture 2

Informativeness and Information Acquisition

1 Blackwell Informativeness

Fix a finite state space Ω and a full-support prior μ_0. An experiment $\sigma = (S, \Pi)$ induces a distribution τ of posterior beliefs. We want to compare experiments without specifying a particular decision problem.

For actions A and payoff u, define the value at posterior μ:

$$V_A^u(\mu) = \max_{a \in A} \sum_{\omega} u(a, \omega)\mu(\omega).$$

This function is convex: it is the maximum of linear functions. For binary states, every action gives a straight line in the probability of the high state. The decision maker chooses the upper envelope.

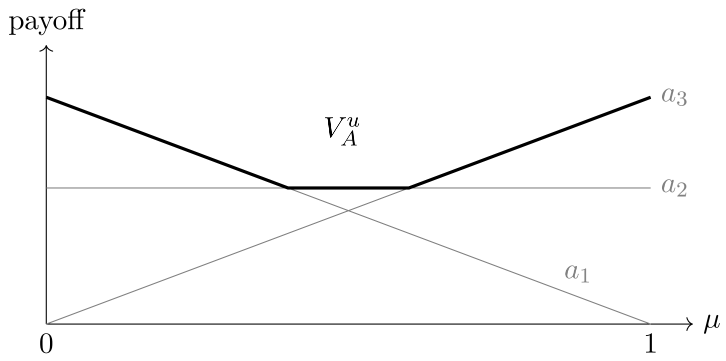

The ex-ante value of the experiment is $\mathbb{E}_\tau[V_A^u(\mu)]$. By Jensen's inequality and Bayes plausibility,

$$\mathbb{E}_\tau[V_A^u(\mu)] \geq V_A^u(\mu_0).$$

Information weakly benefits a decision maker who can ignore it. This statement holds with the decision problem fixed.

Definition. Experiment σ Blackwell dominates σ', written $\sigma \succeq_B \sigma'$, if

$$\mathbb{E}_\tau[V_A^u(\mu)] \geq \mathbb{E}_{\tau'}[V_A^u(\mu)]$$

for every finite decision problem (A, u), at the fixed prior μ_0.

The inequality is weak. Some decision makers choose the same action under both experiments and therefore obtain the same payoff.

1.1 Garbling

Experiment $\sigma' = (S', \Pi')$ is a garbling of $\sigma = (S, \Pi)$ if there is a stochastic matrix $K(s, s')$ such that

$$\Pi'_\omega(s') = \sum_s \Pi_\omega(s)K(s, s'), \qquad K(s, s') \geq 0, \quad \sum_{s'} K(s, s') = 1.$$

Generate s first, then draw s' using K. Conditional on s, the additional randomisation is independent of the state.

Conditional independence and the law of iterated expectations give

$$\mu_{s'} = \mathbb{E}[\mu_s \mid s'].$$

A garbling averages posterior beliefs. Its distribution of posteriors is a mean-preserving contraction of the original distribution.

Proposition 2.1 (Blackwell)**.** *For finite experiments and a fixed full-support prior, the following are equivalent:*

1. $\sigma \succeq_B \sigma'$.
2. σ' *is a garbling of* σ.
3. $\mathbb{E}_\tau[Z(\mu)] \geq \mathbb{E}_{\tau'}[Z(\mu)]$ *for every continuous convex function* Z *on the belief simplex.*

Proof. Garbling implies (3) by conditional Jensen, and (3) implies (1) because V_A^u is continuous and convex. It remains to show (1) implies (2).

Let $\mathcal{C} = \{\Pi K : K \text{ is stochastic}\}$, a compact convex set of matrices with columns indexed by S'. If $\Pi' \notin \mathcal{C}$, there is a matrix D such that

$$\langle D, \Pi' \rangle > \max_{Q \in \mathcal{C}} \langle D, Q \rangle, \qquad \langle D, Q \rangle = \sum_{\omega, s'} D_{\omega s'} Q_{\omega s'}.$$

To see the separation directly, choose $Q^* \in \mathcal{C}$ closest to Π' and set $D = \Pi' - Q^*$. Minimising squared distance along each segment from Q^* to Q gives $\langle D, Q - Q^* \rangle \leq 0$, whereas $\langle D, \Pi' - Q^* \rangle = \|D\|^2 > 0$.

Consider actions S' and payoffs $u(s', \omega) = D_{\omega s'}/\mu_0(\omega)$. Under Π, decision rules are precisely stochastic matrices K; their best payoff is the maximum on the right. Under Π', choosing action s' after signal s' gives the left-hand side. Thus Π' performs strictly better in this decision problem, contradicting (1). $\square$

Full support matters when recovering a kernel in every state. If the prior assigns probability zero to some states, restrict the comparison to its support. We will use only this fixed-prior formulation.

1.2 A binary example

Let the prior put probability 1/2 on each state, L, H. The more informative experiment has kernel

$$\widetilde{\Pi} = \begin{pmatrix} 3/4 & 1/4 \\ 1/4 & 3/4 \end{pmatrix},$$

with rows L, H and columns $\widetilde{s}_1, \widetilde{s}_2$. Its two posteriors on H are $1/4, 3/4$, each with probability 1/2. Garble using

$$K = \begin{pmatrix} 5/6 & 1/6 \\ 1/6 & 5/6 \end{pmatrix}.$$

Then

$$\Pi = \widetilde{\Pi} K = \begin{pmatrix} 2/3 & 1/3 \\ 1/3 & 2/3 \end{pmatrix}.$$

The garbled posteriors are $1/3, 2/3$, again with equal probabilities. Both pairs average to $1/2$; the first pair is more spread out.

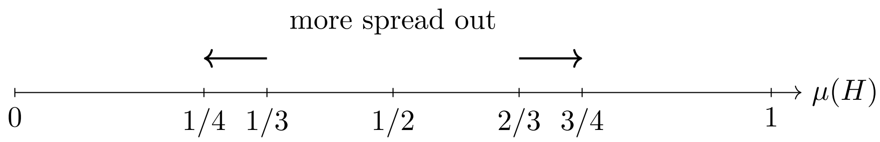

Reference. Blackwell (1953), for the equivalence between comparison of decision problems and garbling.

2 Majorization and Posterior Means

For scalar random variables, the relevant order has a simple form.

Definition. For distributions G, F on $[0,1]$, write $G \preceq_{\text{cx}} F$ if

$$\mathbb{E}_G[\phi(X)] \le \mathbb{E}_F[\phi(X)] \quad \text{for every continuous convex } \phi.$$

Then G is a mean-preserving contraction of F, or F is a mean-preserving spread of G.

Taking $\phi(x) = x$ and $\phi(x) = -x$ shows that the means are equal. In terms of CDFs, the equivalent test is

$$\int_0^t G(x)\,dx \le \int_0^t F(x)\,dx \quad (0 \le t \le 1), \qquad \int_0^1 G(x)\,dx = \int_0^1 F(x)\,dx. \tag{2.1}$$

Here $x_+ = \max\{x, 0\}$. Integrating the indicator of $X \le x$ gives

$$\int_0^t F(x)\,dx = \mathbb{E}_F[(t - X)_+], \qquad \mathbb{E}_F[X] = 1 - \int_0^1 F(x)\,dx.$$

Piecewise-linear convex functions are affine functions plus nonnegative combinations of the hinges $(t - x)_+$: each coefficient is an increase in slope. Thus the CDF inequalities and equal means imply the convex inequalities for these functions. Linear interpolation on successively finer grids approximates any continuous convex function uniformly, proving the stated equivalence.

In the binary-state Blackwell comparison, apply this test to the distributions of the posterior probability. More information means a more dispersed distribution of posterior beliefs.

2.1 Feasible posterior means

Now let the state itself be scalar, $\theta \in [0,1]$, with CDF F_0. After observing a signal, the receiver's posterior mean is

$$M = \mathbb{E}[\theta \mid s].$$

Write G for the distribution of M.

Proposition 2.2 (Posterior-mean feasibility)**.** *A distribution G can be induced as the distribution of posterior means if and only if $G \preceq_{\text{cx}} F_0$. Signals need not be finite here.*

Proof. For every convex ϕ, conditional Jensen gives

$$\mathbb{E}[\phi(M)] = \mathbb{E}[\phi(\mathbb{E}[\theta \mid s])] \le \mathbb{E}[\phi(\theta)].$$

For sufficiency we construct a joint distribution with $\mathbb{E}[\theta \mid M] = M$, then use M as the signal.

Finite distributions. Let G put masses g_j at m_j, and F_0 put masses f_i at x_i. Include $0, 1$ among the x_i, allowing zero masses. For each m_j, choose a distribution w^j on the x_i with mean m_j. Let $\mathcal{S}$ be the set of mixtures $\sum_j g_j w^j$. It is nonempty, compact and convex. We need the vector f to belong to $\mathcal{S}$.

Otherwise, the separating argument in the Blackwell proof gives numbers h_i with $h \cdot f < \min_{s\in\mathcal{S}} h \cdot s$. Define

$$\phi(m) = \min_{\substack{w_i \ge 0,\ \sum_i w_i = 1 \\ \sum_i w_i x_i = m}} \sum_i w_i h_i.$$

This is the lower convex envelope of the points (x_i, h_i): mixing feasible weights proves convexity, and the finite polygon makes it continuous. In particular, $\phi(x_i) \le h_i$. Minimising independently at each m_j gives

$$\sum_i f_i h_i < \sum_j g_j \phi(m_j) \le \sum_i f_i \phi(x_i) \le \sum_i f_i h_i,$$

where the middle inequality is $G \preceq_{\text{cx}} F_0$. This contradiction proves $f \in \mathcal{S}$. The probabilities $\mathbb{P}(M = m_j, \theta = x_i) = g_j w_i^j$ are the required joint distribution.

General distributions. Approximate G by replacing values in each positive-probability interval of a grid of mesh $1/n$ by their conditional mean. This gives a finite $G_n \preceq_{\text{cx}} G$. Approximate F_0 by randomly rounding each value to the two adjacent grid points, with probabilities preserving its mean. This gives $F_0 \preceq_{\text{cx}} F_n$. Both approximations change values by at most $1/n$, and $G_n \preceq_{\text{cx}} F_n$. The finite argument supplies a joint distribution (M_n, θ_n) with these marginals and $\mathbb{E}[\theta_n \mid M_n] = M_n$.

Take a weakly convergent subsequence of these joint distributions on the compact square.[1] Its marginals are G, F_0. For every continuous h, pass to the limit in

$$\mathbb{E}[(\theta_n - M_n)h(M_n)] = 0.$$

The limit satisfies $\mathbb{E}[(\theta - M)h(M)] = 0$, which characterises $\mathbb{E}[\theta \mid M] = M$. Continuous test functions suffice, since they approximate interval indicators in these integrals. □

Full revelation induces $G = F_0$; no information induces a point mass at $\mathbb{E}[\theta]$. Any feasible distribution lies between these two in convex order; the lower comparison follows from Jensen. Equal means alone are not sufficient. A Bernoulli distribution on $\{0, 1\}$ with mean $1/2$ cannot be induced from a uniform state: its second moment is $1/2$, while the state's second moment is $1/3$, violating the convex-order inequality for $\phi(x) = x^2$.

There are two different objects:

- A distribution over full posterior beliefs must average to the prior.

[1] Weak convergence means convergence of integrals of continuous functions. Compactness here follows by taking convergent subsequences of cell probabilities on successively finer finite grids. A diagonal subsequence gives convergence of continuous-function integrals because continuous functions on the square are uniformly continuous.

- A distribution over posterior means must be a contraction of the state distribution.

With more than two states, posterior means need not contain all the information in the posteriors. On states $\{0, 1/2, 1\}$, the point mass at $1/2$ and equal probabilities on $0, 1$ have the same mean but different beliefs.

Reference. Strassen (1965), for the martingale coupling; Dworczak and Kolotilin (2024), Section 4, pp. 1710–1712, and Kleiner et al. (2021), Section 4.4, for the application to posterior means.

2.2 Mean-based persuasion

Suppose the receiver's best reply, including tie-breaking, depends only on $m = \mathbb{E}[\theta \mid s]$. For example, if $u(a, \theta) = b(a) + c(a)\theta$, the expected payoff is $b(a) + c(a)m$, so the set of best replies depends only on m; use a tie rule with the same property. Suppose also that the designer's expected payoff, after the best reply, is a bounded measurable function $W(m)$. The problem becomes

$$\sup_{G \preceq_{\mathrm{cx}} F_0} \int_0^1 W(m)\, dG(m).$$

The designer's payoff need not equal the receiver's value.

This formulation replaces a distribution over beliefs by a distribution over numbers. The objective is linear in G, while feasibility is given by (2.1).

We write a supremum because existence has not been assumed for an arbitrary W. In the examples below an explicit construction attains the bound.

Reference. Kleiner et al. (2021), Theorem 2 and Section 4.4, for the general extreme-point characterization.

2.3 Convex bounds

Let ϕ be continuous and convex, with $\phi(m) \geq W(m)$. Every feasible G satisfies

$$\int W\, dG \leq \int \phi\, dG \leq \int \phi\, dF_0. \tag{2.2}$$

The last term is an upper bound on the designer's payoff.

To attain it, check two things:

1. The induced means are contact points: $\phi(m) = W(m)$.
2. Pooling does not reduce the expected value of ϕ.

The second condition holds if the signal pools only within regions on which ϕ is affine. Revealing a state also preserves ϕ.

The first lecture used affine bounds on posterior beliefs. Here the extra freedom to use a convex bound comes from the convex-order constraint. Finding contact points and affine regions often suggests the experiment.

Example (A threshold decision)**.** Let θ be uniform on $[0, 1]$. The receiver invests if $m \geq 3/4$, and the designer wants investment:

$$W(m) = \mathbf{1}_{\{m \geq 3/4\}}.$$

Pool the upper interval $[1/2, 1]$, whose conditional mean is $3/4$. Reveal states below $1/2$. Investment occurs with probability $1/2$.

To prove optimality, use

$$\phi(m) = 4(m - 1/2)_+.$$

This function is convex and dominates W. Moreover,

$$\int_0^1 \phi(\theta)\, d\theta = \int_{1/2}^1 4(\theta - 1/2)\, d\theta = \frac{1}{2}.$$

The candidate attains the bound. The upper pooled interval is an affine region of ϕ, and its posterior mean $3/4$ is a contact point.

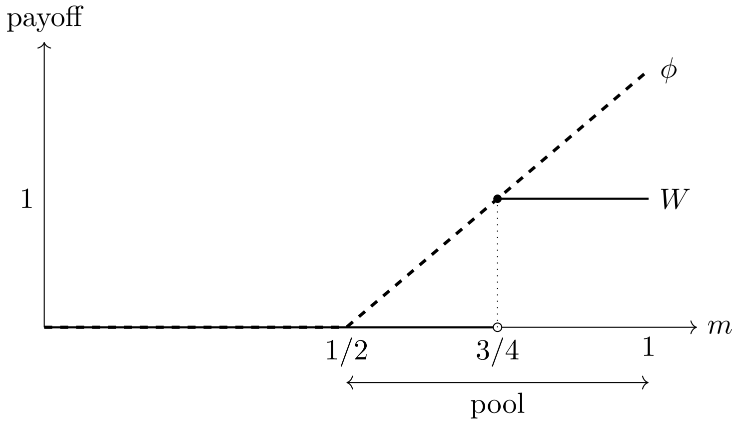

The bound (2.2) is sufficient whenever a candidate attains it. Stronger duality results state when an optimal bound exists; their proofs are not needed for this calculation.

Reference. Dworczak and Martini (2019), Theorem 1, for convex certificates; Kolotilin (2018), for the linear-programming approach; Dworczak and Kolotilin (2024), Sections 3–4, for general and moment persuasion duality.

3 Cost of Information

Return to finite states. Let $C(\mu_0, \tau)$ be the cost of inducing posterior distribution τ. Experiments that induce the same distribution are assigned the same cost; alternatively, use the least costly implementation.

Natural requirements are

$$C(\mu_0, \delta_{\mu_0}) = 0, \qquad C(\mu_0, \tau) \geq 0,$$

and monotonicity:

$$\tau \text{ is more informative than } \tau' \quad \Longrightarrow \quad C(\mu_0, \tau) \geq C(\mu_0, \tau').$$

This is a one-way implication. For example, $C \equiv 0$ is monotone, but equal costs do not imply equal informativeness. At a nondegenerate prior, a receiver paid one for guessing the state correctly obtains one under full revelation and only $\max_\omega \mu_0(\omega) < 1$ without information.

3.1 Shannon's mutual information

The entropy of a belief is

$$H(\mu) = -\sum_{\omega} \mu(\omega) \log \mu(\omega), \qquad 0 \log 0 = 0.$$

We use natural logarithms unless a base is specified. For binary states,

$$H(\mu) = -\mu \log \mu - (1-\mu) \log(1-\mu).$$

It is concave, zero at $0, 1$, and maximal at $1/2$. Indeed,

$$H'(\mu) = \log \frac{1-\mu}{\mu}, \qquad H''(\mu) = -\frac{1}{\mu(1-\mu)} < 0 \quad (0 < \mu < 1).$$

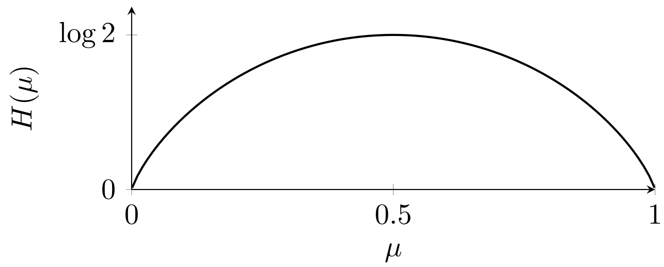

Entropy is related to the number of yes-or-no questions needed to identify a state. For a three-sided die with probabilities $0.5, 0.3, 0.2$, entropy using base-two logarithms is approximately 1.4855 bits. Asking first "Is it the first outcome?" and then, if necessary, distinguishing the other two outcomes takes 1.5 questions on average. Entropy is a lower bound, not in general the exact expected number of questions for a single draw. For the lower bound, let ℓ_i be the number of questions leading to outcome i. A binary decision tree satisfies $\sum_i 2^{-\ell_i} \leq 1$: assign half the incoming mass to each branch, starting from mass one. Concavity of $\log_2$ gives

$$H_2(p) - \sum_i p_i \ell_i = \sum_i p_i \log_2 \frac{2^{-\ell_i}}{p_i} \leq \log_2 \sum_i 2^{-\ell_i} \leq 0.$$

Outcomes of zero probability are omitted.

The expected reduction in entropy is mutual information:

$$I(\omega; s) = H(\mu_0) - \mathbb{E}_\tau[H(\mu)].$$

It is nonnegative by Jensen and zero for no information. More informative experiments reduce expected posterior entropy and therefore increase mutual information. A realised posterior can have higher entropy than the prior; the claim concerns the expectation.

3.2 Posterior-separable costs

The same construction works with any continuous concave uncertainty function H:

$$C(\mu_0, \tau) = H(\mu_0) - \mathbb{E}_\tau[H(\mu)].$$

Using the same function H across priors gives a uniformly posterior-separable cost. Examples include Shannon entropy and $H(\mu) = 1 - \sum_\omega \mu(\omega)^2$. For a scalar state, posterior variance is another concave uncertainty measure: the law of total variance gives $\mathrm{Var}_{\mu_0}(\theta) - \mathbb{E}_\tau[\mathrm{Var}_\mu(\theta)] = \mathrm{Var}_\tau(\mathbb{E}_\mu[\theta]) \geq 0$.

Not every monotone cost has this form. For example,

$$C(\mu_0, \tau) = \big(H(\mu_0) - \mathbb{E}_\tau[H(\mu)]\big)^2$$

is nonnegative, normalised, and Blackwell monotone. It is generally not posterior separable. At a fixed prior, a posterior-separable cost is affine in τ; this squared cost is not.

To see the difference at a nondegenerate prior, take H to be Shannon entropy and mix full revelation with no information, revealing which experiment was selected. If full revelation is selected with probability λ, the squared cost is $\lambda^2 H(\mu_0)^2$, rather than $\lambda H(\mu_0)^2$.

3.3 Costly persuasion

With cost $k[H(\mu_0) - \mathbb{E}_\tau H(\mu)]$, where $k \geq 0$, the designer's value is

$$\sup_{\tau : \mathbb{E}_\tau \mu = \mu_0} \mathbb{E}_\tau[V(\mu) + kH(\mu)] - kH(\mu_0).$$

Thus concavify $V + kH$, then subtract the constant $kH(\mu_0)$. The splitting lemma is unchanged.

In the prosecutor example, take $k = 1$ and binary Shannon entropy. Below $1/2$, the adjusted payoff is $H(\mu)$; at $1/2$ it jumps to $1 + \log 2$. A supporting line is tangent to H at $t < 1/2$ and passes through $(1/2, 1 + \log 2)$:

$$H(t) + (1/2 - t)H'(t) = 1 + \log 2.$$

Substituting the entropy formula reduces the left-hand side to $-\frac{1}{2}\log[t(1-t)]$. Hence $t(1-t) = e^{-2}/4$, giving

$$t = \frac{1 - \sqrt{1 - e^{-2}}}{2} \simeq 0.0351.$$

At prior 0.3, use posteriors $t, 1/2$. The probability on $1/2$ is

$$\frac{0.3 - t}{1/2 - t}.$$

The tangent line dominates H below $1/2$; above $1/2$ it is increasing while $1+H$ is decreasing. It is therefore a valid upper bound everywhere, and proves optimality. After the low signal, the posterior probability of guilt is now positive.

3.4 Rational inattention

A decision maker choosing her own information solves

$$\sup_{\tau : \mathbb{E}_\tau \mu = \mu_0} \mathbb{E}_\tau\left[\max_{a \in A} \mathbb{E}_\mu[u(a, \omega)]\right] - C(\mu_0, \tau).$$

This is the same choice of experiment, with the decision maker's value as the objective. Information is useful, but may not be worth its cost.

Reference. Sims (2003), for rational inattention; Caplin and Dean (2015) and Denti (2022), for information costs; Frankel and Kamenica (2019), for measures of information and uncertainty.

Lecture 3

Bayes Correlated Equilibrium

1 Information Design in Games

The basic game consists of players $i = 1, \dots, I$, a finite state space Ω, a common prior μ_0, finite action sets A_i, and payoffs $u_i(a, \omega)$, where $a \in A = \prod_i A_i$. The designer's payoff is $v(a, \omega)$. The designer has no action in the basic game, or can be represented by a player with a singleton action set.

Players have finite type sets T_i and observe an information structure $S = (T, \pi)$:

$$T = \prod_i T_i, \qquad \pi(t \mid \omega) \in \Delta(T).$$

Player i observes t_i. The basic game and S define a Bayesian game. If every T_i is a singleton, players have no private information.

A decision rule is a mapping

$$\sigma : T \times \Omega \longrightarrow \Delta(A).$$

View it as a mediator who observes the state and all types, then privately recommends an action to each player. The induced joint distribution is

$$\psi(\omega, t, a) = \mu_0(\omega)\pi(t \mid \omega)\sigma(a \mid t, \omega).$$

1.1 Obedience

Definition. A decision rule is a Bayes correlated equilibrium (BCE) if, for every i, t_i, a_i, a_i',

$$\sum_{\omega, t_{-i}, a_{-i}} \psi(\omega, t, a)\big[u_i(a_i, a_{-i}, \omega) - u_i(a_i', a_{-i}, \omega)\big] \geq 0. \tag{3.1}$$

The player knows her type and recommendation. Holding the resulting beliefs fixed, she prefers to obey. To write this as conditional expected utility, divide by $\mathbb{P}(t_i, a_i)$. The common denominator cancels. If that event has probability zero, the inequality is vacuous.

The restrictions are linear. They allow correlation in recommendations, both with the state and across players. When both the state and the initial information are trivial, they reduce to the usual correlated-equilibrium constraints.

1.2 Bayes Nash equilibrium and information expansions

A Bayes Nash strategy profile consists of $\beta_i : T_i \to \Delta(A_i)$. If $\beta_i(a_i \mid t_i) > 0$, action a_i must maximise expected utility conditional on t_i, taking opponents' strategies as given. Conditional on the full type profile, independent private randomisation gives

$$\mathbb{P}(a \mid t) = \prod_i \beta_i(a_i \mid t_i).$$

Actions can still be correlated unconditionally because types are correlated.

An expansion adds signals z_i to the existing types. Its distribution $\pi^*(t, z \mid \omega)$ must preserve the original information:

$$\sum_z \pi^*(t, z \mid \omega) = \pi(t \mid \omega).$$

Players then use strategies $\beta_i(a_i \mid t_i, z_i)$. Integrating out z induces a decision rule on the original game:

$$\sigma(a \mid t, \omega) = \frac{\sum_z \pi^*(t, z \mid \omega) \prod_i \beta_i(a_i \mid t_i, z_i)}{\pi(t \mid \omega)}$$

whenever the denominator is positive. Different expansions can preserve the same original information while correlating the additional signals differently.

Proposition 3.1 (Bergemann–Morris)**.** *A decision rule is a BCE of the original game if and only if it is induced by a Bayes Nash equilibrium of some information expansion.*

Proof. Given a BCE, use the recommended actions themselves as additional private signals: set $\pi^*(t, a \mid \omega) = \pi(t \mid \omega)\sigma(a \mid t, \omega)$. Obedience says that following recommendations is a Bayes Nash equilibrium.

Conversely, in an equilibrium of an expansion, whenever $\beta_i(a_i \mid t_i, z_i) > 0$,

$$\mathbb{E}[u_i(a_i, a_{-i}, \omega) - u_i(a_i', a_{-i}, \omega) \mid t_i, z_i] \geq 0.$$

Multiply by $\mathbb{P}(t_i, z_i)\beta_i(a_i \mid t_i, z_i)$ and sum over z_i. The player's private randomisation is independent of the state and opponents' actions conditional on t_i, z_i. Thus the resulting joint probabilities are exactly those in (3.1). Zero-probability types and unused actions contribute zero. □

An analyst who knows the game and initial information, but not what else players know, obtains the set of possible outcomes by finding BCE. A designer who can supply the additional information chooses her preferred point in that set.

Reference. Bergemann and Morris (2016), Theorem 1.

1.3 The designer's linear program

It is often simplest to work directly with ψ:

$$\begin{aligned} \max_{\psi} \quad & \sum_{\omega,t,a} \psi(\omega, t, a) v(a, \omega) \\ \text{subject to} \quad & \psi(\omega, t, a) \geq 0, \\ & \sum_a \psi(\omega, t, a) = \mu_0(\omega)\pi(t \mid \omega) \quad \text{for every } \omega, t, \\ & \text{the obedience inequalities (3.1).} \end{aligned}$$

The second constraint preserves the distribution of states and initial types. It also normalises the probabilities. With finite sets this is a finite linear program.

2 Examples

2.1 One-player investment

Let $\Omega = \{B, G\}$, with probability 1/2 on each state. The receiver chooses whether to invest:

	B	G
invest	-1	x
not invest	0	0

$0 < x < 1.$

The designer maximises the probability of investment.

With no initial private information, write p_B, p_G for the probability of recommending investment in each state. The joint probabilities are

	B	G
invest	$p_B/2$	$p_G/2$
not invest	$(1-p_B)/2$	$(1-p_G)/2$

Obedience requires

$$-p_B + xp_G \geq 0, \qquad -(1-p_B) + x(1-p_G) \leq 0.$$

The first condition implies the second, since $xp_G - p_B \geq 0 > x - 1$. Thus the problem is

$$\max_{0 \leq p_B, p_G \leq 1} \frac{p_B + p_G}{2} \quad \text{subject to } p_B \leq xp_G.$$

The optimum is

$$p_G = 1, \qquad p_B = x, \qquad \mathbb{P}(\text{invest}) = \frac{1+x}{2}.$$

An upper-bound proof is immediate:

$$p_B + p_G \leq (1+x)p_G \leq 1 + x.$$

The candidate attains both inequalities.

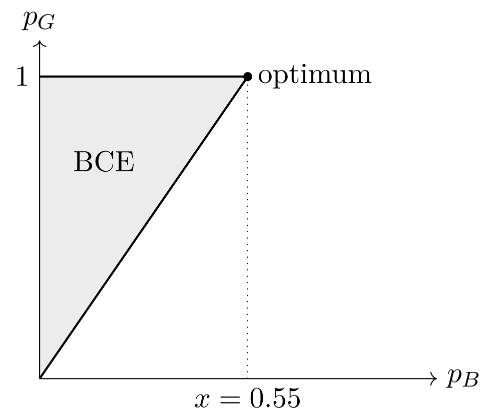

Following an investment recommendation, the receiver is indifferent. Following a recommendation not to invest, the receiver knows the state is bad.

2.2 Investment with private information

The receiver now observes $t \in \{b, g\}$, with precision $q \in [1/2, 1)$:

	B	G
b	q	$1-q$
g	$1-q$	q

The entries are conditional signal probabilities given the state. The posteriors on G are $1-q$ after b and q after g.

Allow the designer to condition recommendations on both the state and the receiver's type, as in the BCE formulation. Write p_{Bt}, p_{Gt} for investment probabilities. If the posterior on G at type t is z, the two obedience inequalities are

$$zxp_{Gt} - (1-z)p_{Bt} \geq 0,$$

$$zx(1-p_{Gt}) - (1-z)(1-p_{Bt}) \leq 0.$$

Equivalently,

$$zxp_{Gt} - (1-z)p_{Bt} \geq \max\{0, zx - (1-z)\}.$$

These conditions do not restrict which precisions q are possible. They restrict the recommendations at each type.

The investment-maximising choice is

$$p_{Gt} = 1, \qquad p_{Bt} = \min\left\{1, \frac{xz}{1-z}\right\}.$$

Indeed, $p_{Gt} \leq 1$ and the investment constraint bound p_{Bt} by this quantity, so the candidate maximises both investment probabilities. If $xz < 1-z$, it attains zero in the combined inequality, whose right-hand side is zero. If $xz \geq 1-z$, it recommends investment with probability one and the no-investment event has probability zero.

Integrating over types gives $p_G = 1$ and

$$\begin{aligned} p_B &= q\,p_{Bb} + (1-q)p_{Bg} \\ &= x(1-q) + \min\{1-q, xq\} \\ &= \begin{cases} x, & q \leq 1/(1+x), \\ (1+x)(1-q), & q \geq 1/(1+x). \end{cases} \end{aligned}$$

For low precision, the maximum investment probability remains $(1+x)/2$. Once a good private signal alone induces investment, further information reduces the designer's ability to induce investment in the bad state. As $q \to 1$, $p_B \to 0$, while $p_G = 1$.

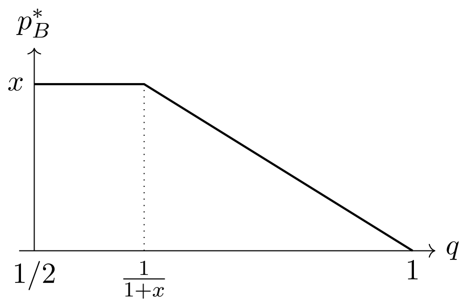

To obtain precision $q' \in [1/2, q]$ from $q > 1/2$, reverse each signal with probability $\rho = (q - q')/(2q - 1)$. The resulting precision is $q(1-\rho) + (1-q)\rho = q'$. Thus more precise signals are Blackwell more informative. They weakly shrink the set of attainable state-action distributions: garble the type in a BCE and sum its obedience inequalities with the garbling probabilities to obtain obedience for each coarser type. The calculation above concerns the designer's optimum; it does not assert that all BCE remain unchanged when the optimum is unchanged.

2.3 Two investors

There are two identical players and no initial private information. The two states still have probability 1/2 each. Not investing gives zero. Investing gives

$$\begin{cases} -1 + \varepsilon a_j, & \omega = B, \\ x + \varepsilon a_j, & \omega = G, \end{cases}$$

where $a_j = 1$ means that the other player invests. Positive ε gives strategic complements; negative ε gives substitutes. For this example take $0 < x < 1/2$. The restriction will keep the substitutes case simple.

The designer maximises the expected number of investors. Average any decision rule with the rule obtained by exchanging player labels. Identical payoffs and linear obedience constraints preserve feasibility; the objective is unchanged. We can therefore restrict attention to symmetric rules. Write p_ω for each player's marginal probability of investing and r_ω for the probability that both invest:

	invest	not invest
invest	r_ω	$p_\omega - r_\omega$
not invest	$p_\omega - r_\omega$	$1 - 2p_\omega + r_\omega$

Nonnegativity requires

$$\max\{0, 2p_\omega - 1\} \le r_\omega \le p_\omega.$$

In particular, p_ω is not the probability that only one player invests. The latter is $2(p_\omega - r_\omega)$.

After multiplying by two, the two distinct obedience inequalities are

$$-p_B + xp_G + \varepsilon(r_B + r_G) \ge 0,$$

$$-(1 - p_B) + x(1 - p_G) + \varepsilon(p_B - r_B + p_G - r_G) \le 0.$$

The designer's objective is $p_B + p_G$. These four variables and linear inequalities give the complete program.

2.4 Complements: public recommendations

Suppose $0 < \varepsilon < (1 - x)/2$. The investment constraint and $r_\omega \le p_\omega$ imply

$$p_B \le xp_G + \varepsilon(p_B + p_G), \qquad p_B \le \frac{x + \varepsilon}{1 - \varepsilon} p_G.$$

Hence the objective is bounded above by $1 + (x + \varepsilon)/(1 - \varepsilon)$. This bound is attained by

$$p_G = r_G = 1, \qquad p_B = r_B = \frac{x + \varepsilon}{1 - \varepsilon} < 1.$$

The no-investment constraint becomes $-(1 - p_B) \leq 0$.

Both players always receive the same recommendation. In the good state both invest. In the bad state the mediator recommends investment to both with probability p_B, and to neither otherwise. A public signal implements the optimum.

2.5 Substitutes: private recommendations

Suppose $-x < \varepsilon < 0$. The investment constraint, $r_B \geq 0$, and $r_G \geq 2p_G - 1$ imply

$$p_B \leq (x + 2\varepsilon)p_G - \varepsilon.$$

Consequently,

$$p_B + p_G \leq (1 + x + 2\varepsilon)p_G - \varepsilon \leq 1 + x + \varepsilon.$$

The coefficient of p_G is positive: $1 + x + 2\varepsilon > 1 - x > 0$. The bound is attained by

$$p_G = r_G = 1, \qquad p_B = x + \varepsilon, \qquad r_B = 0.$$

Since $0 < x + \varepsilon < 1/2$, all probabilities are nonnegative. The no-investment inequality becomes

$$-1 + (x + \varepsilon)(1 + \varepsilon) \leq 0.$$

Indeed, $0 < x + \varepsilon < 1/2$ and $0 < 1 + \varepsilon < 1$. The bad-state recommendations are

	invest	not invest
invest	0	$x + \varepsilon$
not invest	$x + \varepsilon$	$1 - 2x - 2\varepsilon$

while both players are told to invest in the good state. In the bad state, investment recommendations are negatively correlated. They are private. A public recommendation of the asymmetric action profile would reveal the bad state to the player asked to invest.

Thus complements favour common recommendations in this example, while substitutes favour private, negatively correlated recommendations. The conclusion follows from the obedience constraints and the objective; it is not a general ranking of public and private information.

Reference. Bergemann and Morris (2016), for BCE and comparison of information structures in games; Bergemann and Morris (2019), for a survey of information design.

Lecture 4

Market Segmentation and Price Discrimination

1 Information Design in Price Discrimination

A seller has zero marginal cost and posts a price to a buyer whose value has CDF F. The buyer purchases when $v \geq p$. Demand and profit are

$$D_F(p) = \mathbb{P}(v \geq p) = 1 - F(p-), \qquad \pi_F(p) = pD_F(p).$$

The left limit includes buyers whose value equals the price. This matters when F has atoms.

A segmentation is a signal about the buyer's value observed by the seller. The seller prices each segment separately. Equivalently, there is a population of buyers and the seller observes which segment each buyer belongs to.

We consider surplus pairs attainable by a segmentation and a seller best reply. When several prices are optimal, we specify which is selected. This tie-breaking convention affects buyer surplus, though not seller profit.

1.1 Two values

Let values be $0 < L < H$, with probability μ on H. The seller need only consider prices L, H: any lower price can be raised to the next supported value without losing buyers, and prices above H give zero revenue. Put $c = L/H$. Then

$$p(\mu) = \begin{cases} L, & \mu < c, \\ H, & \mu > c, \end{cases} \qquad \pi(\mu) = \max\{L, \mu H\}.$$

Both prices are optimal at $\mu = c$. Buyer surplus is $\mu(H - L)$ at price L, and zero at price H. The maximum possible welfare is

$$w(\mu) = (1 - \mu)L + \mu H.$$

By the splitting lemma, a distribution τ of segment posteriors is feasible if and only if

$$\int_0^1 z \, d\tau(z) = \mu.$$

Given a pricing rule, aggregate profit and buyer surplus are

$$\Pi = \int_0^1 \pi(z) \, d\tau(z), \qquad U = \int_0^1 u(z, p(z)) \, d\tau(z).$$

1.2 Monopolist-optimal information

Full revelation puts probabilities $1 - \mu, \mu$ on posteriors $0, 1$. The seller charges each buyer her value:

$$\Pi = w(\mu), \qquad U = 0.$$

This extracts all available surplus and is optimal for the seller.

1.3 No information

With no signal, the seller earns $\pi(\mu)$. If $\mu < c$, price L induces efficient trade and $U = \mu(H - L)$. If $\mu > c$, price H excludes the low-value buyers and $U = 0$. At $\mu = c$, both outcomes are seller-optimal.

1.4 Buyer-optimal information

Use the lower price at the indifference point c. Buyer surplus as a function of the seller's posterior is

$$u(z) = \begin{cases} z(H - L), & z \le c, \\ 0, & z > c. \end{cases}$$

Concavify this function.

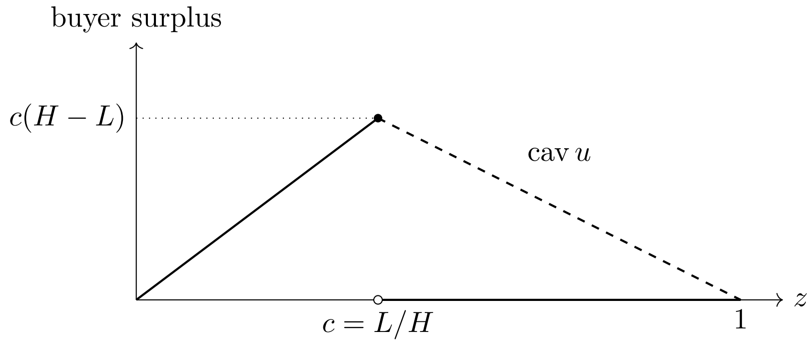

If $\mu \le c$, no information is buyer-optimal. If $\mu > c$, split between c and 1:

$$\mathbb{P}(z = c) = \alpha = \frac{1-\mu}{1-c}, \qquad \mathbb{P}(z = 1) = 1 - \alpha.$$

Charge L in the first segment and H in the second. Then

$$U = \alpha c(H - L) = (1 - \mu)L, \qquad \Pi = \alpha L + (1 - \alpha)H = \mu H.$$

Every buyer trades. The seller obtains the no-information profit and buyers receive the remaining welfare. For either value of μ,

$$U_{\max} = w(\mu) - \pi(\mu).$$

This is also an upper bound for every segmentation: total welfare is at most $w(\mu)$, and the seller can guarantee $\pi(\mu)$ by ignoring the signal.

1.5 Welfare minimisation

The seller can always ignore a signal, so profit cannot be below $\pi(\mu)$. Since buyer surplus is nonnegative, welfare cannot be below $\pi(\mu)$. This bound is attainable.

If $\mu \ge c$, use the same split between c and 1, but charge H in both segments. Buyer surplus is zero and profit is μH.

If $\mu \le c$, use posteriors $0, c$, with probability μ/c on c. Charge H at posterior c, and L at posterior 0. Both segments yield profit L, and buyer surplus is zero. Thus the minimum-welfare outcome is

$$(U, \Pi) = (0, \pi(\mu)).$$

1.6 The surplus triangle

The three outcomes constructed above are

$$(0, \pi(\mu)), \qquad (w(\mu) - \pi(\mu), \pi(\mu)), \qquad (0, w(\mu)).$$

Mixing segmentations, and revealing which one was selected, mixes their surplus pairs. Therefore every point in their convex hull is attainable. Conversely every outcome satisfies

$$U \geq 0, \qquad \Pi \geq \pi(\mu), \qquad U + \Pi \leq w(\mu).$$

These inequalities describe precisely that triangle.

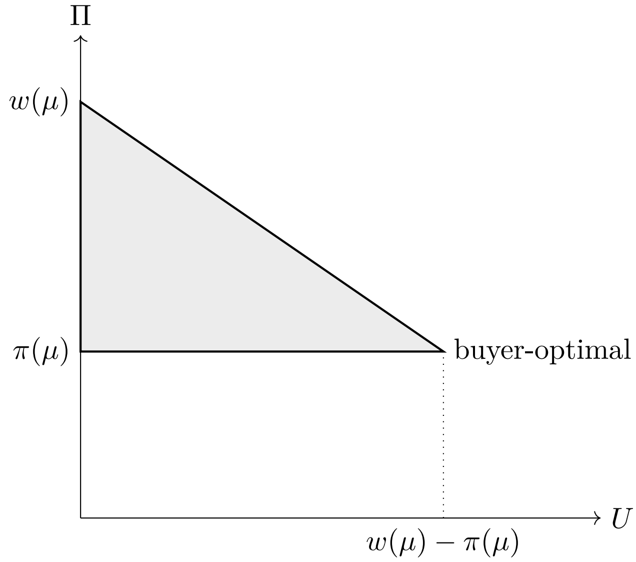

2 Bergemann, Brooks and Morris

The same conclusion holds with many values. We give the finite-value construction.

Let $V = \{v_1, \ldots, v_K\}$, with $0 < v_1 < \cdots < v_K$, and let $X = \Delta(V)$ be the set of markets. A market x gives the proportions of each value. At price v_k, profit is

$$R_k(x) = v_k \sum_{j \geq k} x(v_j).$$

Write x^* for the original market,

$$\pi^* = \max_k R_k(x^*), \qquad w^* = \sum_k v_k x^*(v_k).$$

A segmentation consists of markets x^s and weights α_s satisfying

$$\sum_s \alpha_s x^s = x^*, \qquad \alpha_s \geq 0, \quad \sum_s \alpha_s = 1.$$

For each k, define

$$X_k = \{x \in X : R_k(x) \geq R_j(x) \text{ for every } j\}.$$

It is a convex polytope: the simplex intersected with linear inequalities. We need a convenient set of markets from which to build points in X_k.

2.1 Extremal markets

For a nonempty support $S = \{s_1 < \cdots < s_m\} \subseteq V$, define

$$x^S(s_j) = \begin{cases} s_1(1/s_j - 1/s_{j+1}), & j < m, \\ s_1/s_m, & j = m, \end{cases}$$

and zero probability outside S. The terms telescope, so the probabilities sum to one and

$$\mathbb{P}_{x^S}(v \geq s_j) = \frac{s_1}{s_j}.$$

The seller earns s_1 at every price in S. A price between two supported values gives the same demand as the higher value and strictly less revenue; a price above the support gives zero. Thus precisely the prices in S are optimal. These are the *extremal markets*.

Proposition 4.1 (Extremal-market lemma)**.**

$$X_k = \mathrm{co}\{x^S : v_k \in S\}.$$

Here co *denotes the convex hull, the set of all finite mixtures.*

Every market where v_k is optimal is a mixture of extremal markets where v_k remains optimal. The construction below proves the lemma.

Reference. Bergemann et al. (2015), Lemma 1, pp. 933–934.

Choose a monopoly price $p^* = v_{k^*}$ for the original market. A segmentation into extremal markets all containing p^* is *uniform-profit-preserving* (UPP). It exists by the lemma. Its aggregate profit is

$$\sum_s \alpha_s R_{k^*}(x^s) = R_{k^*}(x^*) = \pi^*.$$

The seller can use any optimal price within each extremal segment without changing this profit.

2.2 The greedy construction

Use unnormalised residual masses $r(v)$, initially $r = x^*$. At each step:

1. Let $S = \{v : r(v) > 0\}$.
2. Construct x^S.
3. Remove a segment of mass
$$\alpha = \min_{v \in S} \frac{r(v)}{x^S(v)}.$$
4. Replace $r(v)$ by $r(v) - \alpha x^S(v)$.

At least one value disappears at each step, so the procedure ends in at most K steps. The residual is always nonnegative. Since each x^S has total mass one, the removed masses sum to one and their weighted markets recover x^*.

Why does it preserve profit? Removing an extremal market subtracts the same revenue, $\alpha \min S$, from every price in the current support. Thus the original optimal price remains optimal in the residual market. It cannot disappear while a nonempty residual remains: a price carrying no mass either sells nothing or can be raised to the next supported value without losing buyers. Neither can be optimal in a nonempty positive-value market. Every extracted support therefore contains p^*.

This also proves the extremal-market lemma. Apply the algorithm to any $x \in X_k$, keeping v_k as the chosen optimal price. All extracted markets have $v_k \in S$. Conversely, each such x^S is in X_k, and mixtures remain there because the revenue inequalities are linear.

The x^S with $v_k \in S$ are the vertices of X_k. To see this, any convex decomposition of x^S within X_k must preserve its zero masses outside S and its equalities $R_j = R_k$ for $v_j \in S$. On support S, these equalities uniquely determine the tails s_1/s_j, and hence determine x^S. Thus no nontrivial decomposition is possible. The convex-hull identity shows there are no other vertices.

Reference. Bergemann et al. (2015), Section II.C, pp. 937–940, for the greedy algorithm.

2.3 Three values

Let $V = \{1,2,3\}$ and $x^* = (1/3, 1/3, 1/3)$. The original revenues at prices $1, 2, 3$ are $1, 4/3, 1$, so $p^* = 2$, $\pi^* = 4/3$, and $w^* = 2$.

First use $x^{\{1,2,3\}} = (1/2, 1/6, 1/3)$. The maximum removable mass is

$$\min\left\{\frac{1/3}{1/2}, \frac{1/3}{1/6}, \frac{1/3}{1/3}\right\} = \frac{2}{3}.$$

The residual masses are $(0, 2/9, 1/9)$. Next use $x^{\{2,3\}} = (0, 1/3, 2/3)$, with mass $1/6$. The residual is $(0, 1/6, 0)$, giving a final segment containing only value 2.

segment	$x(1)$	$x(2)$	$x(3)$	α	lowest price	highest price
$\{1,2,3\}$	1/2	1/6	1/3	2/3	1	3
$\{2,3\}$	0	1/3	2/3	1/6	2	3
$\{2\}$	0	1	0	1/6	2	2

The weighted columns recover $(1/3, 1/3, 1/3)$.

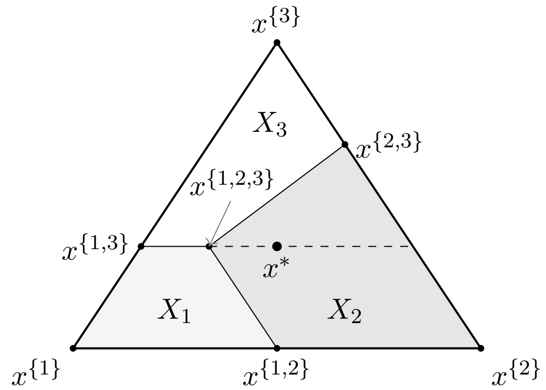

To implement the segmentation as a signal, use $\mathbb{P}(s \mid v) = \alpha_s x^s(v)/x^*(v)$. Here the kernel is

	$\{1,2,3\}$	$\{2,3\}$	$\{2\}$
$v=1$	1	0	0
$v=2$	1/3	1/6	1/2
$v=3$	2/3	1/3	0

Each row sums to one.

Lowest-price selection gives efficient trade:

$$\Pi = \frac{2}{3}(1) + \frac{1}{6}(2) + \frac{1}{6}(2) = \frac{4}{3}, \qquad U = 2 - \frac{4}{3} = \frac{2}{3}.$$

Highest-price selection gives the same profit and zero buyer surplus. Full revelation gives profit 2 and zero buyer surplus.

2.4 The general surplus result

Proposition 4.2 (BBM surplus triangle)**.** *The attainable surplus pairs are exactly*

$$U \geq 0, \qquad \Pi \geq \pi^*, \qquad U + \Pi \leq w^*.$$

Proof. The inequalities follow from voluntary trade, the seller's option to ignore information, and the bound on total welfare.

Take a UPP extremal segmentation. Choosing the lowest supported price in each segment makes every buyer trade, giving $(w^* - \pi^*, \pi^*)$. Choosing the highest supported price makes only highest-value buyers trade; each pays her value, giving $(0, \pi^*)$. Full revelation gives $(0, w^*)$. Mixing these constructions attains every point in the triangle. ☐

The construction preserves the seller's profit and reallocates surplus through information. Lecture 5 will use these binary segments when the platform can also choose the quality of the product assigned to each buyer.

Reference. Bergemann et al. (2015), Proposition 1 and Theorem 1 in Section II.B; Theorem 1B for the extension to a continuum of values.

Lecture 5

Prior Design, Learning, and Matching

1 Prior Information Design

In the previous lecture the distribution of values was fixed and the designer chose what the seller learned. We now let the buyer choose the distribution itself, following Condorelli and Szentes.

A seller has one object and zero cost. The buyer chooses a CDF F supported on $[0,1]$. The seller observes F and posts a price. The buyer observes her realised value and buys when $v \geq p$. The choice of F can represent an observable investment made before trade.

Write

$$D_F(p) = 1 - F(p-), \qquad \Pi(F) = \max_{p\in[0,1]} pD_F(p),$$

and, for a seller-optimal price p,

$$U(F,p) = \mathbb{E}_F[(v-p)_+] = \int_p^1 [1 - F(v)]\, dv.$$

The integral formula follows by integrating the indicator of $v > t$ over $t \in [p,1]$. The demand convention ensures $\limsup_{p'\to p} p'D_F(p') \leq pD_F(p)$: approaching from the left retains buyers at p, while approaching from the right can only lose them. Thus revenue is upper semicontinuous and attains its maximum on $[0,1]$. Its set of maximisers is compact, so a lowest optimal price exists. Buyer surplus is decreasing in price. Also,

$$U(F,p) + pD_F(p) = \mathbb{E}_F[v\mathbf{1}_{\{v\geq p\}}] \leq \mathbb{E}_F[v].$$

When constructing a buyer-optimal outcome, we use the lowest seller-optimal price. We first set the cost of choosing F equal to zero.

1.1 Equal-revenue distributions

For $0 < \pi \leq B \leq 1$, define

$$F_\pi^B(v) = \begin{cases} 0, & 0 \leq v < \pi, \\ 1 - \pi/v, & \pi \leq v < B, \\ 1, & B \leq v \leq 1. \end{cases} \tag{5.1}$$

There is density π/v^2 on (π, B) and an atom of size π/B at B. If $B = \pi$, the distribution is a point mass at π. For $\pi = 0$, define F_0^B as the point mass at zero.

For any $p \in [\pi, B]$,

$$D_{F_\pi^B}(p) = \frac{\pi}{p}, \qquad pD_{F_\pi^B}(p) = \pi.$$

This includes $p = B$, because buyers at the atom buy when indifferent. Prices below π yield revenue below π; prices above B yield zero. Thus every price in $[\pi, B]$ is optimal.

The mean and the buyer's surplus at the lowest price are

$$\mathbb{E}_{F_\pi^B}[v] = \pi + \pi \log(B/\pi), \tag{5.2}$$
$$U(F_\pi^B, \pi) = \pi \log(B/\pi). \tag{5.3}$$

For the mean, integrate survival: demand is one below π and π/p between π and B. At price π, everyone trades, so buyer surplus is the mean minus π.

1.2 The buyer's optimal prior

Proposition 5.1. *With costless distribution choice on* $[0, 1]$, *a buyer-optimal outcome is*

$$F = F^1_{1/e}, \qquad p = 1/e.$$

Both the buyer and seller obtain $1/e$.

Proof. Take any distribution G, let $\pi = \Pi(G) > 0$, and let p be an optimal price. The seller's revenue constraints imply

$$v\mathbb{P}_G(V \geq v) \leq \pi.$$

Hence $G(v) \geq 1 - \pi/v$ for $v \geq \pi$, and F^1_π first-order stochastically dominates G. Also $\pi = pD_G(p) \leq p$. Therefore

$$U(G, p) \leq \mathbb{E}_G[(V - \pi)_+] \leq \mathbb{E}_{F^1_\pi}[(V - \pi)_+] = -\pi \log \pi.$$

The derivative of $-\pi \log \pi$ is $-1 - \log \pi$, with second derivative $-1/\pi < 0$. The maximum is therefore at $\pi = 1/e$, with value $1/e$. The proposed distribution, with price $1/e$, attains this bound. If $\pi = 0$, then $\mathbb{P}(v \geq \varepsilon) = 0$ for every $\varepsilon > 0$; otherwise price ε would earn positive revenue. Thus the value and buyer surplus are zero almost surely. □

The method is to fix the seller's profit first. The constraints $vD(v) \leq \pi$ give an upper envelope for demand. Equal-revenue demand attains this envelope. The remaining choice is one-dimensional.

The buyer's expected value at the optimum is $2/e$. All trade takes place, but investment is inefficient: concentrating value at one would generate welfare one. The loss $1 - 2/e$ comes from the buyer's choice of distribution.

The low price can also be made uniquely optimal. Put $\pi = 1/e$ and move mass $\varepsilon \in (0, \pi)$ from the atom at one to π. Price π still earns π, while any $p \in (\pi, 1]$ earns $p(\pi/p - \varepsilon) = \pi - \varepsilon p < \pi$. Lower prices earn less too. Buyer surplus falls by only $\varepsilon(1 - \pi)$.

Reference. Condorelli and Szentes (2020), "Costless Distributions," Lemma 1 and Theorem 1, for the equilibrium analysis.

2 Distribution Costs and a Majorization Argument

Let $C(F)$ be a finite cost of choosing F. Assume it is lower semicontinuous under convergence in distribution and *increasing in risk*:

$$H \preceq_{\text{cx}} G \implies C(H) \leq C(G).$$

This is a weak inequality. Any continuous cost depending only on the mean satisfies it, since convex order preserves the mean. So does $C(F) = \mathbb{E}_F[c(v)]$ for continuous convex c, by the definition of convex order. Lower semicontinuity means $C(F) \le \liminf_n C(F_n)$ whenever F_n converges in distribution to F.

The following observation gives the main reduction.

Proposition 5.2 (Equal-revenue reduction)**.** *Let G have mean m, monopoly profit $\pi > 0$, and an optimal price p. There is a unique $B \in [\pi, 1]$ such that $H = F_\pi^B$ has mean m. Moreover,*

$$H \preceq_{\mathrm{cx}} G, \qquad U(G,p) \le U(H,\pi) = m - \pi.$$

Proof. First, $\pi = p\mathbb{P}_G(v \ge p) \le \mathbb{E}_G[v\mathbf{1}_{\{v \ge p\}}] \le m$. The previous revenue-envelope argument gives $m \le \mathbb{E}_{F_\pi^1}[v]$. The mean in (5.2) increases continuously from π to $\mathbb{E}_{F_\pi^1}[v]$ as B runs from π to one, giving the required B.

Below B, $H = F_\pi^1 \le G$; above B, $H = 1 \ge G$. Their means agree, so $\int_0^1 (H - G) = 0$. For $t \le B$, integrate $H - G \le 0$ from zero to t. For $t \ge B$, use $\int_0^t (H-G) = -\int_t^1 (H - G) \le 0$. Thus

$$\int_0^t H(v)\,dv \le \int_0^t G(v)\,dv \quad \text{for all } t,$$

with equality at one. This is $H \preceq_{\mathrm{cx}} G$.

Finally, total surplus under G is at most m, so buyer surplus is at most $m - \pi$. Under H, price π induces trade with probability one, attaining that bound. □

For example, take G uniform on $[0,1]$, with mean $1/2$. Revenue is $p(1-p)$, maximised at $p = 1/2$, giving profit $1/4$. Buyer surplus is $\int_{1/2}^1 (v - 1/2)\,dv = 1/8$. The replacement is $F_{1/4}^{e/4}$, since

$$\frac{1}{4} + \frac{1}{4}\log\left(\frac{e/4}{1/4}\right) = \frac{1}{2}.$$

It preserves mean and monopoly profit, reduces risk, and raises buyer surplus to 1/4.

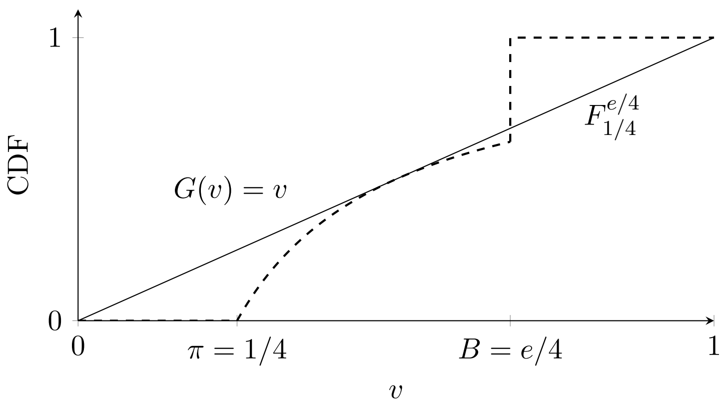

If cost is increasing in risk, the replacement is also weakly cheaper. Therefore an optimal choice can be found by solving

$$\max_{0 \le \pi \le B \le 1} \left\{ \pi \log(B/\pi) - C(F_\pi^B) \right\},$$

interpreting $\pi\log(B/\pi)$ as zero at $\pi = 0$. The objective there is $-C(\delta_0)$. Existence requires checking the boundary too. The map $(\pi, B) \mapsto F_\pi^B$ is continuous in distribution for $\pi > 0$, including $B = \pi$, directly from its CDF. As $\pi \to 0$,

$$0 \le \mathbb{E}_{F_\pi^B}[v] \le \pi(1 - \log\pi) \longrightarrow 0.$$

Thus F_π^B converges to the point mass at zero: for every $\varepsilon > 0$, its probability above ε is at most its mean divided by ε. The surplus also tends to zero. Consequently the objective is upper semicontinuous on the compact parameter triangle $0 \le \pi \le B \le 1$, and attains its maximum.

The argument proves existence of an optimum in this family; it does not establish uniqueness of the optimal distribution.

If cost depends only on the mean, an optimum can be chosen with $B = 1$. For a positive fixed mean, increasing B requires decreasing π. Indeed, $\pi[1 + \log(B/\pi)]$ increases in each variable on $0 < \pi < B$. Thus moving to $B = 1$ increases $m - \pi$ while leaving cost unchanged. A zero-mean choice is already $F_0^1 = \delta_0$.

Reference. Condorelli and Szentes (2020), "Increasing Cost," Lemma 2, Theorem 2, and Remark 1. The short proof above uses the integrated-CDF test from Lecture 2.

3 Roesler–Szentes as a Special Case

The buyer's underlying value now has a fixed CDF F_0 on $[0, 1]$. She chooses an experiment about this value. The seller observes the experiment but not its realisation. Given a signal, the buyer purchases when her posterior mean is at least the price.

Let G be the distribution of posterior means. From Lecture 2, the buyer can choose precisely

$$G \preceq_{\mathrm{cx}} F_0.$$

The seller's demand is $1 - G(p-)$, so the pricing problem is the same as before.

To embed this problem in Condorelli–Szentes, define

$$C(G) = \begin{cases} 0, & G \preceq_{\mathrm{cx}} F_0, \\ 2, & \text{otherwise.} \end{cases}$$

Infeasible choices are never worthwhile: gross buyer surplus is at most one, whereas a feasible uninformative experiment gives zero net surplus. The cost is increasing in risk. If $H \preceq_{\mathrm{cx}} G \preceq_{\mathrm{cx}} F_0$, then H is also free; if G costs two, H cannot cost more. The feasible set is closed: each inequality $\int \phi\, dG \le \int \phi\, dF_0$ for continuous convex ϕ is preserved by convergence in distribution. A cost equal to zero on a closed set and two outside is lower semicontinuous.

The equal-revenue reduction therefore applies. There is an optimal posterior-mean distribution F_π^B. Write $m_0 = \mathbb{E}_{F_0}[v]$. If $m_0 = 0$, the value is zero almost surely. For $m_0 > 0$, the problem becomes

$$\min_{0<\pi\le B\le 1} \pi \quad \text{subject to} \quad \pi[1 + \log(B/\pi)] = m_0, \qquad F_\pi^B \preceq_{\mathrm{cx}} F_0. \tag{5.4}$$

At the optimal low price π, trade is certain and buyer surplus is $m_0 - \pi$.

This is prior design with a feasibility cost. The mean is now fixed and the full majorization constraint must hold. In particular, the unrestricted answer $\pi = 1/e$ need not be feasible or optimal.

Reference. Condorelli and Szentes (2020), "Learning About Valuation" and the discussion following Remark 1, for this reduction. Roesler and Szentes (2017), for buyer-optimal learning and implementation.

3.1 Uniform prior

Let $F_0(v) = v$ on $[0, 1]$, so $m_0 = 1/2$. The mean restriction determines

$$B = \pi \exp\left(\frac{1}{2\pi} - 1\right).$$

For a candidate F_π^B, define the integrated-CDF slack

$$A(t) = \int_0^t [F_0(v) - F_\pi^B(v)]\, dv.$$

Feasibility requires $A(t) \geq 0$ for all t. Direct integration gives

$$A(t) = \begin{cases} t^2/2, & t \leq \pi, \\ t^2/2 - t + \pi + \pi \log(t/\pi), & \pi \leq t \leq B, \\ (1-t)^2/2, & t \geq B. \end{cases}$$

A candidate is obtained by making the middle expression touch zero at an interior minimum t^*:

$$A(t^*) = 0, \qquad A'(t^*) = t^* - 1 + \pi/t^* = 0.$$

Equivalently,

$$\pi = t^*(1 - t^*), \qquad \frac{t^*}{2} = -(1 - t^*) \log(1 - t^*).$$

There is exactly one root in $(1/2, 1)$. With $u = 1 - t$, the equation is $1 - u + 2u \log u = 0$; the left-hand side decreases from 1 to $1/2 - \log 2 < 0$ as u runs from 0 to $1/2$, since its derivative is $1 + 2 \log u < 0$. The resulting values are

$$t^* \simeq 0.71533, \qquad \pi^* \simeq 0.20363, \qquad B^* \simeq 0.87281.$$

The parameters are admissible. The equations imply

$$B^* = t^* \exp\left(\frac{1 - t^*}{2t^*}\right), \qquad \pi^* < t^* < B^* < 1.$$

The last inequality follows from $\log t^* \leq t^* - 1$ and $t^* > 1/2$. In the middle region, $A'(t) = (t^2 - t + \pi^*)/t$. Its two roots are $1 - t^*$, a local maximum, and t^*, a local minimum with value zero. The endpoint values are $A(\pi^*) = (\pi^*)^2/2 > 0$ and $A(B^*) = (1 - B^*)^2/2 > 0$. Hence the slack is nonnegative everywhere.

For a smaller π with the same mean, B increases, because $d \log B/d\pi = (2\pi - 1)/(2\pi^2) < 0$ for $\pi < 1/2$. If $B > 1$, the support constraint fails. Otherwise t^* is still in the middle region, where $\partial A(t^*)/\partial \pi = \log(t^*/\pi) > 0$. Reducing π therefore makes $A(t^*) < 0$. This proves global optimality.

Buyer surplus is approximately 0.29637, compared with 1/8 under full learning. The convex-order test verifies implementability; the tangency condition identifies the best equal-revenue distribution.

4 Buyer-Optimal Matching

There is another choice a platform can make: which product each buyer sees. This changes both the value of a match and what the seller infers. The binary-value case of Condorelli and Szentes gives a simple allocation argument.

4.1 The model

There is a unit mass of buyers, with values $0 < L < H$ for one unit of quality. A fraction $\mu \in (0,1)$ have value H. There is a unit mass of sellers. Each owns one product, with quality $q > 0$ drawn from an atomless CDF Q on $[\underline{q}, \overline{q}]$. A buyer with value v obtains $vq - p$ from buying quality q at price p.

The platform observes buyer values and chooses a matching. Quality is public. After matching, a seller makes a take-it-or-leave-it offer. The seller knows the matching rule, but not the matched buyer's value.

For a complete matching, let $z(q)$ be the probability that a quality-q seller meets a high-value buyer. Feasibility requires

$$0 \le z(q) \le 1, \qquad \int z(q)\, dQ(q) = \mu.$$

These conditions are also sufficient: assign joint masses $z(q)\, dQ(q)$ to high-value buyers and $(1 - z(q))\, dQ(q)$ to low-value buyers. Their marginals are the specified buyer and seller populations. As in Lecture 4, put $c = L/H$. The seller charges qL when $z(q) \le c$, and qH when $z(q) > c$. Use the lower price at a tie. Buyer surplus at this quality is

$$q(H - L)z(q)\mathbf{1}_{\{z(q) \le c\}}.$$

4.2 Allocate the rent-generating mixtures

An extremal mixture with high-value share c makes the seller indifferent between the two prices. At price qL, its buyer surplus is $qc(H - L)$. The same mixture is therefore more valuable at higher quality.

The largest mass of these mixtures that can be formed is

$$m = \min\left\{\frac{\mu}{c}, \frac{1-\mu}{1-c}\right\}.$$

Each unit uses c high-value buyers and $1 - c$ low-value buyers. Place this mixture at the highest-quality sellers, those above

$$q^* = Q^{-1}(1 - m).$$

Here $Q^{-1}(r) = \inf\{q \in [\underline{q}, \overline{q}] : Q(q) \ge r\}$. Since Q is continuous, the mass above q^* is m. Complete the matching with the remaining buyer type below q^*:

$$z^*(q) = \begin{cases} c, & q \ge q^*, \\ 0, & q < q^* \text{ and } \mu \le c, \\ 1, & q < q^* \text{ and } \mu \ge c. \end{cases} \tag{5.5}$$

The choice at the cutoff is irrelevant because Q has no atoms. If $\mu \le c$, the high-value mass is $cm = \mu$. If $\mu \ge c$, it is $cm + 1 - m = \mu$. Thus the matching is feasible in both cases. Every matched pair trades.

Optimality argument. In any matching, let $a(q) = z(q)\mathbf{1}_{\{z(q)\le c\}}$, the joint probability of a high-value buyer and the low price conditional on quality q. Then $0 \le a(q) \le c$, and

$$\int a(q)\,dQ(q) \le \mu, \qquad \frac{1-c}{c}\int a(q)\,dQ(q) \le 1-\mu.$$

The second inequality follows because a seller charging the low price must have at least $(1-c)/c$ low-value buyers per high-value buyer. Buyer surplus is $(H-L)\int q\,a(q)\,dQ(q)$. Let $a^*(q) = c\mathbf{1}_{\{q\ge q^*\}}$, whose total mass cm exhausts the smaller resource bound. For every feasible a,

$$\int q(a-a^*)\,dQ = \int (q-q^*)(a-a^*)\,dQ + q^*\int (a-a^*)\,dQ \le 0.$$

The first integrand is nonpositive on each side of q^*; the second integral is nonpositive by the resource bounds. Thus assigning these mixtures to the highest qualities is optimal. The matching (5.5) attains this bound. □

If $\mu < c$, high-value buyers are scarce and all of them enter the mixed segments. If $\mu > c$, low-value buyers are scarce and all of them enter those segments. The remaining high-value buyers are assigned to low-quality sellers who charge their full value. In this case the matching is stochastically negative assortative: $z^*(q)$ falls with quality.

4.3 An example

Take $L = 1$, $H = 2$, $\mu = 3/4$, and quality uniform on $[1, 2]$. Then $c = 1/2$, $m = 1/2$, and $q^* = 3/2$. The top half of sellers receive a half-high, half-low mixture and charge q. The bottom half receive only high-value buyers and charge $2q$.

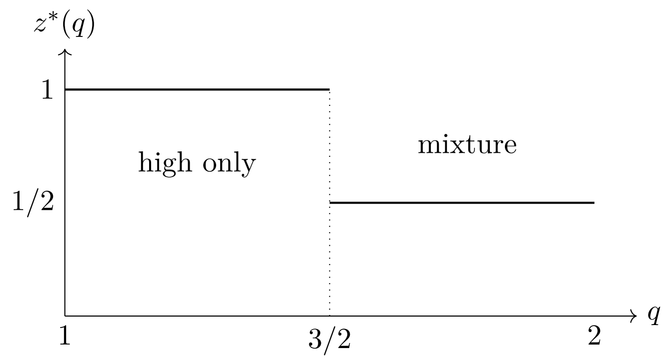

Buyer surplus is

$$\frac{1}{2}\int_{3/2}^{2} q\,dq = \frac{7}{16}.$$

All buyers trade, so total welfare is

$$\int_1^{3/2} 2q\,dq + \int_{3/2}^{2} \frac{3}{2}q\,dq = \frac{41}{16}.$$

Positive assortative matching instead places high-value buyers above quality 5/4. Its welfare is

$$\int_1^{5/4} q\,dq + \int_{5/4}^{2} 2q\,dq = \frac{87}{32}.$$

This assignment maximises total match value: exchanging a high-value buyer at q_1 with a low-value buyer at $q_2 > q_1$ raises value by $(H - L)(q_2 - q_1) > 0$. It also reveals buyer values through quality, so sellers charge each buyer her value and leave zero buyer surplus.

Thus the buyer-optimal platform trades off sorting efficiency against information rents. The method combines the binary segmentation from BBM with an allocation of scarce mixtures across qualities.

Reference. Condorelli and Szentes (2026), Sections 2–3 and Theorem 1; Proposition 1 for the positive-assortative welfare benchmark.

5 Comparison

In BBM the distribution of values is fixed and the seller receives information. In Condorelli–Szentes the buyer chooses the distribution before trade. In Roesler–Szentes that choice is restricted by the fixed prior, giving a majorization-constrained special case. In the matching problem, the platform also chooses which qualities are assigned to the different mixtures of buyers.

The methods are correspondingly different: concavification and decomposition; revenue envelopes and equal-revenue reduction; convex-order feasibility; and allocation by quality. Longer proofs and extensions are in the references. For related surplus bounds in Cournot oligopoly, see Condorelli and Szentes (2022).

References

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2015). "The Limits of Price Discrimination." *American Economic Review*, 105(3), 921–957.

Bergemann, Dirk, and Stephen Morris (2016). "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games." *Theoretical Economics*, 11(2), 487–522.

Bergemann, Dirk, and Stephen Morris (2019). "Information Design: A Unified Perspective." *Journal of Economic Literature*, 57(1), 44–95.

Blackwell, David (1953). "Equivalent Comparisons of Experiments." *Annals of Mathematical Statistics*, 24(2), 265–272.

Caplin, Andrew, and Mark Dean (2015). "Revealed Preference, Rational Inattention, and Costly Information Acquisition." *American Economic Review*, 105(7), 2183–2203.

Condorelli, Daniele, and Balázs Szentes (2020). "Information Design in the Holdup Problem." *Journal of Political Economy*, 128(2), 681–709.

Condorelli, Daniele, and Balázs Szentes (2022). "Surplus Sharing in Cournot Oligopoly." *Theoretical Economics*, 17(3), 955–975.

Condorelli, Daniele, and Balázs Szentes (2026). "Buyer-Optimal Platform Design." *The RAND Journal of Economics*, 57(2), 285–299.

Denti, Tommaso (2022). "Posterior Separable Cost of Information." *American Economic Review*, 112(10), 3215–3259.

Dworczak, Piotr, and Anton Kolotilin (2024). "The Persuasion Duality." *Theoretical Economics*, 19(4), 1701–1755.

Dworczak, Piotr, and Giorgio Martini (2019). "The Simple Economics of Optimal Persuasion." *Journal of Political Economy*, 127(5), 1993–2048.

Frankel, Alexander, and Emir Kamenica (2019). "Quantifying Information and Uncertainty." *American Economic Review*, 109(10), 3650–3680.

Kamenica, Emir, and Matthew Gentzkow (2011). "Bayesian Persuasion." *American Economic Review*, 101(6), 2590–2615.

Kleiner, Andreas, Benny Moldovanu, and Philipp Strack (2021). "Extreme Points and Majorization: Economic Applications." *Econometrica*, 89(4), 1557–1593.

Kolotilin, Anton (2018). "Optimal Information Disclosure: A Linear Programming Approach." *Theoretical Economics*, 13(2), 607–635.

Roesler, Anne-Katrin, and Balázs Szentes (2017). "Buyer-Optimal Learning and Monopoly Pricing." *American Economic Review*, 107(7), 2072–2080.

Sims, Christopher A. (2003). "Implications of Rational Inattention." *Journal of Monetary Economics*, 50(3), 665–690.

Strassen, Volker (1965). "The Existence of Probability Measures with Given Marginals." *Annals of Mathematical Statistics*, 36(2), 423–439.