Information, Persuasion, and Decisionmaking

Part 1: Decisionmaking under uncertainty

Bo Waggoner (University of Pennsylvania → Microsoft NY)
Haifeng Xu (University of Southern California → Harvard)
June 18, 2018, Ithaca, New York.

Background

- Much of economic activity and strategic behavior centers around the **flow of information**

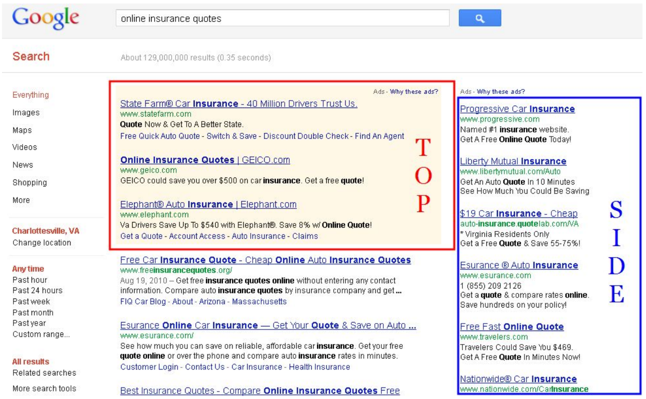

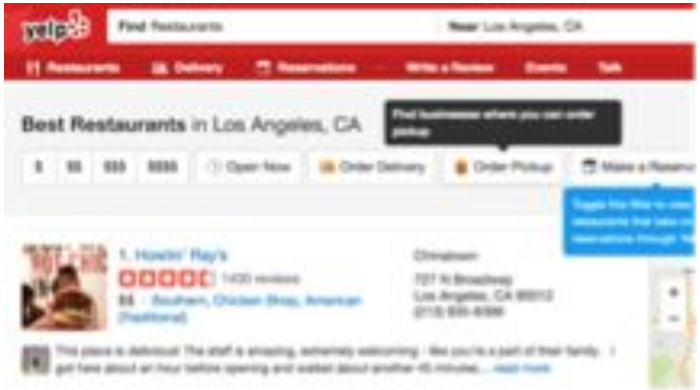

Background

- Much of economic activity and strategic behavior centers around the **flow of information**

- Traditional approaches: information in a game is fixed

- Reality: information can be actively designed/elicited/transferred

- **Recent, fundamental questions:**
  - How to reason about value of information?
  - How does information influence strategic behavior?
  - How to elicit valuable information from strategic sources?
  - How to design information structures to yield desired equilibrium?

Topics Covered in this Tutorial

- Signals as carriers of information, and their properties
- Single-agent decision problems and effect of information
- Bayesian games, equilibrium concepts, and effect of information
- Informational substitutes: definitions, applications, and algorithms
- Persuasion: models, algorithmic study, applications and generalizations
- Open problems and directions

Schedule of the Tutorial

8:30 am - 9:30 am
Part 1: Basics of decisionmaking under uncertainty

(short break)

9:40 am - 10:30 am
Part 2: Informational substitutes and complements

(10:30 am - 11:00 am coffee break)

11:00 am - 12:30 am
Part 3: Algorithmic persuasion

Outline of Part 1

A. Model of information and signals
 - basic properties of signals

B. Model of a single decisionmaker
 - basic properties of decision problems
 - how information impacts decisions
 - Blackwell ordering

C. Bayesian games
 - equilibrium concepts

Part 1A:
Model of information and signals

Notation

• A	set of actions	agent chooses a
• Θ	set of states of the world	nature draws θ
• $u(a, \theta)$	utility function	
• p	prior distribution on Θ	known to agent
• Σ	a signal (also refers to set of realizations)	agent observes $\Sigma = \sigma$
• $\varphi(\sigma, \theta)$	probability of signal σ given state θ	
• p_σ	posterior distribution on Θ given σ	given by Bayes' rule

Basic properties of signals

Probability distributions, signals

- Agent starts with prior belief p in Δ_Θ
- Agent observes signal σ from conditional distribution $\varphi(\sigma, \theta)$
- Agent updates to posterior belief p_σ using Bayes' rule

Fact 1

(1) For all conditional distributions, $\mathbf{E}[\, p_\sigma] = p$.
(On average, the posterior equals the prior.)
(Your current belief is your expectation of your future belief.)

(2) For any set of points $\{p_1,\ldots,p_n\}$ such that p is in their convex hull, there exists a φ inducing this set of posterior beliefs.

Proof.

(1) $\mathbf{E}\ \mathrm{Pr}[\theta\,|\,\sigma] = \mathbf{\Sigma}_\sigma\ \mathrm{Pr}[\sigma]\ \mathrm{Pr}[\theta\,|\,\sigma] = \mathbf{\Sigma}_\sigma \mathrm{Pr}[\theta, \sigma] = \mathrm{Pr}[\theta]$.

(2) Write $p = \mathbf{\Sigma}_\sigma\, \alpha_\sigma\, p_\sigma$ and let $\varphi(\sigma, \theta) = \alpha_\sigma\, p_\sigma(\theta) \,/\, p(\theta)$.

Example: a journey through Ithaca

Probability simplex on Θ = {clear, rain, snow}

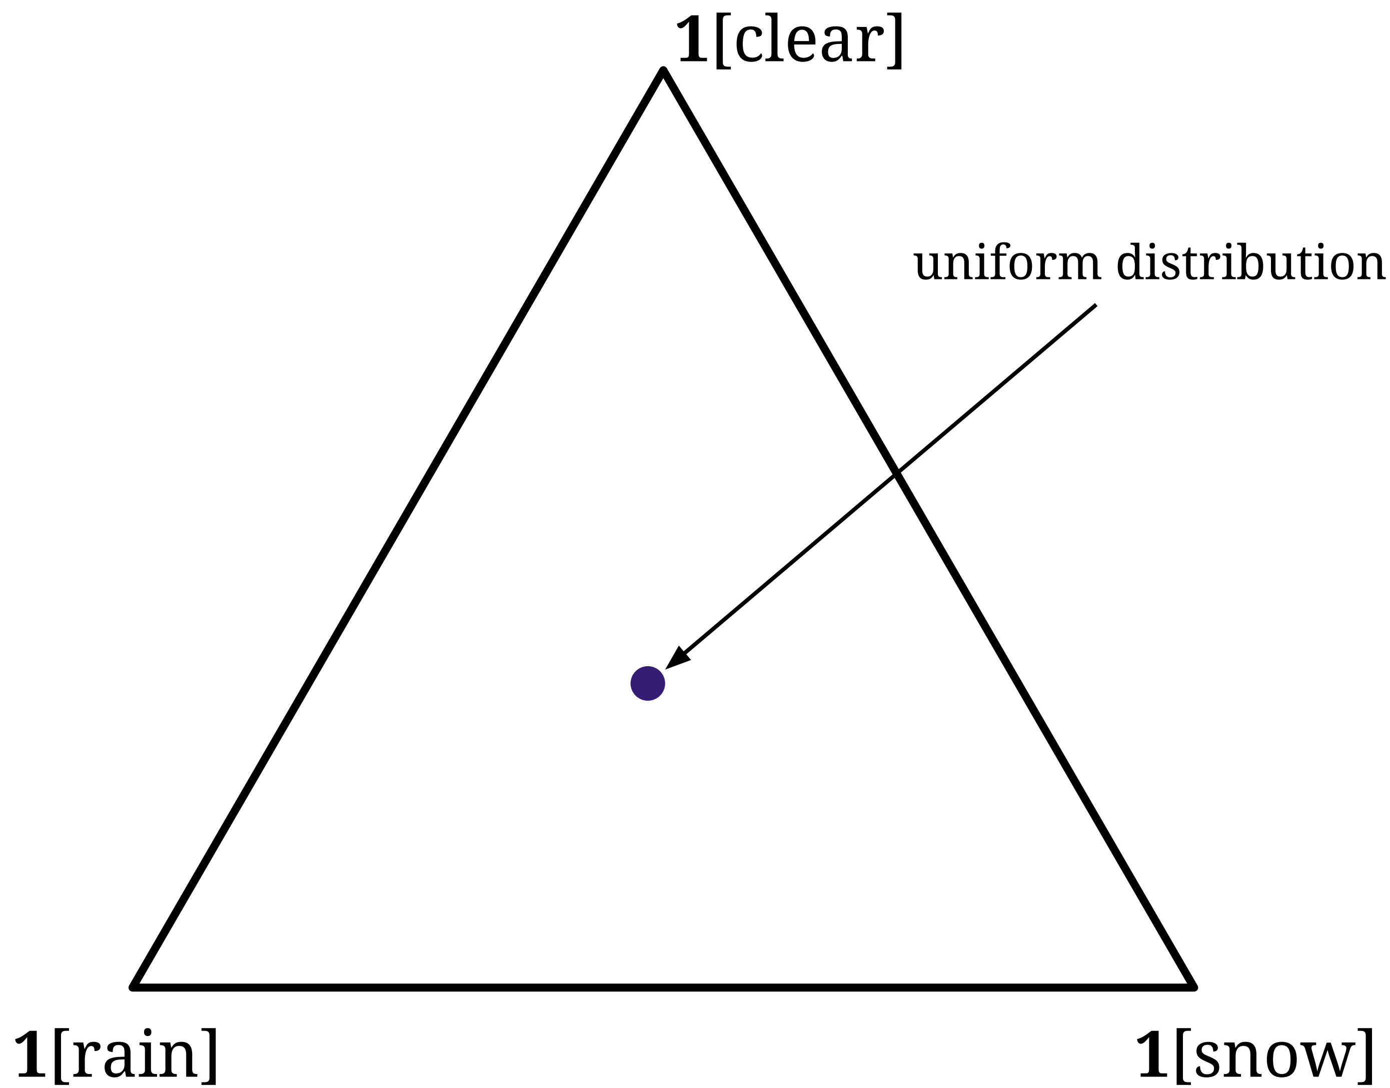

Example: a journey through Ithaca

Probability simplex on Θ = {clear, rain, snow}

Simplest signal: observe θ itself. $\varphi(\theta, \theta) = 1$.

- $p_{\text{rain}} = \mathbf{1}[\text{rain}]$ and so on.
- Fact 1.1: $p = \mathbf{E}[\, p_\theta]$.

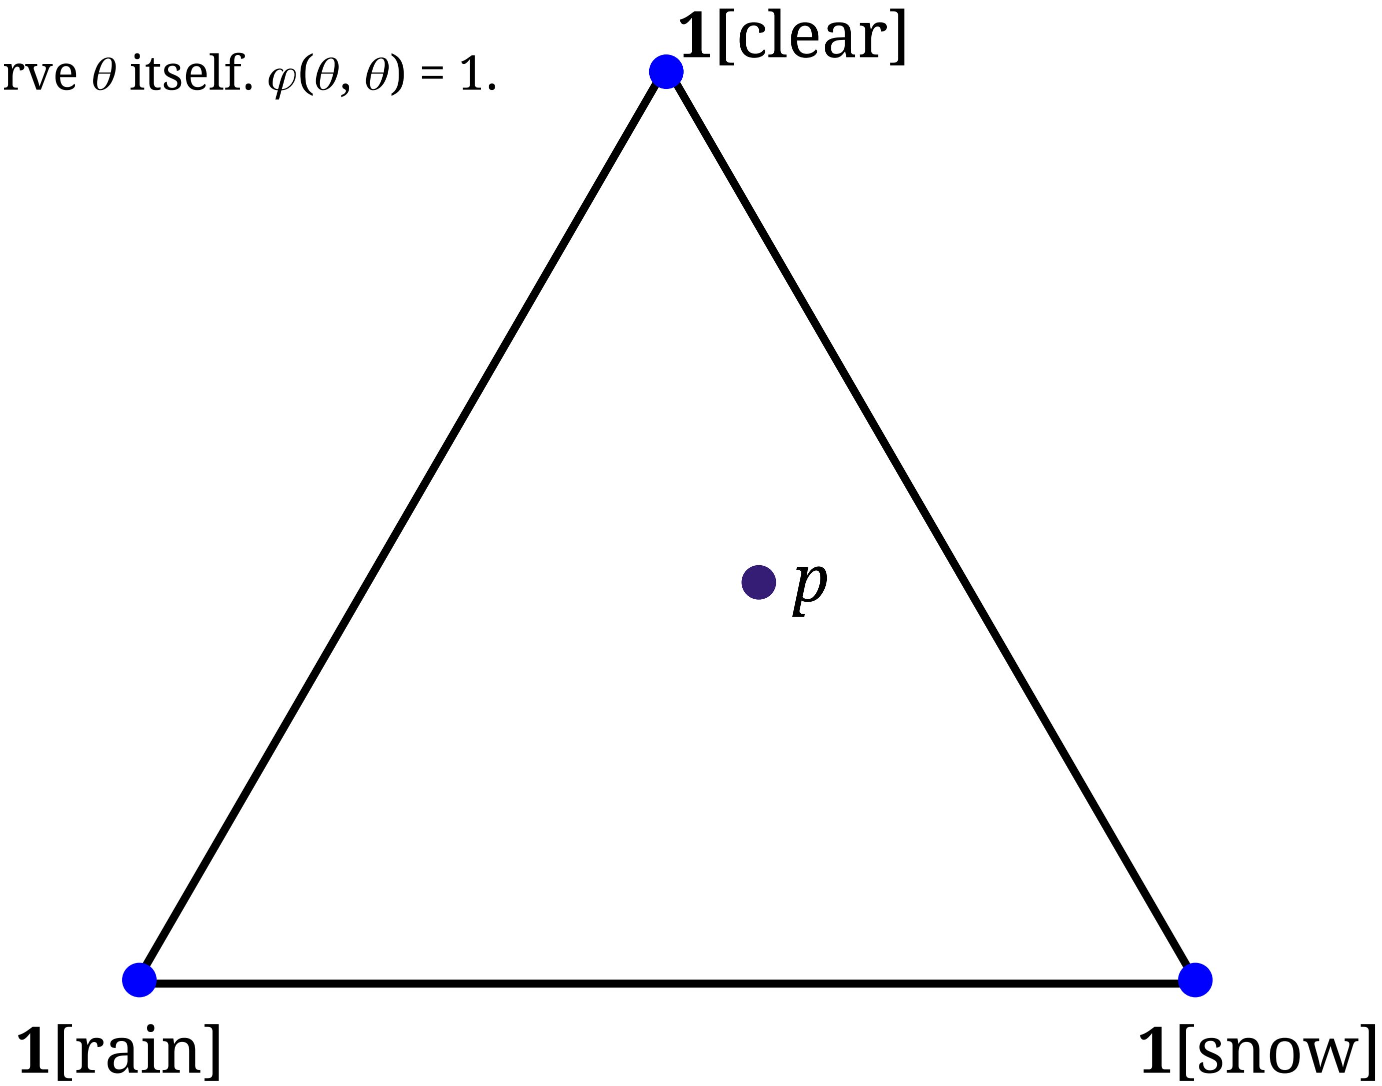

Example: a journey through Ithaca

Probability simplex on Θ = {clear, rain, snow}

Simplest signal: observe θ itself. $\varphi(\theta, \theta) = 1$.

An example where each p_σ is uniform on two of the states.

- Fact 1.2 says this is possible because p is in the convex hull.
- Fact 1.1 implies $\Pr[\sigma = 3]$ is small.

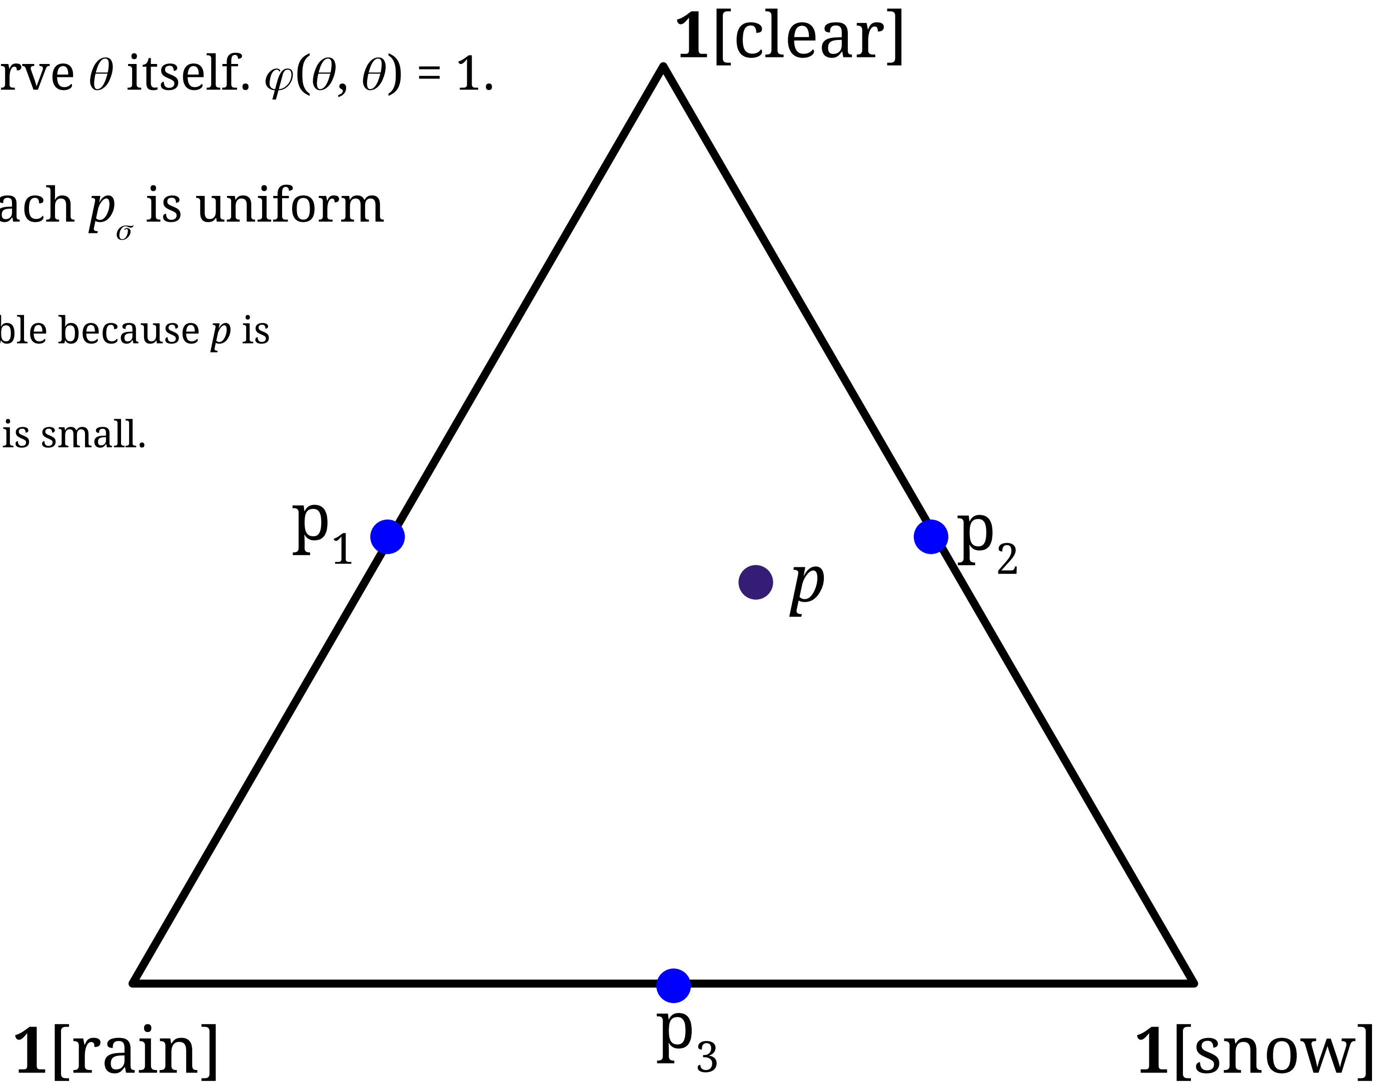

Example: a journey through Ithaca

Probability simplex on Θ = {clear, rain, snow}

Simplest signal: observe θ itself. $\varphi(\theta, \theta) = 1$.

An example where each p_σ is uniform on two of the states.

A generic example with $|\Sigma| = 4$.

- System overdetermined; multiple schemes exist

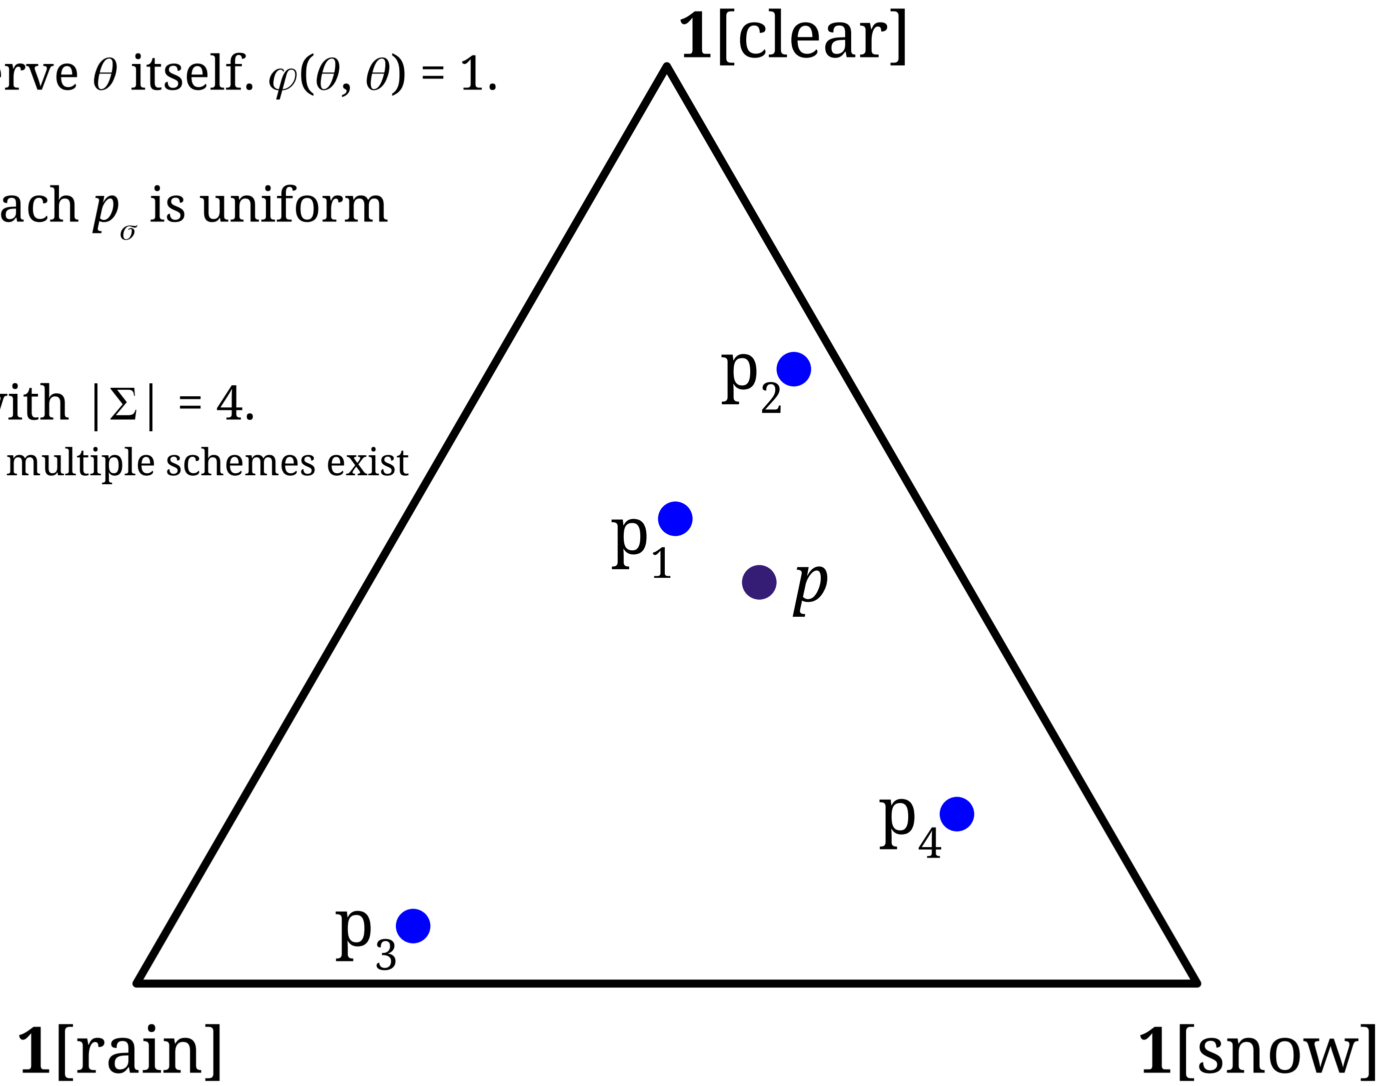

The Ithaca meteorologist

- The prior is p.
- Meteorologist Marsha observes signal $\boldsymbol{\varphi(\sigma, \theta)}$.
- Marsha wants to design a signal $\varphi'(\sigma', \sigma)$.

What is the space of achievable signalling schemes?

- Think of Σ as the new state space!

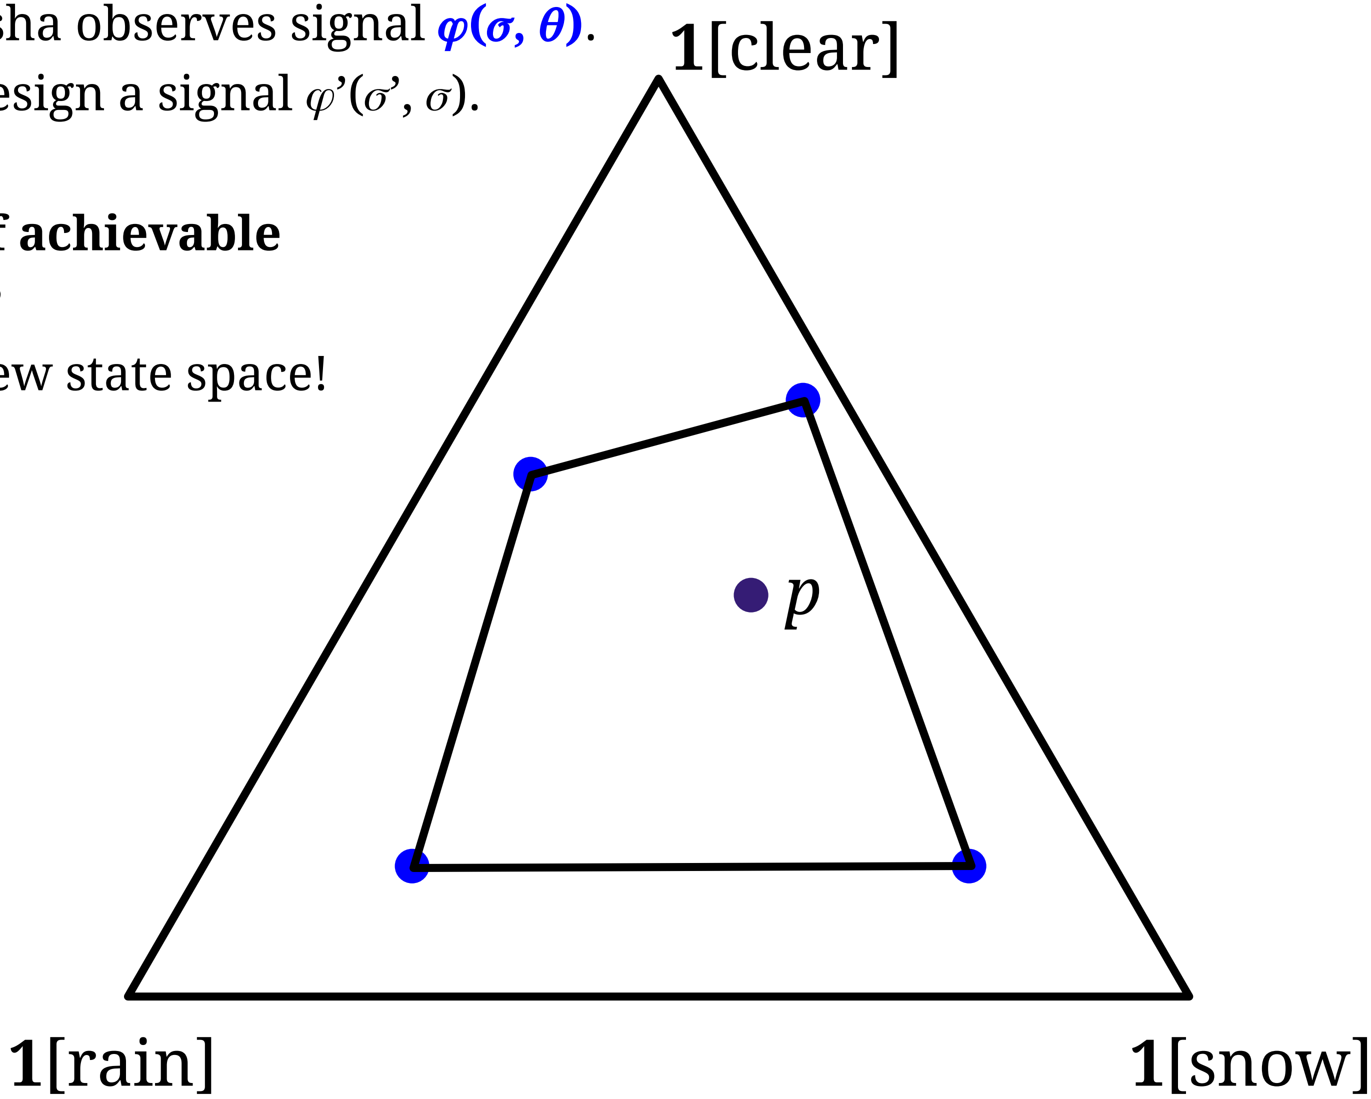

The Ithaca meteorologist

- The prior is p.
- Meteorologist Marsha observes signal $\boldsymbol{\varphi(\sigma, \theta)}$.
- Marsha wants to design a signal $\varphi'(\sigma', \sigma)$.

What is the space of achievable signalling schemes?

- Think of Σ as the new state space!

  (The convex hull of $\{\boldsymbol{p_\sigma}\}$ is the new simplex.)

- Can easily characterize all schemes:
  - $\mathbf{E}\, \boldsymbol{p_{\sigma'}} = p$.
  - p must be in the convex hull of $\{\boldsymbol{p_{\sigma'}}\}$ which must be in the convex hull of $\{\boldsymbol{p_\sigma}\}$.
  - For each $\boldsymbol{\sigma'}$, $\mathbf{E}[\, \boldsymbol{p_\sigma} \mid \boldsymbol{\sigma'}\,] = \boldsymbol{p_{\sigma'}}$.

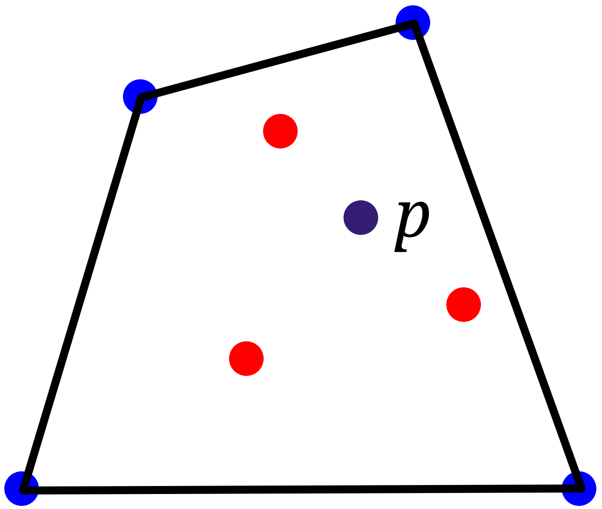

Therefore: It is without much loss to assume that a signaller observes the true state θ.

Part 1B:
Decision problems

Notation

• A	set of actions	agent chooses a
• Θ	set of states of the world	nature draws θ
• $u(a, \theta)$	utility function	
• p	prior distribution on Θ	known to agent
• Σ	a signal (also refers to set of realizations)	agent observes $\Sigma = \sigma$
• $\varphi(\sigma, \theta)$	probability of signal σ given state θ	
• p_σ	posterior distribution on Θ given σ	given by Bayes' rule
• $u(a\,;q)$	$= \mathbf{E}_{\theta \sim q}\, u(a, \theta)$	linear function of q
• $G(q)$	$= \max_a u(a\,;q)$	convex function of q
• $a^*(q)$	$= \operatorname{argmax}_a u(a\,;q)$	optimal action given q

Basics of decision problems

Decision problems and convex functions

- Agent must choose a based on belief q

q may be prior or posterior

- Assume: chooses to maximize expected$_q$ utility
- Write $G(q)$ = “expected utility for optimal$_q$ action”

How to characterize all possible decision problems?

Decision problems and convex functions

- Agent must choose a based on belief q

q may be prior or posterior

- Assume: chooses to maximize expected_q utility
- Write $G(q)$ = "expected utility for optimal_q action"

Fact 2

(1) For every decision problem (A, Θ, u), $G(q) = \max_a u(a\,; q)$ is **convex**.

(2) Every convex $G : \Delta_\Theta \to \mathbf{R}$ is the expected utility function for some decision problem (A, Θ, u).

Proof.

(1) Each $u(\text{a}\,; q)$ is a linear function of q; a max of linear functions is convex.

(2) We can write G as a maximum of linear functions of q.
Assign each linear function to an action a and write it as $u(a\,; q)$.
Define $u(a, \theta) = \text{u}(\text{a}\,; \mathbf{1}[\theta])$ = expected utility for a under belief $\Pr[\theta] = 1$.

Example: a journey through Ithaca

Probability simplex on Θ = {clear, rain, snow}

A = {walk, ride}

Pictured: best-response $a^*(q)$.

For some beliefs q, it is preferable to walk;

for others, it is preferable to ride.

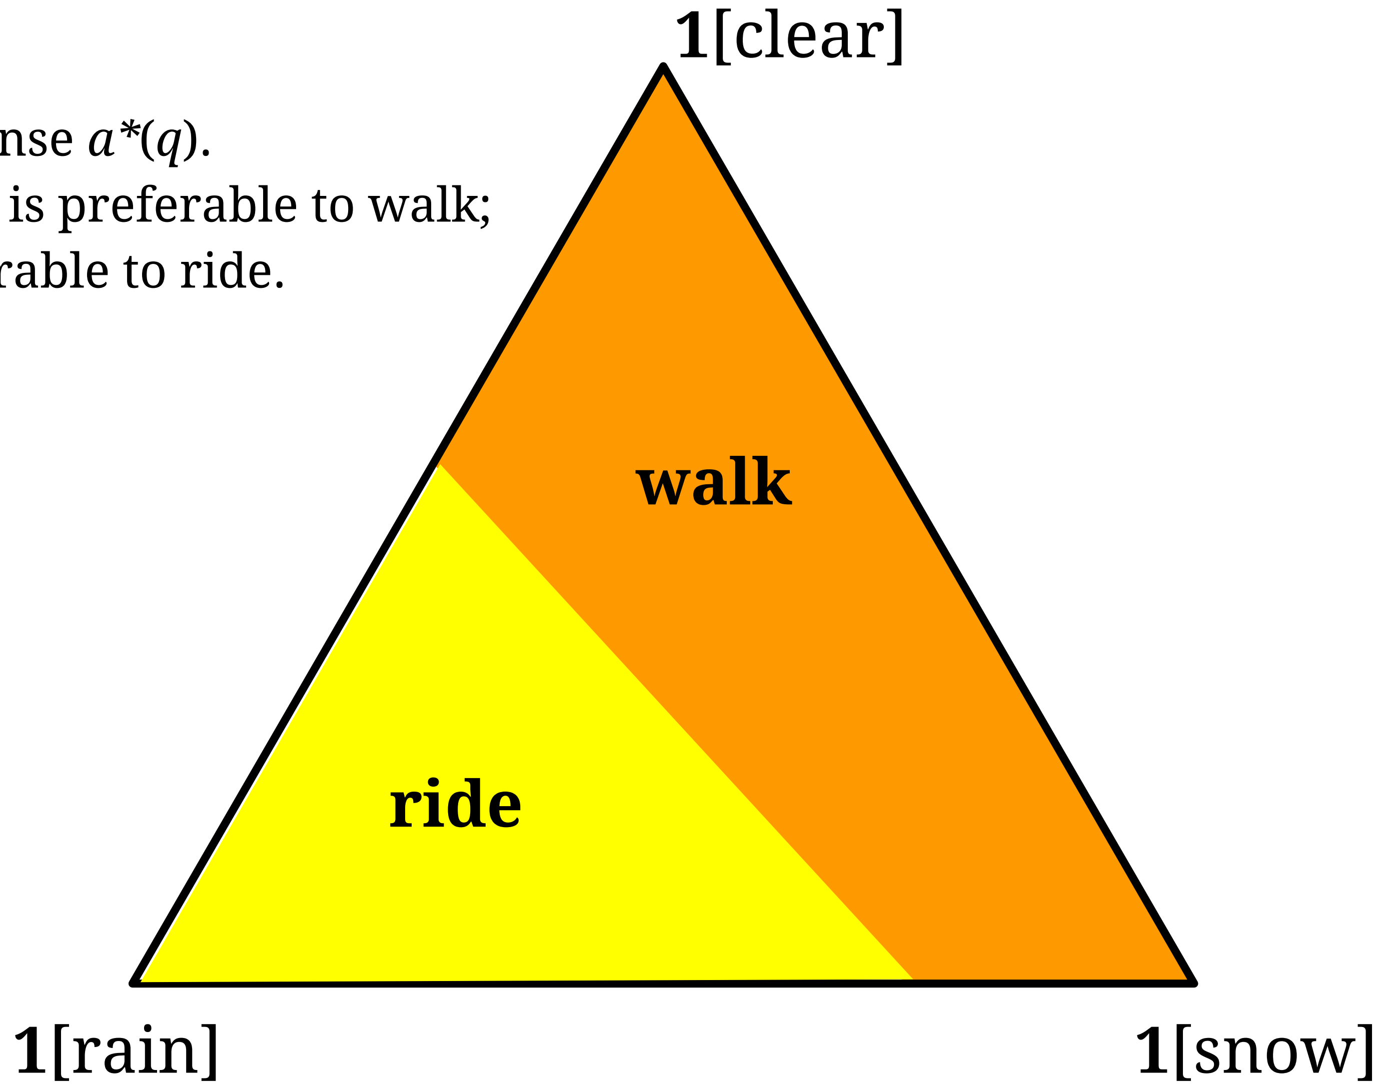

Example: a journey through Ithaca

Probability simplex on Θ = {clear, rain, snow}

A = {walk, ride}

If we ride:

- u(ride, clear) = 7
- u(ride, snow) = 7
- u(ride, rain) = 7

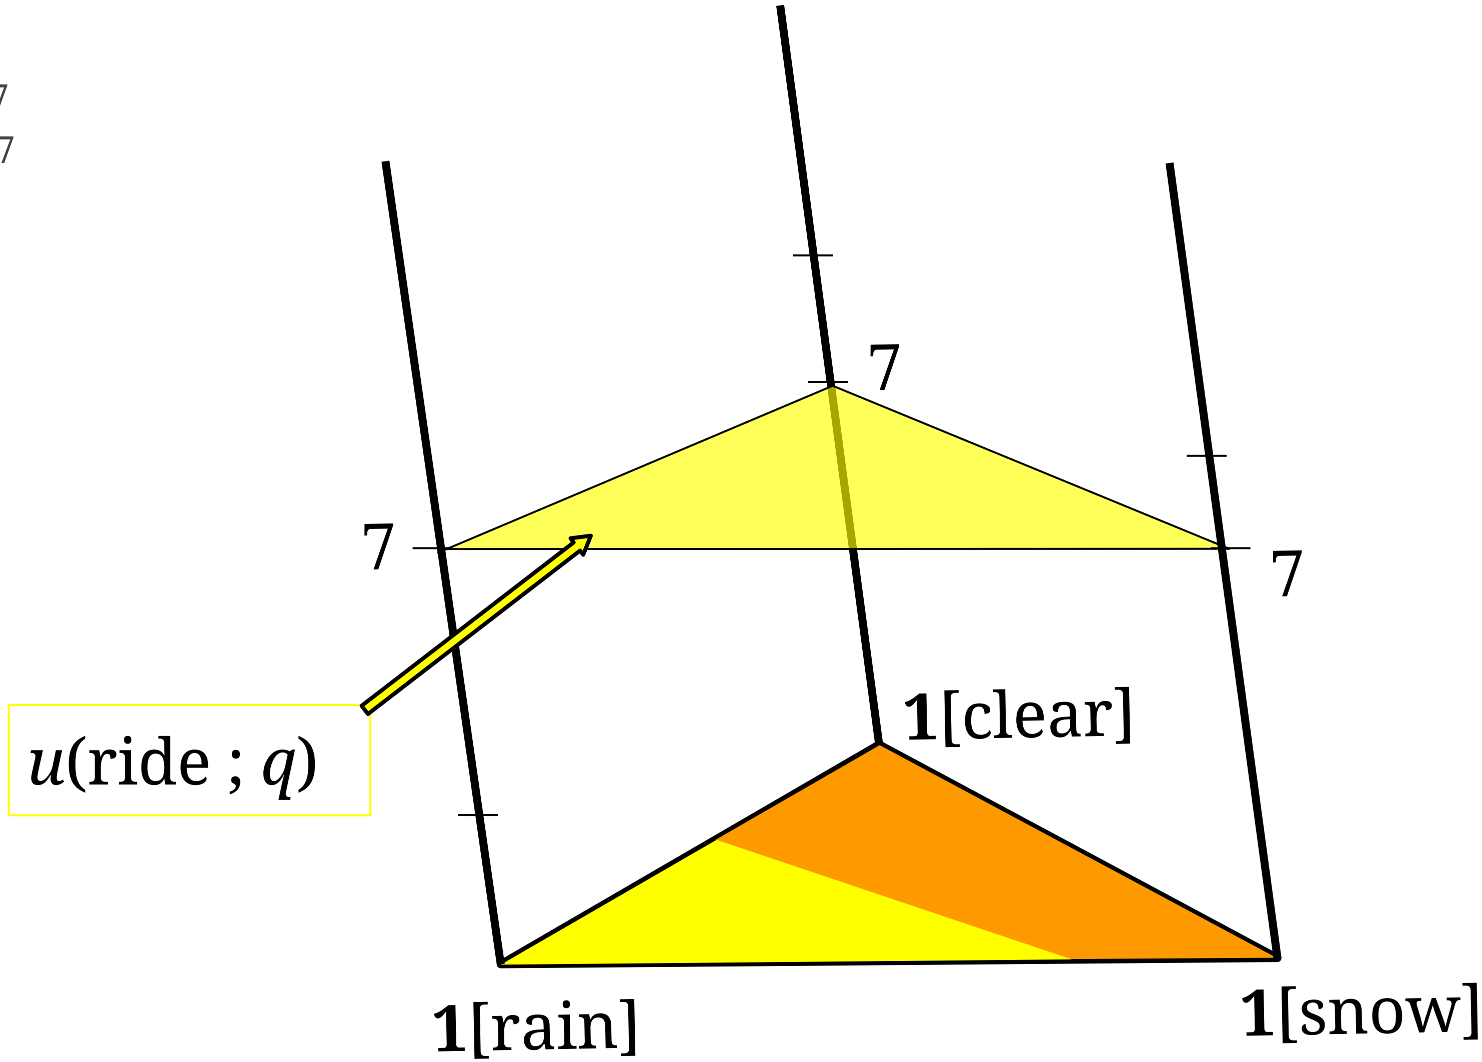

Example: a journey through Ithaca

Probability simplex on Θ = {clear, rain, snow}

A = {walk, ride}

If we ride:

u(ride, clear) = 7

u(ride, snow) = 7

u(ride, rain) = 7

If we walk:

u(walk, clear) = 10

u(walk, snow) = 8

u(walk, rain) = 2

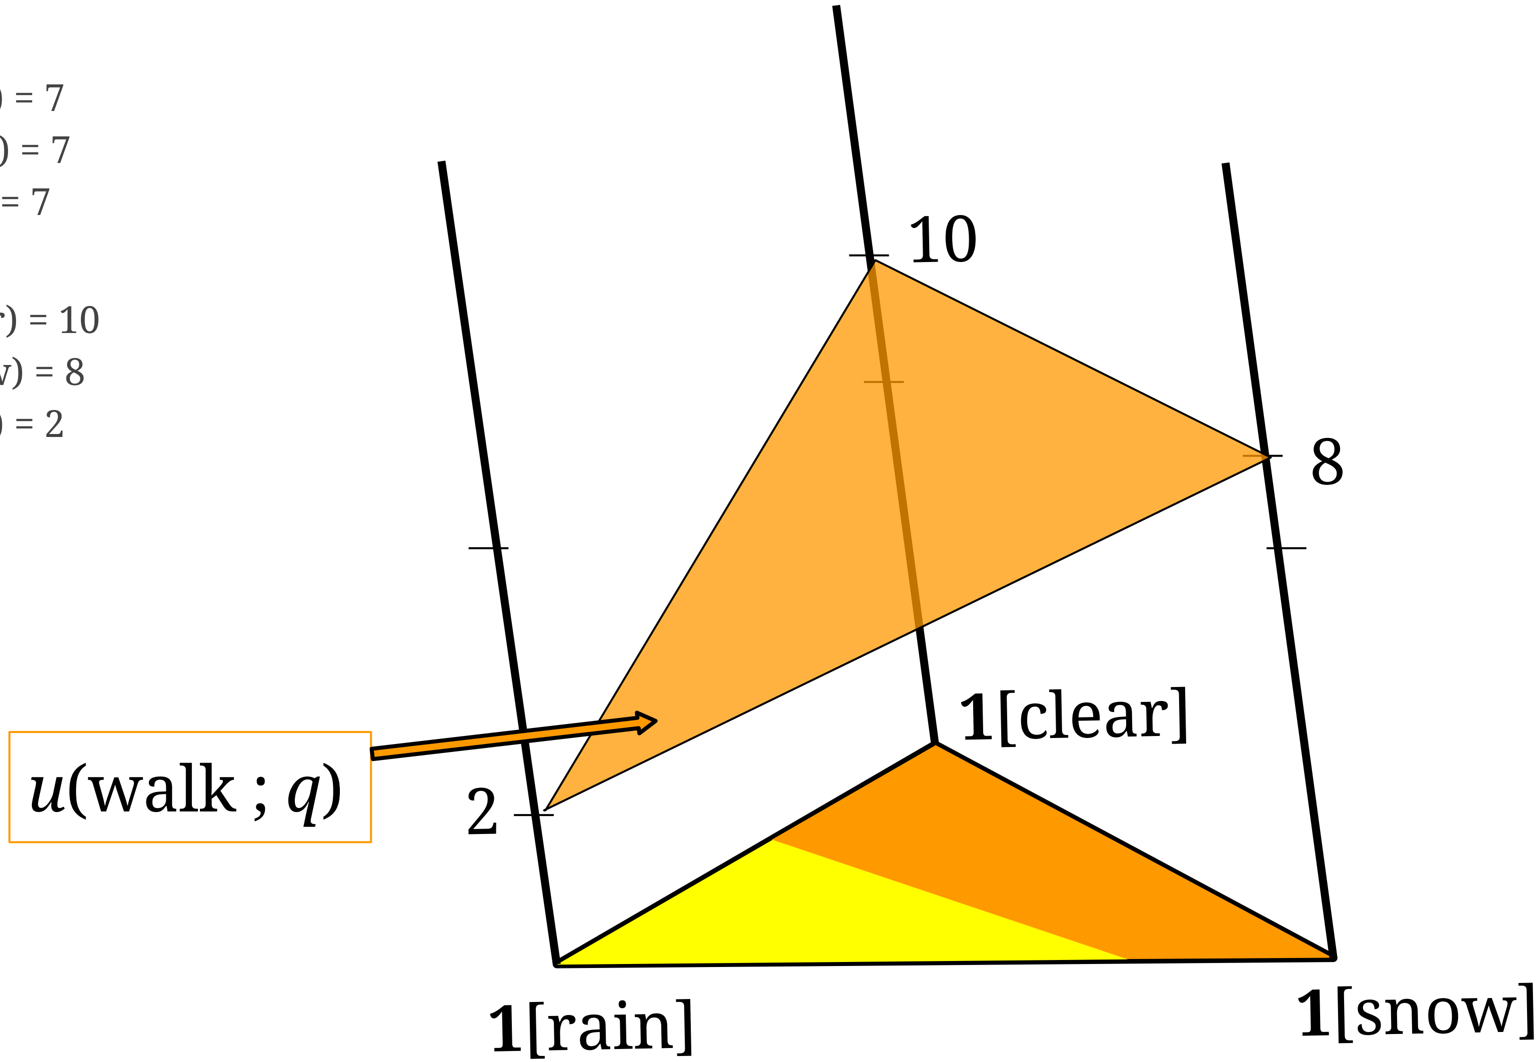

Example: a journey through Ithaca

Probability simplex on Θ = {clear, rain, snow}

A = {walk, ride}

If we ride:

- u(ride, clear) = 7
- u(ride, snow) = 7
- u(ride, rain) = 7

If we walk:

- u(walk, clear) = 10
- u(walk, snow) = 8
- u(walk, rain) = 2

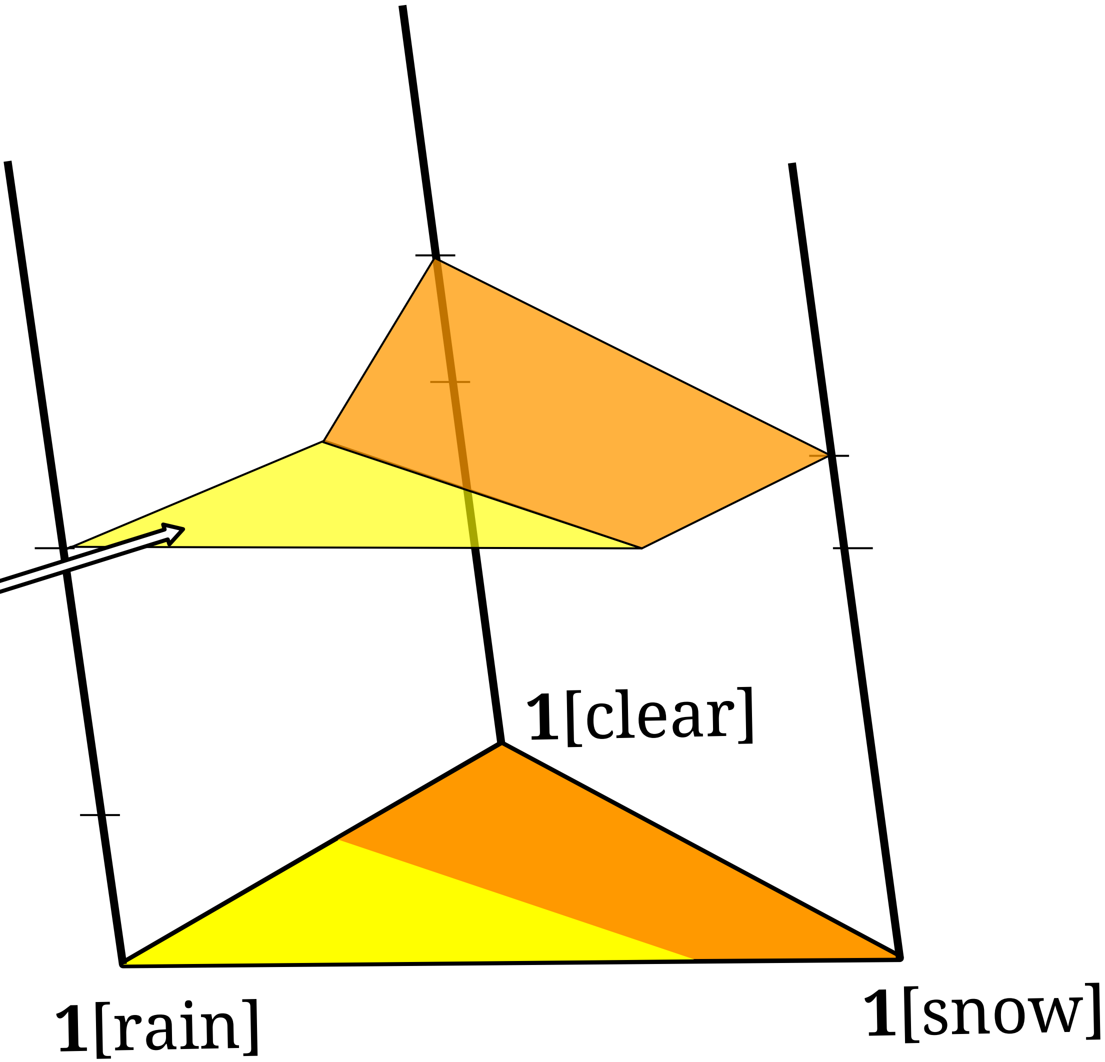

$G(q)$ = expected utility for acting optimally under belief q

Example: proper scoring rules

A decision problem $S(a, \theta)$ with $A = \Delta_\Theta$ is a **proper scoring rule** if $a^*(q) = q$, i.e. it is always optimal to choose one's true belief.

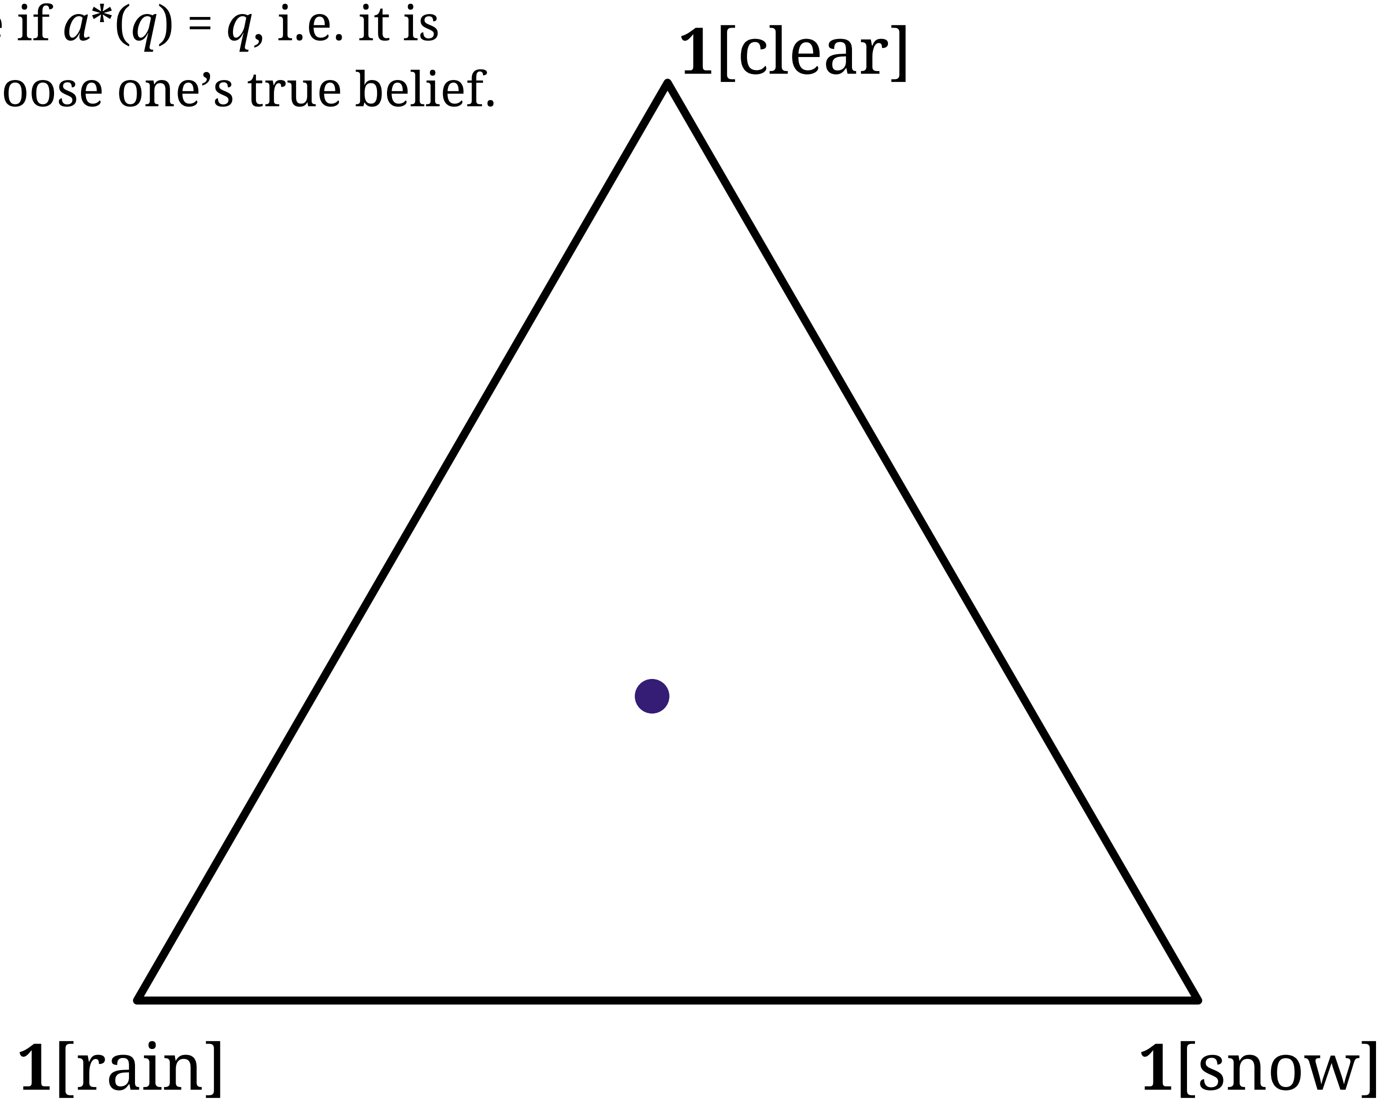

Example: proper scoring rules

A decision problem $S(a, \theta)$ with $A = \Delta_\Theta$ is a **proper scoring rule** if $a^*(q) = q$, i.e. it is always optimal to choose one's true belief.

Solution: take any convex function G.

For each q, there will be a tangent hyperplane.

Strictly convex $\longleftrightarrow$ *strictly* proper (truthfulness is uniquely optimal)

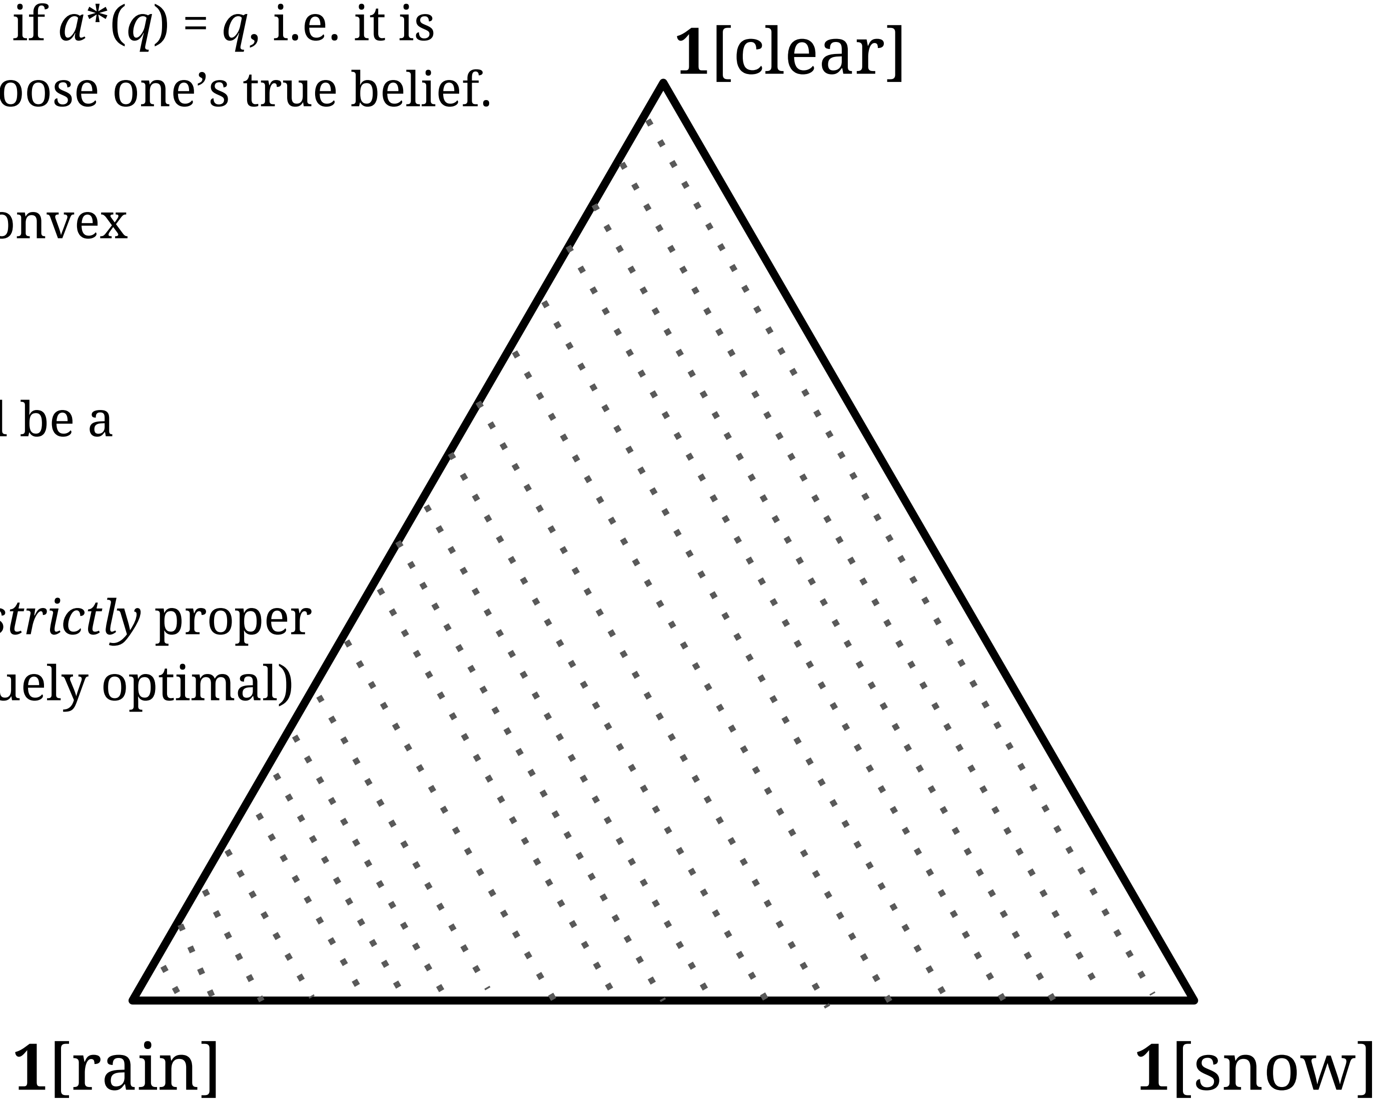

Example: proper scoring rules

A decision problem $S(a, \theta)$ with $A = \Delta_\Theta$ is a **proper scoring rule** if $a^*(q) = q$, i.e. it is always optimal to choose one's true belief.

Solution: take any convex function G.

For each q, there will be a tangent hyperplane.

Example: $S(a, \theta) = \log a(\theta)$.

- $S(a\ ; q) = \mathbf{\Sigma}_\theta\, q(\theta) \log a(\theta)$.
- Optimal action $a = q$.
- $G(q) = \mathbf{\Sigma}_\theta\, q(\theta) \log q(\theta) = -H(q)$.

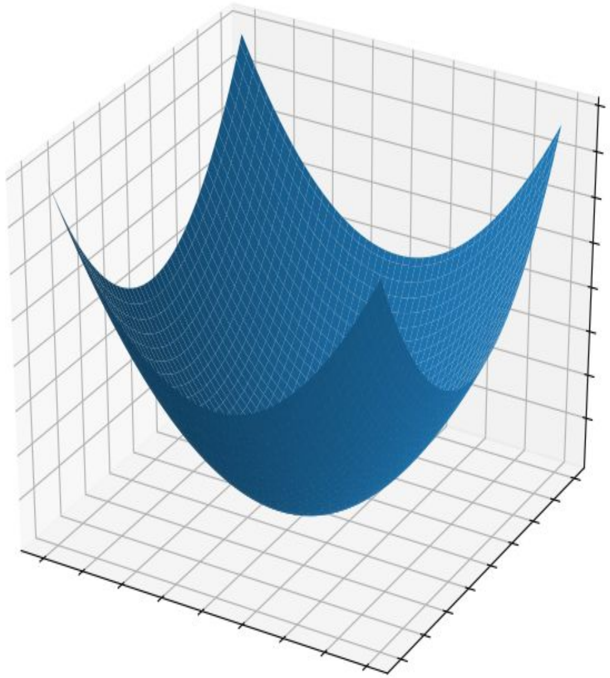

Revelation principle for agents

Fact 3

For every decision problem (A, Θ, u), there is a corresponding proper scoring rule $(\Delta_\Theta, \Theta, S)$ that is utility-equivalent: the expected utility is always equal.

Proof.
Define $S(q, \theta) = u(a^*(q), \theta)$.

In other words: given agent's reported belief q, plug action $a^*(q)$ into the original decision problem.

Truthfulness is an optimal action, and expected utility for any belief is equal in both problems.

Signals and decisionmaking

How signals affect decisionmaking

- Agent begins with prior belief p on Θ.
- She would take action $a^*(p) = \text{argmax}_a\, u(a\,; p)$.
- After receiving signal σ, she takes $a^*(p_\sigma)$.

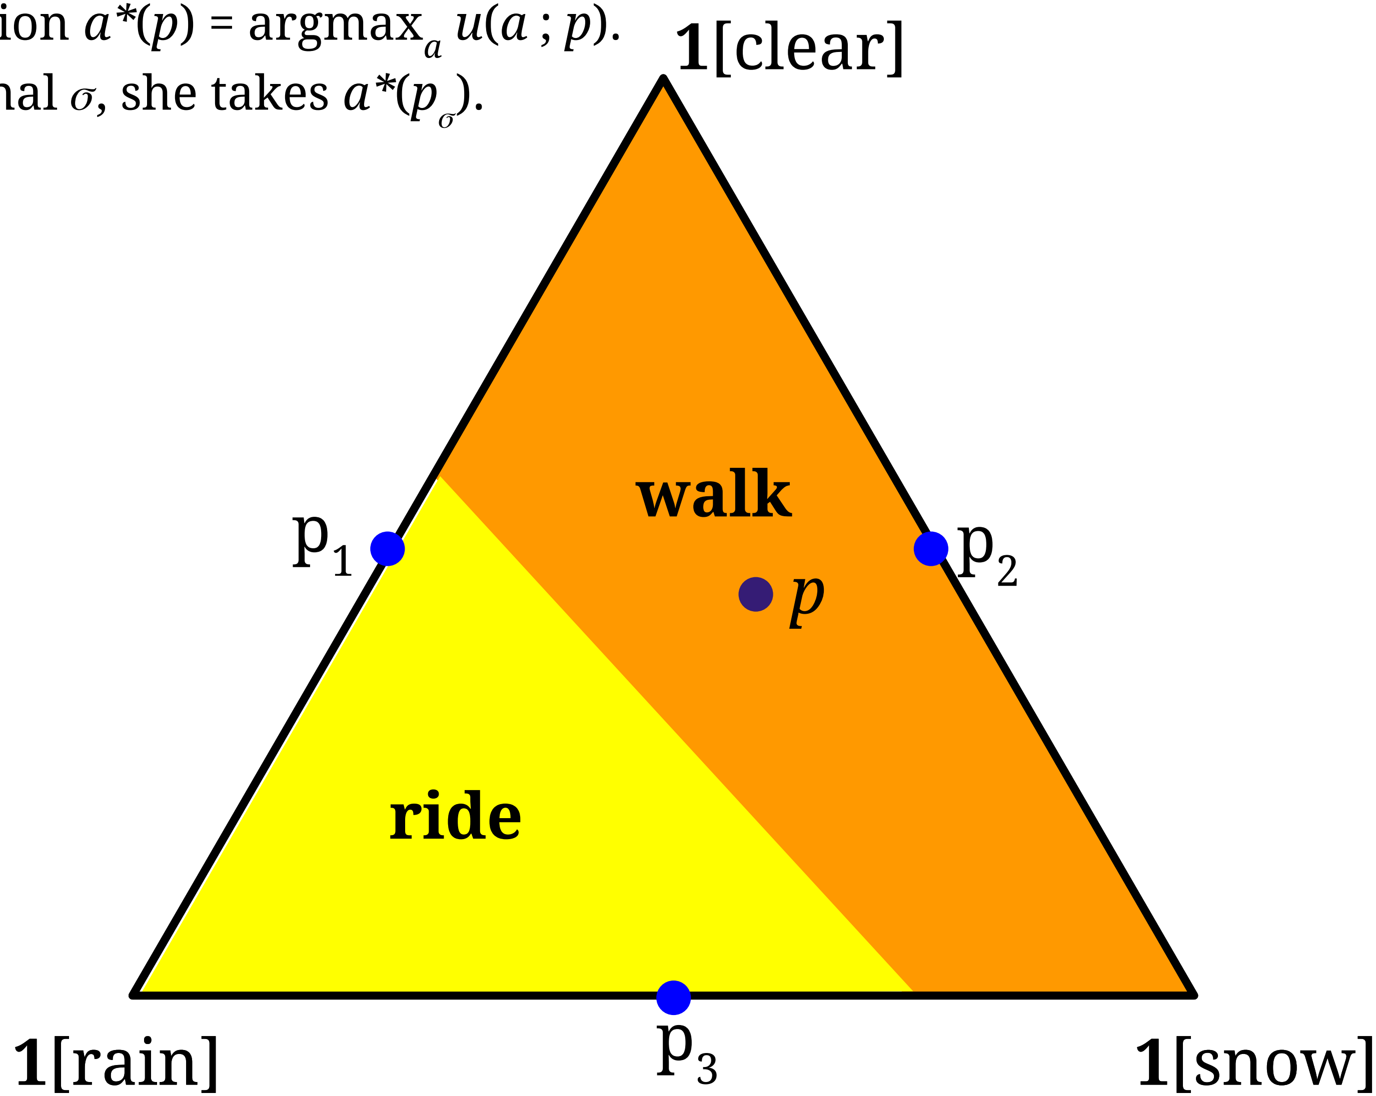

How signals affect decisionmaking

- Agent begins with prior belief p on Θ.
- She would take action $a^*(p) = \text{argmax}_a\ u(a\,;p)$.
- After receiving signal σ, she takes $a^*(p_\sigma)$.
- Expected utility is $V^{u,\varphi}(\Sigma) := \mathbf{E}_{\sigma\sim\varphi} G(p_\sigma)$.

Fact 4

More information always increases expected utility in a decision problem.

Proof.

$V^{u,\varphi}(\Sigma)$

$= \mathbf{E}_\sigma G(p_\sigma)$

$\geq G(\mathbf{E}_\sigma p_\sigma)$ (Jensen's inequality)

$= G(p)$

= exp. utility with no signal.

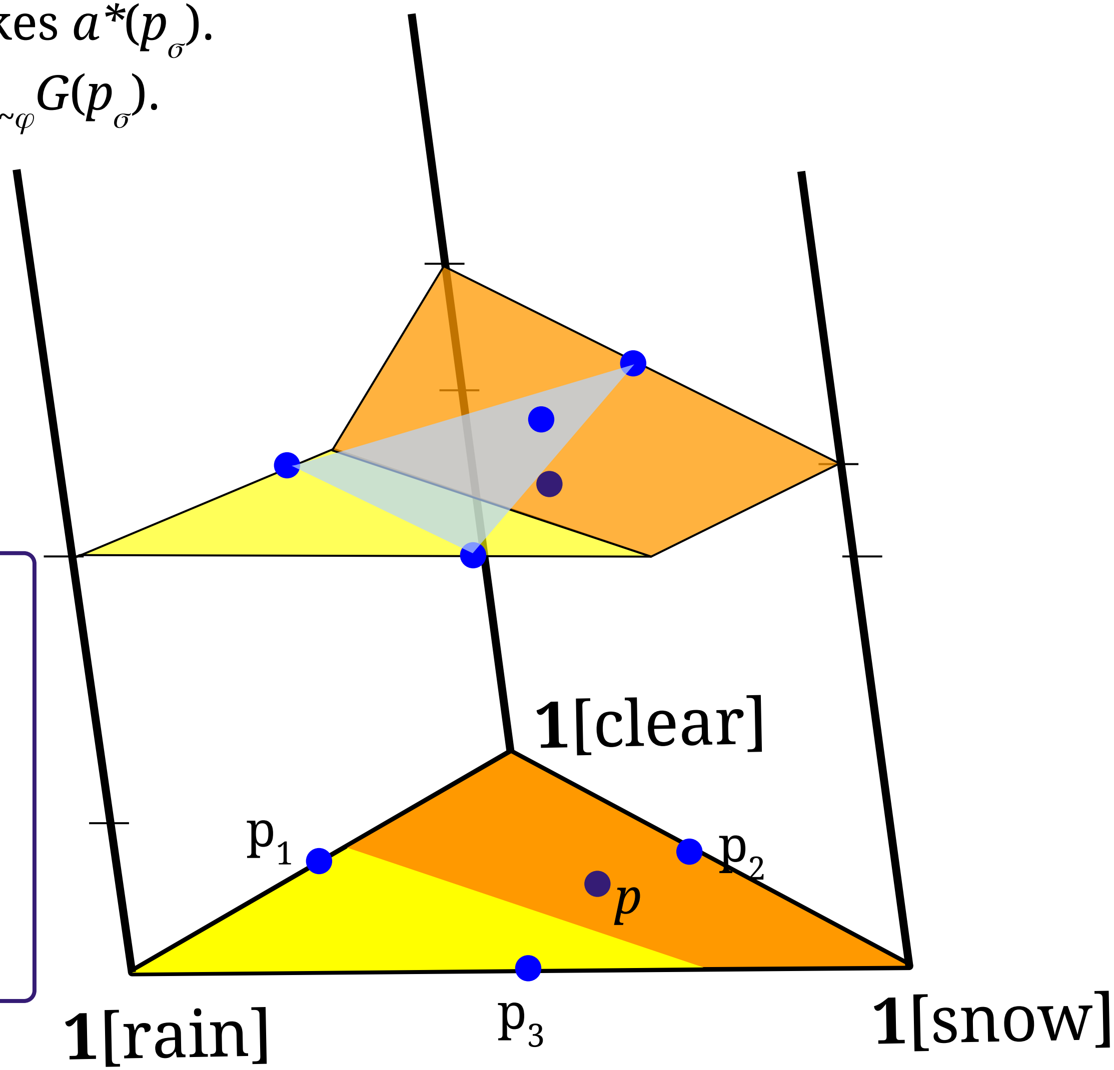

Revelation principle for signallers

φ is **direct** if $\Sigma = A$ (each signal *recommends* a unique action),
and **persuasive** if it is optimal to comply, i.e. $a^*(p_a) = a$.

Fact 5

For every signalling scheme φ in decision problem (A, Θ, u), there is a direct, persuasive φ' that outcome-equivalent: induces the same distribution on a, θ.

Proof.
Merge all signals σ inducing action a, i.e. $a = a^*(p_\sigma)$, into a single signal s.

Then p_s is a convex combination of the $\{p_\sigma\}$, so
$\text{argmax}_{a'}\, \text{u}(a' ; p_s) \;=\; \text{argmax}_{a'}\, \boldsymbol{\Sigma}_\sigma\, \alpha_\sigma\, u(a' ; p_s) \;=\; a.$

Repeat for all actions.

Illustrating the signaller's revelation principle

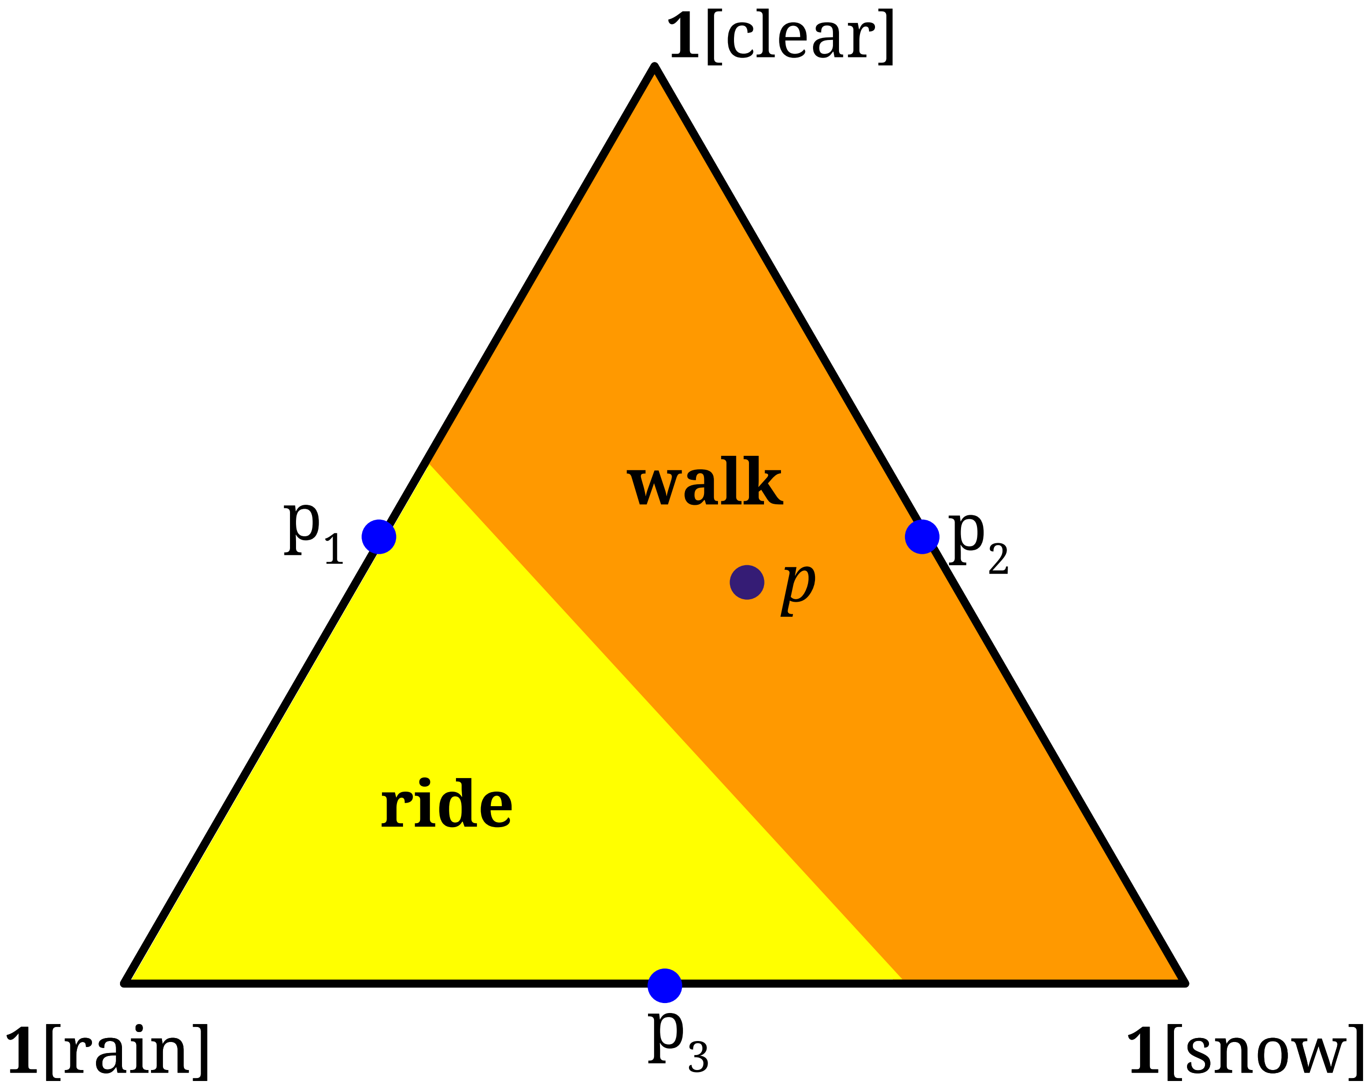

Illustrating the signaller's revelation principle

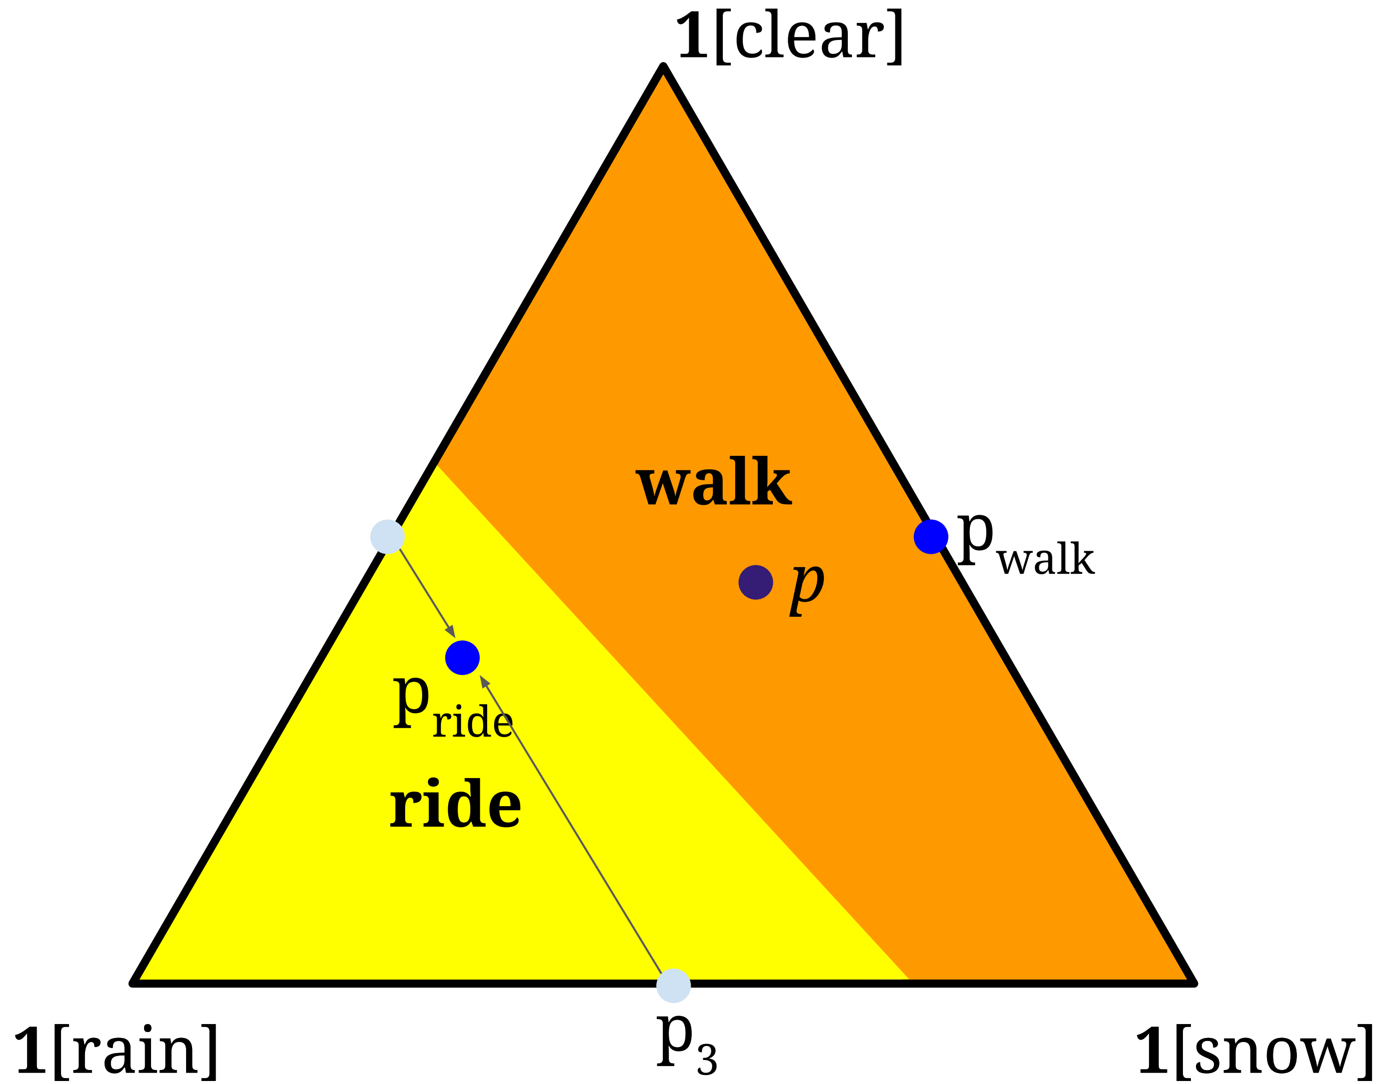

The Blackwell order

Garblings and Blackwell order

- We have a signal **Σ** distributed according to $\boldsymbol{\varphi(\sigma, \theta)}$.
- And we have **Σ'** distributed according to $\boldsymbol{\varphi'(\sigma', \theta)}$.

Say **Σ'** is a **garbling** of **Σ** if it can be simulated given **Σ**, i.e. $\boldsymbol{\sigma'}$ is distributed as a randomized function f($\boldsymbol{\sigma}$). (Σ' is conditionally independent of Θ given Σ.)

Fact (cf. 1.2): **Σ'** is a garbling of **Σ** if and only if each $\boldsymbol{p_{\sigma'}} = \mathbf{E}[\, \boldsymbol{p_{\sigma}} \mid \boldsymbol{\sigma'} \,]$.

Theorem (Blackwell 1953):
Σ' is a garbling of **Σ** if and only if, for all decision problems, $V^{u,\varphi}(\boldsymbol{\Sigma}) \geq V^{u,\varphi}(\boldsymbol{\Sigma'})$.

Proof sketch:
(→) Given **Σ**, we can simulate a draw from **Σ'** and take the optimal action.
(←) If not a garbling, there exists a realization $\boldsymbol{\sigma'}$ that is more informative in some direction than **Σ** is on average. Make a two-action decision problem that rewards this knowledge…

Blackwell proof - an easy version

If $\boldsymbol{p_{\sigma'}}$ is not in the convex hull of { $\boldsymbol{p_{\sigma}}$}, then definitely not a garbling $\Rightarrow$ for some u, **Σ'** is preferable.

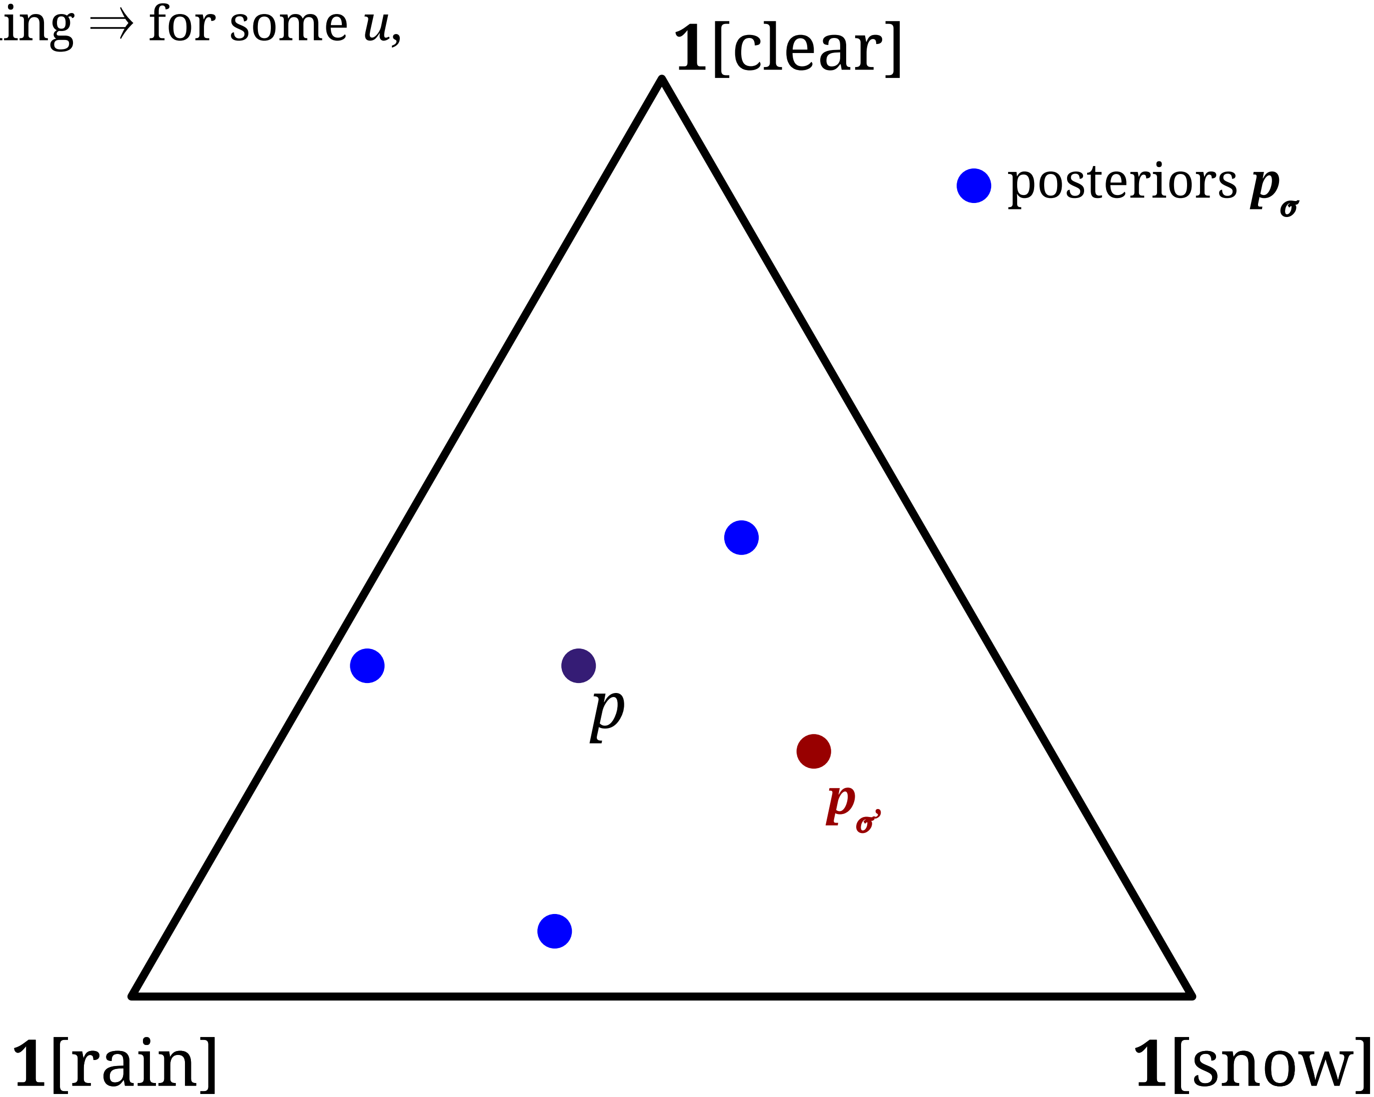

Blackwell proof - an easy version

If $p_{\sigma'}$ is not in the convex hull of { p_σ}, then definitely not a garbling ⇒ for some u, Σ' is preferable.

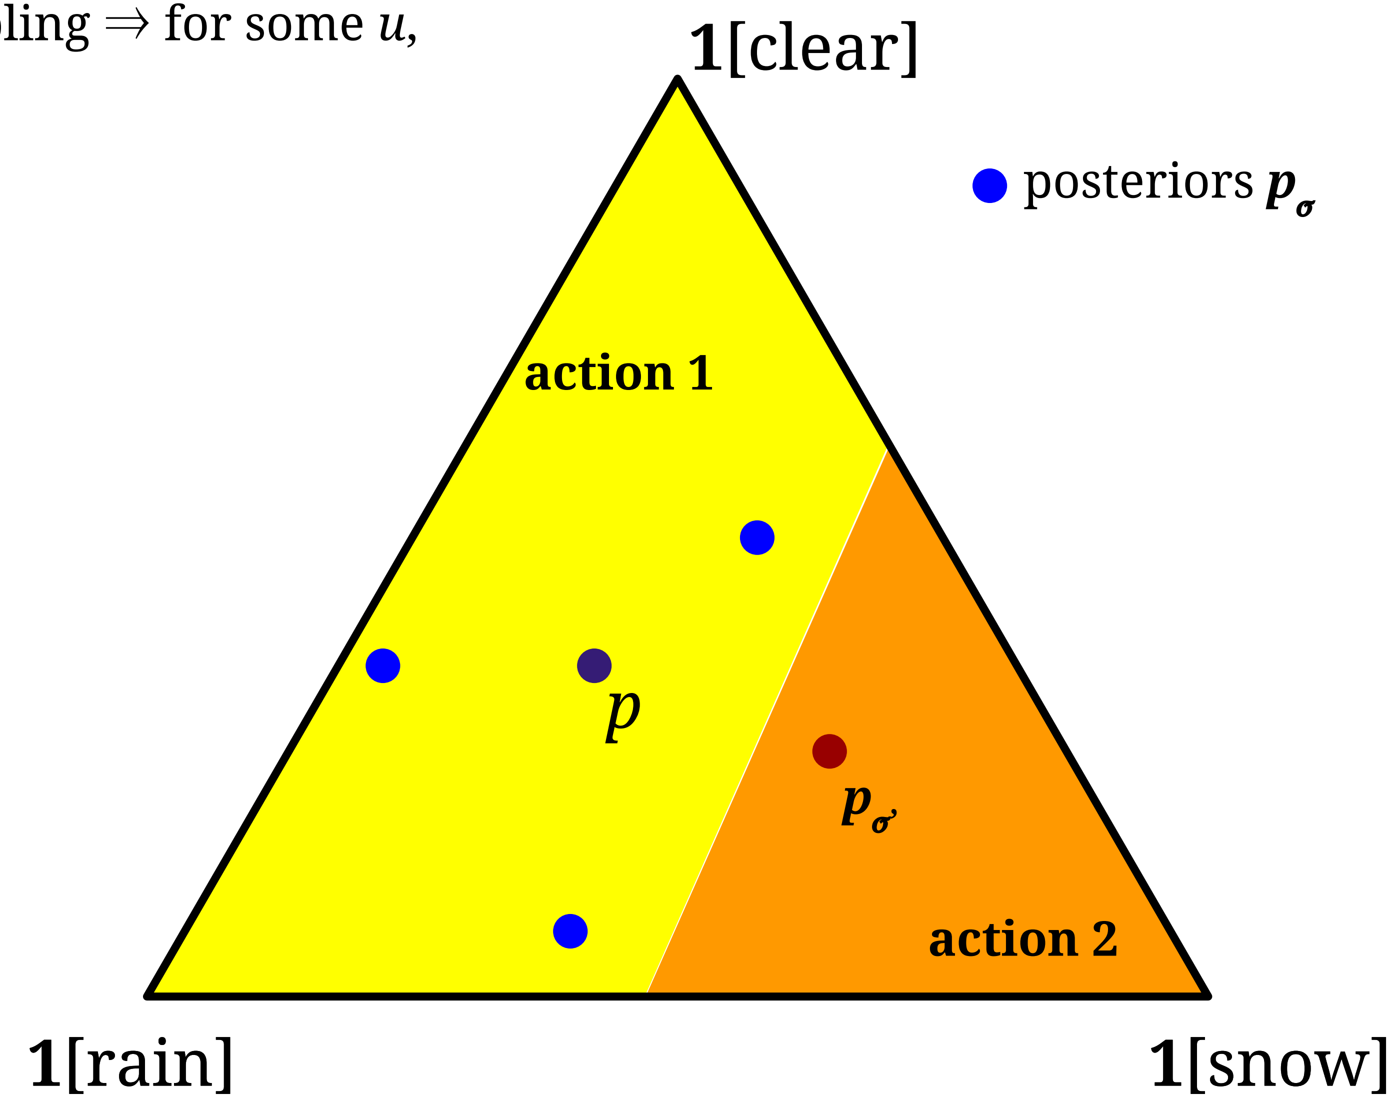

Blackwell proof - an easy version

If $\boldsymbol{p}_{\boldsymbol{\sigma}'}$ is not in the convex hull of { $\boldsymbol{p}_{\boldsymbol{\sigma}}$}, then definitely not a garbling ⇒ for some u,
Σ' is preferable.

$V^{u,\varphi}(\boldsymbol{\Sigma}) = 0$
$V^{u,\varphi}(\boldsymbol{\Sigma'}) > 0.$

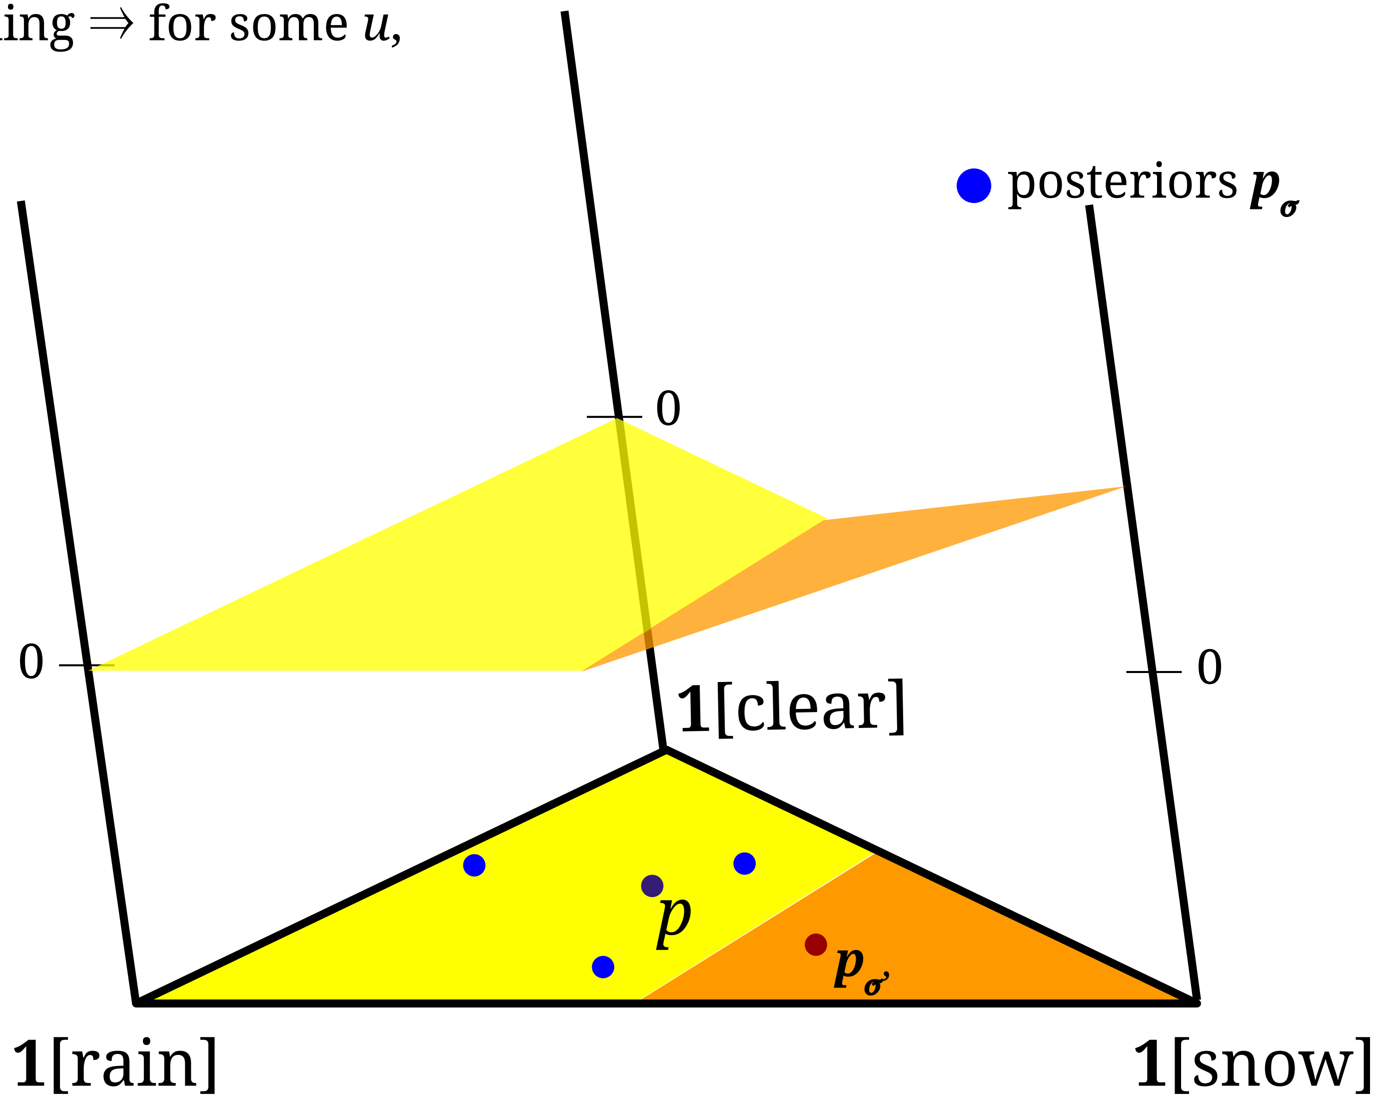

Blackwell proof - general idea

If not a garbling, there exists a realization $\boldsymbol{\sigma'}$ that is more informative in some direction than $\boldsymbol{\Sigma}$ is on average.

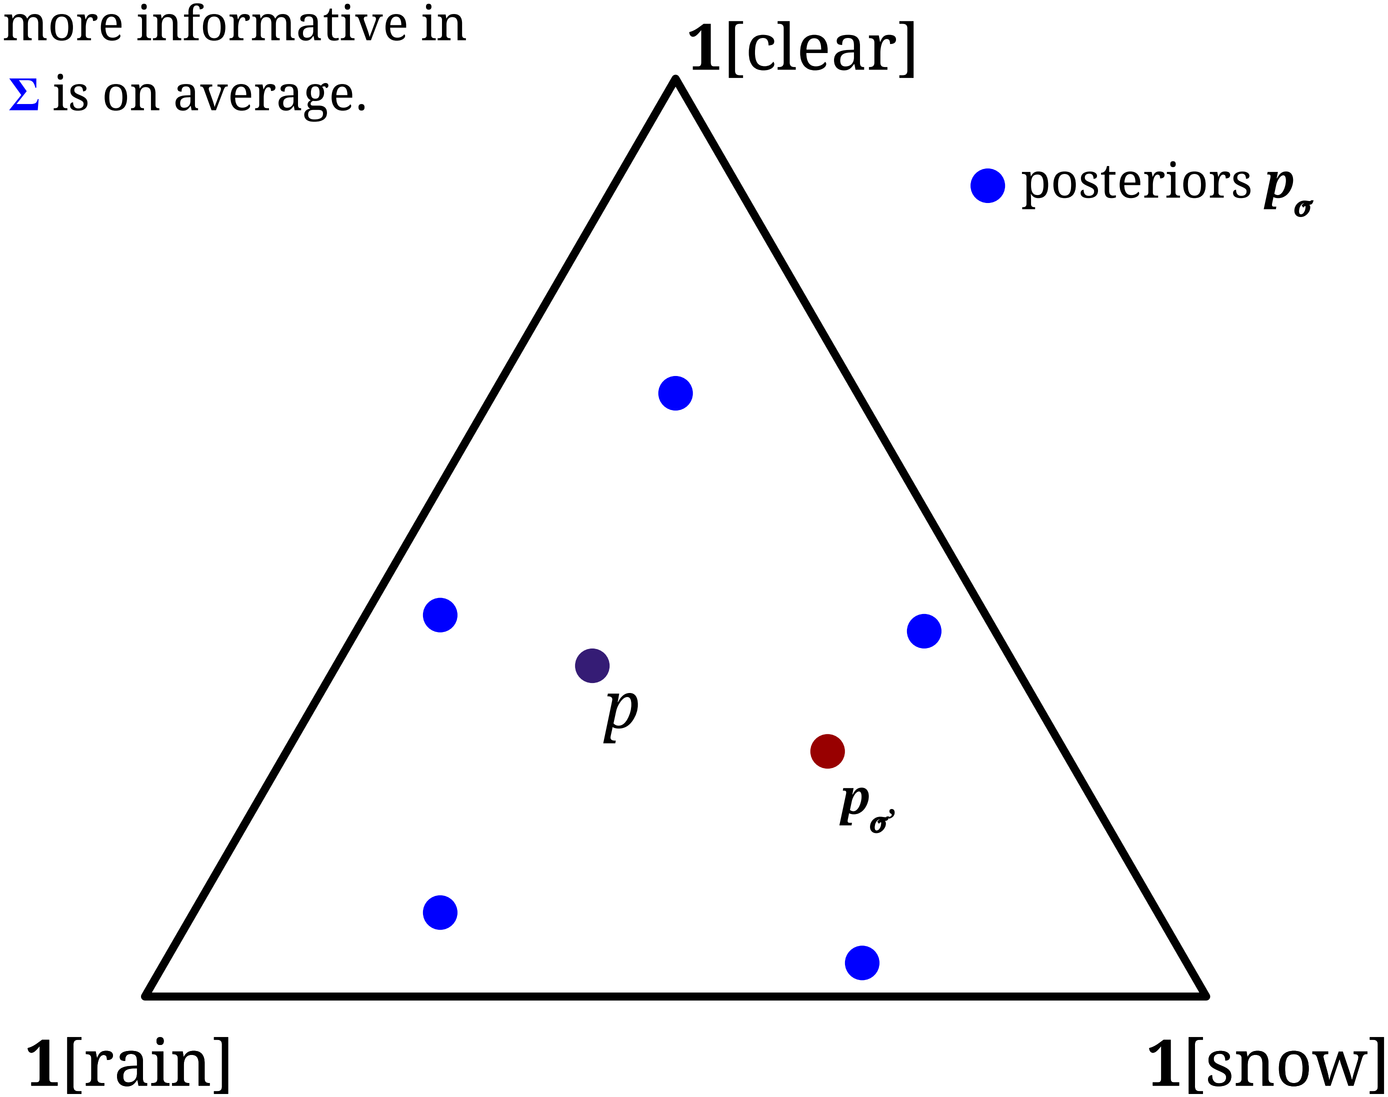

Blackwell proof - general idea

If not a garbling, there exists a realization σ' that is more informative in some direction than Σ is on average.

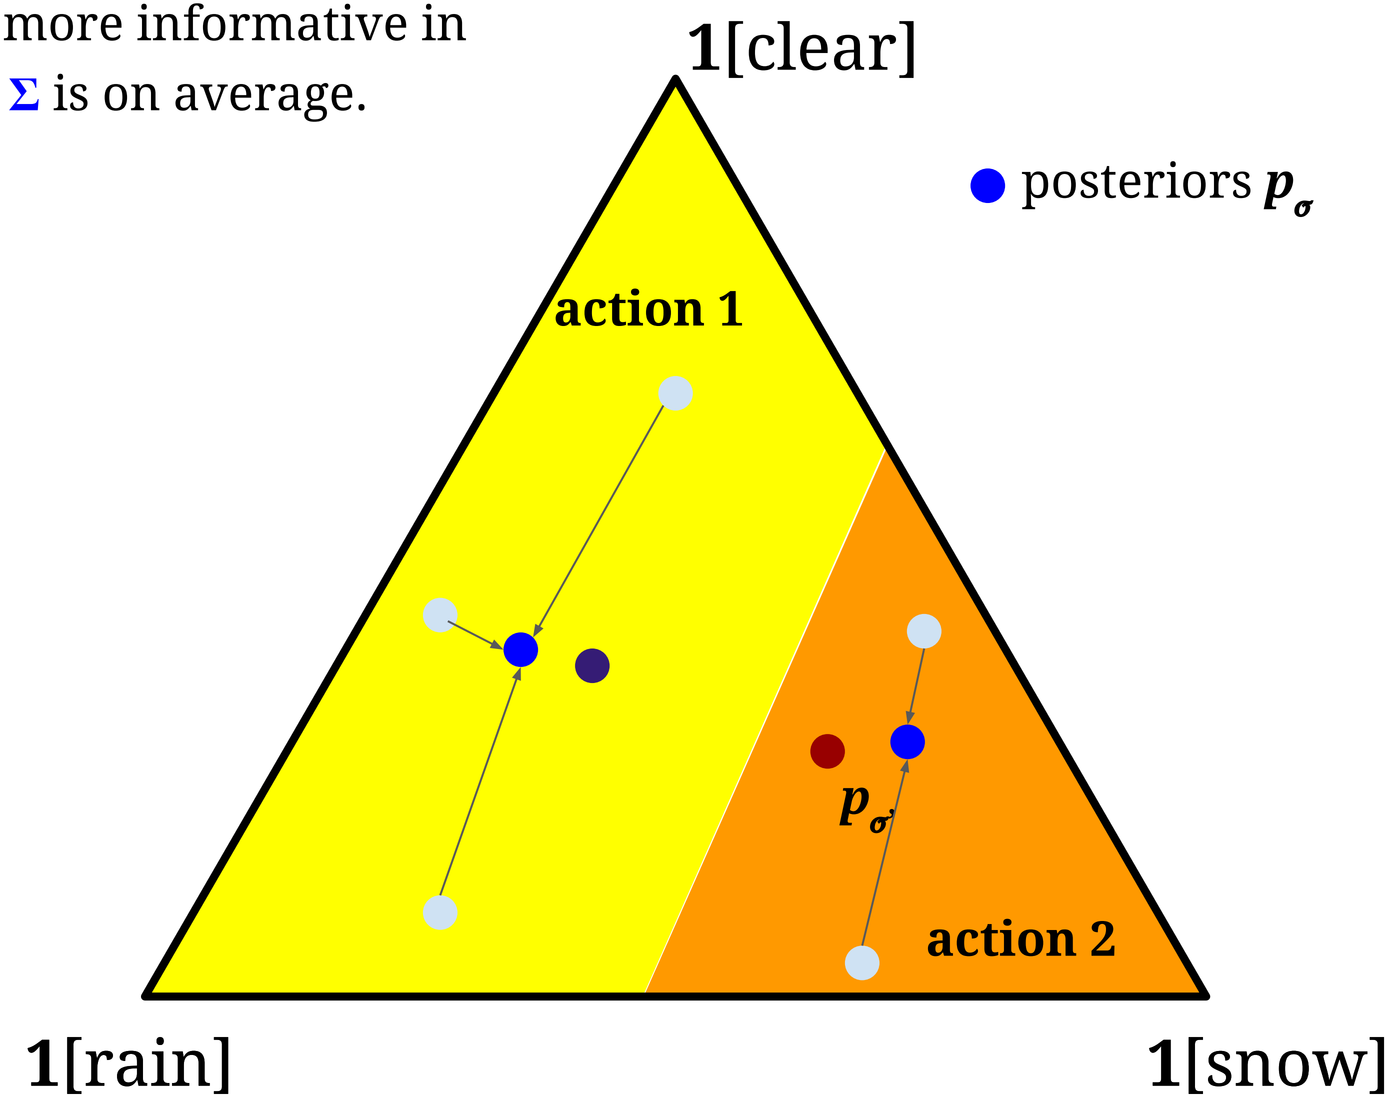

Part 1C:
Bayesian games

The Game Model

The basic game

- $I = \{1,\ldots,n\}$ set of agents — Agent i
- A_i set of actions for agent i — Agent i chooses $a_i \in A_i$
- $A = A_1 \times A_2 \ldots A_n$ — $a \in A$ is an action profile
- $\theta \in \Theta$ — state of nature with prior p
- $u_i((a_i, a_{-i}), \theta)$ — utility function for i

Information structure (Σ, φ) of the game

- Σ_i set of signals for i — agent i observes $\sigma_i \in \Sigma_i$
- $\Sigma = \Sigma_1 \times \Sigma_2 \ldots \Sigma_n$ — $\sigma \in \Sigma$ is a signal profile
- $\varphi(\sigma, \theta)$ — prob. of $\sigma \in \Sigma$ given state θ

Note: a_{-i} denotes the set of all agents' actions except i's
(similar definition for σ_{-i})

Equilibrium Concepts

Bayes Nash Equilibrium (BNE)

- A **strategy** for player i is $\beta_i : \Sigma_i \to \Delta(A_i)$
- $\beta_i\,(a_i|\sigma_i)$= Prob(take a_i when observing σ_i)

$\{\beta_i\}_{i=1,\ldots,n}$ forms a BNE if unilateral deviation is not beneficial for any agent.
Specifically, for any agent i, signal $\sigma_i \in \Sigma_i$, action $a_i \in A_i$ with $\beta_i\,(a_i|\sigma_i) > 0$,

$$\sum_{\sigma_{-i},a_{-i}} p(\theta)\varphi((\sigma_i,\sigma_{-i}),\theta)\left(\Pi_{j\neq i}\beta_j(a_j|\sigma_j)\right)u_i((a_i,a_{-i}),\theta)$$

$$\geq \sum_{\sigma_{-i},a_{-i}} p(\theta)\varphi((\sigma_i,\sigma_{-i}),\theta)\left(\Pi_{j\neq i}\beta_j(a_j|\sigma_j)\right)u_i((a'_i,a_{-i}),\theta)$$

for all $a'_i \in A_i$

Equilibrium Concepts

Bayes Correlated Equilibrium (BCE)

- An action recommendation rule $\pi : \Theta \times \Sigma \to \Delta(A)$
- $\pi\,(a \mid \theta, \sigma\,)$= Prob(recommend action profile a conditioned on θ, σ)

π is a BCE if the recommendation satisfies following obedience constraints:

for any agent i, signal $\sigma_i \in \Sigma_i$, action $a_i \in A_i$,

$$\sum_{\sigma_{-i}, a_{-i}} p(\theta)\varphi((\sigma_i, \sigma_{-i}), \theta)\pi[(a_i, a_{-i})|\theta, (\sigma_i, \sigma_{-i})]u_i((a_i, a_{-i}), \theta)$$

$$\geq \sum_{\sigma_{-i}, a_{-i}} p(\theta)\varphi((\sigma_i, \sigma_{-i}), \theta)\pi[(a_i, a_{-i})|\theta, (\sigma_i, \sigma_{-i})]u_i((a'_i, a_{-i}), \theta)$$

for all $a'_i \in A_i$

A Simple Fact

Fact 6

Any BNE corresponds to a BCE.

Proof.

Follows from definition.

Formally, let $\pi(a \mid \theta, \sigma) = \mathbf{\Pi}_{\mathrm{i}} \beta_i (a_i \mid \sigma_i)$ for each θ.

Comparison of Information Structures

- Goal: compare informativeness of information structures
- Recall: Blackwell order compares informativeness of signaling schemes
  - (Σ', φ') is a **garbling** of (Σ, φ) if they can be coupled such that $\mathbf{\Sigma'}$ is independent of θ conditioned on $\mathbf{\Sigma}$

A generalization of **garbling** for $\sigma=(\sigma_1, \sigma_2, \ldots, \sigma_n) \in \Sigma$:

Information structure (Σ, φ) is **individually sufficient** for (Σ', φ') if they can be coupled such that for any $i = 1,2,\ldots,n$, $\mathbf{\Sigma'}_i$ is independent of θ and $\mathbf{\Sigma}_{-i}$ conditioned on $\mathbf{\Sigma}_i$.

Intuitively, (Σ, φ) is more informative than (Σ', φ')

Comparison of Information Structures

Theorem [Bergemann/Morris 2016]

(Σ, φ) is **individually sufficient** for (Σ', φ') *if and only if* the set of BCE induced by (Σ, φ) is a **subset** of the set of BCE induced by (Σ', φ') for all Bayesian games.

Remarks:

- More information means smaller BCE set (because more constraints)

Obedience constraints in BCE: for any agent i, signal $\sigma_i \in \Sigma_i$, action $a_i \in A_i$,

$$\sum_{\sigma_{-i}, a_{-i}} p(\theta)\varphi((\sigma_i, \sigma_{-i}), \theta)\pi[(a_i, a_{-i})|\theta, (\sigma_i, \sigma_{-i})]u_i((a_i, a_{-i}), \theta)$$

$$\geq \sum_{\sigma_{-i}, a_{-i}} p(\theta)\varphi((\sigma_i, \sigma_{-i}), \theta)\pi[(a_i, a_{-i})|\theta, (\sigma_i, \sigma_{-i})]u_i((a_i', a_{-i}), \theta)$$

for all $a_i' \in A_i$

Comparison of Information Structures

Theorem [Bergemann/Morris 2016]

(Σ, φ) is **individually sufficient** for (Σ', φ') *if and only if* the set of BCE induced by (Σ, φ) is a **subset** of the set of BCE induced by (Σ', φ') for all Bayesian games.

Remarks:

- More information means smaller BCE set (because more constraints)
- Defines a *partial order* over information structures

Prisoner's Dilemma of Incomplete Information

	Cooperate	Defect
Cooperate	-3 -3	0 -9
Defect	-9 0	-6 -6

Prisoner's Dilemma of Incomplete Information

	Cooperate	Defect
Cooperate	$-3 + \theta$, $-3 + \theta$	$-9 + \theta$, 0
Defect	0, $-9 + \theta$	-6, -6

Reward $\theta \sim$ uniform{0, 2, 4.1} to encourage cooperation

- Information structure I: players know θ exactly (full information)
  - A game of complete information is played for each θ
  - Unique BNE and BCE

Prisoner's Dilemma of Incomplete Information

	Cooperate	Defect
Cooperate	$-3 + \theta$, $-3 + \theta$	$-9 + \theta$, 0
Defect	0, $-9 + \theta$	-6, -6

Reward $\theta \sim$ uniform{0, 2, 4.1} to encourage cooperation

- Information structure II: players know whether $\theta = 0$ or not
  - After realization of θ, φ reveals "$\theta \neq 0$" or "$\theta = 0$"

- Unique Bayes Nash Equilibrium:
  - For each player: "$\theta = 0$" $\rightarrow$ defect; "$\theta \neq 0$" $\rightarrow$ cooperate

Prisoner's Dilemma of Incomplete Information

	Cooperate	Defect
Cooperate	$-3 + \boldsymbol{\theta}$ $-3 + \boldsymbol{\theta}$	0 $-9 + \boldsymbol{\theta}$
Defect	$-9 + \boldsymbol{\theta}$ 0	-6 -6

Reward $\boldsymbol{\theta} \sim$ uniform{0, 2, 4.1}
to encourage cooperation

- Information structure II: players know whether $\boldsymbol{\theta} = 0$ or not
  - After realization of $\boldsymbol{\theta}$, φ reveals "$\boldsymbol{\theta} \neq 0$" or "$\boldsymbol{\theta} = 0$"

- This BNE corresponds to a BCE, but there are also other BCEs.
  - π *((*cooperate, cooperate*)* $|\, \theta = 4.1$, "$\boldsymbol{\theta} \neq 0$") = 1
  - π *((*defect, defect*)* $|\, \theta = 2$, "$\boldsymbol{\theta} \neq 0$") = 1
  - π *((*defect, defect*)* $|\, \theta = 0$, "$\boldsymbol{\theta} = 0$") = 1

Prisoner's Dilemma of Incomplete Information

	Cooperate	Defect
Cooperate	$-3+\theta$, $-3+\theta$	$-9+\theta$, 0
Defect	0, $-9+\theta$	-6, -6

Reward $\theta \sim$ uniform{0, 2, 4.1} to encourage cooperation

- Information structure III: players know nothing about θ besides its prior

Exercise: prove there is still a unique BNE, but the set of BCE is even larger than the previous one

A Useful Remark for Bayesian Games

A Bayesian agent
in equilibrium
cannot be misinformed

(There is no lying or misinformation.
Agents can only be more informed or less informed)

Explanation: models assume that players know all prior distributions, and player actions and signals are indeed drawn from these distributions.

Recap and Takeaways

Signals

- Signal $\longleftrightarrow$ set of posterior beliefs $\{p_1,\ldots,p_n\}$ whose expectation is the prior. (amenable to linear programs)

Decision problems

- Decision problem $\longleftrightarrow$ convex function G on Δ_Θ.

Revelation principles

- For agents: can WLOG report belief on Θ; optimal action is simulated.
- For signaller: can WLOG signal a recommended action (it will be followed).

Bayesian Games

- Games with uncertainty and player private information
- Main solution concepts: Bayes-Nash and Bayes-Correlated equilibria
- Information structure affects equilibrium

Notation

• A	set of actions	agent chooses a
• Θ	set of states of the world	nature draws θ
• $u(a, \theta)$	utility function	

• p	prior distribution on Θ	known to agent
• Σ	a signal (also refers to set of realizations)	agent observes $\Sigma=\sigma$
• $\varphi(\sigma, \theta)$	probability of signal σ given state θ	
• p_σ	posterior distribution on Θ given σ	given by Bayes' rule

• $u(a\,;q)$	$= \mathbf{E}_{\theta\sim q}\, u(a, \theta)$	linear function of q
• $G(q)$	$= \max_a u(a\,;q)$	convex function of q
• $a^*(q)$	$= \operatorname{argmax}_a u(a\,;q)$	optimal action given q
• $\mathrm{V}^{\mathrm{u},\varphi}(\Sigma)$	$= \mathbf{E}_{\sigma\sim\varphi} G(p_\sigma)$	exp. utility observing Σ