Convex order and preservation of convexity for Bayesian posterior updates

Yuqiong Wang
yuqw@umich.edu

Joint work with Erhan Bayraktar

Department of Mathematics
University of Michigan

Workshop on Optimal Stopping and Its Applications, Manchester

September 14, 2026

Outline

A motivating example

- A coin has unknown probability $\Theta \in (0,1)$. Two observers, Alice and Bob observes $X_1, X_2 \ldots$ and try to predict the next outcome. Each wrong prediction is penalized by 1.
- Conditioning on Θ, predicting T would have penalization $\mathsf{P}(X_{n+1} = H|\Theta) = \Theta$, while predicting H has $1 - \Theta$.
- After n tosses, $\Pi_n = \mathsf{E}[\Theta \mid \mathcal{F}_n]$ is the predictive probability of heads. The total penalization is $g(\Pi_n) = \min\{\Pi_n, 1 - \Pi_n\}$.
- Assume one is allowed to observe at most one more time with cost c, before making the prediction.

The one-step stopping problem is

$$V^{(1)}(n,z) = \min\left\{g(z), c + \mathsf{E}_{n,z}\left[g(\Pi_{n+1})\right]\right\}.$$

Assume a Uniform prior $\mathrm{Beta}(1,1)$ for Θ. Alice observes one H and one T; Bob observes 10 heads and 10 tails:

$$\mu_2^A = \mathrm{Beta}(2,2), \quad \mu_{20}^B = \mathrm{Beta}(11,11), \quad \Pi_2^A = \Pi_{20}^B = \tfrac{1}{2}.$$

Question: who is better off (has a lower cost)?

A motivating example

- If Alice and Bob each do one more toss, each outcome has probability $1/2$:

$$\Pi_3^A \in \left\{\tfrac{2}{5}, \tfrac{3}{5}\right\}, \quad \Pi_{21}^B \in \left\{\tfrac{11}{23}, \tfrac{12}{23}\right\}.$$

- Thus Alice is better off: $\Pi_3^A \mid \Pi_2^A = \frac{1}{2} \geq_{\text{cx}} \Pi_{21}^B \mid \Pi_{20}^B = \frac{1}{2}$.
- For all $c \geq 0$,

$$V^{(1)}(2, \frac{1}{2}) = \min(\frac{1}{2}, c + \frac{2}{5}) \leq \min(\frac{1}{2}, c + \frac{11}{23}) = V^{(1)}(20, \frac{1}{2}).$$

- If $1/46 < c < 1/10$, Alice buys the observation and Bob stops.

Questions:

- Isn't more information better? By Jensen, $\Pi_m \leq_{\text{cx}} \Pi_n$ for $m < n$. But we consider the conditional one: coming back to the same estimate itself is evidence.
- Is Alice's advantage a consequence of Beta/Bernoulli conjugacy, or does it reflect a general property?

A motivating example

In this work, we show that it is a general property, for one-parameter exponential families, arbitrary priors, and any non-decreasing statistic T.

In particular, we show

- Along a posterior level curve, future updates decrease in convex order as observation time increases. (cf. Ekström-Vaicenavicius, Ekström-Karatzas-Vaicenavicius, Ekström-Wang)
- Posterior transition operators preserve convexity. (cf. Janson-Tysk, Ekström-Tysk)

Together, these yield time monotonicity of the value function for a broad class of stopping problems.

Problem setup

- Conditioning on $\Theta = u$, observations X_i's are i.i.d with density $p_u(x) \propto e^{ux - B(u)}$.
- For $Y_n = X_1 + \cdots + X_n$ and an arbitrary prior μ, define the posterior

$$\mu_{n,Y_n}(\mathrm{d}u) \propto \mathrm{e}^{uY_n - nB(u)}\mu(\mathrm{d}u).$$

- We interpolate this posterior distribution continuously:

$$\mu_{t,y}(\mathrm{d}u) \propto \mathrm{e}^{uy - tB(u)}\mu(\mathrm{d}u), \quad t \geq 0.$$

- For any non-decreasing T, define

$$\Pi_t^T = \mathsf{E}[T(\Theta) \mid \mathcal{F}_t], \quad \Lambda_t(y) = \int T \mathrm{d}\mu_{t,y}(du) =: z.$$

- Since $\partial_y \Lambda_t(y) = \mathrm{Cov}_{\mu_{t,y}}(T(\Theta), \Theta) > 0$, a state z determines a unique posterior $\mu_{t,z} := \mu_{t,y(t,z)}$.

Problem setup

Some possible choices of T:

- **Estimation:** $T(u) = u$. Then Π_n^T is the posterior mean.
- **Sequential testing:** $T(u) = \mathbf{1}_{(\theta_0, \infty)}(u)$. Then $\Pi_n^T = \mathsf{P}(\Theta > \theta_0 \mid \mathcal{F}_n)$.
- **Forecasting:** $T(u) = B'(u)$. Then $\Pi_n^T = \mathsf{E}[X_{n+1} \mid \mathcal{F}_n]$.

Outline

The covariance lemma

Lemma

Let T and $q \geq 0$ be non-decreasing, let G be concave, and suppose $\mathrm{Cov}_\mu(T, G) = 0$. Then

$$\mathrm{Cov}_{q\mu/\int q \mathrm{d}\mu}(T, G) \leq 0.$$

The inequality becomes reversed when G is convex.

- Concavity means that the slope of G decreases as the parameter increases.
- The condition $\mathrm{Cov}_\mu(T, G) = 0$ says that the contributions from smaller and larger parameter values initially balance.
- An increasing q puts relatively more weight on the larger u where G' is smaller.

What $q(u)$ do we choose

Fix a posterior $\mu_{t,z}$ and a threshold $a \in \mathbb{R}$. After the next k observations, define

$$q_a(u) = \mathsf{P}_u(\Pi^T_{t+k} > a).$$

- We evaluate at $\mu_{t,z}$,

$$\mathsf{E}_{t,z}\left[(\Pi^T_{t+k} - a)^+\right] = \int (T(u) - a) q_a(u) \mu_{t,z}(\mathrm{d}u).$$

- We use $(z-a)^+$ because
  - For distributions α and β with the same expected value,

$$\alpha \leq_{\mathrm{cx}} \beta \iff \int (x-a)^+ \alpha(dx) \leq \int (x-a)^+ \beta(dx) \quad \text{for every } a \in \mathbb{R}.$$

  - Together with affine functions they generate every convex function.

Main result 1: convex order comparison

Hold the state fixed: $\Lambda_t(y_t) = z$. Differentiating along the level curve,

$$\frac{\mathrm{d}}{\mathrm{d}t}\mu_{t,z}(\mathrm{d}u) = G_t(u)\,\mu_{t,z}(\mathrm{d}u), \quad G_t(u) = u\,\frac{\partial}{\partial t}y(t,z) - B(u) - \text{constant}.$$

- G_t is concave
- $\mu_{t,z}$ is a probability measure: $\frac{\mathrm{d}}{\mathrm{d}t}\int \mathrm{d}\mu_{t,z} = \int G_t \mathrm{d}\mu_{t,z} = 0$.
- We stay on the z-level: $\frac{\mathrm{d}}{\mathrm{d}t}\int T\mathrm{d}\mu_{t,z} = \int TG_t\mathrm{d}\mu_{t,z} = 0$. Hence $\mathrm{Cov}_{\mu_{t,z}}(T, G_t) = 0$.

With q_a frozen at the later time, the Corollary gives

$$\frac{\mathrm{d}}{\mathrm{d}t}\int (T(u)-a)q_a(u)\,\mu_{t,z}(\mathrm{d}u) \le 0 \implies \text{every } \mathbb{E}[(\Pi_s - a)^+ | \Pi_t = z] \text{ decreases in } t,\ \ t < s.$$

Theorem (1)

For $m < n$, $k \ge 1$, and a state z attainable throughout $[m, n]$,

$$\Pi^T_{m+k} \mid \{\Pi^T_m = z\} \ge_{\mathrm{cx}} \Pi^T_{n+k} \mid \{\Pi^T_n = z\}.$$

Corollary: $t \mapsto \mathrm{Cov}_{\mu_{t,z}}(T(\Theta), \Theta)$ is non-increasing.

Main result 2: preservation of convexity

- Fix t and differentiate the posterior twice with respect to z:

$$\frac{\partial^2}{\partial z^2}\mu_{t,z}(\mathrm{d}u) = Q_z(u)\,\mu_{t,z}(\mathrm{d}u),$$

  where Q_z is quadratic and convex.
- The same two constraints differentiated twice gives $\int Q_z \mathrm{d}\mu_{t,z} = \int TQ_z \mathrm{d}\mu_{t,z} = 0$.

Theorem (2)

Fix t and $k \geq 1$. If f is convex, then $z \mapsto \mathsf{E}_{t,z}[f(\Pi^T_{t+k})]$ is convex.

- For a concave payoff g, consider

$$V(n,z) = \inf_{\tau} \mathsf{E}_{n,z}\left[g(\Pi^T_{n+\tau}) + c\tau\right],$$

$$V_{j+1}(n,z) = \min\left\{g(z),\, c + \mathsf{E}_{n,z}\left[V_j(n+1, \Pi^T_{n+1})\right]\right\}.$$

- Theorem 2 keeps every iterate concave, so Theorem 1 can be applied to it.
- $V(n,\cdot)$ is concave and $V(m,z) \leq V(n,z)$ for $m \leq n$.

Outline

Example 1: Bayesian sequential testing

Fix θ_0 and let $T(u) = \mathbf{1}_{(\theta_0,\infty)}(u)$, so $\Pi_n = \mathsf{P}(\Theta > \theta_0 \mid \mathcal{F}_n)$. With unit loss for a wrong decision and cost c per observation,

$$V(n,\pi) = \inf_{\tau} \mathsf{E}_{n,\pi}\left[\Pi_{n+\tau} \wedge (1 - \Pi_{n+\tau}) + c\tau\right].$$

- Ekström-Wang impose the one-step comparison as Assumption 5.1 and conjectured for time-monotonicity.
- The corollary $t \mapsto \mathrm{Cov}_{\mu_{t,z}}(T(\Theta), \Theta)$ non-increasing recovers their posterior concentration result, that the π-level curves spread out in t.

Example 2: investment under learning

Let $T(u) = u$ and $\widehat{\Theta}_n = \mathsf{E}[\Theta \mid \mathcal{F}_n]$. The decision maker learns about an unknown profitability and chooses both when to stop and whether to invest, discounting at rate $r > 0$:

$$V(n, z) = \sup_{\tau \geq n} \mathsf{E}\left[\mathrm{e}^{-r(\tau - n)}(\widehat{\Theta}_\tau - K)^+ \mid \widehat{\Theta}_n = z\right].$$

With $g(x) = (x - K)^+$.

- g is convex, so Theorem 2 and induction keep $x \mapsto V_j(n, x)$ convex;
- Theorem 1 then gives $V(n, z) \leq V(m, z)$ for $m < n$.

At the same posterior mean, the value of learning is larger to the one that has learned less.

Example (When T is not monotone)

For $T(p) = \mathbf{1}_{\{p=1/2\}}$ with prior $(\frac{1}{4}, \frac{1}{2}, \frac{1}{4})$ on $p \in \{\frac{1}{4}, \frac{1}{2}, \frac{3}{4}\}$,

$$\Pi_1^T \mid \{\Pi_0^T = \tfrac{1}{2}\} \sim \delta_{1/2}, \quad = \Pi_2^T \mid \{\Pi_1^T = \tfrac{1}{2}\} \sim \tfrac{9}{16}\delta_{4/9} + \tfrac{7}{16}\delta_{4/7}.$$

T being monotone is essential.

Comment: the continuous-time analogue

In discrete time, after n observations,

$$\mu_{n,Y_n}(\mathrm{d}u) \propto \mathrm{e}^{uY_n - nB(u)}\mu(\mathrm{d}u).$$

- For general exponential families, replacing n by a real number t initially gives only an analytic interpolation:
$$\mu_{t,y}(\mathrm{d}u) \propto \mathrm{e}^{uy - tB(u)}\mu(\mathrm{d}u).$$
- If one member of the exponential family is infinitely divisible, the family admits a continuous-time realization by a Lévy process L.
- Thus t becomes genuine information time, and both theorems hold for all real times.

Examples: unknown Brownian drift, unknown Poisson intensity, gamma processes.

Not covered: the construction does not cover unknown volatility, arbitrary changes of the Lévy triplet, or multidimensional parameters.

Comment: from observation number to continuous time

Consider Brownian observations with unknown drift:

$$\mathrm{d}Y_t = \Theta \mathrm{d}t + \mathrm{d}W_t, \quad \Pi_t^T = \mathsf{E}[T(\Theta) \mid \mathcal{F}_t^Y].$$

Itô's formula gives

$$\mathrm{d}\Pi_t^T = \partial_y \Lambda_t(Y_t)\mathrm{d}\widehat{W}_t = \mathrm{Cov}(T(\Theta), \Theta \mid \mathcal{F}_t^Y)\mathrm{d}\widehat{W}_t.$$

- The conditional covariance is therefore the instantaneous volatility of the posterior statistic.
- For $T(u) = \mathbf{1}_{\{u>\theta_0\}}$, Ekström-Vaicenavicius show that this volatility decreases with t along each fixed posterior probability level.
- For $T(u) = u$, Ekström-Karatzas-Vaicenavicius establish the corresponding result for the posterior mean.

Example (3)

For $Y_t = \Theta t + W_t$, let

$$\Pi_t = \mathsf{P}\left(\sup_{r\le h}(Y_{t+r} - Y_t) \ge a \mid \mathcal{F}_t^Y\right) = \mathsf{E}[T(\Theta) \mid \mathcal{F}_t^Y],$$

where

$$T(u) = \Phi\left(u\sqrt{h} - \frac{a}{\sqrt{h}}\right) + \mathrm{e}^{2au}\Phi\left(-u\sqrt{h} - \frac{a}{\sqrt{h}}\right).$$

The person with more accumulated information has the less dispersed future revisions.

Summary

- Same estimates today but different history can give different value of information.
- Convexity survives Bayesian updating.
- These yield structural results for the value function of a large family of stopping problems. Convexity and time-monotonicity.
- We do not require conjugate priors. They hold for arbitrary priors and many standard observation models, both discrete time and continuous time.

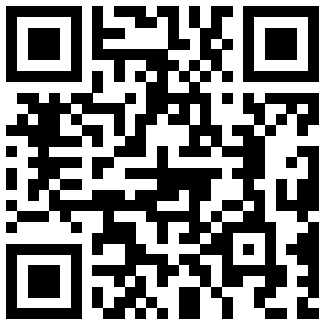

Preprint: arXiv:2609.05065

Thank you!