Bayesian sequential testing and estimation in discrete time

Yuqiong Wang

Partly joint work with Erik Ekström.

Department of Mathematics
Uppsala University

Young Researchers in Stochastic Control and Games,
December 14 2023, Bielefeld

Statistics & Optimal stopping

Classical problem: Testing the unknown drift of a BM.

- Observe the trajectory of a BM with unknown drift:

$$X_t = \Theta t + W_t.$$

  where $\mathbb{P}(\Theta = 1) = \pi = 1 - \mathbb{P}(\Theta = 0)$, $\pi \in (0,1)$.
- Want to test: $H_1 : \Theta = 1$, $H_0 : \Theta = 0$, as accurately as possible.
- Observation is not free: $c > 0$ per unit time of observation.
- Need to test as fast as possible.
- The time to stop observing is part of the decision.

Statistics & Optimal stopping

Solution: optimal stopping in another coordinate.

- The minimised cost V:

$$V = \inf_{\tau,d} \{\mathbb{P}(d=0,\Theta=1) + \mathbb{P}(d=1,\Theta=0) + c\mathbb{E}[\tau]\}. \qquad (1)$$

- Defining the posterior probability process

$$\Pi_t := \mathbb{P}_\pi(\Theta = 1 | \mathscr{F}_t^X),$$

- Problem (1) can be written as

$$V(\pi) = \inf_\tau \mathbb{E}_\pi[c\tau + \Pi_\tau \wedge (1-\Pi_\tau)]$$

where

$$d\Pi_t = \Pi_t(1-\Pi_t)d\tilde{W}_t,$$

- Standard method: explicit solution: Shiryaev (1969).

Application in finance: unknown parameters

- Direct application: testing the unknown drift of a stock (GBM).
- What about an unknown volatility?
  – Not a valid question in continuous time!

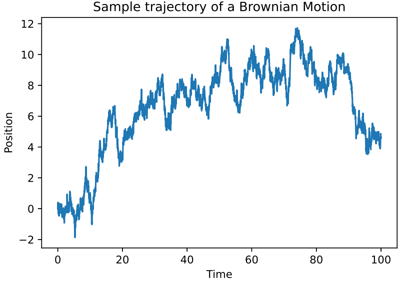

- Financial data is not really continuous.

This motivates us to consider things in **discrete time**.

Natural questions to ask

Question 1: what is the natural class of problems to study in discrete time?
(In terms of composite testing, whole range of posibilities)

- Or, what family of distributions do we consider?

Question 2: can we assign arbitrary distribution to the unknown parameter?

- Explicit solutions?
- Otherwise, what structural properties does the problem exhibit?

What are we building on?

Discrete time: popular in the 60s and 70s

- Studied on a case-by-case basis and rely on conjugate priors: Lindley and Barnett (1965), Moriguti and Robbins (1962).
- Focus on asymptotic behaviour: Schwartz (1962), Bickel (1973), Lai (1988).
- c.f. Sobel (1953), Alvo (1977), Cablio (1977).

Continuous time: many generalisations

- Finite horizon: Gapeev and Peskir (2004), Poisson: Peskir and Shiryaev (2000), multi-dimensional: Ekström and Wang (2022).
- Most literature uses binary priors (c.f. Zhitlukhin and Shiryaev (2011), Ekström and Vaicenavicius (2015), Ekström, Karatzas & Vaicenavicius (2022)).

Discrete-time sequential testing

The general set-up:

- The tester observes $X_1, X_2, \ldots$ sequentially with cost c at each step.
- X_k's are drawn from a one-parameter exponential family depending on r.v. Θ: conditioning on $\Theta = u$, X_k's are independent, and

$$\mathbb{P}(X_1 \in A | \Theta = u) = \int_A e^{ux - B(u)} \nu(dx).$$

- μ: (arbitrary) prior of the unknown parameter Θ

Set-up

- Want to test:

$$H_0: \quad \Theta \leq \theta_0,$$
$$H_1: \quad \Theta > \theta_0,$$

- Let $d = i$ represent H_i is accepted.
- Define the minimal cost

$$V := \inf_{\tau, d} \{\mathbb{P}(d = 1, \Theta \leq \theta_0) + \mathbb{P}(d = 0, \Theta > \theta_0) + c\mathbb{E}[\tau]\}.$$

Set-up

- Define the *posterior probability process* Π

$$\Pi_n := \mathbb{P}(\Theta > \theta_0 | \mathscr{F}_n^X),$$

with $\Pi_0 = \pi = \mu((\theta_0, \infty))$.

- Given a stopping time τ, an optimal decision is

$$d = \begin{cases} 0 & \text{if } \Pi_\tau \leq 1/2, \\ 1 & \text{if } \Pi_\tau > 1/2. \end{cases}$$

- Consequently,

$$V = \inf_\tau \mathbb{E}\left[\Pi_\tau \wedge (1 - \Pi_\tau) + c\tau\right].$$

Markovian embedding

Define $Y_n := \sum_{i=1}^{n} X_i$. For any fixed n, the process Π_n can be written as

$$\Pi_n = q(n, Y_n).$$

One can show that

The function $y \mapsto q(n,y) : \mathbb{R} \to (0,1)$ is an increasing bijection.

We need the **exponential family** for the above to hold!

- At time n, knowing y gives the shape of the posterior.
- Can embed any (n,y) for $n \geq 0$, $y \in \mathbb{R}$.
- We denote by $y(n,\pi)$ the unique value that $q(n,y(n,\pi)) = \pi$, and refer the set $\{(n,y(n,\pi)), n \geq 0\}$ as the $\pi-$*level curve.*

Main result 1: concavity

Define $\mathbb{P}_{n,\pi}(\cdot) := \mathbb{P}(\cdot|\Pi_n = \pi)$. The optimal stopping problem can be written as

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

Concavity of V

- Let $f : [0,1] \to [0,\infty)$ be a concave function. Then $\pi \mapsto \mathbb{E}_{n,\pi}[f(\Pi_{n+1})]$ is concave on $(0,1)$.
- The function $\pi \mapsto V(n,\pi)$ is concave for each fixed $n \geq 0$.

Concavity and its consequence

Introduce now

- The continuation region $\mathscr{C}$:

$$\mathscr{C} := \{(n,\pi) \in \mathbb{N}_0 \times [0,1] : V(n,\pi) < \pi \wedge (1-\pi)\},$$

- The stopping region $\mathscr{D}$ by

$$\mathscr{D} := \{(n,\pi) \in \mathbb{N}_0 \times [0,1] : V(n,\pi) = \pi \wedge (1-\pi)\}.$$

- By standard optimal stopping theory, the stopping time

$$\tau^* := \inf\{k \geq 0 : (n+k, \Pi_{n+k}) \in \mathscr{D}\}$$

  is an optimal strategy.
- Concavity implies that the stopping boundaries are **Two-sided**,

Main result 2: concentration of the posterior

The posterior distribution squeezes in

If $a < \theta_0 < b$, then

$$n \mapsto \mathbb{P}_{n,\pi}(\Theta \leq a) \quad \& \quad n \mapsto \mathbb{P}_{n,\pi}(\Theta > b)$$

are decreasing.

As a consequence, the $\pi-$level curves are spreading out.

Monotonicity in time

An assumption

For any $\pi \in (0,1)$ and $n \geq m \geq 0$, the random variable $\Pi_{m+1}|\{\Pi_m = \pi\}$ dominates $\Pi_{n+1}|\{\Pi_n = \pi\}$ in convex order.

- Assume the above holds. Then $V(n,\pi)$ is non-decreasing in n, and the boundaries are monotone.
- But does this assumption always hold? **A:** We don't know.

Time-monotonicity?

- Holds for some examples with **any** prior (Gaussian w. unknown mean, Bernoulli, Binomial).
- Holds for some other examples for **some families** of priors (Exp, Gaussian w. unknown variance).
- No counter-example is found.

Conjecture: V is non-decreasing in n.

Introduction

Aside from testing, another natural question to ask is

What is the value of the unknown parameter?

- Want to obtain an accurate estimate in the presence of cost.
- We can ask similar questions as in the testing problem. (remind the audience the questions)

Discrete-time sequential estimation

- The basic set-up is the same as in *testing*
- Formulate the stopping problem in another coordinate.

The coordinate

Define the *posterior estimate process*:

$$\hat{\Theta}_n := \mathbb{E}\left[\Theta | \mathscr{F}_n^X\right].$$

Want to minimize:

$$\mathbb{E}\left[(\Theta - \hat{\Theta}_\tau)^2 + c\tau\right]$$

over stopping times.

Markovian embedding

Similarly, we are fine in this coordinate because

$\hat{\Theta}_n = G_n(Y_n)$ is a strictly increasing bijection.

Again, we need the **exponential family** for this to hold!

- Define $\Psi(n, \hat{\Theta}_n) = \mathrm{Var}(\Theta | \mathscr{F}_n^X)$, then V can be written in the θ_0 coordinate:

$$V(n, \theta_0) = \inf_{\tau \in \mathscr{T}} \mathbb{E}_{n,\theta_0}[\Psi(n+\tau, \hat{\Theta}_{n+\tau}) + c\tau].$$

- Note that $M_n := \Psi(n, \hat{\Theta}_n) + \hat{\Theta}_n^2$ is a martingale, we can further write

$$V(n, \theta_0) = \Psi(n, \theta_0) + \inf_{\tau \in \mathscr{T}} \mathbb{E}_{n,\theta_0} \left[\sum_{i=0}^{\tau} \left(c - \left(\hat{\Theta}_{i+1}^2 - \hat{\Theta}_i^2 \right) \right) \right]$$
$$=: \Psi(n, \theta_0) + v(n, \theta_0).$$

Main result: conditions for space-monotonicity

First-order stochastic dominance

If $\theta_0 \leq \tilde{\theta}_0$, then $\mathbb{P}(\hat{\Theta}_n^{\theta_0} \leq a) \geq \mathbb{P}(\hat{\Theta}_n^{\tilde{\theta}_0} \leq a)$, for all $a \in \mathbb{R}$ and all $n \geq 0$.

As a consequence,

Space-monotonicity of v

Assume that for all $k \geq 0$, the mapping

$$\theta_0 \mapsto \mathbb{E}_{k,\theta_0}\left[\hat{\Theta}_{k+1}^2 - \theta_0^2\right]$$

is non-decreasing (non-increasing), then the value function $v(n, \theta_0)$ is non-increasing (non-decreasing) in θ_0 for any fixed $n \geq 0$.

This implies a **one-sided** stopping boundary.

Monotonicity in Space

This is not a general result. It depends on both the *prior* and the *observation*.

Some examples

- Bernoulli observations with any prior: **not monotone**.
- Exponential observations with a gamma prior: **monotone**.
- Gaussian observations with unknown variance and an inverse gamma prior: **monotone**.

What about time-monotonicity? We have some partial results. e.g. . . .

Dynamic pricing under a binary prior

An example in stochastic control: set-up

- Consider a **seller** who offers a product for sale.
- The potential **buyers** arrive in a sequential fashion.
- At time n, the seller offers a price p_n.
- The probability that p_n is accepted is the **demand**, $D(p)$.
- But $D(\cdot)$ is unknown:

$$\mathbb{P}\left(D(\cdot) = D^1(\cdot)\right) = \pi = 1 - \mathbb{P}\left(D(\cdot) = D^0(\cdot)\right).$$

- The seller seeks to maximise the profit:

$$V = \sup_{\{p_n\}_{n\geq 0}} \mathbb{E}\left[\sum_{n=0}^{\infty} e^{-rn} p_n D(p_n)\right].$$

Economic & operations research literatures: incomplete learning, myopic strategy ....

Dynamic pricing under a binary prior

How does it relate to our setting?

- Observe that it is with Bernoulli observations with a Bernoulli prior:

$$\Theta = \begin{cases} 1, & D(\cdot) = D^1(\cdot), \\ 0, & D(\cdot) = D^0(\cdot). \end{cases}$$

- Define the posterior probability process

$$\Pi_n := \mathbb{P}\left(D(\cdot) = D^1(\cdot) \,|\, \mathscr{F}_n\right),$$

- The value can be written as

$$V(\pi) = \sup_{\{p_n\}_{n\geq 0}} \mathbb{E}_\pi \left[\sum_{n=0}^{\infty} e^{-rn} p_n \left(\Pi_n D^1(p_n) + (1 - \Pi_n) D^0(p_n) \right) \right].$$

Dynamic pricing under a binary prior

And clearly satisfies

$$V(\pi) = \sup_p \left\{ e^{-r} \mathbb{E}_\pi \left[V(\Pi_1^p) \right] + p \left(\pi D^1(p) + (1-\pi) D^0(p) \right) \right\}.$$

A monotone sequence can then be constructed to find a fixed point, which coincides with V.

- Used convexity of $\mathbb{E}_\pi[f(\Pi_1)]$, for f convex.
- The prior distribution can be relaxed to an arbitrary prior.
- The observation can be relaxed.

This opens up doors to general problems of "exploration-exploitation type".

Summary

To summarise the talk:

- We study the Bayesian sequential **testing** and **estimation** problems in discrete time.
- The unknown parameter is taken from the **exponential family**
- The prior can be arbitrary
- In general, no explicit solutions. We are after structural properties.
- The problems we study open up doors to control problems with learning and earning features.

Thank you for your attention!

Contact: yuqiong.wang@math.uu.se