Quickest Detection With Post-Change Distribution Uncertainty

Heng Yang

City University of New York, Graduate Center

Olympia Hadjiliadis

City University of New York, Brooklyn College and Graduate Center

Michael Ludkovski

University of California Santa Barbara

June 24th, 2015

Content of Topics

Introduction

- This quickest detection problem is concerned with detecting an abrupt change in the distribution of a stochastic process.
- In this work, we assume that there is uncertainty about the post-change distribution.
- This problem is different from the estimation of the change point when given the information that the change has already happened in a time window.

Contribution

- We develop a novel family of composite stopping times that combines multiple CUSUM stages along with the CUSUM reaction period (CRP).
- We design such composite stopping times to establish *third-order* optimality for the problem of detecting a Wiener disorder with uncertain post-change drift.
- We also use such composite stopping times to analyze the question of distinguishing the different values of post-change drift and show an asymptotic identification result.

Mathematical Setup

- On a probability space $(\Omega, \mathcal{F}, P)$, we observe the sequential process $\{Z_t\}_{t\geq 0}$ with $Z_0 = 0$.
- Suppose that this process changes from a standard Brownian motion to a Brownian motion with drift m at a *fixed but unknown* change time $\tau \geq 0$, that is:

$$dZ_t := \begin{cases} dW_t & t < \tau \\ mdt + dW_t & t \geq \tau \end{cases}$$

where W is a standard Brownian motion.

Post-Change Drift

- We have no probability structure on the change time τ.
- Suppose that the post-change drift m comes from a Bernoulli distribution, which is independent of the pre-change Brownian motion process $\{W_s\}_{s<\tau}$:

$$m = \begin{cases} m_1 & \text{with likelihood } p \\ m_2 & \text{with likelihood } 1-p \end{cases}$$

  where $0 < m_1 < m_2$ and $0 < p < 1$ are known constants.
- We can generalize to the discrete range $\{m_1, m_2, \ldots, m_N\}$.

Filtration

To facilitate our analysis, we introduce the measures and filtration.

- Assume that the probability space supports a Bernoulli random variable U that is independent of $\{Z_t\}_{t=0}^{\infty}$.
- Let $\mathcal{G}_t$ correspond to the $\sigma-$algebra generated by both $\{Z_s\}_{s \leq t}$ and U. Note that $\mathcal{G}_0 = \sigma(U)$.
- This enlargement of the natural filtration of Z is to enable randomization.

Measures

- $P_\tau^{m_i}$ is the measure defined on this filtration, under which we have $dZ_t = dW_t$ for $t < \tau$, and $dZ_t = m_i dt + dW_t$ for $t \geq \tau$, where $\tau \in [0, \infty)$ and $i = 1, 2$.
- P_∞ represents the measure under which there is no change, and so $dZ_t = dW_t$ for all t.
- The uncertainty regarding m is modeled by a probability measure

$$P_\tau^m = pP_\tau^{m_1} + (1-p)P_\tau^{m_2}.$$

Detection Delay Criterion And False Alarm

- Our objective is to find an stopping time to detect the change point that balances the trade off between a small detection delay and the penalty on the frequency of false alarms.

Detection Delay Criterion And False Alarm

- Our objective is to find an stopping time to detect the change point that balances the trade off between a small detection delay and the penalty on the frequency of false alarms.
- For any $\mathcal{G}$-stopping time R, the false alarm can be represent by $E_{\infty}[R]$.

Detection Delay Criterion And False Alarm

- Our objective is to find an stopping time to detect the change point that balances the trade off between a small detection delay and the penalty on the frequency of false alarms.
- For any $\mathcal{G}$-stopping time R, the false alarm can be represent by $E_\infty[R]$.
- The worst detection delay of R given the post-change drift m_i is of Lorden's form:

$$J_i(R) := \sup_{\tau \geq 0} \text{esssup}_\omega \, E_\tau^{m_i}[(R-\tau)^+ | \mathcal{G}_\tau].$$

- Since m is unknown, we average over J_i's according to the distribution of m to define the detection delay criterion:

$$J(R) := pJ_1(R) + (1-p)J_2(R).$$

The Detection Problem

- We would like to find a stopping time $T \in \mathcal{S}$ to detect τ, where

$$\mathcal{S} = \{\mathcal{G} - \text{stopping time T: } E_\infty[T] \geq \gamma\}$$

for a constant $\gamma > 1$.

- Given γ, m_1, m_2, p, our quickest detection problem is

$$\inf_{T \in \mathcal{S}} J(T)$$

The Λ-CUSUM Stopping Time

A Λ-CUSUM stopping time (Cumulative Sum) is a stopping time with the parameters Λ and $K > 0$ defined as

$$T(\xi, \Lambda, K) := \inf\left\{ t > 0 : V_t - \inf_{0 \le s \le t} V_s \ge K \right\}$$

where

$$V_t = \Lambda \xi_t - \frac{1}{2}\Lambda^2 t$$

and $\xi := \{\xi_t\}_{t \ge 0}$ is a stochastic process.

The Λ-CUSUM Stopping Time

- When $m_1 = m_2 = \Lambda$ is a constant, the log-form of maximum likelihood ratio becomes

$$\log \sup_{0 \le \tau \le t} \frac{dP_\tau^\Lambda}{dP_\infty}\bigg|_{\mathcal{G}_t} = V_t - \inf_{0 \le \tau \le t} V_\tau$$

  where $V_t = \Lambda Z_t - \frac{1}{2}\Lambda^2 t$.
- The CUSUM-form stopping time provides an optimal solution in the detection problem in this case.

Detection Problem Without Post-change Uncertainty

Theorem (Beibel'96; Shiryaev'96; Moustakides'04)

When the post-change drift is a constant m, for any stopping time $T \in \mathcal{S}$, we have

$$J(T) \geq J(T_K^m)$$

, where T_K^m is the CUSUM stopping time satisfying $E_\infty[T_K^m] = \gamma$.

- In our work, we suppose that the post-change drift has uncertainty, that is $0 < m_1 < m_2$. We can show that the Λ-CUSUM stopping time is no longer optimal stopping time or even asymptotic optimal stopping time.

Construction Of A Composite Stopping Time

- We consider a composite stopping time, namely T_{com}, that is constructed in two stages.

- In the first stage, we apply a CUSUM stopping time

$$T_{\nu}^{\lambda} := T(Z, \lambda, \nu)$$

, where $0 < \lambda < 2m_1$ and $\nu > 0$.

- We define the CUSUM reaction period (CRP) with respect to the first stage.

CUSUM Reaction Period (CRP)

- Up to the first stage, there is the last reset time (last running minimum position):

$$\rho(Z,\lambda,\nu) := \sup\left\{t \in [0, T_\nu^\lambda) : V_t = \inf_{s\le t} V_s\right\}.$$

- The time length between the last reset and the end of the first stage is

$$S(Z,\lambda,\nu) := T(Z,\lambda,\nu) - \rho(Z,\lambda,\nu).$$

- Denote $S_\nu^\lambda = S(Z,\lambda,\nu)$ and we call it as the CUSUM reaction period (CRP) with respect to T_ν^λ.

Motivation of CRP

Conditional on the case that the change has happened before the beginning of CRP S_ν^λ, CRP has the following pattern.

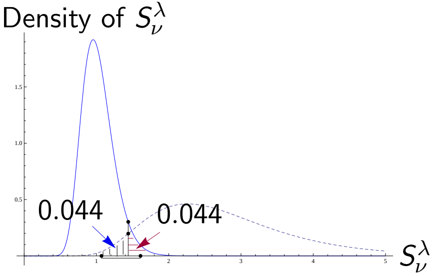

Figure: The parameters are $\lambda = 1$ and $\nu = 5$. The solid curve (left) gives the density function under the case $\{m = 5, \tau < \rho_\nu^\lambda\}$; the dashed curve (right) gives the density under the case $\{m = 2, \tau < \rho_\nu^\lambda\}$.

Construction Of A Composite Stopping Time (Cont.)

- To continue the second stage, we define the shift function θ_u of the process for a time $u \geq 0$ as follows

$$\theta_u(Z) := \{Z_{t+u} - Z_u\}_{t\geq 0}.$$

Note that θ_s makes the path re-start at time s from 0.

- In the second stage, we define $\mathcal{G}$-stopping times $T_{h_1}^{\mu_1}$ and $T_{h_2}^{\mu_2}$:

$$T_{h_1}^{\mu_1} := T(\theta_{T_\nu^\lambda}(Z), \mu_1, h_1) \quad \text{and} \quad T_{h_2}^{\mu_2} := T(\theta_{T_\nu^\lambda}(Z), \mu_2, h_2),$$

where μ_i's are constants satisfying $0 < \mu_i < 2m_1$ and $h_i > 0$ are the second-stage thresholds for $i = 1, 2$.

Construction Of A Composite Stopping Time (Cont.)

- To connect the first and second stages, for a parameter $b > 0$, when $S_\nu^\lambda \geq b$, we continue the second step and stop at $T_{h_1}^{\mu_1}$; when $S_\nu^\lambda < b$, we continue the second step and stop at $T_{h_2}^{\mu_2}$.
- So the composite stopping time T_{com} is defined as

$$T_{com} = T_\nu^\lambda + 1_{\{S_\nu^\lambda \geq b\}} T_{h_1}^{\mu_1} + 1_{\{S_\nu^\lambda < b\}} T_{h_2}^{\mu_2},$$

where $1_{\{S_\nu^\lambda \geq b\}}$ and $1_{\{S_\nu^\lambda < b\}}$ are indicator functions.

Idea Of T_{com}

The diagram below illustrates the construction of T_{com}.

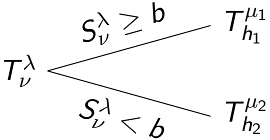

A Lower Bound Of Detection Delay

Lemma (Lower Bound of Delay)

For any $\mathcal{G}$-stopping time $T \in \mathcal{S}$, we have the lower bound on the detection delay $J(T)$ as

$$J(T) \geq LB(\gamma) := \frac{2p}{m_1^2} g\left(f^{-1}(\frac{m_1^2\gamma}{2})\right) + \frac{2(1-p)}{m_2^2} g\left(f^{-1}(\frac{m_2^2\gamma}{2})\right)$$

where

$$g(x) = e^{-x} + x - 1 \quad \textit{and} \quad f(x) = e^{x} - x - 1, \quad \textit{for } x > 0.$$

Third-Order Asymptotic Optimality

Theorem (Asymptotic Optimality of T_{com})

The composite $\mathcal{G}$-stopping time the form T_{com} defined above, satisfying the frequency of false alarm constraint $E_\infty[T_{com}] = \gamma$, is asymptotically optimal of third order in the sense that

$$\lim_{\gamma \to \infty} [J(T_{com}) - LB(\gamma)] = 0$$

when its parameters $\lambda, \nu, b, h_i, \mu_i$ satisfy appropriate conditions.

The Behavior of CRP

Conditional on the case that the change has happened before the recording of CRP S_ν^λ, we have the following result.

Lemma

For any parameter $\lambda \in (0, 2m_1)$ and $b, \nu > 0$ such that $b/\nu = l$ is a positive constant that satisfies

$$l > \frac{2}{\lambda(2m_1 - \lambda)},$$

there exists a positive constant $L = L(m_1, \lambda, l)$ such that

$$\lim_{\nu \to \infty} P_\tau^{m_1}\left(S_\nu^\lambda \geq b \mid \tau < \rho_\nu^\lambda\right) e^{-L\nu} = 0.$$

The Behavior of CRP (Cont.)

When there is no change during the recording of CRP, we have the following result.

Lemma

For parameters $\lambda, b, \nu > 0$ such that b/ν is a positive constant , we have

$$\lim_{\nu \to \infty} P_\infty(S_\nu^\lambda < b) = 0.$$

Thus, the event $\{S_\nu^\lambda < b\}$ can help to distinguish the two cases.

Identification Of Post-Change Distribution

- In our work, we suppose that the post-change drift has uncertainty, that is $0 < m_1 < m_2$.
- We may want not only to detect the change point, but also to identify the post-change distribution.
- We want to find the additional information about the post-change drift based on the composite stopping time.
- This requirement is different from the estimation problem solely.

An Identification Function

- For the composite rule T_{com}, it is possible to construct an identification function $\delta_{T_{com}} \in \mathcal{G}_{T_{com}}$ taking values in $\{m_1, m_2\}$ to serve the purpose of post-change identification of the drift.
- To this effort, we denote

$$S_{h_2} := S(Z, \mu_2, h_2)$$

as CRP in the second stage $T_{h_2}^{\mu_2}$, and denote the CRP $S_\nu = S_\nu^\lambda$ in the first stage.

- The identification function $\delta_{T_{com}}$ is defined as follows

$$\delta_{T_{com}} := \begin{cases} m_1 & \text{if } \{S_\nu \geq b_\nu\} \cup \{S_{h_2} \geq b_{h_2}, S_\nu < b_\nu\}; \\ m_2 & \text{if } \{S_{h_2} < b_{h_2}, S_\nu < b_\nu\}. \end{cases}$$

for positive constants b_ν and b_{h_2}.

Asymptotic Identification Performance

Proposition

If we choose $\mu_2 = m_1$ *and a constant* l *such that*

$$\frac{2}{m_1(2m_2 - m_1)} < l < \frac{2}{m_1^2} \quad \textit{and let} \quad b_{h_2} = l\ h_2.$$

then, as $\gamma \to \infty$, *we have*

$$\lim_{\gamma \to \infty} P_{\tau}^{m_i}\left(\delta_{T_{com}} \neq m_i \middle| \ \tau < \rho_{\nu}^{\lambda}\right) = 0 \quad \textit{for } i = 1, 2$$

- As $\gamma \to \infty$, the condition is equivalent to $\{\tau < \infty\}$.

The Behavior of CRP

The identification performance is based on the CRP in the second stage.

Lemma

For any parameter $\mu_2 \in (0, 2m_1)$, when we choose $b_{h_2}/h_2 = l$ to be a positive constant that satisfies

$$l < \frac{2}{\mu_2(2m_1 - \mu_2)},$$

there exists a positive constant $L_1 = L_1(m_1, \mu_2, l)$ such that

$$\lim_{h_2 \to \infty} P_\tau^{m_1}\left(S_{h_2} < b_{h_2} \mid \tau < \rho_{h_2}^{\mu_2}\right) e^{-L_1 h_2} = 0.$$

The Behavior of CRP (Cont.)

Lemma

For any parameter $\mu_2 \in (0, 2m_2)$, when we choose $b_{h_2}/h_2 = l$ to be a positive constant that satisfies

$$l > \frac{2}{\mu_2(2m_2 - \mu_2)},$$

there exists a positive constant $L_2 = L_2(m_2, \mu_2, l)$ such that

$$\lim_{h_2 \to \infty} P_\tau^{m_2}\left(S_{h_2} \geq b_{h_2} \mid \tau < \rho_{h_2}^{\mu_2}\right) e^{-L_2 h_2} = 0.$$

Numerical Illustrations

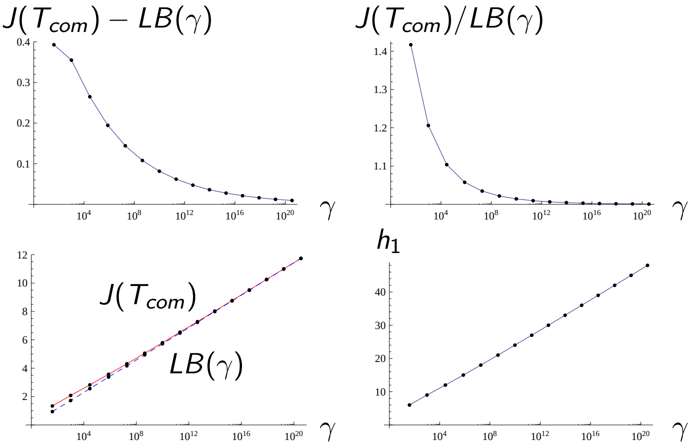

Figure: Performance of T_{com} in the case $m_1 = 2$, $m_2 = 5$, $p = 0.4$. All graphs are plotted against γ with a logarithmic x-axis.

Numerical Illustrations (Cont.)

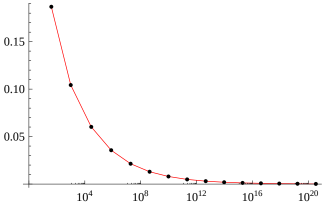

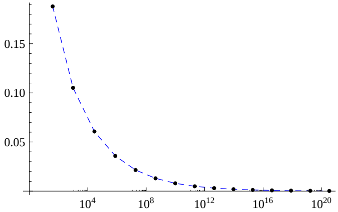

Left: $P_\tau^{m_1}\left(\delta_{T_{com}} \neq m_1 \middle| \tau < \rho_\nu^\lambda\right)$. Right: $P_\tau^{m_2}\left(\delta_{T_{com}} \neq m_2 \middle| \tau < \rho_\nu^\lambda\right)$.
All graphs are plotted against γ with a logarithmic x-axis.

Future Work

- It is interesting to consider the case where the post-change drift has a more complicated struction. For example, the post-change drift is a piece-wise continuous function $\mu(t)$.

- We may let the post-change drift has a distribution which is unknown but in a family of the given distributions. So we can consider Bayesian estimation with respect to the posterior distribution of the post-change drift.

- The applications that satisfy such models are interesting. The efficient and robust algorithms will be needed.

Reference

- M. Beibel, A note on Ritov's Bayes approach to the minimax property of the CUSUM procedure, *Ann. Statist.*, Vol. 24, No. 4, pp. 1804-1812, 1996.
- G. Lorden, Procedures for reacting to a change in distribution, *The Annals of Mathematical Statistics*, Vol. 42, No. 6, 1897-1908, 1971.
- G. V. Moustakides, Optimality of the CUSUM procedure in continuous time, *Annals of Statistics*, Vol. 32, No. 1, pp. 302-315, 2004.
- A.N. Shiryaev, Minimax optimality of the method of cumulative sums (CUSUM) in the case of continuous time, *Russian Mathem. Surv.* 51, 750-751, 1996.
- H. Zhang and O. Hadjiliadis, Drawdowns and the speed of a market crash, *Methodol. Comput. Appl. Probab.* Vol. 14, no. 3 pp. 739-752, 2012.

Thank You!