

UC Berkeley

Proceedings of the PowerUp Conference

Title

Outage Identification from Electricity Market Data: Quickest Change Detection Approach

Permalink

<https://escholarship.org/uc/item/2h39k5f1>

Journal

Proceedings of the PowerUp Conference, 1(1)

Authors

Hoseinpour, Milad

Shekhar, Shubhanshu

Dvorkin, Vladimir

Publication Date

2026-09-10

DOI

None/powerup_proceedings.69128

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NoDerivatives License, available at <https://creativecommons.org/licenses/by-nd/4.0/>

Peer reviewed

PowerUp 2026

BOULDER, COLORADO
SEPTEMBER 9 – 11, 2026

Outage Identification from Electricity Market Data: Quickest Change Detection Approach

Milad Hoseinpour, Shubhanshu Shekhar, Vladimir Dvorkin

University of Michigan

Suggested Citation

M. Hoseinpour, S. Shekhar, and V. Dvorkin, "Outage Identification from Electricity Market Data: Quickest Change Detection Approach," in *Proceedings of the PowerUp Conference*, 2026.

BibTeX Entry

```
@inproceedings{hoseinpour2026outageidentification,  
  author = {Hoseinpour, Milad and Shekhar, Shubhanshu and Dvorkin, Vladimir},  
  title = {Outage Identification from Electricity Market Data: Quickest Change  
    Detection Approach},  
  booktitle = {Proceedings of the PowerUp Conference},  
  year = {2026},  
}
```

Outage Identification from Electricity Market Data: Quickest Change Detection Approach

Milad Hoseinpour, Shubhanshu Shekhar, and Vladimir Dvorkin

Abstract—Power system outages expose market participants to significant financial risk unless promptly detected and hedged. We develop an outage identification method from public market signals grounded in the parametric quickest change detection (QCD) theory. Parametric QCD operates on stochastic data streams, distinguishing pre- and post-change regimes using the ratio of their respective probability density functions. To derive the density functions for normal and post-outage market signals, we exploit multi-parametric programming to decompose complex market signals into parametric random variables with a known density. These densities are then used to construct a QCD-based statistic that triggers an alarm as soon as the statistic exceeds an appropriate threshold. Numerical experiments on a stylized PJM testbed demonstrate rapid line outage identification from public streams of electricity demand and price data.

Index Terms—electricity markets, multi-parametric programming, outage identification, quickest change detection.

I. INTRODUCTION

Power grids are facing increasing threats from extreme weather and aging infrastructure. Slow anomaly detection incurs enormous costs: power outages alone cost the U.S. economy as much as \$150 billion annually [1]. The outages caused by the 2021 Winter Storm Uri led to billions in damages in Texas, with days-long blackouts, and surging prices up to \$9,000/MWh [2]. Early warning signs appeared in public price signals, yet delayed detection and hedging drove major utilities into bankruptcy [3]. While subsequent reforms mandated faster outage reporting by ERCOT [4], market participants (utilities, traders, etc.) require independent detection to act preemptively rather than await official reports.

System operators across the U.S. publish information on marginal prices, aggregated generator and load statistics, binding constraints, along with other operational and market data. This has given rise to modern data analytics companies, such as Grid Status [5] and Electricity Maps [6], which provide real-time access to market data streams, leading to the question: *How to leverage public operational and market data for independent and fast anomaly detection in power systems?*

The theory of Quickest Change Detection (QCD) offers the right formalism to answer this question [7]. QCD involves continuously monitoring a data stream, e.g., electricity prices, to raise an alarm as soon as an abrupt distributional change is detected. A Cumulative Sum (CuSum) is a powerful QCD statistic that accumulates evidence over time and compares it against a threshold, providing a theoretically optimal framework for detecting abrupt distributional changes [8]. These

properties have led to the development of numerous variants of CuSum informed by the network topology [9]–[12]. However, topology alone may not fully characterize the cause of distributional shifts. This is particularly relevant to electricity markets governed by optimization engines with specific objectives and constraints. For example, a transmission line outage in power systems triggers an economic re-dispatch optimization, whose solution must respect generation cost structure, along with power flow and generation limits in the post-outage topology. This raises the question: *How to integrate optimization objectives and constraints within the design of CuSum statistics for faster outage detection?*

Contributions: We propose a CuSum-based method for detection and identification of critical events (e.g., line or generator outage) in power systems using publicly available market data streams. Our method is based on the statistical relationship between observable market data (e.g., demand, prices, and aggregated dispatch statistics) and the underlying market-clearing structure, that is leveraged in a CuSum scheme accumulating evidence of structural changes in real-time. We contribute to anomaly detection in power systems by:

- 1) Developing probabilistic models of real-time market data streams in normal and post-outage states using multi-parametric programming [13]–[15]. The models decompose stochastic signals into simpler streams that exhibit distinct characteristics in normal and post-outage states, facilitating outage detection and identification.
- 2) Designing a CuSum-based statistic that leverages these models to identify structural changes, such as line and generator outages, from streams of market-clearing data. The proposed statistic tests whether the observed stream originates from the normal or post-outage model, accumulating statistical evidence when an outage occurs and remaining quiescent otherwise.
- 3) Identifying and exploring detection thresholds for the proposed statistic, enabling the control of false alarm rates and detection delays.

We apply the method to line outage identification using electricity demand and prices from the stylized PJM testbed.

Related Work: Outage detection is central to power system operations. A substantial body of work has focused on outage identification from streams of Phasor Measurement Unit (PMU) data, with early contributions relating observed streams to those induced by models under varying topologies [16], [17]. The combinatorial nature of line outage identification subsequently led to mixed-integer formulations, as in [18]. The computational burden was later alleviated by framing the problem as sparse recovery problem in [19] and by leveraging related methods from the signal processing domain [20]. More

The authors are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109, USA. E-mail: {miladh, shubhan, dvorkin}@umich.edu.

Supported by the Propelling Original Data Science (PODS) Program, Michigan Institute for Data and AI in Society (MIDAS). Project ID: 2521

recently, detection has been offloaded to machine learning models, e.g., based on convolutional neural networks [21] and decision trees [22]. Despite strong empirical validation, these methods lack awareness of the fundamental performance metrics in stochastic settings: detection delay and false alarm rate. The two are intrinsic to QCD schemes, which have been applied to line outage detection from PMU measurements in both quasi-steady-state [9] and transient [10] settings. Existing QCD schemes, however, rely on PMU measurements available only to system operators, rendering them unsuitable for *independent* analytics from streams of market-clearing data.

The use of public market data for inference was explored in [23], [24] using electricity prices for topology identification. In a similar spirit, inverse optimization techniques have been applied to recover market structure, including transmission constraints [25], supply curves [26], [27], and electricity demands [28]. While able to extract structural information, these methods assume *static* market structure. Hence, they do not provide definitive outage detection, as they fail to distinguish between pre- and post-outage samples in data streams. When applied to outage identification (e.g., inferring that a faulted line has zero transfer capacity), the inertia induced by pre-outage samples leads to significant detection delays.

By contrast, the proposed QCD statistic posits nominal and candidate post-outage market structures (e.g., one per line outage) and acts on public market signals to raise an alarm as soon as possible after a structural change occurs. We focus on parametric QCD [29], which requires explicit probabilistic densities of observed signals. This is not restrictive for market-clearing inputs such as electricity demand and renewable power generation, which are known to be well-approximated by parametric distributions [30]. However, the densities of dispatch statistics and electricity prices are not readily available, as market-clearing optimization is a nonlinear transformation of input parameters. To address this challenge, we exploit multi-parametric programming [13]–[15], which decomposes market-clearing optimization into critical regions, each admitting an affine mapping from inputs to market outcomes. The affine structure of critical regions permits closed-form derivation of output densities from known input distributions, thus enabling parametric QCD on public market signals.

Paper Organization: Next Section II details probabilistic models of market-clearing data streams, which are then leveraged in Section III to design the CuSum detection statistics. Section IV reports numerical results, and Section V concludes. All auxiliary information is relegated to the Appendix A.

II. PROBABILISTIC MODELS OF REAL-TIME MARKET

In this section, we derive probabilistic models of observed data streams from system operators' dashboards, including locational marginal prices (LMPs) and aggregated dispatch statistics. We begin with a perturbed market-clearing problem that is repeatedly solved under stochastic demand realizations, and then leverage multi-parametric programming theory [13] to obtain tractable models that allow us to characterize the probability density functions of observations.

A. Perturbed Market-Clearing Optimization Problem

We consider electricity market-clearing problem, which is solved every 5 minutes for real-time realizations of demand $1+\xi$, modeled by the mean vector 1 perturbed by random vector ξ . In response to new demand, system operators compute the cost-optimal dispatch $\mathbf{p}$ of conventional generators subject to generation and power flow constraints. When conventional generation is insufficient to meet demand, system operators resort to load shedding, denoted by ℓ . The perturbed market-clearing is posed as the following convex problem [31]:

$$\underset{\mathbf{p}, \ell}{\text{minimize}} \quad \frac{1}{2} \mathbf{p}^\top \mathbf{C} \mathbf{p} + \mathbf{c}^\top \mathbf{p} + \frac{1}{2} \ell^\top \mathbf{S} \ell + \mathbf{s}^\top \ell \quad (1a)$$

$$\text{subject to} \quad \mathbf{1}^\top (\mathbf{M}_p \mathbf{p} - \mathbf{M}_l (1 + \xi - \ell)) = 0, \quad (1b)$$

$$|\mathbf{F}(\mathbf{M}_p \mathbf{p} - \mathbf{M}_l (1 + \xi - \ell))| \leq \bar{\mathbf{f}}, \quad (1c)$$

$$\underline{\mathbf{p}} \leq \mathbf{p} \leq \bar{\mathbf{p}}, \quad (1d)$$

$$\mathbf{0} \leq \ell \leq 1 + \xi, \quad (1e)$$

where objective function (1a) minimizes quadratic costs of generation, with parameters $\mathbf{C}$ and $\mathbf{c}$, and load shedding, parameterized by $\mathbf{S}$ and $\mathbf{s}$. Constraint (1b) enforces power balance and (1c) introduces power flow limits using a power transfer distribution factor (PTDF) matrix $\mathbf{F}$. The incidence matrices $\mathbf{M}_p$ and $\mathbf{M}_l$ relate generators and demands to their corresponding buses, respectively. Generator limits are enforced by (1d), while constraints (1e) bound load shedding between zero and real-time demand realizations.

For further derivation, it is convenient to rewrite the perturbed problem (1) in a compact quadratic programming (QP) form. Let $\mathbf{x} := [\mathbf{p}^\top, \ell^\top]^\top$ denote the vector of decision variables. Then, the problem can be written as

$$\underset{\mathbf{x}}{\text{minimize}} \quad \frac{1}{2} \mathbf{x}^\top \mathbf{Q} \mathbf{x} + \mathbf{q}^\top \mathbf{x} \quad (2a)$$

$$\text{subject to} \quad \mathbf{A} \mathbf{x} \leq \mathbf{B} \xi + \mathbf{b} \quad : \mu, \quad (2b)$$

where vector μ stacks the duals, from which the LMPs $\lambda = \Lambda \mu$ are obtained using auxiliary matrix Λ . Parameters $\mathbf{Q}$, $\mathbf{q}$, $\mathbf{A}$, $\mathbf{B}$, $\mathbf{b}$, and Λ are defined in Appendix A.

B. Mapping Load Perturbations to Market Outcomes

The uncertainty of demand ξ is often described by known probability distributions, e.g., Normal, Log-Normal, or Weibull [30]. However, characterizing the distributions of aggregated dispatch $\mathbf{1}^\top \mathbf{x}$ and prices λ via density of ξ is challenging since these quantities depend nonlinearly on ξ through optimization (2). We overcome this challenge using multi-parametric programming [13], providing an adaptive affine mapping between load perturbations and market outcomes.

Consider some realization of the demand perturbation $\hat{\xi}$, and let $\bar{\mathbf{A}}, \bar{\mathbf{B}}, \bar{\mathbf{b}}$ contain those entries of $\mathbf{A}, \mathbf{B}, \mathbf{b}$ that correspond to the active constraints at $\hat{\xi}$, and let $\tilde{\mathbf{A}}, \tilde{\mathbf{B}}, \tilde{\mathbf{b}}$ denote the remaining entries corresponding to inactive constraints. The stationarity and primal feasibility conditions of (2) yield

$$\mathbf{Q} \mathbf{x}^* + \mathbf{q} + \tilde{\mathbf{A}}^\top \tilde{\mu}^* = \mathbf{0}, \quad (3a)$$

$$\tilde{\mathbf{A}} \mathbf{x}^* - \tilde{\mathbf{B}} \hat{\xi} - \tilde{\mathbf{b}} = \mathbf{0}, \quad (3b)$$

at optimality. Solving (3a) for $\mathbf{x}^*$ and substituting into (3b), we obtain an affine relationship between the dual solution and the demand perturbation:

$$\tilde{\mu}^*(\hat{\xi}) = \mathbf{D}\hat{\xi} + \mathbf{d}, \quad (4a)$$

with auxiliary parameters $\mathbf{D}$ and $\mathbf{d}$ defined as

$$\mathbf{D} := -\left(\tilde{\mathbf{A}}\mathbf{Q}^{-1}\tilde{\mathbf{A}}^\top\right)^{-1}\tilde{\mathbf{B}}, \quad (4b)$$

$$\mathbf{d} := -\left(\tilde{\mathbf{A}}\mathbf{Q}^{-1}\tilde{\mathbf{A}}^\top\right)^{-1}\left(\tilde{\mathbf{b}} + \tilde{\mathbf{A}}\mathbf{Q}^{-1}\mathbf{q}\right). \quad (4c)$$

Substituting (4a) into (3a) and solving for $\mathbf{x}^*$, the corresponding primal solution also becomes affine in the perturbation

$$\mathbf{x}^*(\hat{\xi}) = \mathbf{P}\hat{\xi} + \mathbf{p}, \quad (5a)$$

with auxiliary parameters $\mathbf{P}$ and $\mathbf{p}$ defined as

$$\mathbf{P} := -\mathbf{Q}^{-1}\tilde{\mathbf{A}}^\top\mathbf{D}, \quad (5b)$$

$$\mathbf{p} := -\mathbf{Q}^{-1}\mathbf{q} - \mathbf{Q}^{-1}\tilde{\mathbf{A}}^\top\mathbf{d}. \quad (5c)$$

Taken together, for some realization $\hat{\xi}$ of the random perturbation, the primal-dual solution depends affinely on $\hat{\xi}$:

$$\begin{bmatrix} \mathbf{x}^* \\ \tilde{\mu}^* \end{bmatrix}(\hat{\xi}) = \begin{bmatrix} \mathbf{P} \\ \mathbf{D} \end{bmatrix} \hat{\xi} + \begin{bmatrix} \mathbf{p} \\ \mathbf{d} \end{bmatrix}. \quad (6)$$

An important result from [13] is that the affine mapping (6) is valid not only for a particular realization $\hat{\xi}$, but for all ξ from the polyhedral neighborhood of $\hat{\xi}$, referred to as *critical region*. The critical region $\mathcal{R}$ of $\hat{\xi}$ is defined by the primal feasibility conditions on inactive constraints and the dual feasibility conditions on active constraints:

$$\mathcal{R} = \{\xi \mid (\overline{\mathbf{A}}\mathbf{P} - \overline{\mathbf{B}})\xi < \overline{\mathbf{b}} - \overline{\mathbf{A}}\mathbf{p}, \mathbf{D}\xi + \mathbf{d} \geq \mathbf{0}\}. \quad (7)$$

A well-posed QP (2) has a finite number of critical regions [13], each having its own affine mapping of load perturbations to optimization results. Hence, observing perturbation $\hat{\xi}$, we pick the corresponding critical region and map the density of ξ to the density of market outcomes via (6) from that region.

The set $\mathcal{CR}$ of critical regions is computed using Alg. 1. It takes market-clearing parameters and the set Ξ of perturbation samples as inputs. For each sample, it solves the QP problem (2) and computes the associated region, unless the sample belongs to one of the regions found at previous iterations.

Algorithm 1 Computing critical regions of problem (2)

Input: Market-clearing parameters $\mathbf{Q}$, $\mathbf{q}$, $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{b}$, and load perturbation set Ξ

Output: Set of critical regions $\mathcal{CR}$

- 1: Initialize $\mathcal{CR} \leftarrow \emptyset$
 - 2: **for** each load perturbation $\hat{\xi} \in \Xi$ **do**
 - 3: **if** $\hat{\xi}$ lies in an existing region $\mathcal{R}_i \in \mathcal{CR}$ **then**
 - 4: **skip** ▷ Already covered by an existing region
 - 5: Solve the QP problem (2) at $\xi = \hat{\xi}$
 - 6: Identify active and inactive constraints
 - 7: Compute the critical region $\mathcal{R}_i$ via (7) and add to $\mathcal{CR}$
 - 8: **return** $\mathcal{CR}$
-

C. Probabilistic Models of Electricity Market Data Streams

The multi-parametric decomposition of problem (2) allows us to derive densities of market-clearing outcomes reported in system operators' dashboards through adaptive transformations of the demand perturbation density. When perturbations belong to critical region $\mathcal{R}_i$, market outcomes follow region-specific affine relationship with this density [32, Theorem 2.1.8]. For example, LMPs can be expressed using (6) as

$$\lambda_i(\xi) = \tilde{\mathbf{A}}_i \tilde{\mu}^*(\xi) = \tilde{\mathbf{A}}_i \mathbf{D}_i \xi + \tilde{\mathbf{A}}_i \mathbf{d}_i, \quad (8)$$

where $\mathbf{D}_i$ and $\mathbf{d}_i$ are parameters of the critical region $\mathcal{R}_i$, and matrix $\tilde{\mathbf{A}}_i$ discards the columns of the full matrix $\tilde{\mathbf{A}}$ corresponding to inactive power flow constraints in region i . When a critical region changes over time, expression (8) adapts parameters $\mathbf{D}_i$, $\mathbf{d}_i$, and $\tilde{\mathbf{A}}_i$ to a new critical region.

Before deriving the exact density, we note that the classical CuSum method assumes i.i.d. observations [33], an assumption that rarely holds in market data streams due to demand seasonality (e.g., daily morning and evening peaks). To satisfy CuSum assumptions, we act on the increments of observed signals, thereby producing approximately stationary observations with tractable distributions. The increment of LMPs are

$$\Delta\lambda_t = \lambda_i(\xi_t) - \lambda_i(\xi_{t-1}), \quad (9)$$

computed at two subsequent time steps.

Suppose that the increments of demand perturbations follow a Wiener process, i.e., $\xi_t - \xi_{t-1} \sim \mathcal{N}(\mathbf{0}, \Sigma)$, with zero-mean and known covariance matrix Σ estimated from operational history. Then, the mean and covariance of LMP increments $\Delta\lambda_t$ from the critical region $\mathcal{R}_i$ are given by

$$\mathbb{E}[\Delta\lambda_t] = \tilde{\mathbf{A}}_i \mathbf{D}_i \mathbb{E}[\xi_t - \xi_{t-1}] = \mathbf{0}, \quad (10)$$

$$\text{Cov}[\Delta\lambda_t] = \text{Cov}\left[\tilde{\mathbf{A}}_i \mathbf{D}_i (\xi_t - \xi_{t-1})\right] = \tilde{\mathbf{A}}_i \mathbf{D}_i \Sigma (\tilde{\mathbf{A}}_i \mathbf{D}_i)^\top.$$

Accordingly, within each critical region $\mathcal{R}_i$, the incremental observation $\Delta\lambda_t$ follows a zero-mean multivariate Gaussian distribution with probability density function proportional to

$$f_i(\Delta\lambda_t) \propto \exp\left(-\frac{1}{2}\Delta\lambda_t^\top (\tilde{\mathbf{A}}_i \mathbf{D}_i \Sigma (\tilde{\mathbf{A}}_i \mathbf{D}_i)^\top)^{-1} \Delta\lambda_t\right). \quad (11)$$

With a finite set $\mathcal{CR} = \{\mathcal{R}_i\}_{i=1}^I$ of critical regions, the probabilistic model of real-time LMPs is given by a finite collection of probability density functions (11).

In what follows, we will narrow down the scope of observed market signals to LMPs alone. This is without much loss of generality, as we can derive the densities for other published statistics (total generation by fuel type, emissions, power flows in tie-lines, etc.) in a similar way, as long as they admit an affine transformation of the primal-dual solution in (6). For example, the density for increments of the total generator dispatch $g_i(\xi) = \mathbf{1}^\top \mathbf{x}_i^*(\xi)$ takes the following form

$$f_i(\Delta g_t) \propto \exp\left(-\frac{1}{2}\Delta g_t (\mathbf{1}^\top \mathbf{P}_i \Sigma (\mathbf{1}^\top \mathbf{P}_i)^\top)^{-1} \Delta g_t\right), \quad (12)$$

which is analogous to (11), except that it adapts the primal parameters $\mathbf{P}_i$ of critical region $\mathcal{R}_i$ to demand realizations.

III. CUSUM-BASED OUTAGE IDENTIFICATION FROM STREAMS OF MARKET-CLEARING DATA

Consider a stream $\Delta\lambda_1, \Delta\lambda_2, \Delta\lambda_3, \dots$ of LMP increments as a stochastic process whose statistical properties are determined by the underlying market-clearing structure. Under normal operating conditions, the system evolves according to a nominal structure, and the distribution of $\Delta\lambda_t$ remains stationary over time. A structural change, such as a line or generator outage, alters the feasible set of the market-clearing problem, leading to an abrupt shift in LMP distribution. Formally, let $\mathcal{A}$ denote the set of possible structural changes in the market-clearing problem. The family of distributions $\{f^a\}_{a \in \mathcal{A} \cup \{0\}}$ is defined such that $a = 0$ corresponds to the nominal structure, and $a \in \mathcal{A}$ corresponds to a particular structural change. We assume the existence of an unknown change (outage) point $T \in \mathbb{N} \cup \{\infty\}$ such that

$$\Delta\lambda_t \sim \begin{cases} f^0, & t \leq T, \\ f^a, & t > T, \end{cases} \quad a \in \mathcal{A}. \quad (13)$$

Our goal is twofold: (1) promptly detect the structural change from the streams of LMP observation, and (2) identify the specific outage underlying the change. That is, we need to distinguish between the null hypothesis $\mathbb{H}_0 : T = \infty$ (no change), and a set of alternative $\mathbb{H}_1 : T < \infty$ hypotheses using streams of market observations. Toward this goal, in this section we develop CuSum statistics where the log-likelihood ratios are computed using probabilistic market models (11) induced by time-varying critical regions.

A. CuSum Design using Probabilistic Market Models

Consider two variants of market-clearing problem (2), which are a nominal problem modeling normal operation:

$$\underset{\mathbf{x}}{\text{minimize}} \quad \frac{1}{2} \mathbf{x}^\top \mathbf{Q}_0 \mathbf{x} + \mathbf{q}_0^\top \mathbf{x} \quad (14a)$$

$$\text{subject to} \quad \mathbf{A}_0 \mathbf{x} \leq \mathbf{B}_0 \boldsymbol{\xi} + \mathbf{b}_0 \quad : \mu_0, \quad (14b)$$

and an alternative problem corresponding to an outage:

$$\underset{\mathbf{x}}{\text{minimize}} \quad \frac{1}{2} \mathbf{x}^\top \mathbf{Q}_a \mathbf{x} + \mathbf{q}_a^\top \mathbf{x} \quad (15a)$$

$$\text{subject to} \quad \mathbf{A}_a \mathbf{x} \leq \mathbf{B}_a \boldsymbol{\xi} + \mathbf{b}_a \quad : \mu_a. \quad (15b)$$

A line outage is modeled by excluding the entries associated with the faulted line from $\mathbf{A}_0$, $\mathbf{B}_0$, and $\mathbf{b}_0$, yielding $\mathbf{A}_a$, $\mathbf{B}_a$, and $\mathbf{b}_a$. A generation outage is modeled by removing the relevant entry in $\mathbf{x}$ along with related problem data. The nominal and alternative problems have distinct critical regions and thus distinct LMP densities (11), facilitating outage detection.

We propose a two-stage CuSum design. At the first (offline) stage, before observing market signals, we compute the critical regions for the nominal and alternative market structures, denoted by $\mathcal{CR}_0 = \{\mathcal{R}_i^0\}_{i=1}^I$ and $\mathcal{CR}_a = \{\mathcal{R}_j^a\}_{j=1}^J$, respectively, using Alg. 1. At the second (online) stage, at each time step t we observe LMP increments $\Delta\lambda_t$ and the corresponding demand perturbation $\boldsymbol{\xi}_t$. Based on realizations of $\boldsymbol{\xi}_t$, we identify the critical region $\mathcal{R}_i^0 \in \mathcal{CR}_0$ for the nominal

problem and $\mathcal{R}_j^a \in \mathcal{CR}_a$ for all alternative problems. The log-likelihood ratio for testing whether increments come from the nominal or alternative problem is computed as:

$$\log \left(\frac{f_j^a(\Delta\lambda_t)}{f_i^0(\Delta\lambda_t)} \right) = \log \left(\frac{\exp \left(-\frac{1}{2} \Delta\lambda_t^\top \left(\tilde{\Lambda}_j^a \mathbf{D}_j^a \boldsymbol{\Sigma} (\tilde{\Lambda}_j^a \mathbf{D}_j^a)^\top \right)^{-1} \Delta\lambda_t \right)}{\exp \left(-\frac{1}{2} \Delta\lambda_t^\top \left(\tilde{\Lambda}_i^0 \mathbf{D}_i^0 \boldsymbol{\Sigma} (\tilde{\Lambda}_i^0 \mathbf{D}_i^0)^\top \right)^{-1} \Delta\lambda_t \right)} \right). \quad (16)$$

This statistic has the following intuition. If the alternative hypothesis is false (no outage), the expected value of (16) is in the negative domain. If the alternative hypothesis is true, the expected value of (16) is positive instead. This intuition is captured in the design of the following CuSum statistics:

$$w_t^a = \max \left\{ 0, w_{t-1}^a + \log \left(\frac{f_j^a(\Delta\lambda_t)}{f_i^0(\Delta\lambda_t)} \right) \right\}, \quad \forall a \in \mathcal{A}, \quad (17)$$

computed for every hypothesized outage. The CuSum value w_t^a evolves over time as follows. For every alternative hypothesis, the log-likelihood ratio is more likely to be negative when the hypothesis is false. Applying the max operator clips the statistic at zero, preventing accumulation of large negative values that would delay detection after a true outage occurs. In contrast, if the alternative hypothesis is true, the log-likelihood ratios are consistently positive, leading to sustained linear growth of w_t^a , providing evidence that the outage has occurred.

B. Adapting CuSum for Bounded Demand Perturbations

The log-likelihood ratio (17) assumes unbounded perturbation increments $\boldsymbol{\xi}_t - \boldsymbol{\xi}_{t-1}$. In practice, however, they are bounded by load limits and tend to temporally remain constant during peak-hour operation, thus violating the underlying Gaussian assumptions. Here, we propose an adaptive CuSum, which computes the likelihood only over the components of $\boldsymbol{\xi}_t$ that are strictly within the interior of its feasible domain.

Let $\mathbf{S}_t$ denote the selection matrix that retains only those components of perturbation $\boldsymbol{\xi}_t$ that remain strictly within the feasible domain at two subsequent times steps. The adapted perturbation increment and its covariance matrix are

$$\hat{\boldsymbol{\xi}}_t - \hat{\boldsymbol{\xi}}_{t-1} = \mathbf{S}_t (\boldsymbol{\xi}_t - \boldsymbol{\xi}_{t-1}) \implies \hat{\boldsymbol{\Sigma}} = \mathbf{S}_t \boldsymbol{\Sigma} \mathbf{S}_t^\top. \quad (18a)$$

The adaptive covariance of LMP increments becomes

$$\text{Cov}[\Delta\lambda_t] = (\mathbf{S}_t \tilde{\Lambda}_j^a \mathbf{D}_j^a) \hat{\boldsymbol{\Sigma}} (\mathbf{S}_t \tilde{\Lambda}_j^a \mathbf{D}_j^a)^\top, \quad \forall a \in \mathcal{A} \cup \{0\}.$$

Then, accommodating bounded perturbation increments, the log-likelihood ratio in (16) is computed using the adaptive covariance, while the CuSum recursion (17) remains the same.

C. Threshold For Outage Detection and Identification

Running CuSum statistics in parallel for all candidate events, an outage is declared at time τ as soon as one of the statistics $\{w_t^a\}_{a \in \mathcal{A}}$ exceeds a prescribed threshold $\eta > 0$:

$$\tau = \inf \{t \geq 1 : \max_{a \in \mathcal{A}} w_t^a \geq \eta\}. \quad (19)$$

The identity of the detected outage is determined by the index a of the statistic that exceeds the threshold first. As with

any sequential detection procedure acting on stochastic data streams, the choice of the threshold η governs a trade-off between detection delay and the rate of false alarms.

Rather than explicitly specifying the probability of false alarms over an infinite horizon, we adopt an average run length (ARL)-based approach for threshold selection. Specifically, under nominal operation ($\mathbb{H}_0$), the threshold is chosen to achieve a desired in-control ARL:

$$\text{ARL}_0 = \mathbb{E}_{\mathbb{H}_0}[\tau], \quad (20)$$

representing the expected number of samples until a false alarm is declared [34]. This metric is estimated by evaluating the CuSum statistics on LMP streams generated from the nominal market-clearing optimization, and the threshold η is adjusted so that the resulting ARL_0 matches a prescribed target, e.g., one false alarm per week of operation.

Once the threshold is fixed, the performance of the CuSum detector in post-outage operation ($\mathbb{H}_1$) is characterized in terms of the expected detection delay:

$$\mathbb{E}_{\mathbb{H}_1}[\tau - T \mid \tau \geq T], \quad (21)$$

i.e., the expected number of post-outage samples until detection is declared. Consistent with classical sequential analysis, the detection delay scales inversely with the Kullback–Leibler divergence $D(f^a \| f^0)$ between post-outage and nominal LMP distributions [34]. Hence, outages inducing smaller distributional shifts are expected to have longer detection delays.

D. Outage Identification Algorithm

Algorithm 2 summarizes the outage identification procedure. It takes as input the critical regions from the nominal and hypothesized post-outage optimization problems, the demand perturbation increment covariance, and a detection threshold. It continuously outputs $a^* = 0$ unless the outage is detected, reporting outage index $a^* \in \mathcal{A}$. At every time step, it finds the critical regions of the nominal and each alternative optimization, then updates the parallel CuSum statistics. When one of the hypothesized statistics exceeds the threshold, the algorithm declares an outage associated with that statistic.

Algorithm 2 CuSum-Based Outage Identification

Input: Sets $\mathcal{CR}_0, \mathcal{CR}_{a \in \mathcal{A}}$, covariance Σ , threshold η

Output: Identification result $a^* \in \mathcal{A} \cup \{0\}$

- 1: Initialize $w_0^a \leftarrow 0$ for all $a \in \mathcal{A}$, $a^* \leftarrow 0$
- 2: **for** $t = 1, 2, \dots$ **do**
- 3: Observe perturbation ξ_t and LMP increment $\Delta \lambda_t$
- 4: Find the critical region $\mathcal{R}_i^0 \in \mathcal{CR}_0$ (nominal)
- 5: **for all** $a \in \mathcal{A}$ **do**
- 6: Find the critical region $\mathcal{R}_j^a \in \mathcal{CR}_a$ (alternative)
- 7: Update the CuSum statistic

$$w_t^a \leftarrow \max \left\{ 0, w_{t-1}^a + \log \left(\frac{f_j^a(\Delta \lambda_t)}{f_i^0(\Delta \lambda_t)} \right) \right\}$$

- 8: **if** $\max_{a \in \mathcal{A}} w_t^a \geq \eta$ **then**
 - 9: **return** $a^* \leftarrow \arg \max_{a \in \mathcal{A}} w_t^a$
-

IV. NUMERICAL EXPERIMENTS

The CuSum statistic is tested for line outage identification in the PJM network [35] (Fig. 1). We show how it distinguishes demand variation from outages during price spikes, then evaluate statistical performance across multiple scenarios.

A. Settings

The stochastic demand perturbation is applied to loads 1 and 2 as a Wiener process with a zero drift and a standard deviation of 8 MW, illustrated in Fig. 2 (left), over a horizon of 1,000 samples. Each sample corresponds to a market-clearing outcome computed by solving problem (2) every 5 minutes. We model a fault on the line connecting buses 1 and 5, which is isolated by the system operator between samples 499 and 500. The outage alters the multi-parametric geometry of the market-clearing problem, as depicted in Fig. 1 for the nominal (middle) and the post-outage (right) market operation. They respectively consist of 18 and 9 critical regions, which are notably different. When the same stochastic process (blue trajectories in Fig. 1) develops over time, it activates distinct critical regions of the two problems, yielding an abrupt shift in the LMP streams after $T = 500$, as shown in Fig. 2 (middle). The proposed CuSum statistic is applied to distinguish whether these streams originate from critical regions of the nominal or post-outage problem. Detailed settings, data, and code are provided in the accompanying online repository [36].

B. Distinguishing Line Outages from Demand Variations

We first illustrate identification for a specific demand trajectory in Fig. 2 (left), where the outage at $T = 500$ coincides with a sudden demand growth. Visual inspection of demand and price trajectories might mistakenly attribute the price spike to demand growth rather than the outage. We show that the proposed identification method resolves this ambiguity.

We select the detection threshold $\eta = 50$, whose rationale is discussed in the next subsection. Figure 2 (right) depicts six CuSum statistics, each hypothesizing the outage of a specific line. Observe that, before the outage, statistics vary but hardly sustain consistent evidence accumulation. This is because the expected log-likelihood ratios are negative for all candidate hypotheses before the outage, keeping the statistics closer to zero. Line outage at $T = 500$, however, triggers three out of the six statistics to suddenly grow; this is because the outage of the corresponding three lines leads to statistically similar impacts on LMPs. The statistic in green, corresponding to the true line outage, grows most rapidly and crosses the threshold first with a detection delay of only 5 samples in this particular scenario, i.e., 25 minutes in the real-time market timescale. This result confirms that the price spike cannot be attributed to demand variation alone, and is in fact consistent with a transmission line outage between buses 1 and 5.

C. Performance Analysis of the CuSum Statistic

Table I provides the trade-offs between performance metrics for various threshold choices η . The threshold is selected offline using a horizon of $t_{\max} = 5,000$ samples from nominal




Fig. 1. PJM system and the corresponding critical regions of the real-time market-clearing problem in the nominal and post-outage operation.




Fig. 2. Demand (left) and LMP (middle) data streams and the dynamics of the six CuSum statistics (right). The outage occurs at $T = 500$ (red dashed line).

TABLE I
CuSum PERFORMANCE ACROSS 1,000 DEMAND TRAJECTORIES FOR VARYING DETECTION THRESHOLDS η .

Offline stage (before operation)			Online stage (during operation)				
η	ARL $\mathbb{E}_{\mathbb{H}_0}[\tau]$	Probability of false alarm [%] $\mathbb{P}_{\mathbb{H}_0}[a^* \in \mathcal{A} \mid \tau < t_{\max}]$	Detection delay $\tau - T$ if $\tau \geq T$		Probability of [%]		
			average	median	False detection $\mathbb{P}_{\mathbb{H}_1}[a^* \in \mathcal{A} \mid \tau < T]$	Successful detection $\mathbb{P}_{\mathbb{H}_1}[a^* \in \mathcal{A} \mid \tau \geq T]$	Successful identification $\mathbb{P}_{\mathbb{H}_1}[a^* = a^{\text{true}} \mid \tau \geq T]$
10	871.2	90.5	56.1	16	46.6	53.4	54.5
20	1793.0	59.1	76.6	33	12.2	87.7	65.6
30	2319.2	34.6	114.8	49	2.4	97.4	69.4
40	2732.9	15.3	146.1	62	0.2	99.4	73.1
50	3008.4	7.6	172.7	78	0.0	99.4	76.4
60	3045.8	3.0	200.2	93	0.0	99.4	78.7

operation. Larger thresholds increase the ARL and hence reduce false alarm probability. For instance, selecting $\eta = 50$ triggers a false alarm once every $\approx 3,000$ samples, or once per 250 hours in real-time markets with 5-minute resolution.

The online performance is evaluated for the same line outage event. Larger ARLs inevitably increase detection delays, with median values ranging from 16 samples (1.3h) to 93 samples (7.7h). The probability of false detection is minimized for $\eta \geq 40$, achieving near-certain detection. However, detecting a distributional shift is simpler than identifying which outage caused the shift, as shown by the probability of successful identification. This is because multiple transmission line outages produce similar LMP distributions, yet the true outage can still be identified with probability up to 78.7%, but at the cost of increased detection delay. These results show that η can be naturally calibrated via cross-validation on known (historical) outages to balance ARL with detection delay.

V. CONCLUSION

We developed probabilistic models of real-time electricity markets and showed how they enable QCD schemes to detect and identify critical power system events from market data streams. A careful calibration of the detection threshold led to detection rates over 97%, while the probability of false alarm approached zero. Although identification accuracy is lower, the true outage identification rate still remains close to 79%. We believe that incorporating more data streams (generation by type, tie-line power flows) would improve the performance.

Our approach assumes known market parameters (e.g., generation costs), enabling precise estimation of critical regions and parametrization of probability density functions for market data streams. Future work will address parameter uncertainty and explore alternative, nonparametric QCD methods. We expect these methods may have longer detection delays, which we plan to estimate and optimize.

VI. AI USAGE DISCLOSURE

The authors used AI tools to assist with narrative polishing, including spell-checking and streamlining arguments, as well as for coding and debugging. None of the narrative was directly produced by AI; it was used solely in an advisory role and did not replace the authors. All models and ideas are the sole intellectual property of the authors, and no AI model contributed to their development.

APPENDIX

A. QP Parameters

The cost coefficients in (2a) are given by

$$\mathbf{Q} = \begin{bmatrix} \mathbf{C} & \mathbf{0} \\ \mathbf{0} & \mathbf{S} \end{bmatrix}, \quad \mathbf{q} = \begin{bmatrix} \mathbf{c} \\ \mathbf{s} \end{bmatrix}.$$

The constraint parameters in (2b) are

$$\mathbf{A} = \begin{bmatrix} -\mathbf{1}^\top \mathbf{M}_p & -\mathbf{1}^\top \mathbf{M}_l \\ \mathbf{F} \mathbf{M}_p & \mathbf{F} \mathbf{M}_l \\ -\mathbf{F} \mathbf{M}_p & -\mathbf{F} \mathbf{M}_l \\ \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \\ -\mathbf{I} & \mathbf{0} \\ \mathbf{0} & -\mathbf{I} \end{bmatrix}, \mathbf{B} = \begin{bmatrix} -\mathbf{1}^\top \mathbf{M}_l \\ \mathbf{F} \mathbf{M}_l \\ -\mathbf{F} \mathbf{M}_l \\ \mathbf{M}_l^\top \mathbf{M}_l \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix}, \mathbf{b} = \begin{bmatrix} \mathbf{1}^\top \mathbf{M}_l \mathbf{l} \\ \mathbf{F} \mathbf{M}_l \mathbf{l} + \bar{\mathbf{f}} \\ -\mathbf{F} \mathbf{M}_l \mathbf{l} + \bar{\mathbf{f}} \\ \bar{\mathbf{p}} \\ \mathbf{l} \\ -\bar{\mathbf{p}} \\ \mathbf{0} \end{bmatrix}.$$

The auxiliary matrix to compute LMPs from the dual variables of (2b) is defined as

$$\mathbf{\Lambda} = [\mathbf{1} - \mathbf{F}^\top \mathbf{F}^\top \mathbf{0}],$$

yielding the standard LMP definition

$$\boldsymbol{\lambda} = \mathbf{\Lambda} \boldsymbol{\mu} = \mu_0 \mathbf{1} - \mathbf{F}^\top (\bar{\boldsymbol{\mu}} - \boldsymbol{\mu}),$$

where μ_0 is the dual variable associated with the power balance constraint (1b), and $\bar{\boldsymbol{\mu}}$ and $\boldsymbol{\mu}$ are the dual variables of the upper and lower power flow limits in (1c), respectively.

REFERENCES

- [1] The Pew Charitable Trusts, "America's electric grid: Growing cleaner, cheaper and stronger," October 2015. <https://tinyurl.com/6ssea7ss>; Accessed: 2025-05-09.
- [2] FERC, NERC, and Regional Entity, "The February 2021 cold weather outages in Texas and the South Central United States," *FERC and NERC and Regional Entity Staff, Tech. Rep.*, 2021.
- [3] F. Li and C. Heilbrun, "How extreme weather events have bankrupted utility players and changed the electric grid," 2024. FTI Consulting.
- [4] ERCOT, "Discussion on FERC, NERC, and NERC regional entity staff final report," Dec 2021. Available: https://www.ercot.com/files/docs/2021/12/06/14_Revised_12_6_2021_ERCOT_Discussion_on_FERC,_NERC_and_NERC_Regional_Entity_Staff_Final_Report.pdf.
- [5] GridStatus.io, "Gridstatus: Real-time power grid status and market information." <https://www.gridstatus.io/>.
- [6] Electricity Maps, "Live electricity map (15-minute intervals)." https://app.electricitymaps.com/map/live/fifteen_minutes.
- [7] V. V. Veeravalli and T. Banerjee, "Quickest change detection," in *Academic press library in signal processing*, vol. 3, pp. 209–255, Elsevier, 2014.
- [8] E. S. Page, "Continuous inspection schemes," *Biometrika*, vol. 41, no. 1/2, pp. 100–115, 1954.
- [9] Y. C. Chen, T. Banerjee, A. D. Dominguez-Garcia, and V. V. Veeravalli, "Quickest line outage detection and identification," *IEEE Transactions on Power Systems*, vol. 31, no. 1, pp. 749–758, 2015.
- [10] G. Rovatsos, X. Jiang, A. D. Dominguez-Garcia, and V. V. Veeravalli, "Statistical power system line outage detection under transient dynamics," *IEEE Transactions on Signal Processing*, vol. 65, no. 11, pp. 2787–2797, 2017.
- [11] F. Nizam, S. Chaki, S. Al Mamun, and M. S. Kaiser, "Attack detection and prevention in the cyber physical system," in *2016 international conference on computer communication and informatics (ICCCI)*, pp. 1–6, IEEE, 2016.
- [12] V. Raghavan and V. V. Veeravalli, "Quickest change detection of a markov process across a sensor array," *IEEE transactions on information theory*, vol. 56, no. 4, pp. 1961–1981, 2010.
- [13] P. Tøndel, T. A. Johansen, and A. Bemporad, "An algorithm for multi-parametric quadratic programming and explicit MPC solutions," *Automatica*, vol. 39, no. 3, pp. 489–497, 2003.
- [14] S. Taheri, M. Jalali, V. Kekatos, and L. Tong, "Fast probabilistic hosting capacity analysis for active distribution systems," *IEEE Transactions on Smart Grid*, vol. 12, no. 3, pp. 2000–2012, 2020.
- [15] S. Taheri, V. Kekatos, and S. Veeramachaneni, "Strategic generation investment in energy markets: A multiparametric programming approach," *IEEE Transactions on Power Systems*, vol. 37, no. 4, pp. 2590–2600, 2021.
- [16] J. E. Tate and T. J. Overbye, "Line outage detection using phasor angle measurements," *IEEE Transactions on Power Systems*, vol. 23, no. 4, pp. 1644–1652, 2008.
- [17] J. E. Tate and T. J. Overbye, "Double line outage detection using phasor angle measurements," in *2009 IEEE Power & Energy Society General Meeting*, pp. 1–5, IEEE, 2009.
- [18] R. Emami and A. Abur, "External system line outage identification using phasor measurement units," *IEEE Transactions on Power Systems*, vol. 28, no. 2, pp. 1035–1040, 2012.
- [19] H. Zhu and G. B. Giannakis, "Sparse overcomplete representations for efficient identification of power line outages," *IEEE Transactions on Power Systems*, vol. 27, no. 4, pp. 2215–2224, 2012.
- [20] G. B. Giannakis, V. Kekatos, N. Gatsis, S.-J. Kim, H. Zhu, and B. F. Wollenberg, "Monitoring and optimization for power grids: A signal processing perspective," *IEEE Signal Processing Magazine*, vol. 30, no. 5, pp. 107–128, 2013.
- [21] W. Li, D. Deka, M. Chertkov, and M. Wang, "Real-time faulted line localization and pmu placement in power systems through convolutional neural networks," *IEEE Transactions on Power Systems*, vol. 34, no. 6, pp. 4640–4651, 2019.
- [22] C. Hannon, D. Deka, D. Jin, M. Vuffray, and A. Y. Likhov, "Real-time anomaly detection and classification in streaming PMU data," in *2021 IEEE Madrid PowerTech*, pp. 1–6, IEEE, 2021.
- [23] V. Kekatos, G. B. Giannakis, and R. Baldick, "Grid topology identification using electricity prices," in *2014 IEEE PES general meeting—conference & exposition*, pp. 1–5, IEEE, 2014.
- [24] V. Kekatos, G. B. Giannakis, and R. Baldick, "Online energy price matrix factorization for power grid topology tracking," *IEEE Transactions on Smart Grid*, vol. 7, no. 3, pp. 1239–1248, 2015.
- [25] J. R. Birge, A. Hortaçsu, and J. M. Pavlin, "Inverse optimization for the recovery of market structure from market outcomes: An application to the miso electricity market," *Operations Research*, vol. 65, no. 4, pp. 837–855, 2017.
- [26] L. Mitridati and P. Pinson, "A Bayesian inference approach to unveil supply curves in electricity markets," *IEEE Transactions on Power Systems*, vol. 33, no. 3, pp. 2610–2620, 2017.
- [27] Z. Liang and Y. Dvorkin, "Data-driven inverse optimization for marginal offer price recovery in electricity markets," in *Proceedings of the 14th ACM International Conference on Future Energy Systems*, pp. 497–509, 2023.
- [28] V. Dvorkin, P. Van Hentenryck, J. Kazempour, and P. Pinson, "Differentially private distributed optimal power flow," in *2020 59th IEEE Conference on Decision and Control (CDC)*, pp. 2092–2097, IEEE, 2020.
- [29] L. Xie, S. Zou, Y. Xie, and V. V. Veeravalli, "Sequential (quickest) change detection: Classical results and new directions," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 2, pp. 494–514, 2021.
- [30] L. A. Roald, D. Pozo, A. Papavasiliou, D. K. Molzahn, J. Kazempour, and A. Conejo, "Power systems optimization under uncertainty: A review of methods and applications," *Electric Power Systems Research*, vol. 214, p. 108725, 2023.
- [31] S. Chatzivasilas, "Lecture notes on optimal power flow (OPF)," *arXiv preprint arXiv:1811.00943*, 2018.
- [32] G. Casella and R. Berger, *Statistical Inference*. Duxbury Pacific Grove, 2002.
- [33] M. Basseville and I. V. Nikiforov, "Detection of abrupt changes: theory and applications," *Journal of the Royal Statistical Society-Series A Statistics in Society*, vol. 158, no. 1, p. 185, 1995.

- [34] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential analysis: Hypothesis testing and changepoint detection*. CRC press, 2014.
- [35] S. Babaeinejadsarookolae *et al.*, “The power grid library for benchmarking AC optimal power flow algorithms,” *arXiv preprint arXiv:1908.02788*, 2019.
- [36] M. Hoseinpour, “Code: outage identification from electricity market data: Quickest change detection approach.” <https://github.com/miladpour/paper-cusum-detection>, 2026.

I. REVIEW #78A

A. Paper summary

This interesting paper analyzes market signals to detect specific outages by detecting likely changes in the probability distributions governing those signals. The market signals are from an OPF that changes structure according to each outage. The outages are detected in a small power system.

B. Comments for authors

My impression is that the paper is interesting, the analysis is credible, and the results shown to be quite promisingly practical for a small system.

As is reasonable for the scope of a conference paper, the test power system is small, but I do think that the paper should briefly discuss thoughts on how the method might scale to larger systems, as it is obvious that this is more challenging because there are more different outages, and there is perhaps or probably a problem that it is harder to distinguish different outages or deal with multiple outages.

C. Summarize your recommendation in one to two sentences

A very good paper suitable for PowerUp

— Anonymous reviewer

II. REVIEW #78B

A. Paper summary

This paper proposes using market data for identifying outages in the power system.

B. Comments for authors

The idea behind this paper is interesting, and one can imagine that energy traders would benefit directly from the proposed solution. However, it is not clear how the timing of market-clearing data releases would align with detection efforts. This approach may be more suitable for forensic analysis or post-event review.

C. Summarize your recommendation in one to two sentences

The paper is novel and interesting; however, the practical value of the proposed method is not clear.

— Anonymous reviewer

III. REVIEW #78C

A. Paper summary

Power system outages alter the economics of dispatch, and their signature appears in public streams of locational marginal prices (LMPs). The paper asks whether this signal is sufficient for independent, fast outage identification, and proposes a CuSum-based detection method grounded in explicit probabilistic models of LMP streams. The key technical ingredient is multi-parametric programming, which decomposes the market-clearing optimization into regions where LMPs depend affinely on demand, yielding tractable Gaussian densities that power the likelihood ratios in the CuSum. Results on a small PJM testbed are promising but the method's scalability to realistic systems remains an open question.

B. Comments for authors

Multi-parametric programming decomposes the market-clearing optimization into critical regions with affine LMP-demand mappings, yielding closed-form Gaussian densities that make parametric QCD tractable. This is a reasonable technical chain. The development is clear and the adaptive CuSum extension for bounded perturbations is a practical addition. However, the “public data” framing overstates accessibility since setting up the method requires full knowledge of generation costs, PTDF matrix, and generator limits, which are not generally public. The validation on a small, stylized case seems insufficient to support claims of practical applicability, particularly given that the number of MPP critical regions can grow exponentially with system size. Finally, the paper does not establish that the median detection delay is operationally useful. Essentially, the value proposition for market participants rests on independent detection being faster than official outage reporting but it is not clear whether this is sufficient for market participants to act and hedge effectively.

C. Summarize your recommendation in one to two sentences

The paper presents a technically interesting idea. The validation is limited to a small stylized testbed and several key assumptions particularly around data accessibility and scalability of multi-parametric programming are not adequately addressed.

— Anonymous reviewer