Sequential Analysis: Testing, Estimation, Detection, and Decision

Syllabus

Aaditya Ramdas

September 21, 2026

Contents

1 Course at a glance

Course title. Sequential Analysis: Testing, Estimation, Detection, and Decision
Format. Ten weeks; twenty 80-minute meetings selected from thirty candidate lectures; primarily whiteboard lectures
Expected enrollment. 20 students
Audience. Master's and PhD students in Statistics, with mathematically prepared students from computer science, economics, electrical engineering, operations research, and related fields welcome
Catalog number. 223/323
Instructor. Aaditya Ramdas `aramdas@stanford.edu`
TA. Reese Mae Feldmeier `rfeld@stanford.edu`
Website & live schedule. https://web.stanford.edu/∼aramdas/26F-SeqAnalysis/

Important scheduling note for students. This is a rough, deliberately overcomplete syllabus. The sequence below contains thirty candidate lectures for less than twenty class meetings. At least ten lectures will therefore be omitted, combined, or assigned as guided reading, depending on pace, student preparation, and industry interests.

Sequential analysis asks what changes when data arrive over time and the analyst may use the data observed so far to decide **whether to continue, what to measure next, and when to stop**. These freedoms invalidate many fixed-sample guarantees if ignored, but they also create opportunities: a well-designed sequential procedure can stop early, adapt intelligently, detect changes quickly, and report valid uncertainty at every time.

The course develops a common mathematical language for these problems. The central objects are stopping times, likelihood-ratio martingales, nonnegative supermartingales, e-values and e-processes, confidence sequences, and information inequalities. Classical sequential probability ratio tests lead naturally to modern testing by betting. The same ideas are then reused in A/B testing, sequential change detection, election auditing, multiple testing, multi-armed bandits, and optimal stopping.

The emphasis is on **fundamental ideas rather than encyclopedic coverage**. Every meeting centers on at least one formal result, with a proof chosen to fit on a whiteboard. More advanced theorems are stated precisely and interpreted, even when their full proofs lie beyond the course.

2 Learning objectives

By the end of the course, students should be able to do at least half of the following:

1. formulate sequential problems using filtrations, stopping times, predictable processes, and adapted decisions;
2. construct likelihood-ratio martingales and apply optional-stopping and maximal inequalities safely;
3. derive and analyze the sequential probability ratio test and information lower bounds;
4. design e-values and e-processes using conditional betting factors, exponential supermartingales, mixtures, and universal inference;
5. invert sequential tests to obtain confidence sequences and reason about time-uniform coverage;
6. construct fixed-width sequential confidence sets and distinguish terminal asymptotic coverage from simultaneous coverage;
7. build valid analyses for adaptively monitored and adaptively designed experiments;
8. formulate sequential experimental design as maximin information acquisition under predictable control;
9. analyze change detectors, finite-population audits, and offline or online multiple testing;
10. prove representative regret and identification guarantees for bandit algorithms;
11. formulate Robbins–Monro stochastic approximation and SGD and verify their conditional-noise and step-size assumptions;
12. solve finite-horizon stopping problems with the Snell envelope and analyze the secretary problem;
13. use conformal evidence, group-sequential tests, quickest-detection criteria, contextual policy evaluation, post-deployment inference, and Bayesian decisions when those flex lectures are selected; and
14. recognize the shared information-theoretic and conditional-expectation structure behind testing, estimation, design, optimization, detection, and adaptive sampling.

3 Prerequisites

Students should be comfortable with graduate probability at the level of conditional expectation, martingales, modes of convergence, and Radon–Nikodym derivatives, together with mathematical statistics at the level of likelihood ratios, exponential families, and confidence intervals. Familiarity with convexity and Kullback–Leibler divergence is useful. Measure-theoretic details will be stated carefully, but the course will favor discrete-time arguments that can be proved directly at the board.

A short prerequisite list:

- filtrations, adapted and predictable processes;
- stopping times and stopped processes;
- conditional Jensen and tower properties;
- likelihood ratios and change of measure;
- Hoeffding's lemma and sub-Gaussian variables; and
- binary relative entropy.

4 Course policies

4.1 AI policy

The use of AI for solving homework problems is strongly discouraged (it's akin to sending your robot to the gym so that you can get fitter). It is ok to use it to brainstorm project ideas, but write the project proposal yourself, and think deeply about everything you write in the report (I have no doubt that AI knows about filtrations). Regardless, all submitted work in this class must be written by you and not by AI (it is often much clearer when you write it, and even if you use AI, put it away and write stuff down yourself so that you actually find your gaps by seeing where you struggle to put down the math).

4.2 Electronics policy

The use of electronics in class is strongly discouraged (no laptops, phones, etc.). If you must use it, sit in the last row, in a corner, so that you do not distract others. The instructor will provide Latex notes on the website, and you can take written notes as well.

4.3 Late days

You have a total of 5 late days to be used across all the deliverables (3 homeworks, project proposal, final report) to use as you see fit. Eg: one late day for each, or, 5 late days for the final report (but no late days before that).

4.4 Access and Accommodations

Stanford University is committed to providing an inclusive and accessible educational environment for all students. Students seeking disability-related accommodations should contact the Office of Accessible Education (OAE) as early as possible to initiate the accommodation process. Once accommodations are approved, students should provide their accommodation letter to the instructor and discuss implementation within the course. Early communication supports timely access to course materials, instruction, and assessments. For more information, please visit: `oae.stanford.edu`.

4.5 Intended 80-minute meeting

Time	Board activity
0–10 min	motivating problem, recap, and one diagnostic question
10–28 min	definitions and setup
28–58 min	proof of the central result
58–72 min	worked example or algorithm
72–80 min	limitations, takeaway, and an exit exercise

4.6 Scope choices

The main sequence deliberately includes statistical inference, sequential design, and sequential decision-making, but it does not attempt a deep treatment of every subfield. Ten flex lectures place sequential conformal methods, fixed-width confidence sets, group-sequential design, sequential experimental design, minimax quickest change detection, online FDR, contextual-bandit policy evaluation, stochastic approximation and SGD, inference after adaptive deployment, and Bayesian sequential decisions beside the core material they extend. The calendar cannot accommodate all thirty candidates, so selection and compression are part of the course design rather than a promise of encyclopedic coverage.

5 Weekly schedule

The first class will discuss the syllabus, requirements, expectations, etc., so the rest of the numbering should be understood as starting from Thu, Sep 24. And it does not include the in-class midterm.

5.1 Week 1 — Why stopping changes inference

5.1.1 Lecture 1. Stopping times, martingales, and safe optional sampling

- sequential experiments and filtrations;
- stopping times, predictable indicators, and stopped processes;
- bounded optional stopping;
- gambler's ruin and a first warning about optional stopping.

Proof in class: if (M_t) is an integrable martingale and $\tau \le N$ is a stopping time, then $\mathbb{E}[M_\tau] = \mathbb{E}[M_0]$.

5.1.2 Lecture 2. Likelihood-ratio martingales and Ville's inequality

- sequential likelihood ratios and change of measure;
- nonnegative supermartingales;
- Ville's maximal inequality;
- anytime-valid tests and always-valid p-values.

Proof in class: for a nonnegative supermartingale (E_t),

$$\mathbb{P}\left(\sup_{t\ge 0} E_t \ge 1/\alpha\right) \le \alpha\,\mathbb{E}[E_0].$$

Suggested reading: Wald (1947); Ville (1939); Ramdas et al. (2023).

5.2 Week 2 — Classical sequential testing

5.2.1 Lecture 3. The sequential probability ratio test

- simple versus simple testing;
- upper and lower likelihood-ratio boundaries;
- exact conservative type-I and type-II error bounds;
- Bernoulli and Gaussian examples;
- Wald's approximate boundary calibration and overshoot.

Proof in class: the SPRT with upper boundary A and lower boundary B has type-I error at most $1/A$ and type-II error at most B.

5.2.2 Lecture 4. Expected sample size and information lower bounds

- Wald's identities and the log-likelihood random walk;
- data processing for relative entropy;
- lower bounds on expected stopping time;
- statement and interpretation of Wald–Wolfowitz optimality.

Proof in class: any sequential test with errors (α, β) satisfies

$$\mathbb{E}_0 \tau \, \mathrm{KL}(P_0, P_1) \geq \mathrm{kl}(\alpha, 1 - \beta),$$

under standard integrability conditions.

Suggested reading: Wald (1947); Wald and Wolfowitz (1948); Siegmund (1985).

5.3 Week 3 — Testing by betting

5.3.1 Lecture 5. E-values, e-processes, and Kelly growth

- e-values as nonnegative evidence with null mean at most one;
- conditional e-factors and multiplicative capital;
- e-processes and optional continuation;
- calibrating e-values to p-values;
- likelihood ratios as log-growth-optimal bets.

Proof in class: products of conditional e-factors form an e-process; among all e-values, the likelihood ratio maximizes expected log wealth under a fixed alternative.

5.3.2 Lecture 6. Constructing e-processes for means

- linear betting for bounded observations;
- predictable plug-in bets;
- Hoeffding's lemma;
- exponential supermartingales for sub-Gaussian means;
- one-sided versus two-sided tests.

Proof in class: if $X_t \in [a, b]$ and $\mathbb{E}[X_t \mid \mathcal{F}_{t-1}] \leq \mu$, then

$$\exp\left\{ \lambda \sum_{i=1}^{t} (X_i - \mu) - \frac{\lambda^2 (b-a)^2 t}{8} \right\}$$

is a nonnegative supermartingale for every $\lambda \geq 0$.

Suggested reading: Ramdas et al. (2023); Waudby-Smith and Ramdas (2024); Grünwald, Heide, and Koolen (2024).

5.4 Week 4 — Confidence sequences and conformal evidence

5.4.1 Lecture 7. Test inversion and time-uniform confidence sets

- fixed-time confidence intervals versus confidence sequences;
- inversion of a family of e-processes;
- nested confidence sequences and anytime coverage;
- Gaussian-mean and Bernoulli-mean examples.

Proof in class: inverting level-α sequential tests produces a $(1-\alpha)$ confidence sequence.

5.4.2 Lecture 8. Mixtures, stitching, and variance adaptation

- why one tuning parameter cannot be optimal at every time;
- method of mixtures;
- normal-mixture boundary in closed form;
- stitching over geometric epochs;
- statement of empirical-Bernstein/Freedman confidence sequences.

Proof in class: a nonnegative mixture of test supermartingales is a test supermartingale; integrate a Gaussian prior over the exponential parameter to derive a curved boundary.

5.4.3 Lecture 8A. Sequential conformal prediction and conformal evidence

Flex candidate. This lecture may be taught, compressed into a reading session, or omitted.

- exchangeability and online conformity scores;
- randomized conformal ranks and finite-sample validity;
- prediction sets obtained by inversion;
- betting functions that turn conformal p-values into an e-process;
- the distinction between marginal predictive coverage and time-uniform evidence.

Proof in class: randomized ranks are uniform under exchangeability, and predictable betting factors with unit integral multiply into a nonnegative martingale.

Suggested reading: Darling and Robbins (1967); Robbins (1970); Howard et al. (2020); Howard et al. (2021).

5.4.4 Lecture 8B. Fixed-width sequential confidence sets

Flex candidate. This lecture contrasts terminal precision guarantees with confidence-sequence coverage.

- Chow–Robbins stopping for an unknown mean and variance;
- oracle fixed-width sample size and variance-adaptive stopping;
- almost-sure termination and first-order sample-size efficiency;
- random-time studentization and asymptotic terminal coverage;
- confidence-sequence width stopping and multivariate confidence sets.

Proof in class: if $S_n^2 \to \sigma^2$ and the pilot grows slowly, the Chow–Robbins stopping time satisfies $N_d/(z^2\sigma^2/d^2) \to 1$ almost surely.

Suggested reading: Y. S. Chow and Robbins (1965); Yuan Shih Chow, Robbins, and Siegmund (1971); compare Howard et al. (2021).

5.5 Week 5 — Composite models and adaptive experiments

5.5.1 Lecture 9. Composite nulls and universal inference

- why generalized likelihood ratios are not automatically e-values;
- domination by simple-null likelihood ratios;
- split likelihood and universal inference;
- nuisance parameters and safe model fitting;
- optional continuation across studies.

Proof in class: if $\bar{q}(x^t)$ is any sequential density and $\mathcal{P}_0$ is the null family, then

$$E_t = \frac{\bar{q}(X^t)}{\sup_{P \in \mathcal{P}_0} p(X^t)}$$

is an e-process for $\mathcal{P}_0$.

5.5.2 Lecture 10. Anytime-valid A/B testing under adaptive allocation

- continuous monitoring and peeking;
- sequential randomization and predictable treatment probabilities;
- inverse-propensity-weighted martingale differences;
- confidence sequences for treatment effects;
- practical issues: delayed outcomes, multiple metrics, guardrails, and power.

Proof in class: under sequential randomization,

$$Z_t = \frac{A_t Y_t}{\pi_t} - \frac{(1 - A_t) Y_t}{1 - \pi_t}$$

has conditional mean equal to the contemporaneous treatment effect.

5.5.3 Lecture 10A. Group-sequential testing and alpha spending

Flex candidate. This lecture may replace part of the adaptive-experiment or multiple-testing material.

- planned looks in canonical information time;
- joint Gaussian/Brownian null laws;
- cumulative versus local type-I error;
- alpha-spending functions and conservative Bonferroni construction;
- Pocock- and O'Brien–Fleming-style boundaries;
- finite-look calibration versus open-ended e-process validity.

Proof in class: if boundaries are calibrated so that the cumulative crossing probability at information time t_k is $a(t_k)$, then the final type-I error is $a(1) = \alpha$.

Suggested reading: Johari, Pekelis, and Walsh (2022); Howard et al. (2021).

5.5.4 Lecture 10B. Sequential design of experiments and controlled sensing

Flex candidate. This lecture adds adaptive control of the next measurement to adaptive stopping.

- finite hypotheses, predictable sensing actions, and controlled likelihoods;
- likelihood-ratio martingales under adaptive action selection;
- information rates and maximin KL experiment allocation;
- Chernoff-style policies, forced exploration, and identifiability;
- diagnostic selection, sensor scheduling, and information versus reward.

Proof in class: predictable action selection preserves pairwise likelihood-ratio martingales; a Ville-plus-union-bound threshold controls the probability of a wrong terminal declaration.

Suggested reading: Chernoff (1959); Kiefer and Sacks (1963).

5.6 Week 6 — Sequential change detection

5.6.1 Lecture 11. CUSUM and Shiryaev–Roberts statistics

- change-point model and detection delay;
- log-likelihood CUSUM;
- Shiryaev–Roberts mixture statistic;
- recursive implementations;
- false-alarm metrics and average run length.

Proof in class: derive the CUSUM and Shiryaev–Roberts recursions from maxima and sums over candidate change times.

5.6.2 Lecture 12. E-detectors and nonparametric false-alarm control

- e-processes started at every possible change time;
- sums and maxima of restarted evidence;
- definition of an e-detector;
- average-run-length control;
- bounded-mean-shift detector without a post-change parametric model.

Proof in class: if M_t is an e-detector, then $T = \inf\{t : M_t \geq 1/\alpha\}$ satisfies $\mathbb{E}_\infty T \geq 1/\alpha$.

5.6.3 Lecture 12A. Advanced quickest change detection

Flex candidate. This lecture deepens Week 6 when change detection is a course priority.

- Lorden and Pollak worst-case delay criteria;
- average-run-length constraints and false-alarm calibration;
- CUSUM's least favorable post-change state;
- information lower bounds of order $\log(\gamma)/\mathrm{KL}(P_1, P_0)$;
- first-order minimax optimality, equalizers, head starts, and unknown post-change models.

Proof in class: every detector with no-change average run length at least γ has first-order worst-case delay at least $\log(\gamma)/I$, and CUSUM attains this scale under standard renewal conditions.

Suggested reading: Page (1954); Roberts (1966); Shin, Ramdas, and Rinaldo (2024).

5.7 Week 7 — Finite populations, many hypotheses, and online FDR

5.7.1 Lecture 13. Sampling without replacement and risk-limiting election audits

- random-permutation filtration for a finite population;
- conditional remaining means;
- betting martingales without replacement;
- assorters and reduction of election assertions to mean tests;
- risk limit as a sequential type-I guarantee.

Proof in class: under a candidate finite-population mean m, the product

$$\prod_{t=1}^{n}\{1+\lambda_t(X_t-m_t)\}, \qquad m_t=\frac{Nm-\sum_{i<t}X_i}{N-t+1},$$

is a martingale for every admissible predictable betting sequence.

5.7.2 Lecture 14. Multiple testing with e-values

- false discovery rate;
- the e-BH step-up rule;
- arbitrary dependence;
- combining evidence across streams and cautions about global stopping.

Proof in class: e-BH controls FDR at level α under arbitrary dependence when each true-null e-value has expectation at most one.

5.7.3 Lecture 14A. Online false discovery control and alpha-investing

Flex candidate. This lecture extends multiple testing from a fixed family to an indefinite stream of confirmatory hypotheses.

- online FDR and modified FDR;
- conditional super-uniformity and predictable test levels;
- alpha wealth, spending, and rewards after discoveries;
- an accounting proof of modified-FDR control;
- e-LOND and fixed-time FDR control under arbitrary dependence;
- operational issues including alpha-death, retries, shared users, and stopping horizons.

Proof in class: nonnegative alpha wealth bounds cumulative testing levels by $\alpha\{R(t)+\eta\}$, while the e-LOND rejection threshold gives a direct fixed-time FDR proof from null e-value expectations.

Suggested reading: Stark (2020); Stark (2023); Waudby-Smith and Ramdas (2020); Waudby-Smith, Stark, and Ramdas (2021); Wang and Ramdas (2022).

5.8 Week 8 — Adaptive allocation, policy evaluation, and identification

5.8.1 Lecture 15. Stochastic bandits and UCB

- regret and the exploration–exploitation tradeoff;
- upper confidence bound algorithms;
- concentration under adaptive sampling;
- logarithmic instance-dependent regret.

Proof in class: for bounded rewards, UCB1 pulls a suboptimal arm a only $O(\log n/\Delta_a^2)$ times in expectation.

5.8.2 Lecture 15A. Contextual bandits and off-policy evaluation

Flex candidate. This lecture is recommended when personalization, recommender systems, or experimentation platforms are central applications.

- contexts, actions, potential rewards, and stochastic policies;
- linear contextual UCB and contextual regret;
- logging policies, overlap, and importance weights;
- inverse-propensity-score policy evaluation under adaptive logging;
- predictable doubly robust scores;
- effective sample size, weight clipping, and post-selection bias.

Proof in class: conditioning on the observed context and logged action probabilities makes inverse-propensity scoring unbiased; adding a predicted reward and an importance-weighted residual preserves the same target value.

5.8.3 Lecture 15B. Stochastic approximation and stochastic gradient descent

Flex candidate. This lecture connects classical noisy root finding to modern online optimization.

- Robbins–Monro recursion and martingale-difference noise;
- step-size conditions $\sum a_t = \infty$ and $\sum a_t^2 < \infty$;
- squared-distance Lyapunov proof of convergence;
- SGD, streaming quantile calibration, and projected updates;
- iterate averaging, constant-step tracking, and feedback-loop failures.

Proof in class: strong monotonicity plus square-summable noise energy makes the squared error an almost supermartingale and forces convergence to the root.

Suggested reading: Robbins and Monro (1951); Polyak and Juditsky (1992); Bottou, Curtis, and Nocedal (2018); Lai (2003).

5.8.4 Lecture 16. Best-arm identification with confidence sequences

- fixed-confidence versus fixed-budget goals;
- successive elimination;
- simultaneous per-arm confidence sequences;
- correctness and gap-dependent sample complexity;
- connection to information lower bounds.

Proof in class: on a simultaneous confidence event, the best arm is never eliminated and every suboptimal arm is eventually eliminated; union bounds yield δ-correctness.

5.8.5 Lecture 16A. Inference after adaptive deployment

Flex candidate. This lecture separates the inferential consequences of adaptive allocation, adaptive stopping, and adaptive target selection.

- inverse-propensity martingale scores for adaptively assigned arms;
- finite-sample confidence sequences under minimum exploration;

- selection-safe reporting for an adaptively chosen arm and stopping time;
- augmented scores and predictable variance stabilization;
- delayed outcomes, population drift, policy resets, and logging integrity;
- finite-sample anytime-valid versus asymptotic fixed-time guarantees.

Proof in class: bounded inverse-propensity scores generate an exponential supermartingale, and a simultaneous confidence event remains valid after selecting both an arm and a stopping time.

Suggested reading: Auer, Cesa-Bianchi, and Fischer (2002); Even-Dar, Mannor, and Mansour (2006); Kaufmann, Cappé, and Garivier (2016); Lattimore and Szepesvári (2020).

Additional reading for the flex lectures: Javanmard and Montanari (2018); Dudik, Langford, and Li (2011); Wang, Agarwal, and Dudik (2017); Hadad et al. (2021).

5.9 Week 9 — Optimal stopping

5.9.1 Lecture 17. The Snell envelope

- finite-horizon stopping problems;
- backward induction;
- the smallest dominating supermartingale;
- first-contact stopping rule;
- examples with selling and search.

Proof in class: the Snell envelope dominates the reward process, is the smallest such supermartingale, and its first contact time is optimal.

5.9.2 Lecture 18. The secretary problem and the $1/e$ rule

- random order and relative ranks;
- threshold strategies;
- exact success probability;
- asymptotic optimization;
- what changes with recall, costs, or unknown horizon.

Proof in class: skipping the first r applicants and then taking the next record succeeds with probability

$$\frac{r}{n}\sum_{j=r}^{n-1}\frac{1}{j},$$

which is maximized at $r/n \to e^{-1}$ and converges to e^{-1}.

5.9.3 Lecture 18A. Bayesian sequential decisions

Flex candidate. This lecture may replace or follow the secretary problem when decision theory is emphasized.

- posterior odds as prior odds times accumulated likelihood ratio;
- terminal misclassification loss and per-observation cost;
- finite-horizon Bellman recursion;
- threshold structure and relation to the SPRT;
- one-step Bernoulli example;

- Bayes-optimality versus frequentist operating characteristics.

Proof in class: backward induction yields the stop/continue Bellman equation, with immediate Bayes risk compared against sampling cost plus expected continuation value.

Suggested reading: Yuan Shih Chow, Robbins, and Siegmund (1971); Dynkin (1963); Gilbert and Mosteller (1966).

5.10 Week 10 — Limits and synthesis

5.10.1 Lecture 19. Law-of-the-iterated-logarithm boundaries

- why repeatedly using fixed-time intervals fails;
- the law of the iterated logarithm;
- unavoidable $\sqrt{t \log \log t}$ scale;
- stitching as a near-optimal construction;
- practical finite-sample versus asymptotic tradeoffs.

Proof in class: a fixed-z boundary $c\sqrt{t}$ is crossed infinitely often under a mean-zero finite-variance model, while geometrically stitched boundaries attain the iterated-logarithm order.

5.10.2 Lecture 20. A unifying adaptive information inequality

- chain rule for stopped log-likelihood ratios;
- data processing through the terminal decision;
- lower bounds for testing, best-arm identification, and detection;
- course synthesis: design evidence, calibrate error, optimize information.

Proof in class: for models P, Q, a stopping time τ, and $A \in \mathcal{F}_\tau$,

$$\mathbb{E}_P\left[\sum_{t=1}^{\tau} \mathrm{KL}\big(P_t(\cdot \mid \mathcal{F}_{t-1}), Q_t(\cdot \mid \mathcal{F}_{t-1})\big)\right] \geq \mathrm{kl}(P(A), Q(A)).$$

Suggested reading: Robbins (1970); Lai and Robbins (1985); Kaufmann, Cappé, Garivier (2016).

6 Assessment and course work

The following weighting is a draft designed for a small, advanced class.

Component	Weight	Purpose
3 problem sets	30%	proofs, calculations, and small simulations
Final project	40%	theoretical, methodological, applied sequential analysis
In-class midterm	30%	in Week 7

Problem sets will be released approximately every two weeks and will be due roughly middle of 3rd week, middle of 5th week and start of 9th week. Students may discuss ideas, but submitted proofs and code must be their own and collaborators must be acknowledged. The instructor should replace this paragraph with official university and department policy language before publication.

6.1 Final project

Students may work alone or in pairs. A project should formulate a genuinely sequential question, specify the filtration and operating characteristics, justify the validity of the proposed method, and evaluate power, width, regret, delay, or sample size. A polished project must have at least two out of the following three parts:

- **theory:** prove or sharpen a result for a model not covered in class;
- **methodology:** design an e-process, confidence sequence, detector, or adaptive algorithm and compare it with a baseline; or
- **application:** conduct a careful sequential analysis of a real or realistically simulated problem.

Possible domains include online experiments, public-health surveillance, reliability monitoring, election auditing, adaptive clinical experimentation, anomaly detection, best-arm identification, model comparison, or sequential forecasting.

Milestones are a one-page proposal at the end of Week 4 (5 out of 40), a five-minute board pitch in Week 7 (5 out of 40), and a final project report (30 out of 40).

7 Reading strategy

There is no single required textbook. Ideally, you would get used to consulting old and modern papers and texts, like

1. classical monographs, especially Wald (1947) and Siegmund (1985);
2. probability and optimal-stopping references such as Yuan Shih Chow, Robbins, and Siegmund (1971) and Lattimore and Szepesvári (2020); and
3. modern papers on time-uniform inference, betting, and e-processes, especially Howard et al. (2020), Howard et al. (2021), Ramdas et al. (2023), and Waudby-Smith and Ramdas (2024).

8 Notation and conventions

- $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, P)$ is a filtered probability space in discrete time.
- A process indexed by $t-1$ and chosen before seeing X_t is called predictable.
- "Martingale" includes integrability; "nonnegative supermartingale" may take the value $+\infty$ only when explicitly stated.
- $\mathrm{kl}(p,q) = p\log(p/q) + (1-p)\log((1-p)/(1-q))$ is binary relative entropy.
- $\mathrm{KL}(P,Q) = \mathbb{E}_P \log(dP/dQ)$ when $P \ll Q$.
- A level-α sequential test controls the probability of **ever** rejecting a true null, unless a different error criterion is stated.
- Confidence sequences use closed or open endpoints according to convenience; boundary events can be handled by a harmless convention or randomization.

9 Core reference list

Auer, Peter, Nicolò Cesa-Bianchi, and Paul Fischer. 2002. "Finite-Time Analysis of the Multiarmed Bandit Problem." *Machine Learning* 47: 235–56. https://doi.org/10.1023/A:1013689704352.

Bottou, Léon, Frank E. Curtis, and Jorge Nocedal. 2018. "Optimization Methods for Large-Scale Machine Learning." *SIAM Review* 60 (2): 223–311. https://doi.org/10.1137/16M1080173.

Chernoff, Herman. 1959. "Sequential Design of Experiments." *The Annals of Mathematical Statistics* 30 (3): 755–70. https://doi.org/10.1214/aoms/1177706205.

Chow, Y. S., and Herbert Robbins. 1965. "On the Asymptotic Theory of Fixed-Width Sequential Confidence Intervals for the Mean." *The Annals of Mathematical Statistics* 36 (2): 457–62. https://doi.org/10.1214/aoms/1177700156.

Chow, Yuan Shih, Herbert Robbins, and David Siegmund. 1971. *Great Expectations: The Theory of Optimal Stopping*. Boston: Houghton Mifflin.

Darling, D. A., and Herbert Robbins. 1967. "Confidence Sequences for Mean, Variance, and Median." *Proceedings of the National Academy of Sciences* 58 (1): 66–68. https://doi.org/10.1073/pnas.58.1.66.

Dynkin, E. B. 1963. "The Optimum Choice of the Instant for Stopping a Markov Process." *Soviet Mathematics Doklady* 4: 627–29.

Even-Dar, Eyal, Shie Mannor, and Yishay Mansour. 2006. "Action Elimination and Stopping Conditions for the Multi-Armed Bandit and Reinforcement Learning Problems." *Journal of Machine Learning Research* 7: 1079–1105.

Gilbert, John P., and Frederick Mosteller. 1966. "Recognizing the Maximum of a Sequence." *Journal of the American Statistical Association* 61 (313): 35–73. https://doi.org/10.1080/01621459.1966.10502008.

Grünwald, Peter, Rianne de Heide, and Wouter M. Koolen. 2024. "Safe Testing." *Journal of the Royal Statistical Society: Series B* 86 (5): 1091–1128. https://doi.org/10.1093/jrsssb/qkae011.

Howard, Steven R., Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. 2020. "Time-Uniform Chernoff Bounds via Nonnegative Supermartingales." *Probability Surveys* 17: 257–317. https://doi.org/10.1214/18-PS321.

———. 2021. "Time-Uniform, Nonparametric, Nonasymptotic Confidence Sequences." *The Annals of Statistics* 49 (2): 1055–80. https://doi.org/10.1214/20-AOS1991.

Johari, Ramesh, Leo Pekelis, and David J. Walsh. 2022. "Always Valid Inference: Continuous Monitoring of A/B Tests." *Operations Research* 70 (3): 1806–21. https://doi.org/10.1287/opre.2021.2135.

Kaufmann, Emilie, Olivier Cappé, and Aurélien Garivier. 2016. "On the Complexity of Best-Arm Identification in Multi-Armed Bandit Models." *Journal of Machine Learning Research* 17 (1): 1–42.

Kiefer, Jack, and Jerome Sacks. 1963. "Asymptotically Optimum Sequential Inference and Design." *The Annals of Mathematical Statistics* 34 (3): 705–50.

Lai, Tze Leung. 2003. "Stochastic Approximation." *The Annals of Statistics* 31 (2): 391–406. https://doi.org/10.1214/aos/1051027873.

Lai, Tze Leung, and Herbert Robbins. 1985. "Asymptotically Efficient Adaptive Allocation Rules." *Advances in Applied Mathematics* 6 (1): 4–22. https://doi.org/10.1016/0196-8858(85)90002-8.

Lattimore, Tor, and Csaba Szepesvári. 2020. *Bandit Algorithms*. Cambridge University Press. https://doi.org/10.1017/9781108571401.

Page, E. S. 1954. "Continuous Inspection Schemes." *Biometrika* 41 (1/2): 100–115. https://doi.org/10.1093/biomet/41.1-2.100.

Polyak, Boris T., and Anatoli B. Juditsky. 1992. "Acceleration of Stochastic Approximation by Averaging." *SIAM Journal on Control and Optimization* 30 (4): 838–55. https://doi.org/10.1137/0330046.

Ramdas, Aaditya, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. 2023. "Game-Theoretic Statistics and Safe Anytime-Valid Inference." *Statistical Science* 38 (4): 576–601. https://doi.org/10.1214/23-STS894.

Robbins, Herbert. 1970. "Statistical Methods Related to the Law of the Iterated Logarithm." *The*

Annals of Mathematical Statistics 41 (5): 1397–1409. https://doi.org/10.1214/aoms/1177696786.

Robbins, Herbert, and Sutton Monro. 1951. "A Stochastic Approximation Method." *The Annals of Mathematical Statistics* 22 (3): 400–407. https://doi.org/10.1214/aoms/1177729586.

Roberts, S. W. 1966. "A Comparison of Some Control Chart Procedures." *Technometrics* 8 (3): 411–30. https://doi.org/10.1080/00401706.1966.10490374.

Shin, Jaehyeok, Aaditya Ramdas, and Alessandro Rinaldo. 2024. "E-Detectors: A Nonparametric Framework for Sequential Change Detection." *New England Journal of Statistics in Data Science* 2 (2): 229–50. https://doi.org/10.51387/24-NEJSDS57.

Siegmund, David. 1985. *Sequential Analysis: Tests and Confidence Intervals*. New York: Springer. https://doi.org/10.1007/978-1-4757-1862-1.

Stark, Philip B. 2020. "Sets of Half-Average Nulls Generate Risk-Limiting Audits: SHANGRLA." *Electronic Voting* 34: 319–36. https://doi.org/10.1007/978-3-030-30625-0_17.

———. 2023. "ALPHA: Audit That Learns from Previously Hand-Audited Ballots." *The Annals of Applied Statistics* 17 (1): 641–79. https://doi.org/10.1214/22-AOAS1646.

Ville, Jean. 1939. "Étude Critique de La Notion de Collectif." *Monographies Des Probabilités* 3.

Wald, Abraham. 1947. *Sequential Analysis*. New York: John Wiley & Sons.

Wald, Abraham, and Jacob Wolfowitz. 1948. "Optimum Character of the Sequential Probability Ratio Test." *The Annals of Mathematical Statistics* 19 (3): 326–39. https://doi.org/10.1214/aoms/1177730197.

Wang, Ruodu, and Aaditya Ramdas. 2022. "False Discovery Rate Control with e-Values." *Journal of the Royal Statistical Society: Series B* 84 (3): 822–52. https://doi.org/10.1111/rssb.12489.

Waudby-Smith, Ian, and Aaditya Ramdas. 2020. "Confidence Sequences for Sampling Without Replacement." *Advances in Neural Information Processing Systems* 33: 20204–14.

———. 2024. "Estimating Means of Bounded Random Variables by Betting." *Journal of the Royal Statistical Society: Series B* 86 (1): 1–27. https://doi.org/10.1093/jrsssb/qkad009.

Waudby-Smith, Ian, Philip B. Stark, and Aaditya Ramdas. 2021. "RiLACS: Risk-Limiting Audits via Confidence Sequences." *Proceedings of the 2021 ACM Conference on Advances in Financial Technologies*, 132–43. https://doi.org/10.1145/3479722.3480981.