

Slippage-Aware FPTAS Selection for Time-Zone Event-Based Trading: A Two-Market Perspective

Felix Reichel

June 13, 2026

Abstract

This theoretical perspective develops a formal theory of asynchronous scheduled-news trading between two venues, denoted here as Market 1 and Market 2. The motivating event may be a press release, an earnings release, a scheduled corporate statement, an index-rebalance notice, a macro release, or any pre-announced information arrival whose release time lies after instruments on one venue are no longer directly executable while another venue or proxy remains tradable during that time window; the motivating empirical background includes event-study evidence for earnings and corporate announcements as well as macro-policy announcement effects [2, 3, 35, 32, 4, 21]. The timing wedge creates a finite decision problem rather than a pure arbitrage opportunity: any trader must choose which candidate event trades to execute when expected return, slippage, instrument and leg count, margin, borrow, latency and venue access are all limited resources. We model candidate trades as items, prove that the implementable single-budget selection problem can be modeled exactly as a 0–1 knapsack problem and give fully polynomial-time approximation schemes (FPTAS) for slippage-adjusted net profit maximization and for slippage minimization s.t. a target expected return. We also give immediate formal non-arbitrage, robustness, heterogeneity, basis-risk, estimation-error and scheduled-news generalization results.

JEL Classification: G14; G12; G15; C61; C63.

MSC 2020 Classification: 90C27; 68W25; 90C39; 60K25; 91G10; 91G80.

Keywords: Asynchronous Markets; Scheduled News; Event-Based Trading; Market Microstructure; Execution Latency; Queueing Delay; Stochastic Network Calculus; Basis Risk; Robust Optimization; Knapsack Approximation; Time-Zone Trading; Statistical Arbitrage; iSLIP.

Contents

1 Two-market Timing Model	2
1.1 Scheduled Information	2
1.2 Candidate Event Items	3
2 Payoffs, Slippage and Admissibility	4
3 Queueing-Theoretic Latency Envelopes	5
4 Why the Timing Wedge Is Not Arbitrage	7
5 A Profit-Maximizing FPTAS	8
6 A Slippage-Minimizing FPTAS	9
7 General Scheduled-News Events	9
8 Heterogeneity and Basis-Risk	10
9 Estimation, Rounding Error and Out-of-Sample Degradation	11
10 Conclusions	14
A Main-text Omitted Proofs	19

1 Two-market Timing Model

1.1 Scheduled Information

Let $\mathcal{M} = \{1, 2, P\}$ denote Market 1, Market 2 and a tradable proxy. A day contains ordered times

$$\tau_2^c(t) < \tau_N(t) < \tau_1^c(t), \tag{1}$$

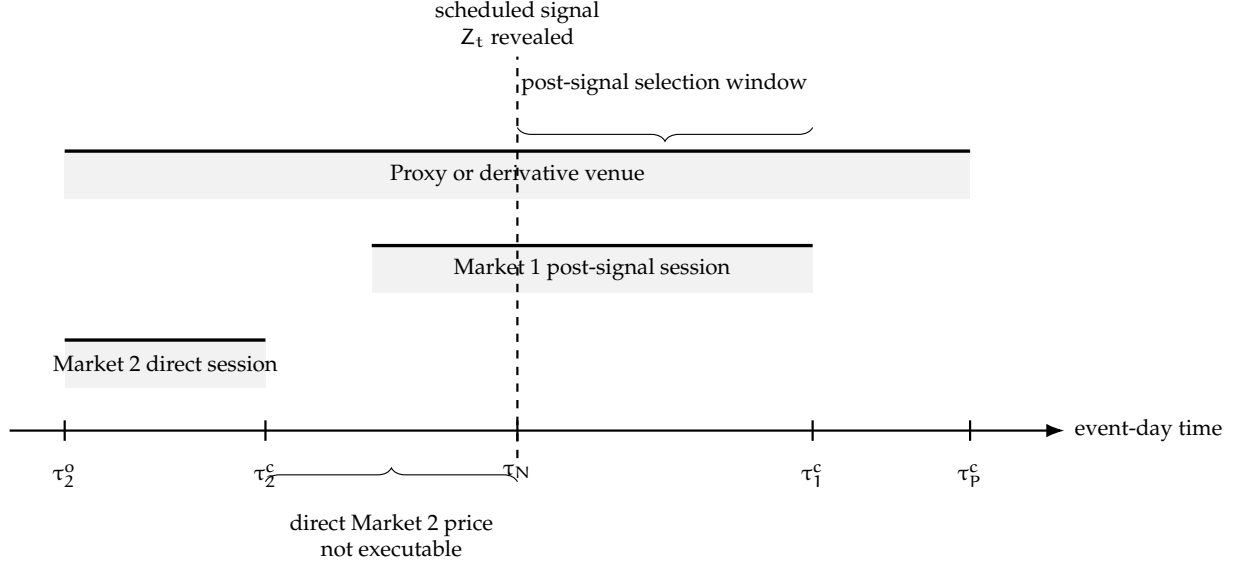


Figure 1: Generic asynchronous scheduled-news timing. The diagram deliberately uses Market 1 and Market 2 labels: the theory applies to any pair of venues with staggered tradability around a scheduled signal.

where $\tau_2^c(t)$ is the Market 2 direct close, $\tau_N(t)$ is a scheduled news release and $\tau_1^c(t)$ is the Market 1 close; the notation abstracts from the standard event-study view of dated information arrivals [8, 32]. The proxy venue P is tradable on a set $\mathcal{T}_P(t)$ such that $\tau_N(t) \in \mathcal{T}_P(t)$.

Definition 1.1 (Scheduled-news Event). On a pre-filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_\tau)_{\tau \geq 0}, \mathbb{P})$, a scheduled-news event is a pair $(\tau_N(t), Z_t)$ for which

$$Z_t \notin \mathcal{F}_{\tau_2^c(t)}, \quad Z_t \in \mathcal{F}_{\tau_N(t)}. \quad (2)$$

The RV formally denoted as Z_t can represent any scheduled signal.

Definition 1.2 (Post-signal Tradability). A claim is post-signal tradable if an admissible order can be submitted and *potentially filled* after $\tau_N(t)$ and before the relevant venue close. A *stale direct claim* is a claim whose last direct execution price occurs before that.

Assumption 1.1 (Asynchronous Executability). Market 2 direct instruments are not directly tradable on $(\tau_2^c(t), \tau_N(t)]$, while at least one instrument in Market 1 or venue P is post-signal tradable.

1.2 Candidate Event Items

Let \mathcal{J} be a finite set of elementary post-signal tradable instruments. An instrument can be a direct Market 1 instrument, a proxy for Market 2 exposure, or any supported substitute. Let \mathcal{D} be a finite set of event-cycle

dates or intraday windows.

Definition 1.3 (Event Item). An item $i \in \mathcal{I}$ is a feasible trade package

$$i = (d_i, \theta_i, h_i, o_i), \quad (3)$$

where $d_i \in \mathcal{D}$ is a scheduled-news window, $\theta_i \in \mathbb{R}^{|\mathcal{J}|}$ is a signed vector of instrument weights, h_i is a holding rule mapping entry and exit times to payoff and o_i is an admissible order protocol.

Example 1.1 (Generic Item Classes). A finite item library may contain: (i) a proxy for Market 2 exposure; (ii) a long–short spread using a Market 1 option against a Market 2 proxy; (iii) a market-neutral basket; (iv) a delayed follow-through trade; or (v) a company-news trade using an issuer option and an industry proxy. The theory treats all of these as finite items with estimated net scores and associated costs.

2 Payoffs, Slippage and Admissibility

The decomposition below follows the standard market-microstructure and optimal-execution view that spreads, market impact, order handling, inventory or financing frictions and routing constraints must be treated as first-order trading costs [22, 34, 24, 19, 28, 1, 6, 16].

For each item i define the gross random return R_i , execution loss L_i and net return as the random variable

$$X_i = R_i - L_i. \quad (4)$$

The ex-ante net score and expected slippage are given formally as

$$\mu_i = \mathbb{E}[X_i \mid \mathcal{F}_0], \quad c_i = \mathbb{E}[L_i \mid \mathcal{F}_0]. \quad (5)$$

Here μ_i and c_i are *ex-ante* magnitudes conditioned on \mathcal{F}_0 ; separating such planned (ex-ante) quantities from their realized (ex-post) outcomes is consistent with the sequence-analysis tradition of the Stockholm School [31]. It is precisely this gap that the estimation and degradation bounds of Section A quantify. A selection vector is $x \in \{0, 1\}^n$ for $n = |\mathcal{I}|$.

Definition 2.1 (Feasible Selection). Given nonnegative resources $a_{\kappa i}$ and budgets B_κ , the feasible set is

$$\mathcal{X} = \{x \in \{0, 1\}^n : Ax \leq B\}. \quad (6)$$




Figure 2: Slippage-aware selection pipeline. Wide spacing and fixed-width nodes are used to avoid label overlap.

Resources may include margin, capital-hours, gross notional, maximum options, borrow availability, latency class, broker access, order-protection or venue-rule constraints and risk budget [38, 33].

Definition 2.2 (Slippage Decomposition). For item i with options $\ell \in \mathcal{L}_i$ and signed sizes $q_{i\ell}$,

$$L_i = \sum_{\ell \in \mathcal{L}_i} (S_{i\ell}|q_{i\ell}| + M_{i\ell}(q_{i\ell}) + F_{i\ell}|q_{i\ell}| + B_{i\ell}|q_{i\ell}|) + \Gamma_i, \quad (7)$$

where S is half-spread cost, M is market impact, F is explicit fee or tax, B is financing or borrow cost and Γ_i is synchronization, latency and partial-fill loss.

Assumption 2.1 (Convex Impact and Bounded Order Size). For every option, $M_{i\ell}$ is nonnegative, convex, satisfies $M_{i\ell}(0) = 0$ and is nondecreasing in $|q|$. Admissible orders satisfy $|q_{i\ell}| \leq \bar{q}$.

Lemma 2.1 (Finite Slippage Envelope). For every finite item library with bounded order size, there are computable constants $\bar{c}_i < \infty$ such that $0 \leq L_i \leq \bar{c}_i$ for every admissible execution of item i .

Proof omitted; see Appendix A, Proof of Lemma 2.1.

Proposition 2.1 (Option Parsimony Weakly Lowers Deterministic Slippage). Suppose items i and j have the same expected gross return and the same resource usage, but i uses a strict subset of the options of j with equal sizes on common options. If all cost components in (7) are nonnegative, then $\mathbb{E}[L_i] \leq \mathbb{E}[L_j]$, with strict inequality if any omitted option has positive expected cost.

Proof omitted; see Appendix A, Proof of Proposition 2.1.

3 Queueing-Theoretic Latency Envelopes

The latency term introduced in the slippage term (7) can be made explicit by modeling the execution path as a finite queueing network. For a time-zone event trade, an order may pass through a strategy engine, broker

risk checks, a smart order router, exchange gateways and a limit-order-book queue before acknowledgement or fill. The relevant object is therefore not only average latency, but the tail probability that an item misses its economically relevant execution deadline. Stochastic network calculus and related queueing bounds, as in Ciucu’s delay-envelope approach and later sharp or tandem-queue refinements, provide a natural way to translate this tail risk into an item-level penalty or feasibility constraint [10, 11, 12, 13].

Definition 3.1 (Execution Path). For item i , let

$$\mathcal{Q}_i = (1, \dots, m_i) \quad (8)$$

be the ordered queueing path used by its order protocol o_i . Node r may represent internal order generation, broker processing, routing, an exchange gateway, queue position in the book, or fill acknowledgement. Let $D_{i,r} \geq 0$ be the delay at node r and define the end-to-end delay

$$D_i = \sum_{r=1}^{m_i} D_{i,r}. \quad (9)$$

Definition 3.2 (Latency Envelope). A nonincreasing function $\eta_i : \mathbb{R}_+ \rightarrow [0, 1]$ is a latency envelope for item i if

$$\mathbb{P}(D_i > d \mid \mathcal{F}_0) \leq \eta_i(d) \quad (10)$$

for all $d \geq 0$. The envelope may be estimated from historical routing data, broker logs, queue-position models, stochastic network calculus, or tandem-queue bounds.

Definition 3.3 (Deadline Admissibility). Let $\Delta_i > 0$ be the latest useful execution delay for item i after the scheduled signal is observed and let $\alpha_i \in [0, 1]$ be a tolerated miss probability. The item is deadline-admissible if

$$\eta_i(\Delta_i) \leq \alpha_i. \quad (11)$$

Definition 3.4 (Queueing-Adjusted Profit). Let $\kappa_i \geq 0$ denote the economic loss from missing the useful execution window for item i . Define the queueing penalty and queueing-adjusted score by

$$\lambda_i = \kappa_i \eta_i(\Delta_i), \quad p_i^Q = p_i - \lambda_i. \quad (12)$$

Additionally, deadline-inadmissible items can be removed from the library before selection.

Proposition 3.1 (Queueing-Adjusted Selection Preserves FPTAS). *Suppose each item is assigned either a deterministic deadline-admissibility indicator or a deterministic queueing penalty λ_i satisfying (12). After deleting*

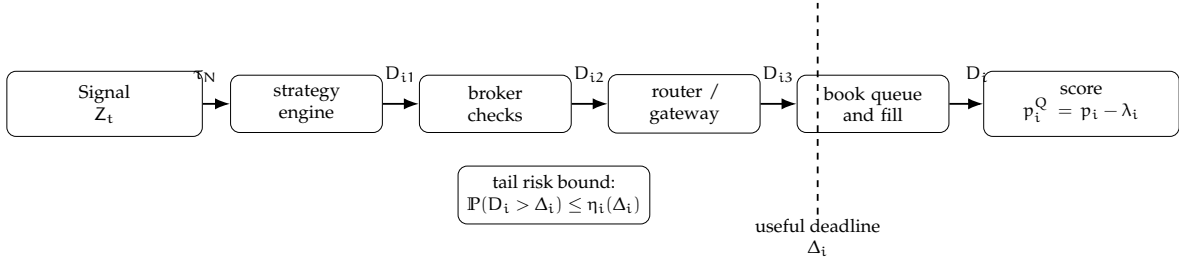


Figure 3: Queueing-theoretic latency envelope. Execution delay is treated as an end-to-end queueing delay and encoded as either a feasibility filter or a penalty in the item score.

deadline-inadmissible items and replacing p_i by p_i^Q , the single-budget selection problem remains a 0–1 knapsack problem and admits the same profit FPTAS.

Proof omitted; see Appendix A, Proof of Proposition 3.1.

Remark 3.1 (Why This Is Useful For Event Trading). Scheduled-news periods are precisely the periods in which routing and book queues can become bursty. A mean-latency adjustment can understate the probability of a late or partial fill. A latency envelope instead penalizes the right tail of D_i , so thin-window strategies, multi-option spreads and proxy trades with slow acknowledgement can be downweighted even when their expected gross return is high.

4 Why the Timing Wedge Is Not Arbitrage

The distinction between a positive expected event edge and a statewise arbitrage is consistent with the asset-pricing and limits-of-arbitrage literature: informational efficiency does not eliminate all predictable or conditional effects, but a risk-bearing strategy is not a pure arbitrage merely because its estimated mean is positive [17, 18, 26, 30, 23, 37, 14, 15, 39, 20].

Definition 4.1 (Pure Arbitrage). A self-financing terminal payoff Y is a pure arbitrage if $Y \geq 0$ a.s.; $\mathbb{P}(Y > 0) > 0$ and initial wealth is zero.

Definition 4.2 (Statistical Event Edge). A family of event strategies is a statistical event edge if its expected net return is positive under a maintained data-generating law, while realized losses remain possible for finite samples.

Proposition 4.1 (No Pure Arbitrage From Asynchronous News Alone). *Assume post-signal tradable prices are adapted to (\mathcal{F}_τ) and satisfy a frictional no-arbitrage pricing rule on their actual trading domains. If stale direct Market 2 prices are not executable after τ_2^c , then the timing relation (1) and information jump (2) do not imply pure arbitrage.*

Proof omitted; see Appendix A, Proof of Proposition 4.1.

Corollary 4.1 (Statistical, Not Risk-Free). *Any implementable rule selected by the FPTAS below is best interpreted as a statistical event strategy unless additional state-wise payoff inequalities are proven.*

5 A Profit-Maximizing FPTAS

The algorithmic step uses the classical profit-scaling approach for knapsack and fully polynomial approximation schemes, with the terminology aligned with standard approximation-algorithm texts [25, 29, 27, 40, 41].

Let $p_i = \hat{\mu}_i - \hat{c}_i$ be an estimated net profit score. Negative-score items are discarded. With one scalar resource w_i and budget B , the selection problem is

$$\text{OPT}(B) = \max_{x \in \{0,1\}^n} \left\{ \sum_{i=1}^n p_i x_i : \sum_{i=1}^n w_i x_i \leq B \right\}. \quad (13)$$

Proposition 5.1 (Exact Reduction). *The single-budget scheduled-news selection problem can be formulated exactly as a 0–1 knapsack problem.*

Proof omitted; see Appendix A, Proof of Proposition 5.1.

Definition 5.1 (FPTAS). An algorithm is a fully polynomial-time approximation scheme (FPTAS) for a maximization problem if, for each $\varepsilon \in (0, 1)$, it returns a feasible solution with value at least $(1 - \varepsilon)$ times optimum and runs in time polynomial in input size and $1/\varepsilon$.

Algorithm 1: Profit Scaling. Given $P_{\max} = \max_i p_i$ and $K = \varepsilon P_{\max}/n$. Define floored integer profits $\tilde{p}_i = \lfloor p_i/K \rfloor$. The dynamic program stores

$$D(j, v) = \min\{\text{weight needed to reach scaled profit } v \text{ using items } 1, \dots, j\}. \quad (14)$$

Return the feasible set with largest v satisfying $D(n, v) \leq B$.

Lemma 5.1 (Scaling Error). *For every item i , $K\tilde{p}_i \leq p_i < K(\tilde{p}_i + 1)$. For every subset S , $0 \leq p(S) - K\tilde{p}(S) < nK$.*

Proof omitted; see Appendix A, Proof of Lemma 5.1.

Theorem 5.1 (Profit FPTAS Guarantee). *Algorithm 1 returns a feasible set S_ε such that*

$$p(S_\varepsilon) \geq (1 - \varepsilon)\text{OPT}(B). \quad (15)$$

Its running time is polynomial in n and $1/\varepsilon$.

Proof omitted; see Appendix A, Proof of Theorem 5.1.

6 A Slippage-Minimizing FPTAS

A trader may prefer a target-return rule: achieve at least P_0 expected net score while minimizing expected slippage,

$$\text{iSLIP}(P_0) = \min_{x \in \{0,1\}^n} \left\{ \sum_{i=1}^n c_i x_i : \sum_{i=1}^n p_i x_i \geq P_0, \sum_{i=1}^n w_i x_i \leq B \right\}. \quad (16)$$

Lemma 6.1 (Feasibility Equivalence). *Problem (16) is feasible iff $\text{OPT}(B) \geq P_0$.*

Proof omitted; see Appendix A, Proof of Lemma 6.1.

Theorem 6.1 (Bicriteria Slippage FPTAS). *Assume integer weights are bounded by B , or a polynomial discretization of the resource budget. For every $\varepsilon \in (0, 1)$, there is a polynomial-time algorithm returning x^ε such that*

$$\sum_i w_i x_i^\varepsilon \leq B, \quad \sum_i p_i x_i^\varepsilon \geq (1 - \varepsilon)P_0, \quad (17)$$

and whose slippage is minimum among selections achieving the corresponding scaled profit target in the dynamic program.

Proof omitted; see Appendix A, Proof of Theorem 6.1.

Remark 6.1 (Economic Interpretation). The Bicriteria relaxation is appropriate for execution. It says that one may give up an arbitrarily small fraction of target expected return to reduce transaction-cost intensity with a polynomial-time rule.

7 General Scheduled-News Events

The preceding construction does not depend on a policy announcement. It applies to any scheduled event satisfying the filtration jump and tradability assumptions.

Definition 7.1 (Event Class). An event class $g \in \mathcal{G}$ consists of a release calendar, a signal variable $Z_{g,t}$, a pair of ordered venue times satisfying (1), an item library \mathcal{I}_g and an admissibility map \mathcal{A}_g .

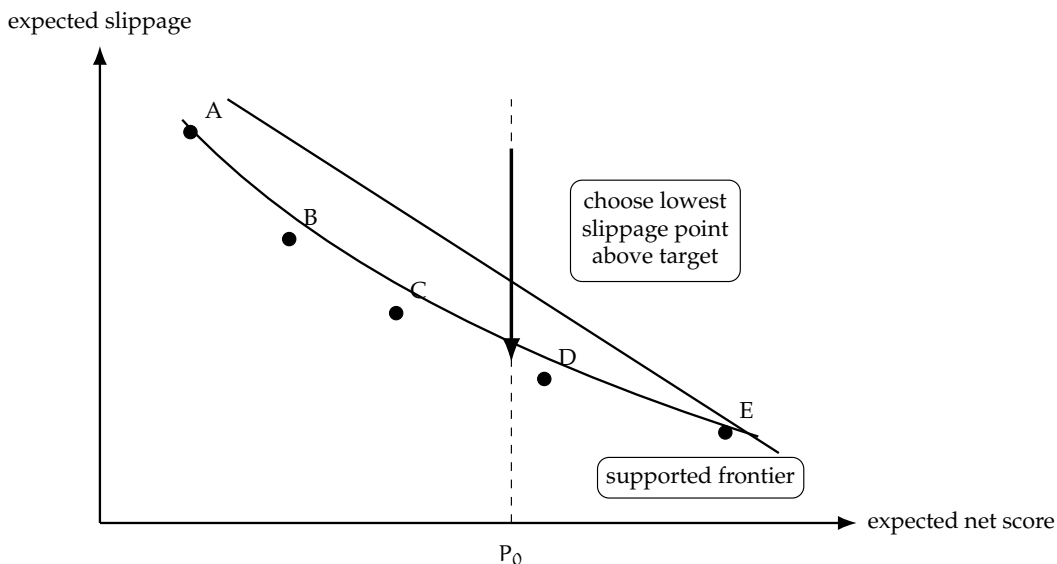


Figure 4: Slippage minimization under a target net score. The target-return formulation can prefer a lower-option item even when a larger spread has higher gross expected payoff.

Example 7.1 (Scheduled News). For a company-news event, $Z_{g,t}$ may be an earnings surprise, guidance update, merger decision, regulatory filing, product approval, or management announcement. Items may include issuer options, peer baskets, sector proxies, options with maturity constraints, or index-proxy hedges. If the release occurs when one venue is closed and another proxy remains tradable, the same item-selection problem arises.

Proposition 7.1 (Event-class invariance). *For any event class g with finite \mathcal{I}_g and a single scalar resource budget, slippage-adjusted trade selection reduces to (13) and admits the profit FPTAS.*

Proof omitted; see Appendix A, Proof of Proposition 7.1.

8 Heterogeneity and Basis-Risk

The robust-score formulation uses the standard robust-optimization device of optimizing against an uncertainty set, with box-type penalties paralleling classical robustness budgets and uncertainty-set reductions [5, 7].

Market 2 exposure is often heterogeneous. A proxy may be simply a noisy substitute for several underlying submarkets, sectors, or issuer groups.

Definition 8.1 (Response Vector and Factor Loadings). Let $q = 1, \dots, Q$ index sub-exposures within Market

2. The event response vector is $Y \in \mathbb{R}^Q$. Item i has loading vector $\beta_i \in \mathbb{R}^Q$ and its Market 2 component is $\beta_i^\top Y$.

Definition 8.2 (Robust Profit). Let \mathcal{U} be a compact set of plausible response means $m = \mathbb{E}[Y]$. The robust item profit is

$$p_i^{\text{rob}} = \inf_{m \in \mathcal{U}} \{\alpha_i + \beta_i^\top m\} - c_i, \quad (18)$$

where α_i collects all other expected components.

Lemma 8.1 (Box Ambiguity). *If $\mathcal{U} = \{m : \|m - \hat{m}\|_\infty \leq \rho\}$, then*

$$p_i^{\text{rob}} = \alpha_i + \beta_i^\top \hat{m} - \rho \|\beta_i\|_1 - c_i. \quad (19)$$

Proof omitted; see Appendix A, Proof of Lemma 8.1.

Corollary 8.1 (Robust FPTAS). *Replacing p_i by p_i^{rob} preserves the one-budget knapsack structure and the profit FPTAS guarantee.*

Proof omitted; see Appendix A, Proof of Corollary 8.1.

9 Estimation, Rounding Error and Out-of-Sample Degradation

The score p_i used by the selector is not observed directly. It is estimated from finite samples, market-impact observations, spread data, queueing measurements, proxy regressions and possibly covariance-sensitive risk summaries. Therefore two errors must be separated: statistical estimation error and numerical rounding error. The former is an out-of-sample problem; the latter is a finite-precision problem. The statistical component is precisely the out-of-sample prediction error, a selection and validation issue closely related to prediction-with-noisy-scores arguments in statistical learning [9]; leave-one-out cross-validation admits finite-sample accuracy guarantees even when the feature dimension is comparable to or larger than the sample size [42]. Following the covariance-estimation comparison in [36], Gram-type raw-moment estimators, Welford updates and Chan–Golub–LeVeque merges are algebraically equivalent in exact arithmetic but can have sharply different floating-point behavior.

Definition 9.1 (Exact, Estimated and Floating-Point Scores). Let p_i denote the true population score of item i . Let $\hat{p}_i^{\mathcal{A}}$ denote the exact-arithmetic score produced by an estimation routine $\mathcal{A} \in \{\text{Gram}, \text{Welford}, \text{CGL}\}$

and let $\hat{p}_i^{A,\text{fl}}$ denote the score actually computed in floating-point arithmetic. The total score error then yields

$$\Delta_i^A = \left| \hat{p}_i^{A,\text{fl}} - p_i \right| \leq \underbrace{\left| \hat{p}_i^A - p_i \right|}_{\text{statistical error}} + \underbrace{\left| \hat{p}_i^{A,\text{fl}} - \hat{p}_i^A \right|}_{\text{rounding error}}. \quad (20)$$

Definition 9.2 (Uniform Statistical and Rounding Envelopes). The estimated scores have statistical envelope δ and rounding envelope ρ^A if

$$\max_i |\hat{p}_i^A - p_i| \leq \delta, \quad \max_i |\hat{p}_i^{A,\text{fl}} - \hat{p}_i^A| \leq \rho^A. \quad (21)$$

Then the total uniform score error is

$$\Delta^A = \delta + \rho^A. \quad (22)$$

Definition 9.3 (Score Sensitivity to Covariance Estimation). Suppose an item score depends on a covariance or second-moment summary Σ_i through a locally Lipschitz map ϕ_i :

$$p_i = \phi_i(\Sigma_i, \zeta_i),$$

where ζ_i collects all non-covariance inputs. Let L_i^ϕ satisfy

$$|\phi_i(\Sigma, \zeta_i) - \phi_i(\Sigma', \zeta_i)| \leq L_i^\phi \|\Sigma - \Sigma'\|_{\max}. \quad (23)$$

A covariance rounding bound therefore induces a score-rounding bound.

Proposition 9.1 (Rounding Envelopes for Score Construction). *Assume the covariance component of item i is estimated from t_i observations using $\mathcal{A} \in \{\text{Gram}, \text{Welford}, \text{CGL}\}$ in double precision with unit roundoff $\varepsilon_{\text{mach}}$. Let X_i bound the absolute raw coordinates, let $\bar{x}_i = \|\mu_i\|_\infty$ denote the raw mean scale and let σ_{ik}^2 denote the relevant diagonal covariance entries. Then, up to absolute constants and the score sensitivity L_i^ϕ ,*

$$\rho_i^{\text{Gram}} \lesssim L_i^\phi \left(X_i^2 \varepsilon_{\text{mach}} + \frac{\bar{x}_i^2}{t_i - 1} \varepsilon_{\text{mach}} \right), \quad (24)$$

$$\rho_i^{\text{Welford}} \lesssim L_i^\phi \max_{k,l} (\sigma_{ik} \sigma_{il}) \varepsilon_{\text{mach}}, \quad (25)$$

$$\rho_i^{\text{CGL}} \lesssim L_i^\phi \max_{k,l} (\sigma_{ik} \sigma_{il}) \varepsilon_{\text{mach}} \log_2 t_i. \quad (26)$$

Consequently, Gram-type raw-moment estimation is sensitive to the raw data scale, whereas Welford and balanced CGL are essentially shift-stable [36].

Proof omitted; see Appendix A, Proof of Proposition 9.1.

Proposition 9.2 (Cancellation Penalty for Raw-Moment Score Estimation). *Suppose the data used to estimate an item-level covariance component satisfy*

$$\mathbf{y}_s = \mathbf{m} + \mathbf{z}_s, \quad \|\mathbf{m}\|_2 = c, \quad \mathbb{E}[\mathbf{z}_s] = 0, \quad \|\Sigma_z\|_2 = \sigma^2. \quad (27)$$

For a D-dimensional covariance summary, Gram-type raw-moment estimation can contribute a floating-point covariance error of order

$$\|\hat{\Sigma}^{\text{Gram}} - \Sigma\|_{\text{F}} \gtrsim dc^2 \varepsilon_{\text{mach}}, \quad (28)$$

whereas Welford satisfies

$$\|\hat{\Sigma}^{\text{Welford}} - \Sigma\|_{\text{F}} = O(d\sigma^2 \varepsilon_{\text{mach}}), \quad (29)$$

independently of the shift magnitude c. A balanced CGL implementation has the same centered scale with an additional logarithmic merge-depth factor. Thus, large raw offsets in prices, notionals, or feature levels should be centered, or estimated with Welford/CGL, before scores are passed to the FPTAS [36].

Proof omitted; see Appendix A, Proof of Proposition 9.2.

Theorem 9.1 (Degradation Bound with Statistical and Rounding Error). *Let $\hat{\mathcal{S}}$ be feasible and $(1 - \varepsilon)$ -optimal under the floating-point estimated scores $\hat{p}_i^{\text{A,fl}}$. Let \mathcal{S}^* be truly optimal under the population scores p_i . If every feasible set contains at most r items and*

$$\max_i |\hat{p}_i^{\text{A,fl}} - p_i| \leq \Delta^{\text{A}},$$

then

$$p(\hat{\mathcal{S}}) \geq (1 - \varepsilon)p(\mathcal{S}^*) - (2 - \varepsilon)r\Delta^{\text{A}}. \quad (30)$$

In particular, under (22),

$$p(\hat{\mathcal{S}}) \geq (1 - \varepsilon)p(\mathcal{S}^*) - (2 - \varepsilon)r(\delta + \rho^{\text{A}}). \quad (31)$$

Proof omitted; see Appendix A, Proof of Theorem 9.1.

Corollary 9.1 (Numerically Stable Validation Requirement). *An approximate combinatorial selector is useful only when the estimated edge dominates both the statistical validation penalty and the finite-precision rounding penalty. For raw price-level or notional-level features with large shifts, a Gram-type raw-moment estimator may inflate ρ^{A} enough to dominate the FPTAS approximation error. Centered Welford or balanced CGL estimation is therefore preferable whenever item scores depend on second moments or covariance-sensitive risk penalties.*

10 Conclusions

This paper treats scheduled-news trading between asynchronous venues as a constrained selection problem. A signal Z_t may arrive after Market 2 has closed but before Market 1, or a proxy venue, has stopped trading. That timing gap matters, but it is not an arbitrage by itself. The direct Market 2 price is no longer executable and the remaining trades must pass through spreads, impact, latency, basis risk and limited trading resources.

Available post-signal trades are modeled as a finite item library, with each item carrying a net score, resource demand, slippage term, execution protocol and possible proxy exposure. Queueing delay enters the same item representation: a trade can be removed if its deadline risk is too large, or its score can be reduced by a deterministic latency penalty. With one scalar resource, the resulting selection problem is exactly a 0–1 knapsack problem. Profit scaling therefore gives an FPTAS for the slippage- and latency-adjusted objective. The target-return version gives the corresponding approximation scheme for minimizing slippage subject to a required expected score.

The same structure survives the main extensions. Basis-risk uncertainty and heterogeneous Market 2 exposure can be handled through robust item scores. Statistical estimation error and floating-point rounding error enter as explicit loss terms in the selected portfolio’s true value. This is especially important when scores use covariance or second-moment estimates. In large-shift data, raw Gram-style moment estimation can magnify cancellation error, while centered Welford updates and balanced Chan–Golub–LeVeque merges keep the arithmetic at the covariance scale.

The main conclusion is that the timing gap is useful only after it is made executable. Once execution, proxy mismatch, latency and estimation error are accounted for, the problem is not a free lunch but a finite, noisy, resource-constrained choice problem. The value of the framework is that these frictions remain visible while the final selection step still has a clean approximation guarantee.

Copyright Statement

F. Reichel, 2026. All rights reserved. Original work of the author; any AI use was auxiliary (editing, typesetting, wording) and does not affect authorship or ownership. The author expressly reserves all text- and data-mining rights, including use of this work as input for training, fine-tuning, benchmarking, or evaluating AI systems, under Art. 4(3) of Directive (EU) 2019/790. Any reuse requires prior written consent

and attribution to the author.

References

- [1] Almgren, R. and Chriss, N. (2001). Optimal execution of portfolio transactions. *The Journal of Risk*, 3(2), 5–39. doi:10.21314/JOR.2001.041.
- [2] Ball, R. and Brown, P. (1968). An empirical evaluation of accounting income numbers. *Journal of Accounting Research*, 6(2), 159–178.
- [3] Beaver, W. H. (1968). The information content of annual earnings announcements. *Journal of Accounting Research*, 6, 67–92.
- [4] Bernanke, B. S. and Kuttner, K. N. (2005). What explains the stock market’s reaction to Federal Reserve policy? *Journal of Finance*, 60(3), 1221–1257.
- [5] Ben-Tal, A., El Ghaoui, L. and Nemirovski, A. (2009). *Robust Optimization*. Princeton University Press.
- [6] Bertsimas, D. and Lo, A. W. (1998). Optimal control of execution costs. *Journal of Financial Markets*, 1(1), 1–50.
- [7] Bertsimas, D. and Sim, M. (2004). The price of robustness. *Operations Research*, 52(1), 35–53.
- [8] Campbell, J. Y., Lo, A. W. and MacKinlay, A. C. (1997). *The Econometrics of Financial Markets*. Princeton University Press.
- [9] Cesa-Bianchi, N. and Lugosi, G. (2006). *Prediction, Learning and Games*. Cambridge University Press.
- [10] Ciucu, F. (2007). Network calculus delay bounds in queueing networks with exact solutions. In *Managing Traffic Performance in Converged Networks: 20th International Teletraffic Congress, ITC 2007*, 495–506. Springer. doi:10.1007/978-3-540-72990-7_45.
- [11] Ciucu, F., Poloczek, F. and Schmitt, J. (2013). Sharp bounds in stochastic network calculus. *ACM SIGMETRICS Performance Evaluation Review*, 41(1), 101–112. doi:10.1145/2494232.2465746.
- [12] Ciucu, F., Mehri, S. and Rizk, A. (2024). On ultra-sharp queueing bounds. In *Proceedings of IEEE INFOCOM*, 1980–1988.
- [13] Ciucu, F. and Mehri, S. (2025). On the distribution of sojourn times in tandem queues. *Proceedings of the ACM on Measurement and Analysis of Computing Systems*, 9(2), 1–34. doi:10.1145/3727140.
- [14] Delbaen, F. and Schachermayer, W. (1994). A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300, 463–520.

- [15] Duffie, D. (2001). *Dynamic Asset Pricing Theory*. Princeton University Press.
- [16] Easley, D., Lopez de Prado, M. M. and O'Hara, M. (2012). Flow toxicity and liquidity in a high-frequency world. *Review of Financial Studies*, 25(5), 1457–1493.
- [17] Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2), 383–417.
- [18] Fama, E. F. (1991). Efficient capital markets II. *Journal of Finance*, 46(5), 1575–1617.
- [19] Glosten, L. R. and Milgrom, P. R. (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1), 71–100.
- [20] Gromb, D. and Vayanos, D. (2010). Limits of arbitrage: The state of the theory. *Annual Review of Financial Economics*, 2, 251–275.
- [21] Gürkaynak, R. S., Sack, B. and Swanson, E. T. (2005). Do actions speak louder than words? The response of asset prices to monetary policy actions and statements. *International Journal of Central Banking*, 1(1), 55–93.
- [22] Harris, L. (2003). *Trading and Exchanges: Market Microstructure for Practitioners*. Oxford University Press.
- [23] Harrison, J. M. and Kreps, D. M. (1979). Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory*, 20(3), 381–408.
- [24] Hasbrouck, J. (2007). *Empirical Market Microstructure: The Institutions, Economics and Econometrics of Securities Trading*. Oxford University Press.
- [25] Ibarra, O. H. and Kim, C. E. (1975). Fast approximation algorithms for the knapsack and sum of subset problems. *Journal of the ACM*, 22(4), 463–468.
- [26] Jensen, M. C. (1978). Some anomalous evidence regarding market efficiency. *Journal of Financial Economics*, 6(2–3), 95–101.
- [27] Kellerer, H., Pferschy, U. and Pisinger, D. (2004). *Knapsack Problems*. Springer.
- [28] Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335.
- [29] Lawler, E. L. (1979). Fast approximation algorithms for knapsack problems. *Mathematics of Operations Research*, 4(4), 339–356.
- [30] Lo, A. W. (2004). The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management*, 30(5), 15–29.

- [31] Jonung, L. (Ed.) (1991). *The Stockholm School of Economics Revisited*. Cambridge University Press.
- [32] MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of Economic Literature*, 35(1), 13–39.
- [33] New York Stock Exchange. (2026). *Trading Information, Order Types and Auction Procedures*. Official NYSE documentation. Retrieved June 13, 2026.
- [34] O’Hara, M. (1995). *Market Microstructure Theory*. Blackwell.
- [35] Patell, J. M. and Wolfson, M. A. (1984). The intraday speed of adjustment of stock prices to earnings and dividend announcements. *Journal of Financial Economics*, 13(2), 223–252.
- [36] Reichel, F. (2026). *2B or Not 2B: A Tale of Three Algorithms for Streaming: Covariance Estimation after Welford and Chan–Golub–LeVeque*. arXiv:2605.00247v1 [stat.CO].
- [37] Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 13(3), 341–360.
- [38] Securities and Exchange Commission. (2005). Regulation NMS; Final Rule, Exchange Act Release No. 34-51808, including Rule 611, the Order Protection Rule.
- [39] Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, 52(1), 35–55.
- [40] Vazirani, V. V. (2001). *Approximation Algorithms*. Springer.
- [41] Williamson, D. P. and Shmoys, D. B. (2011). *The Design of Approximation Algorithms*. Cambridge University Press.
- [42] Rahnema Rad, K., Zhou, W. and Maleki, A. (2020). Error bounds in estimating the out-of-sample prediction error using leave-one-out cross validation in high-dimensions. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics (AISTATS)*, PMLR 108, 4067–4077. arXiv:2003.01770.

A Main-text Omitted Proofs

This appendix collects the proofs omitted from the main text.

Proof of Lemma 2.1 (Finite Slippage Envelope)

Proof. The item library and option set are finite. Half-spreads, fees, borrow or financing schedules and synchronization penalties are finite under admissibility. Convex impact is continuous on the compact interval $[-\bar{q}, \bar{q}]$ and is therefore bounded. Summing finitely many bounded nonnegative terms gives \bar{c}_i . \square

Proof of Proposition 2.1 (Option Parsimony Weakly Lowers Deterministic Slippage)

Proof. Expected slippage is a sum of nonnegative option-level and synchronization terms. Item i contains the common terms and omits at least one nonnegative term from j . Therefore its sum is weakly smaller and strictly smaller when an omitted expected term is positive. \square

Proof of Proposition 4.1 (No Pure Arbitrage From Asynchronous News Alone)

Proof. A pure arbitrage requires a self-financing admissible trade with nonnegative payoff in all states and a positive payoff in at least one state. The stale direct Market 2 price cannot be executed after the signal is observed. The trader can only use Market 1 instruments, proxy instruments, or derivative instruments whose post-signal prices may already reflect Z_t . Any proxy trade has basis risk, liquidity risk and state-dependent payoff. Therefore the timing wedge alone does not create a statewise nonnegative payoff. It can create an estimated positive expected edge, but not pure arbitrage. \square

Proof of Proposition 5.1 (Exact Reduction)

Proof. Map each candidate trade to a knapsack item, its estimated net score to profit p_i and its scarce resource usage to weight w_i . A feasible trading schedule is exactly a subset whose total weight is at most B and the objective is the sum of selected profits. This is (13). \square

Proof of Lemma 5.1 (Scaling Error)

Proof. The itemwise inequality follows from the floor definition. Summing the upper rounding error, which is less than K per selected item, over at most n items gives the subset inequality. \square

Proof of Theorem 5.1 (Profit FPTAS Guarantee)

Proof. Let S^* be optimal. Exact dynamic programming in scaled profits gives $\tilde{p}(S_\epsilon) \geq \tilde{p}(S^*)$ among feasible sets. Therefore

$$p(S_\epsilon) \geq K\tilde{p}(S_\epsilon) \geq K\tilde{p}(S^*) > p(S^*) - nK. \quad (32)$$

Because $P_{\max} \leq \text{OPT}(B)$ whenever a positive-profit item is feasible, $nK = \epsilon P_{\max} \leq \epsilon \text{OPT}(B)$. Therefore the returned value is at least $(1 - \epsilon)\text{OPT}(B)$. The total scaled profit is at most $nP_{\max}/K = n^2/\epsilon$, so the dynamic program has polynomial state space in n and $1/\epsilon$. \square

Proof of Lemma 6.1 (Feasibility Equivalence)

Proof. If (16) is feasible, a budget-feasible set reaches P_0 , so the maximum budget-feasible profit is at least P_0 . Conversely, if $\text{OPT}(B) \geq P_0$ an optimizer of (13) is feasible for (16). \square

Proof of Theorem 6.1 (Bicriteria Slippage FPTAS)

Proof. Scale profits by $K = \epsilon P_{\max}/n$ and define $\tilde{P}_0 = \lceil P_0/K \rceil$. A dynamic program over item index, scaled profit and resource state minimizes total slippage cost. The scaling lemma implies that reducing the scaled target by at most n units loses at most $nK = \epsilon P_{\max}$. If $P_{\max} \leq P_0$, the true profit is at least $(1 - \epsilon)P_0$. The dynamic program minimizes slippage exactly over the scaled target class and budget states. The number of scaled profit states is $O(n^2/\epsilon)$ and the resource state space is polynomial given assumption, so the algorithm is polynomial. \square

Proof of Proposition 7.1 (Event-Class Invariance)

Proof. The proof does not use the economic identity of the signal. It uses only finiteness of the item library, deterministic estimated scores, nonnegative resource weights and a scalar budget. These properties hold by definition for event class g . \square

Proof of Lemma 8.1 (Box Ambiguity)

Proof. The infimum of $\beta_i^\top m$ over an ℓ_∞ box is achieved by shifting each coordinate against the sign of β_{iq} . Therefore $\inf_m \beta_i^\top m = \beta_i^\top \hat{m} - \rho \sum_q |\beta_{iq}|$. Substitution gives the result. \square

Proof of Corollary 8.1 (Robust FPTAS)

Proof. After robustification, each item still has a deterministic scalar profit score. Therefore the scaled-profit dynamic program applies unchanged. \square

Proof of the Degradation Bound (Statistical Error Only)

Proof. For any feasible S with at most r items, $|\hat{p}(S) - p(S)| \leq r\delta$. Let \hat{S}^* maximize estimated score. Since \hat{S} is $(1 - \varepsilon)$ -optimal under estimates,

$$\hat{p}(\hat{S}) \geq (1 - \varepsilon)\hat{p}(\hat{S}^*) \geq (1 - \varepsilon)\hat{p}(S^*). \quad (33)$$

Using the uniform error twice,

$$p(\hat{S}) \geq \hat{p}(\hat{S}) - r\delta \quad (34)$$

$$\geq (1 - \varepsilon)\hat{p}(S^*) - r\delta \quad (35)$$

$$\geq (1 - \varepsilon)(p(S^*) - r\delta) - r\delta, \quad (36)$$

which rearranges to the claim. \square

Proof of Proposition (Constraint Encoding)

Proof. If a restriction makes an item impossible, remove it. If it makes an item more expensive or resource intensive, increase its slippage score, weight, or constraint coefficients. The resulting problem remains a finite item-selection problem. \square

Proof of Proposition 9.1 (Rounding Envelopes for Score Construction)

Proof. The covariance-estimation bounds in [36] distinguish between raw-moment estimation and centered updating. For a Gram-type raw-moment estimator, each covariance entry is obtained from a difference between a raw second-moment term and an outer product of raw sums. In finite precision, this produces a rounding contribution at the raw data scale. Thus, for item i , if X_i bounds the absolute raw coordinates and $\bar{x}_i = \|\mu_i\|_\infty$ denotes the raw mean scale, the entrywise covariance rounding error satisfies, up to absolute constants,

$$\|\hat{\Sigma}_i^{\text{Gram,fl}} - \hat{\Sigma}_i^{\text{Gram}}\|_{\max} \lesssim X_i^2 \varepsilon_{\text{mach}} + \frac{\bar{x}_i^2}{t_i - 1} \varepsilon_{\text{mach}}. \quad (37)$$

By contrast, Welford updates are formed from residuals around the running mean. Their arithmetic therefore takes place at centered covariance scale rather than raw data scale. Therefore

$$\|\hat{\Sigma}_i^{\text{Welford,fl}} - \hat{\Sigma}_i^{\text{Welford}}\|_{\max} \lesssim \max_{k,l}(\sigma_{ik}\sigma_{il}) \varepsilon_{\text{mach}}. \quad (38)$$

For a balanced Chan–Golub–LeVeque merge tree, the same centered-scale bound applies at each merge level and the tree depth contributes a logarithmic factor:

$$\|\hat{\Sigma}_i^{\text{CGL,fl}} - \hat{\Sigma}_i^{\text{CGL}}\|_{\max} \lesssim \max_{k,l}(\sigma_{ik}\sigma_{il}) \varepsilon_{\text{mach}} \log_2 t_i. \quad (39)$$

Now suppose that the item score depends on the covariance summary through the locally Lipschitz map ϕ_i in (23). Then

$$|\hat{p}_i^{\mathcal{A},\text{fl}} - \hat{p}_i^{\mathcal{A}}| \leq L_i^\phi \|\hat{\Sigma}_i^{\mathcal{A},\text{fl}} - \hat{\Sigma}_i^{\mathcal{A}}\|_{\max}. \quad (40)$$

Applying this inequality to the three covariance rounding bounds above gives (24), (25) and (26). This proves the claimed score-level rounding envelopes. \square

Proof of Proposition 9.2 (Cancellation Penalty for Raw-Moment Score Estimation)

Proof. Write the observations used in the covariance component as

$$\mathbf{y}_s = \mathbf{m} + \mathbf{z}_s, \quad \|\mathbf{m}\|_2 = c, \quad \mathbb{E}[\mathbf{z}_s] = 0, \quad \|\Sigma_z\|_2 = \sigma^2. \quad (41)$$

A Gram-type covariance estimator computes a normalized difference between a raw second-moment matrix and an outer product of raw sums. Before normalization, both terms contain leading contributions of order

$$t_i^2 m m^\top. \quad (42)$$

In exact arithmetic, these leading raw-location terms cancel, leaving the centered covariance contribution generated by the z_s terms. In floating-point arithmetic, however, the large raw-location terms are rounded before cancellation. The residual rounding error is therefore controlled by the raw shift scale rather than the centered covariance scale.

Entrywise, the cancellation-sensitive contribution is of order

$$|m_k m_l| \varepsilon_{\text{mach}}. \quad (43)$$

For a d -dimensional covariance summary, summing these entrywise contributions in Frobenius scale gives the lower-envelope behavior

$$\|\hat{\Sigma}^{\text{Gram}} - \Sigma\|_F \gtrsim d c^2 \varepsilon_{\text{mach}}, \quad (44)$$

up to absolute constants.

For Welford, the update is based on residuals relative to the running mean. Consequently, the dominant arithmetic is performed on quantities of centered scale z_s , not on the shifted raw level $m + z_s$. The resulting Frobenius-scale error is therefore

$$\|\hat{\Sigma}^{\text{Welford}} - \Sigma\|_F = O(d \sigma^2 \varepsilon_{\text{mach}}), \quad (45)$$

which is independent of the shift magnitude c . A balanced Chan–Golub–LeVeque implementation uses centered block summaries and centered merge corrections, so it has the same shift-stable scale, with only an additional logarithmic factor from the merge depth. This proves the stated cancellation comparison. \square

Proof of Theorem 9.1 (Degradation Bound with Statistical and Rounding Error)

Proof. Given

$$\hat{p}^{\mathcal{A}, \Pi}(S) = \sum_{i \in S} \hat{p}_i^{\mathcal{A}, \Pi}, \quad p(S) = \sum_{i \in S} p_i. \quad (46)$$

Given assumption,

$$\max_i |\hat{p}_i^{A,fl} - p_i| \leq \Delta^A. \quad (47)$$

Since every feasible set contains at most r items, every feasible set S satisfies

$$p(S) - r\Delta^A \leq \hat{p}^{A,fl}(S) \leq p(S) + r\Delta^A. \quad (48)$$

Additionally,

$$p(S) \geq \hat{p}^{A,fl}(S) - r\Delta^A, \quad (49)$$

and

$$\hat{p}^{A,fl}(S) \geq p(S) - r\Delta^A. \quad (50)$$

Because \hat{S} is $(1 - \varepsilon)$ -optimal under the floating-point estimated scores,

$$\hat{p}^{A,fl}(\hat{S}) \geq (1 - \varepsilon)\hat{p}^{A,fl}(S^*). \quad (51)$$

Using (49) for \hat{S} gives

$$p(\hat{S}) \geq \hat{p}^{A,fl}(\hat{S}) - r\Delta^A. \quad (52)$$

Combination yields,

$$p(\hat{S}) \geq (1 - \varepsilon)\hat{p}^{A,fl}(S^*) - r\Delta^A. \quad (53)$$

Now apply (50) to S^* :

$$\hat{p}^{A,fl}(S^*) \geq p(S^*) - r\Delta^A. \quad (54)$$

Therefore,

$$p(\hat{S}) \geq (1 - \varepsilon)(p(S^*) - r\Delta^A) - r\Delta^A. \quad (55)$$

Expanding,

$$p(\hat{S}) \geq (1 - \varepsilon)p(S^*) - (1 - \varepsilon)r\Delta^A - r\Delta^A. \quad (56)$$

Therefore

$$p(\hat{S}) \geq (1 - \varepsilon)p(S^*) - (2 - \varepsilon)r\Delta^A. \quad (57)$$

Finally, if $\Delta^A = \delta + \rho^A$, then

$$p(\hat{S}) \geq (1 - \varepsilon)p(S^*) - (2 - \varepsilon)r(\delta + \rho^A), \quad (58)$$

which proves the expanded bound.



AI-assisted verification statement. In line with the transparency expectations of the EU AI Act (Regulation (EU) 2024/1689) and the European Commission's guidelines on the responsible use of generative AI in research, we disclose that Claude Opus 4.8 (Anthropic), with SymPy 1.14.0, NumPy and SciPy, was used solely as an assistive tool: to symbolically verify the algebraic identities in this appendix and to run automated falsification tests of the remaining claims (all these 34 checks passed). These tests provide evidence and regression guards, not formal proofs of the universally generalized claims in this paper. The AI system is not an author; the authors directed and reviewed all outputs and retain sole and full responsibility for the correctness of the results.