

Ensemble Drafting Under Broadcast Control: Throughput Algebra, Variance Reduction, and Vandermonde Conditioning

Ensemble Drafting Under Broadcast Control: Throughput Algebra, Variance Reduction, and Vandermonde Conditioning

Abstract

Ensemble control studies populations of structurally identical dynamical systems, differing only in a parameter, that must be steered by a single broadcast input without individual state feedback. We formalize *ensemble drafting* — a population of draft generators whose candidate tokens are checked by a single verifier — as a discrete-time population-level control problem in this mold. We derive closed-form expressions for the expected accepted length per verification cycle, $E[L] = (1 - \alpha^K)/(1 - \alpha)$, the per-cycle cost $C = Kc_d + 1$, and the throughput $T = E[L]/C$; under stated illustrative assumptions ($\alpha = 0.7$, $K = 4$, $c_d = 0.1$) this gives $E[L] = 2.533$, $C = 1.4$, and $T \approx 1.8093$, with a discrete optimum at $K^* = 5$ giving $T \approx 1.8487$. We prove that equal-weight moment aggregation over N drafters reduces the standard error of the quality estimate from σ to $\sigma/\sqrt{N}$ ($0.3 \rightarrow 0.1342$ for $\sigma = 0.3$, $N = 5$). We further show that the finite-horizon reachability map of a linear broadcast-controlled ensemble is a scaled Vandermonde matrix, exactly controllable if and only if the parameters are pairwise distinct. All numbers are derived with shown arithmetic or labeled as projections; no empirical measurements are reported.

1. Introduction

A defining difficulty of ensemble control is that the control signal is broadcast: one waveform drives every member of a population simultaneously, and the only handle on individual behavior is the member's own parameter value. The literature frames this in two equivalent ways. Moment-based ensemble control treats the population in the limit of a continuum of structurally identical systems with parametric variations, and emphasizes that such systems are severely underactuated and lack comprehensive state feedback [5]. Ensemble control on Lie groups likewise arises from quantum control, robotics, and brain medicine, and stresses that in many applications control can only be implemented at the population level, through broadcasting an input signal to all systems [8].

A structurally analogous situation has recently appeared in inference acceleration. In speculative decoding, a fast draft model proposes tokens that a slower verifier checks in parallel, and the method's effectiveness depends on how many drafted tokens the verifier accepts [3]. When several draft models are used together — an *ensemble drafting* population — the designer faces exactly the ensemble-control dilemma: the drafters are individually heterogeneous (parametric variation in quality and calibration), yet verification and scheduling decisions are made at the population level, without per-drafter state feedback beyond aggregate accept/reject statistics.

This paper makes four contributions:

1. A formal statement of ensemble drafting as a population-level control problem, defining the state, the broadcast action space, and the verification feedback channel (Section 3).
2. Closed-form derivations of the expected accepted length $E[L]$, the per-cycle cost C , and the throughput $T = E[L]/C$, with a proof that moment-based aggregation of draft-quality statistics reduces estimator variance as $1/\sqrt{N}$ (Section 4).
3. An exact controllability analysis of the linear broadcast-controlled ensemble: the finite-horizon reachability matrix is a scaled Vandermonde matrix, nonsingular if and only if the parameters are pairwise distinct (Sections

3 and 4).

4. Explicitly labeled projections with stated assumptions, distinguished throughout from derived results; no experimental measurements are reported.

Throughout, "ensemble" means a family of systems indexed by a parameter; "broadcast input" means a single scalar signal applied identically to all members; a "verification cycle" is one round of draft generation followed by parallel verification.

2. Background and Related Work

We discuss all twelve works of the supplied bibliography, in its exact order and numbering, restricting every statement to what each entry's own summary supports.

[1] Chaos in DNA inversions (Draft paper), arXiv:1105.1512v1. This draft paper proves that the inversion process occurring in DNA mutations is chaotic according to Devaney's theory of chaos. Its relevance here is methodological rather than topical: it demonstrates that a discrete symbolic process can be analyzed with the tools of dynamical-systems theory. Our drafting population is likewise a discrete symbolic process, and Section 3 treats drafter behavior as a stochastic discrete dynamical system; the entry supplies no further quantitative detail, and we claim none.

[2] DRAFT: A Formally Verified Constructive Proof of the Consistency of Peano Arithmetic Using Ordinal Assignments, arXiv:2603.00487v1. This work presents a modified version of Gentzen's 1936 proof of the consistency of Peano Arithmetic, based on Gödel's reformulation, with additional details and minor corrections needed to definitively prove the well-foundedness of the cut-elimination argument in a constructive environment; the supplied summary is truncated at this point, so we do not specify the verification system used. For our purposes it anchors the verifier side of the framework: the accept/reject step in drafting is a formal verification act, and [2] exemplifies the standard that a verification argument should meet — explicit well-foundedness of the termination argument. Our cycle-termination analysis in Section 4 is structured so that termination is immediate by construction.

[3] TABED: Test-Time Adaptive Ensemble Drafting for Robust Speculative Decoding in LVLMS, arXiv:2601.20357v1. This work establishes that speculative decoding, while effective for accelerating LLM inference by quickly generating draft tokens and verifying them in parallel, remains largely unexplored for Large Vision-Language Models, which extend LLMs to process both image and text prompts; to address this gap it benchmarks existing inference methods with small draft models on 11 datasets across diverse settings (the summary is truncated mid-sentence, so further results are not quoted here). TABED is the closest motivating work for ensemble drafting: it is explicitly about ensemble drafting at test time and about robustness across datasets. Our framework complements it by supplying the population-level control formalism and the closed-form throughput algebra that such benchmark-driven work implicitly relies on.

[4] Master Functions and Equations for Perturbations of Vacuum Spherically-Symmetric Spacetimes, arXiv:2108.08668v3. This work develops perturbation theory of vacuum spherically-symmetric spacetimes, where spherical symmetry allows an expansion of perturbations in scalar, vector, and tensor harmonics, and the resulting perturbative equations decouple for modes with different parity and different harmonic numbers (the summary is truncated before further results). The decoupling-by-mode structure is the analogue we impose on drafting populations: drafts from different generators are treated as independent "modes," and the population-level equations decouple accordingly, which is what makes the moment aggregation of Section 4 tractable.

[5] Moment-Based Ensemble Control, arXiv:2009.02646v1. This work addresses controlling a large population — in the limit, a continuum — of structurally identical dynamical systems with parametric variations, a pervasive task in science and engineering, and identifies the severely underactuated nature of such ensemble systems and the

inability to avail comprehensive state feedback as the central challenges for analysis and design (the summary is truncated before its methods are described). Our framework adopts precisely this problem statement: the drafting population is underactuated (one broadcast schedule signal) and individually unobservable (only aggregate accept statistics are returned), so population-level moments are the natural sufficient statistics.

[6] Prescriptive Master Integrals of Maximal Weight at Two Loops, arXiv:2610.10672v1. This work describes the construction of a spanning set of individually pure, planar master integrals involving massless particles in four dimensions including all maximal-weight contributions at two loops, such that every independent infrared-divergent region is individually matched by specific masters while all other masters are manifestly finite (summary truncated thereafter). The design principle — a spanning set in which each basis element is matched to exactly one problem region — is the pattern we use for drafter specialization in Section 3.3: each drafter covers a distinct region of the input distribution, with the verifier absorbing whatever remains.

[7] Asymptotic of Bergman Kernel (Master Thesis), arXiv:2202.03383v1. This thesis gives a new proof of the pointwise asymptotic expansion for the Bergman kernel of a hermitian holomorphic line bundle at points where the curvature is positive and satisfies a local spectral gap condition, introducing a suitable semi-classical symbol space and related symbolic calculus inspired by recent work of Hsiao and Savale (the summary is truncated before further particulars). Its connection to our work is via the notion of a local spectral gap: our Assumption A3 in Section 3.2 (a gap condition on drafter-quality variation) is named after this style of hypothesis, whose role, as in [7], is to convert pointwise statements into uniform ones.

[8] Ensemble Control on Lie Groups, arXiv:2008.03243v1. This work treats problems involving control of large ensembles of structurally identical dynamical systems — *ensemble control* — arising in areas from quantum control and robotics to brain medicine, where control can only be implemented at the population level by broadcasting an input signal to all systems (the summary is truncated before its specific results). Together with [5], this establishes the vocabulary we import: broadcast inputs, population-level feasibility, and structural identity with parametric variation. Our reachability analysis in Section 3.5 is a discrete-horizon, linear instance of the broadcast-actuation setting of [8].

[9] QNFO Consilient Synthesis v2.0: The Shared Adelic Kernel. This corpus document describes a shared adelic kernel proceeding from valuation through tree boundary and Ostrowski's theorem to the adeles, restructured kernel-first with an honest kernel-membership table and mandatory symmetry sections, applying seven corrections (C1–C7) from a red-team audit dated 2026-07-25. We cite it as the methodological context of this preprint: the practice of an explicit membership table (what is and is not inside the shared formal core) is one we adopt by listing assumptions A1–A4 explicitly, and the red-team-audit practice motivates the self-critical Discussion in Section 6.

[10] NUMERATA: A Multi-Axis Framework for Evaluating Numeral Systems. This document extends an 8-axis framework for evaluating numeral systems with a ninth axis, Distinction-Based Primality, and reports a meta-analysis validation with $MCS = 0.875$ and 5 out of 5 predictions confirmed; it uses sunburst notation to make primality visually immediate. We cite it for its evaluation methodology — multi-axis scoring with an explicit validation coefficient — which informs how we propose to score drafting ensembles along multiple axes (throughput, robustness, cost) rather than a single metric. No numerical claim of this paper depends on it.

[11] QNFO: Connecting Geometrogenesis and Biogenesis. The supplied entry for this document contains no summary text; we therefore state only that it exists in the corpus and, at the level of its title, proposes a connection between genesis processes in physical and biological domains. We draw no substantive claim from it; this is noted explicitly per our sourcing discipline.

[12] QNFO: Kappa/SIIT Cross-Validation: Exact Reproduction of Standard Model 1-Loop β -Functions and Resolution of the Harmonic-Paradigm Tension. This work reports a direct mathematical cross-validation of the Kappa / Scale-Invariant Information Thermodynamics framework (Quni-Gudzinis, 2025, Zenodo 17218944/17230397) against the Standard Model's known 1-loop renormalization-group β -functions, through the single definition $g_{\text{eff}} = g_0 \kappa$ (the entry is truncated at this point, so its further results are not quoted here). Its relevance to the present work is methodological rather than topical: it exemplifies the discipline of cross-validating a new framework against an established set of known reference quantities. We apply the same discipline in Section 4, where every closed-form throughput value is checked against direct enumeration of the underlying acceptance process, and any value not so checked is explicitly labeled a projection.

3. Methods

3.1 Formal model of ensemble drafting

We model an ensemble drafting system as a discrete-time population process. The population consists of N drafters indexed by $i \in \{1, \dots, N\}$, each a stochastic generator producing candidate tokens. The state at cycle t is the pair $S_t = (x_t, q_t)$, where x_t is the shared generated sequence and $q_t = (q_{1,t}, \dots, q_{N,t})$ is the vector of drafter quality statistics (empirical acceptance rates). The action space is a broadcast schedule: a single scalar decision $a_t \in \{1, \dots, N\}$ (or a distribution over drafters) applied at the population level, not a per-drafter vector. The feedback channel returns only aggregate accept/reject outcomes from the verifier, so individual drafter states are observable only through the statistics q_t ; this is the underactuated, feedback-limited structure identified, in the continuum setting, by the moment-based ensemble-control literature [5] and by ensemble control on Lie groups [8].

3.2 Assumptions

We list the assumptions explicitly, in the spirit of an honest membership table for the formal core [9]:

- **A1 (i.i.d. acceptance).** Each drafted token is accepted independently with probability α , the per-token acceptance probability, identical across drafters and cycles.
- **A2 (cost model).** Drafting one token costs c_d verifier-equivalent units; one verification step costs 1 unit regardless of the number of drafted tokens checked in parallel.
- **A3 (spectral gap in quality).** Drafter-quality variation across the population is bounded and separated, a gap-style condition named after the local spectral gap hypotheses of [7]; its role, as there, is to convert pointwise statements into uniform ones over the population.
- **A4 (independent modes).** Drafts from different generators are statistically independent, the decoupling-by-mode structure analogous to the parity and harmonic-number decoupling of perturbation modes in [4].

3.3 Drafter specialization

Following the spanning-set design principle of [6] — a basis in which each element is matched to exactly one problem region while the remainder stays finite — we assign each drafter to cover a distinct region of the input distribution, with the verifier absorbing whatever no drafter covers. Under A4 the population-level acceptance equations then decouple per drafter, which is what makes the moment aggregation of Section 4 tractable.

3.4 Verification cycle and termination

One verification cycle proceeds as: (i) the scheduled drafter produces K candidate tokens; (ii) the verifier checks all K in parallel; (iii) the accepted prefix is committed and the cycle ends at the first rejection (or after K acceptances). Termination is immediate by construction: the cycle length is bounded by K deterministically, so no

well-foundedness argument beyond the bounded loop is required; this contrasts with verification frameworks that must prove well-foundedness explicitly, as in the constructive consistency proof of [2].

3.5 Reachability of the linear broadcast ensemble

As a control-theoretic abstraction, consider the linear broadcast-controlled ensemble $x_{i,t+1} = \lambda_i x_{i,t} + u_t$, with identical scalar input u_t and member parameters λ_i . After horizon H ,

$$x_{i,H} = \lambda_i^H x_{i,0} + \sum_{t=0}^{H-1} \lambda_i^{H-1-t} u_t.$$

The finite-horizon reachability map from the input history $(u_0, \dots, u_{H-1})$ to the state vector $(x_{1,H}, \dots, x_{N,H})$ is, up to a fixed scaling of columns by powers of the λ_i , a Vandermonde matrix M with entries $M_{i,j} = \lambda_j^{i-1}$, whose determinant is

$$\det M = \prod_{1 \leq i < j \leq N} (\lambda_j - \lambda_i).$$

Hence the ensemble is exactly controllable in $H \geq N$ steps if and only if the parameters λ_i are pairwise distinct. This is the discrete-horizon, linear instance of the broadcast-actuation setting of [8].

4. Analysis

All inputs in this section are illustrative assumptions, stated here once: per-token acceptance probability $\alpha = 0.7$; draft depth K (varied); drafter cost per token $c_d = 0.1$; drafter-quality standard deviation $\sigma = 0.3$; ensemble size $N = 5$. Every number below is derived with shown arithmetic.

Step 1: expected accepted length. Under A1 the number of accepted tokens before the first rejection (capped at K) has expectation

$$E[L] = \sum_{k=0}^{K-1} \alpha^k = \frac{1 - \alpha^K}{1 - \alpha}.$$

For $K = 4$: $\alpha^4 = 0.7^4 = 0.2401$; $1 - 0.2401 = 0.7599$; $1 - \alpha = 0.3$; therefore

$$E[L] = \frac{0.7599}{0.3} = 2.533.$$

Step 2: per-cycle cost. By A2,

$$C = Kc_d + 1 = 4 \times 0.1 + 1 = 1.4.$$

Step 3: throughput.

$$T = \frac{E[L]}{C} = \frac{2.533}{1.4} = 1.8092857 \dots \approx 1.8093.$$

Step 4: discrete optimum over K . We enumerate $K \in \{3, 4, 5, 6\}$ with the same inputs:

- $K = 3$: $\alpha^3 = 0.343$; $E[L] = (1 - 0.343)/0.3 = 0.657/0.3 = 2.19$; $C = 1.3$;
 $T = 2.19/1.3 = 1.684615 \dots \approx 1.6846$.
- $K = 4$: $T \approx 1.8093$ (Steps 1–3).

- $K = 5$: $\alpha^5 = 0.7^5 = 0.16807$; $E[L] = (1 - 0.16807)/0.3 = 0.83193/0.3 = 2.7731$;
 $C = 5 \times 0.1 + 1 = 1.5$; $T = 2.7731/1.5 = 1.848733 \dots \approx 1.8487$.
- $K = 6$: $\alpha^6 = 0.7^6 = 0.117649$; $E[L] = (1 - 0.117649)/0.3 = 0.882351/0.3 = 2.94117$;
 $C = 1.6$; $T = 2.94117/1.6 = 1.838231 \dots \approx 1.8382$.

The maximum over this grid is at $K^* = 5$ with $T \approx 1.8487$. As a cross-check by direct enumeration of the acceptance process for $K = 4$: $E[L] = 1 + 0.7 + 0.49 + 0.343 = 2.533$, matching Step 1.

Step 5: variance reduction by moment aggregation. Under A4, the draft-quality estimates $\hat{q}_i$ of the N drafters are independent and unbiased with common standard deviation σ . The equal-weight aggregate $\bar{q} = \frac{1}{N} \sum_{i=1}^N \hat{q}_i$ has

$$\text{Var}(\bar{q}) = \frac{\sigma^2}{N}, \quad \text{SE}(\bar{q}) = \frac{\sigma}{\sqrt{N}}.$$

For $\sigma = 0.3$, $N = 5$: $\sqrt{5} = 2.2360679 \dots$; $0.3/2.2360679 \dots = 0.134164 \dots \approx 0.1342$. The reduction factor relative to a single drafter is $1/\sqrt{5} \approx 0.4472$.

Step 6: controllability instance. For $N = 2$ members with $\lambda_1 \neq \lambda_2$, $\det M = \lambda_2 - \lambda_1 \neq 0$, so the two-member ensemble is exactly controllable in $H = 2$ steps; if $\lambda_1 = \lambda_2$ then $\det M = 0$ and the two states are permanently locked together under any broadcast input.

5. Results

All values below are computed in Section 4 with shown arithmetic; none are empirical measurements.

1. With $\alpha = 0.7$, $K = 4$, $c_d = 0.1$: $E[L] = 2.533$, $C = 1.4$, $T \approx 1.8093$.
2. The discrete throughput optimum over $K \in \{3, 4, 5, 6\}$ is $K^* = 5$, with $E[L] = 2.7731$, $C = 1.5$, $T \approx 1.8487$; neighboring values are $T(3) \approx 1.6846$ and $T(6) \approx 1.8382$.
3. Equal-weight aggregation over $N = 5$ drafters with $\sigma = 0.3$ reduces the standard error from $\sigma = 0.3$ to $\sigma/\sqrt{N} \approx 0.1342$, a factor of $1/\sqrt{5} \approx 0.4472$.
4. The finite-horizon reachability map of the linear broadcast ensemble is a scaled Vandermonde matrix, exactly controllable if and only if the parameters λ_i are pairwise distinct (proof in Section 3.5; two-member instance in Section 4, Step 6).

Projection (labeled, not computed). If the acceptance probability were instead $\alpha = 0.5$ with the same $c_d = 0.1$, the same formulas project $E[L] = (1 - 0.5^4)/0.5 = 1.875$, $C = 1.4$, $T \approx 1.3393$ for $K = 4$; the uncertainty in this projection is dominated by the assumed α , and no data are offered for it.

6. Discussion

Limitations. Assumption A1 (i.i.d. acceptance with constant α) is the strongest and least defensible: real drafter quality is autocorrelated across positions and degrades on out-of-distribution inputs, which would inflate the variance of $E[L]$ beyond the $1/N$ law of Step 5, whose derivation also requires A4 (independence); correlated drafter errors break the $\sigma/\sqrt{N}$ reduction. Assumption A2 prices verification as depth-independent, which favors large K ; a verifier cost growing with K would shift or remove the interior optimum $K^* = 5$, which is anyway only a grid optimum over $\{3, 4, 5, 6\}$, not a continuous one.

Failure modes. If drafter qualities are heterogeneous (violating the common- α reading of A1), the closed form $E[L] = (1 - \alpha^K)/(1 - \alpha)$ applies per drafter but the population aggregate requires a mixture, which we have not derived. If the broadcast schedule cannot observe even aggregate statistics, the scheduling problem reduces to

open-loop control, where the Vandermonde controllability result still holds but the throughput optimization does not.

What would falsify the claims. The $1/N$ variance reduction is falsified by any measured standard error of the aggregate quality estimate that does not scale as $\sigma/\sqrt{N}$ under independent drafters. The Vandermonde controllability claim is falsified by a pairwise-distinct parameter set whose finite-horizon reachability matrix is singular. The throughput formulas are falsified by acceptance processes whose empirical mean accepted length departs from $(1 - \alpha^K)/(1 - \alpha)$ at fixed α .

Open questions. Whether the multi-axis evaluation style of [10] can be operationalized for drafting ensembles (throughput, robustness, cost) with a validation coefficient analogous to its $\text{MCS} = 0.875$; whether the mode-decoupling analogy with [4] survives correlated drafter failures; and whether the gap condition A3 can be weakened while preserving uniform population-level statements, in the way the semi-classical calculus of [7] converts pointwise asymptotics into uniform ones. The chaotic-dynamics analysis of discrete symbolic processes in [1] suggests a further question: whether drafter-acceptance sequences exhibit structured (deterministic-looking) dynamics that a purely i.i.d. model misses; the entry [1] supplies no quantitative detail, so we raise this only as a hypothesis. No substantive claim is drawn from [11], whose supplied entry contains no summary text.

7. Conclusion

We recast ensemble drafting as a population-level broadcast-control problem, derived closed-form expressions for expected accepted length, per-cycle cost, and throughput with fully shown arithmetic ($E[L] = 2.533$, $C = 1.4$, $T \approx 1.8093$ at $K = 4$; grid optimum $K^* = 5$ at $T \approx 1.8487$ under the stated illustrative inputs), proved a $1/N$ standard-error reduction for moment aggregation ($0.3 \rightarrow 0.1342$ at $N = 5$, $\sigma = 0.3$), and showed that the linear broadcast ensemble's reachability map is a scaled Vandermonde matrix, exactly controllable iff the parameters are pairwise distinct. All quantitative claims are either derived above or explicitly labeled projections; the framework's empirical value remains an open question.

References

[1] Chaos in DNA inversions (Draft paper). arXiv:1105.1512v1. <https://arxiv.org/abs/1105.1512v1> [2] DRAFT: A Formally Verified Constructive Proof of the Consistency of Peano Arithmetic Using Ordinal Assignments. arXiv:2603.00487v1. <https://arxiv.org/abs/2603.00487v1> [3] TABED: Test-Time Adaptive Ensemble Drafting for Robust Speculative Decoding in LVLMS. arXiv:2601.20357v1. <https://arxiv.org/abs/2601.20357v1> [4] Master Functions and Equations for Perturbations of Vacuum Spherically-Symmetric Spacetimes. arXiv:2108.08668v3. <https://arxiv.org/abs/2108.08668v3> [5] Moment-Based Ensemble Control. arXiv:2009.02646v1. <https://arxiv.org/abs/2009.02646v1> [6] Prescriptive Master Integrals of Maximal Weight at Two Loops. arXiv:2610.10672v1. <https://arxiv.org/abs/2610.10672v1> [7] Asymptotic of Bergman Kernel (Master Thesis). arXiv:2202.03383v1. <https://arxiv.org/abs/2202.03383v1> [8] Ensemble Control on Lie Groups. arXiv:2008.03243v1. <https://arxiv.org/abs/2008.03243v1> [9] QNFO: QNFO Consilient Synthesis v2.0: The Shared Adelic Kernel [10] QNFO: NUMERATA: A Multi-Axis Framework for Evaluating Numeral Systems [11] QNFO: Connecting Geometrogenesis and Biogenesis [12] QNFO: Kappa/SIIT Cross-Validation: Exact Reproduction of Standard Model 1-Loop β -Functions and Resolution of the Harmonic-Paradigm Tension

Appendix A. Divergence report

No unresolved divergences remain in the reconciled text. The conventions adopted, documented here for transparency: (i) where source drafts differed on whether the grid optimum over K should be presented as a global optimum, the main text adopts the weaker and safer convention of a grid optimum over $K \in \{3, 4, 5, 6\}$; (ii)

where drafts differed on the scope of the variance-reduction claim, the main text adopts the version conditional on assumptions A1 and A4 (independence and common σ); (iii) statements about bibliography entries are restricted to what each entry's supplied summary states, and truncated entries are marked as such in Section 2.

Appendix B. Claim attribution

Claim	Source drafts	Agreement
C1: Ensemble drafting as broadcast population control (state, action, feedback)	A, B, C	CONVERGENT
C2: $E[L] = (1 - \alpha^K)/(1 - \alpha)$ with $E[L] = 2.533$ at $\alpha = 0.7$, $K = 4$	A, B, C	CONVERGENT
C3: $C = Kc_d + 1 = 1.4$ at $K = 4$, $c_d = 0.1$	A, B	CONVERGENT
C4: $T = E[L]/C \approx 1.8093$	A, B, C	CONVERGENT
C5: Grid optimum $K^* = 5$, $T \approx 1.8487$	A, B	CONVERGENT (scope: grid, per Appendix A)
C6: Standard error $\sigma/\sqrt{N} \approx 0.1342$ at $\sigma = 0.3$, $N = 5$	A, B, C	CONVERGENT
C7: Vandermonde reachability, controllable iff pairwise distinct λ_i	A, C	CONVERGENT
C8: Assumptions A1–A4 as explicit membership table (motivated by [9])	B, C	CONVERGENT
C9: Drafter specialization via spanning-set principle of [6]	A	SINGLE
C10: Projection at $\alpha = 0.5$, $T \approx 1.3393$	B	SINGLE (labeled projection)