

Contents

The Prufrock Protocol — DRAFT Formal Specification	1
1. Domain	2
1.1 Protocol constants [P]	2
1.2 Experiment variables [E]	2
1.3 Derived quantities [D]	3
2. The Contribution	3
3. The Ledger	5
4. The Schedule	5
4.1 Interruption timing	5
4.2 Prompt selection	5
4.3 Forfeit handling [P]	7
5. The Seed	7
5.1 Seed parameters [E]	7
5.2 Seed as genesis contributions	8
6. Form Constraints	8
6.1 Form constraint schema [E]	8
6.2 Form validity [E]	9
7. Experiment Configuration	10
7.1 Configuration summary	10
7.2 Protocol constants summary	10
8. Properties	11
9. Experiment Definitions	12
9.1 The Sonnet Trial	12
9.2 The Ghazal Trial	13
9.3 The Sestina Trial	13
10. Notes on Implementation	13

The Prufrock Protocol — DRAFT Formal Specification

This document gives a mathematical specification of The Prufrock Protocol. It describes the abstract objects, the validity rules they must satisfy, and the properties the protocol is intended to guarantee. The protocol is parameterised: specific experiments (the sonnet trial, the ghazal trial, etc.) are concrete instantiations obtained by fixing the experiment configuration variables.

Throughout, variables are annotated as follows:

- **[P] — Protocol constant.** Fixed by the protocol specification. Not configurable per experiment.
- **[E] — Experiment variable.** Set per experiment definition. Different experiments may choose different values.
- **[D] — Derived.** Computed from other variables; not set directly.

1. Domain

1.1 Protocol constants [P]

These are fixed for all experiments and define the protocol’s identity.

- (Sign, Verify) [P] — a digital signature scheme, assumed EUF-CMA-secure. The protocol fixes Ed25519 over Curve25519. All implementations must use this scheme; it is not a per-experiment choice.
- $H : \{0, 1\}^* \rightarrow \{0, 1\}^{256}$ [P] — a collision-resistant cryptographic hash. The protocol fixes SHA-256. Not configurable.
- $\mathcal{T} = \mathbb{R}^+$ [P] — the global wall-clock domain (UTC). All timestamps are UTC; local time is derived from participant timezone.
- \mathcal{G} [P] — a coarse-grained geographic space. The protocol fixes H3 resolution 5 ($\sim 252 \text{ km}^2$) as the default. Resolution may be adjusted by the protocol governance body but is not a per-experiment variable, since changing it would affect privacy guarantees across experiments.
- **Prompt anonymity** [P] — During the response phase, the participant’s view of $\Theta_{p,d,j}^*$ does not include the public key, name, or any attributable identity of the contribution authors. Author identity is preserved cryptographically in σ and the ledger but is never shown live to a responding participant. Participants respond to lines, not to people. This invariant is structural and applies to all experiments uniformly.

1.2 Experiment variables [E]

These are set per experiment and together define an **experiment configuration** \mathcal{E} .

- \mathcal{P} [E] — a finite set of **participants**, with $|\mathcal{P}| = n$. Each $p \in \mathcal{P}$ holds a keypair (sk_p, pk_p) under the protocol’s signature scheme and declares a daily availability window $W_p \subseteq [0, 24)$ in local time.
- $n = |\mathcal{P}|$ [E] — **cohort size**. The number of participants in a single cohort.
- c [E] — **cohort count**. The number of parallel cohorts running the same seed. Cohorts are isolated unless cross-cohort visibility is enabled.
- $R \in \mathcal{T}$ [E] — the **response window**: maximum elapsed time between prompt delivery and signed reply.
- \mathcal{L} [E] — the space of admissible **contributions**, defined by the experiment’s response medium and bounded by the response unit and length rules. \mathcal{L} may be text ($\subset \Sigma^*$), audio (a bytestream space), image, video, or a tagged union (mixed media). The hash $H(x)$ commits to the bytes of the contribution regardless of medium.
- ResponseMedium [E] — the medium of the contribution. One of {text, audio, image, video, m}. Default: text. Determines the type-space of \mathcal{L} and the storage representation of x . The protocol’s authentication, lineage, and tamper-evidence properties hold uniformly across all media.
- u [E] — the **response unit**: the granularity of a single contribution. For textual responses, one of {line, couplet, stanza, free_block}; for non-textual responses, one of {single_take, fragment, free_block}, where single_take denotes one continuous capture (an audio recording, a single photograph) and fragment denotes

a bounded clip. The interpretation of u depends on ResponseMedium.

- k [E] — **prompt depth**: the number of prior contributions shown to a participant at each interruption.
- m [E] — **poem length**: the total number of rounds (contributions per participant) that constitute a completed poem.
- D [E] — **duration**: the number of calendar days the experiment runs.
- r [E] — **rounds per day**: the number of interruptions per participant per day.
- λ [E] — **interruption rate**: governs the timing distribution of interruptions within availability windows. Interpretation depends on the interruption mode.
- Δ [E] — **minimum inter-arrival interval**: the minimum time between successive interruptions for a single participant.
- Mode_τ [E] — **interruption mode**: one of {uniform, poisson, scheduled}.
- Sel [E] — **selection rule**: the function that determines which prior contribution(s) are presented as the prompt, and which participant is assigned to extend which poem. One of {chosen_rotation, adjacent_pair_rotation, random_peer, weighted_peer, ...}
- \mathcal{F} [E] — **form constraint** (optional): a set of structural rules (rhyme scheme, refrain, end-word rotation, stanza structure, volta position) that the experiment imposes. $\mathcal{F} = \emptyset$ for free-form experiments.
- PromptType [E] — **prompt content type**: the form of the disruption presented to a participant at each interruption. Common values: contributions (the prompt is k prior contributions from the ledger), external_form (the prompt is a text drawn from outside the ledger — e.g. a haiku, an aphorism, a fragment of news), media (the prompt is an image, sound, or short video), location_anchor (the prompt is a coordinate or coarse-grained place description), none (no content; the bare interruption is the prompt and the participant responds only to their own present moment). The default is contributions. The protocol’s authentication and lineage rules apply uniformly: any prompt artefact, whether a prior contribution or an externally introduced disruption, must itself be signed and hashed and is recorded in Θ accordingly.

1.3 Derived quantities [D]

- $N_{\text{total}} = n \times m$ [D] — total contributions per cohort in a completed experiment.
- $T(\ell)$ [D] — **termination predicate**: true when a chain has reached length m . Defined as $T(\ell_j) \iff \text{depth}(\ell_j) = m$.
- Cross-cohort total: $c \times N_{\text{total}}$ [D].

2. The Contribution

A **contribution** is the protocol’s unit of authenticated expression. (The term “line” is used when $u = \text{line}$; the general term is “contribution”.)

A contribution is a tuple

$$\ell = (x, \Theta, a, t_p, t_r, g, \sigma)$$

where:

Field	Domain	Class	Meaning
x	\mathcal{L}	—	the contributed text (one response unit)
Θ	$(\{0, 1\}^{256})^k \cup \{\perp\}$	—	ordered tuple of hashes of the k [E] parent contributions, or \perp for a genesis contribution
a	$\{pk_p : p \in \mathcal{P}\}$	—	the contributor's public key
t_p	\mathcal{T}	—	timestamp of the prompt
t_r	\mathcal{T}	—	(interruption) timestamp of the signed response
g	\mathcal{G}	[P]	contributor's coarse-grained location at t_r , at protocol-fixed resolution
σ	signature space	[P]	$\text{Sign}_{sk_p}(x \parallel \Theta \parallel a \parallel t_p \parallel t_r \parallel g)$ using protocol-fixed Ed25519

The parent reference Θ is an ordered tuple of k hashes. When $k = 1$, Θ is a singleton; when $k > 1$, the contribution records all k parent hashes in the order they were presented to the participant.

A contribution ℓ is **valid**, written $V(\ell)$, iff all of the following hold:

1. **Authenticity** [P]. $\text{Verify}_a(\sigma, x \parallel \Theta \parallel a \parallel t_p \parallel t_r \parallel g) = \text{true}$.
2. **Punctuality** [P, E]. $t_r \in [t_p, t_p + R]$, where R is the experiment's response window **[E]**.
3. **Lineage** [P]. Either $\Theta = \perp$, or for every $h \in \Theta$, $\exists \ell'$ already in the ledger with $H(\ell') = h$ and $V(\ell')$.
4. **Unit conformance** [E]. x conforms to the experiment's response unit u and any applicable form constraint \mathcal{F} .
5. **Self-exclusion** [P]. For every $h \in \Theta$: the contribution ℓ' with $H(\ell') = h$ has $a(\ell') \neq a(\ell)$. A participant is never prompted with their own contribution.

Invalid contributions are not admitted to the ledger.

3. The Ledger

The **poem-ledger** \mathcal{L}^* is the append-only set of all valid contributions committed under the protocol. The ledger structure is a **protocol constant** [P] — it is always an authenticated, append-only log regardless of experiment.

Define the parent relation $\prec \subseteq \mathcal{L}^* \times \mathcal{L}^*$ by

$$\ell' \prec \ell \iff H(\ell') \in \Theta(\ell).$$

When $k = 1$ (single-parent prompt), (\mathcal{L}^*, \prec) is a **directed forest**. When $k > 1$, it is a **directed acyclic graph** (DAG), since each contribution may have multiple parents.

A **chain** is a maximal path $\ell_0, \ell_1, \dots, \ell_m$ where $\Theta(\ell_0) = \perp$ and for each $i > 0$, $H(\ell_{i-1}) \in \Theta(\ell_i)$, selecting a single lineage thread through the DAG.

A **poem** is a maximal chain rooted at a genesis seed (one or more contributions with $\Theta = \perp$) and extended by valid contributions according to the experiment’s selection rule. A poem has length m [E] when complete. Authorship of a poem is determined by which participants contribute to it, as governed by Sel [E]: under single-spine rules (chosen_rotation), one participant may dominate a poem; under dispersed rules (adjacent_pair_rotation), authorship is distributed across many participants and no single poet authors more than a small fraction of any one poem.

4. The Schedule

The protocol’s **schedule** governs when participants are interrupted and which contributions are presented as their prompt.

4.1 Interruption timing

For each participant p on each experiment day $d \in \{1, \dots, D\}$ [E], the protocol generates r [E] interruption times within $W_p^{(d)}$ (the participant’s availability window for day d):

- If $\text{Mode}_\tau = \text{uniform}$ [E]: $\tau_{p,d,j} \sim \text{Uniform}(W_p^{(d)})$ for $j \in \{1, \dots, r\}$, subject to Δ [E].
- If $\text{Mode}_\tau = \text{poisson}$ [E]: arrival times are drawn from a Poisson process with rate λ [E] on $W_p^{(d)}$, truncated to r arrivals, subject to Δ .
- If $\text{Mode}_\tau = \text{scheduled}$ [E]: interruption times are fixed in advance (e.g. at the midpoint of $W_p^{(d)}$).

4.2 Prompt selection

Each interruption produces a triadic moment. **The participant’s situated experience** at $\tau_{p,d,j}$ — their inner and outer state, the texture of where and when they are

— is the **thesis**. The protocol introduces an **antithesis** in the form of $\Theta_{p,d,j}^*$: a disruption whose content type is governed by PromptType **[E]**. The participant’s signed contribution ℓ records the **synthesis**. The Schedule (Section 4.1) determines when the antithesis arrives; the Selection Rule Sel **[E]** determines, for the contribution case, which prior content is drawn forward; the prompt type determines the form the disruption takes.

When PromptType = contributions, the prompt $\Theta_{p,d,j}^*$ is a tuple of k **[E]** prior contributions selected according to Sel **[E]**:

- If Sel = chosen_rotation: a pre-committed permutation $\phi : \{1, \dots, n\} \rightarrow \mathcal{P}$ **[D]** determines that on day d , the contributions of participant $\phi(d - 1)$ from day $d - 1$ form the prompt. If $k > 1$ and only one contribution exists from the chosen participant, the prompt is filled by traversing the lineage (grandparent contributions). All participants extend the same single spine; this rule produces a single dominant authorship lineage per poem.
- If Sel = adjacent_pair_rotation: the protocol maintains n parallel poems, each nominally associated with one participant via the seed phase. Each interruption assigns participant p to extend exactly one non-nominal poem P_q (where q identifies a poem p does not seed). The prompt $\Theta_{p,d,j}^*$ is the most recent two contributions of P_q — i.e., $(\ell_{m_q-1}^{(P_q)}, \ell_{m_q}^{(P_q)})$ where m_q is the current length of P_q at the moment of the interruption. Pairs are therefore always chronologically adjacent and drawn from the same poem. The (day, participant) \rightarrow poem assignment matrix $\mathcal{A} : \{1, \dots, D\} \times \mathcal{P} \rightarrow \{P_1, \dots, P_n\} \cup \{\perp\}$ is committed in advance as a near-balanced combinatorial design satisfying: (i) no fixed points — $\mathcal{A}(d, p) \neq P_p$ for all d during the response phase; (ii) per-participant balance — each participant’s response-day assignments are distributed across as many distinct non-nominal poems as possible, minimising repeated exposure to lines from the same poem. This rule disperses authorship across the cohort rather than concentrating it on a single spine.
- If Sel = random_peer: k contributions are selected uniformly at random from the most recent round, excluding p ’s own contributions.
- If Sel = weighted_peer: selection weights contributions by criteria (e.g. least-prompted participant, geographic diversity).
- If Sel = round_robin: a deterministic rotation selects the prompting participant(s).

In all modes, the **self-exclusion invariant** [P] holds: p is never shown their own prior contributions as the prompt.

When PromptType \neq contributions, the prompt is drawn from a registered, signed source pool defined by the experiment (a curated set of haiku, a media library, a stream of news fragments, etc.). Each such artefact is itself hashed and signed, and its hash is recorded in Θ exactly as a contribution hash would be — preserving lineage and tamper-evidence (validity rule 3) regardless of prompt content type.

4.3 Forfeit handling [P]

If no valid contribution is committed by p within R [E] of $\tau_{p,d,j}$, the slot is recorded as a **forfeit**. Forfeit handling is a protocol constant:

- The forfeit is permanently recorded in the ledger (absence is data).
- If $Sel = chosen_rotation$ and the forfeiting participant was the chosen poet for the next round, the protocol falls back to the most recent non-forfeited chosen poet's contribution.
- The forfeiting participant's poem will have fewer than m contributions; the poem is still valid but marked incomplete.

5. The Seed

Every experiment begins with a **seed**: one or more contributions not authored by any participant, drawn from existing works. The seed is the genesis of the ledger.

5.1 Seed parameters [E]

Parameter	Class	Description
Seed source	[E]	corpus (from existing literary work), manual (hardcoded), generated, or participant_seeded (each participant authors s seed contributions for their nominally associated poem during a seeding phase)
Seed work type	[E]	The form of the source work (sonnet, ghazal, epic, free verse, etc.)
Seed line extraction	[E]	Which part of the source work to extract: last_n, first_n, volta, couplet, envoi, random, custom
Seed line count s	[E]	Number of contributions extracted as the initial prompt. Must satisfy $s \leq k$.
Seed minimum age	[E]	Minimum years since publication (for rights confidence)
Seed language	[E]	ISO 639-1 code

Parameter	Class	Description
Seed rights status	[E]	Minimum tier: global_pd_confident, source_pd_only, or needs_review

5.2 Seed as genesis contributions

Each extracted seed line is encoded as a genesis contribution ℓ_0 with:

- x = the extracted text
- $\Theta = \perp$
- a = a distinguished “seed” public key (not a participant’s key)
- $t_p = t_r$ = the historical publication date (or experiment start date if unknown)
- g = the historical location of composition (if known) or the publication location
- σ = signed by the experiment organiser’s key, attesting provenance

Under participant_seeded mode, each genesis contribution is instead signed by the seeding participant’s key sk_p , with t_p and t_r recorded as the actual interruption and response times, and g recorded as the participant’s coarse location at t_r (just like any response contribution). The $\Theta = \perp$ marker still distinguishes the seed as a genesis. Each participant authors s seed contributions for their nominal poem P_p during the protocol’s seeding phase, before the response phase begins.

6. Form Constraints

When an experiment specifies a form constraint $\mathcal{F} \neq \emptyset$ **[E]**, contributions must satisfy additional validity rules beyond the protocol’s base rules.

6.1 Form constraint schema **[E]**

$$\mathcal{F} = (f_{\text{name}}, S, \rho, \mathcal{R}, \mathcal{W}, \mu, v, \mathcal{X})$$

Field	Domain	Description
f_{name}	string	Named form (sonnet, ghazal, sestina, haiku, tanka, qasida, renga, free, ...)
S	\mathbb{Z}^+ list $\cup \{\perp\}$	Stanza structure: number of contributions per stanza, or \perp if unstructured
ρ	string $\cup \{\perp\}$	Rhyme scheme (e.g. “ABAB CDCD EFEF GG”), or \perp if unenforced

Field	Domain	Description
\mathcal{R}	RefrainRules $\cup \{\perp\}$	Refrain rules (for ghazal radif, villanelle refrains, etc.)
\mathcal{W}	EndWordRules $\cup \{\perp\}$	End-word rotation rules (for sestina)
μ	string $\cup \{\perp\}$	Meter (advisory, e.g. "iambic pentameter"), or \perp
v	$\mathbb{Z}^+ \cup \{\perp\}$	Volta position (contribution number of expected turn), or \perp
\mathcal{X}	map of string \rightarrow typed value	Extensions. Arbitrary form rules outside the canonical schema. The schema fields above are anglophone-canonical and not exhaustive; \mathcal{X} allows any form to express its constraints. Examples: <pre> {"syllable_pattern": [5,7,5]} for haiku; {"syllable_pattern": [5,7,5,7,7]} for tanka; {"kigo_required": true} for haiku seasonal markers; {"parallelism": "synonymous"} for Hebrew biblical verse; {"alliteration_per_line": 3} for Old English alliterative verse; {"image_aspect_ratio": "1:1"} for image-medium experiments; {"audio_max_seconds": 30} for audio-medium experiments. </pre>

6.2 Form validity [E]

When $\mathcal{F} \neq \emptyset$, the unit conformance rule (validity rule 4) additionally requires:

- The contribution’s position in the poem’s stanza structure S is consistent.
- If $\rho \neq \perp$: the contribution’s end-sound is consistent with the rhyme scheme at its position.
- If $\mathcal{R} \neq \perp$: the contribution includes the required refrain at the specified position.
- If $\mathcal{W} \neq \perp$: the contribution’s end-word matches the rotation pattern at its position.

Form constraints are **advisory** in the simulation (scored, not hard-rejected) and **validated** in a live experiment (flagged for review if violated, but not automatically rejected — human judgement applies to poetic form).

7. Experiment Configuration

An **experiment configuration** \mathcal{E} is the complete set of experiment variables [E] that, together with the protocol constants [P], fully determines the behaviour of a run.

7.1 Configuration summary

Variable	Symbol	Class	Sonnet trial value
Cohort size	n	[E]	14
Cohort count	c	[E]	2
Response window	R	[E]	4 hours
Response unit	u	[E]	line
Response medium	ResponseMedium	[E]	text
Prompt depth	k	[E]	2
Prompt type	PromptType	[E]	contributions
Poem length	m	[E]	14
Duration	D	[E]	14 days
Rounds per day	r	[E]	1
Interruption mode	Mode _{τ}	[E]	uniform
Min inter-arrival	Δ	[E]	n/a (1 round/day)
Selection rule	Sel	[E]	adjacent_pair_rotation
Form constraint	\mathcal{F}	[E]	sonnet (advisory)
Seed source	—	[E]	corpus
Seed work type	—	[E]	sonnet
Seed extraction	—	[E]	volta (lines 12-13)
Seed line count	s	[E]	2
Seed min age	—	[E]	100 years
Seed language	—	[E]	en
Cross-cohort visibility	—	[E]	false

7.2 Protocol constants summary

Constant	Value	Rationale
Signature scheme	Ed25519 / Curve25519	EUFCMA-secure, widely implemented, deterministic
Hash function	SHA-256	Collision-resistant, standard
Clock domain	UTC	Unambiguous global reference
Geographic resolution	H3 level 5 (~252 km ²)	Balances analytical utility against participant privacy
Ledger structure	Authenticated append-only log	Tamper-evidence without requiring consensus mechanism
Self-exclusion	Always enforced	Structural invariant: you are never prompted with yourself
Prompt anonymity	Always enforced	Structural invariant: participants respond to lines, not to people. Author identity is preserved in the ledger but never shown live.
Forfeit recording	Always recorded	Absence is data; the ledger does not pretend forfeits didn't happen

8. Properties

The protocol is designed to satisfy the following properties. Properties hold for all valid experiment configurations.

(P1) Authenticity [P]. For every $\ell \in \mathcal{L}^*$ with $V(\ell)$: ℓ was committed by the holder of sk_a during $[t_p, t_p + R]$. (From EUFCMA security of Ed25519 and validity rule 1.)

(P2) Temporal monotonicity [P]. For all $\ell, \ell' \in \mathcal{L}^*$ with $\ell' \prec \ell$: $t_r(\ell) > t_r(\ell')$. (A descendant cannot precede its ancestor in time.)

(P3) Tamper-evidence [P]. Modification of any contribution $\ell^* \in \mathcal{L}^*$ invalidates every descendant in the DAG. (From collision-resistance of SHA-256 propagated through $\Theta(\cdot)$.)

(P4) Non-equivocation [P]. No participant can sign two distinct contributions for the same prompt slot $(p, \tau_{p,d,j})$ without detection, since both signatures bind the same t_p and a .

(P5) Self-exclusion [P]. No contribution in \mathcal{L}^* has a parent authored by the same participant. (*Enforced by validity rule 5.*)

(P6) Liveness [P, E]. For experiments with $\text{Mode}_\tau \in \{\text{uniform, poisson}\}$: as $|\mathcal{L}^*| \rightarrow \infty$, the empirical distribution of contribution times converges to a uniform sample over $\bigcup_p W_p$, conditional on λ , Δ , and r .

(P7) Compression-fidelity trade [P]. The protocol preserves: contributor identity, prompt and response timestamps, coarse location, and the dialectical lineage of inheritance. It does not preserve: the contributor’s full mental state, surrounding context, or alternatives considered. This unrecorded space is the protocol’s **negative space** — the same negative space that poetic form is built to honour.

9. Experiment Definitions

9.1 The Sonnet Trial

Set $n = 14$, $c = 2$, $R = 4$ hours, $u = \text{line}$, $k = 2$, $m = 14$, $D = 14$, $r = 1$, $\text{Mode}_\tau = \text{uniform}$, $\text{Sel} = \text{adjacent_pair_rotation}$, seed source = corpus, $s = 2$, seed work type = sonnet, seed extraction = volta (lines 12-13), seed min age = 100 years.

The trial produces $n = 14$ poems per cohort, each of length $m = 14$ lines (counting only generated lines; the 2 shared seed lines provide context but are not counted toward the sonnet form), with $c = 2$ cohorts running in parallel (poets, laypeople), totalling $2 \times 14 \times 14 = 392$ generated contributions.

Shared seed. Two adjacent lines are extracted from a single existing sonnet’s volta (lines 12-13 of a public-domain source poem) and serve as the genesis pair (s_1, s_2) for **every** poem in the cohort. All 14 poems begin from the same starting condition; the seeds are not part of any single poem’s authorship — they are shared external input.

Day 1 — first generated line. Each participant p is interrupted once at $\tau_{p,1} \sim \text{Uniform}(W_p^{(1)})$ and shown the shared pair (s_1, s_2) . The participant commits a new line, which becomes line 1 of their nominally associated poem P_p . After day 1, each of the 14 poems contains $(s_1, s_2, \ell_1^{(P_p)})$ where $\ell_1^{(P_p)}$ was authored by participant p .

Days 2-14 — dispersed rotation. On each subsequent day $d \in \{2, \dots, 14\}$, each participant is interrupted once and shown the current tail pair of one non-nominal poem P_q ($q \neq p$), as assigned by the matrix \mathcal{A} . On day 2 the tail pair of P_q is $(s_2, \ell_1^{(P_q)})$; from day 3 onward it is $(\ell_{d-2}^{(P_q)}, \ell_{d-1}^{(P_q)})$. The participant commits the next line of P_q , advancing it by 1. Each day produces 14 new contributions, one per poem.

Assignment matrix. $\mathcal{A} : \{2, \dots, 14\} \times \mathcal{P} \rightarrow \{P_1, \dots, P_{14}\}$ is a 13-row Latin square on the 14 poems with no fixed points: no participant ever writes for their own nominal poem during the rotation, and each participant writes for each of the 13 other poems exactly once across days 2-14. The matrix is committed via hash-commitment before day 1 and revealed at experiment close.

Authorship dispersion. Each completed poem of 14 lines comprises one line from its nominal participant (line 1, day 1) and 13 lines from the 13 other participants (lines 2-14, days 2-14, one each). No participant authors more than $1/14 \approx 7\%$ of any one poem. Each participant’s 14 contributions (1 nominal + 13 rotated) are distributed across all 14 poems in the cohort, exactly once each.

Form constraint \mathcal{F} : sonnet (advisory). $S = [4, 4, 4, 2]$, $\rho = \perp$ (unenforced), $v = 9$.

9.2 The Ghazal Trial

Set $n = 7$, $c = 1$, $R = 6$ hours, $u = \text{couplet}$, $k = 2$, $m = 7$, $D = 7$, $r = 1$, $\text{Sel} = \text{round_robin}$.

Form constraint \mathcal{F} : ghazal. $S = [2, 2, 2, 2, 2, 2, 2]$, $\rho = \text{"AA BA CA DA EA FA GA"}$, $\mathcal{R} = \text{radif}$ (refrain phrase extracted from seed’s final couplet, position: end of second line).

9.3 The Sestina Trial

Set $n = 6$, $c = 1$, $R = 8$ hours, $u = \text{stanza}$, $k = 1$, $m = 7$ (6 stanzas + envoi), $D = 7$, $r = 1$, $\text{Sel} = \text{permutation}$.

Form constraint \mathcal{F} : sestina. $S = [6, 6, 6, 6, 6, 6, 3]$, $\mathcal{W} = \text{end-word spiral rotation } [6, 1, 5, 2, 4, 3]$.

10. Notes on Implementation

- The signature scheme (Ed25519) and hash function (SHA-256) are **protocol constants** [P]. Implementations must not substitute these without a protocol version change.
- The ledger may be realised on any authenticated append-only data structure: a Merkle DAG, a Certificate-Transparency-style log, or a blockchain. The protocol requires only append-only authenticated storage, not a specific consensus mechanism. The choice of ledger backend is an **implementation decision**, not an experiment variable.
- The chosen-poet rotation ϕ (when $\text{Sel} = \text{chosen_rotation}$) should be committed via a hash-commitment scheme published before day 1 and revealed at experiment close, so participants can verify the rotation was not adapted in flight.
- Geographic precision (H3 resolution) is a **protocol constant** [P] to ensure consistent privacy guarantees across experiments. If a future experiment requires finer or coarser resolution, this should be proposed as a protocol version change, not a per-experiment override.
- Form constraints \mathcal{F} are validated differently in simulation (scored) versus live experiments (flagged for human review). The protocol does not automatically reject contributions that violate form — poetic licence is honoured — but violations are recorded.