Persona-Based Pre-Validation Screening Protocol v1.0

The Human Imperative • Mars Readiness Program

Persona-Based Pre-Validation Screening Protocol

for the MCSI, ESI, and Arcadia Index

An Item-Construction Screening Study Using AI Persona Simulation

Study Protocol — Version 1.0 — 17 July 2026

Dinu-Stefan Teodorescu, Ph.D.

Developed in collaboration with Claude (Anthropic) and DeepSeek

Companion to The Human Imperative: Awaken and Flourish in the Age of AI in the Symbiocene

1. Purpose and Honest Framing

This protocol describes a structured screening study of the item pools of three instruments developed alongside The Human Imperative: the Mars Colonist Symbiocene Inventory (MCSI), the Earth Symbiocene Inventory (ESI), and the Arcadia Index. The study uses biographically rich AI personas as simulated respondents in order to detect construction faults in the items before any human administration takes place.

What this study is: an item-construction debugger and pre-validation screen. It is designed to catch mis-keyed or reverse-scored items that behave backwards, items that statistically cohere with the wrong subscale, double-barreled or ambiguous wordings that produce chaotic response patterns, subscales whose items fail to hang together, and scoring logic that assigns archetypes incorrectly.

What this study is not: validation. No result from this study establishes reliability, validity, or norms for human use. The method is a filter that can catch broken items but cannot certify sound ones. Every claim made from its results must carry that limitation explicitly, in the study report and in any outreach material that cites it.

The framing mirrors the epistemics of the training manuals themselves: conventionally, this is real methodological work worth doing with discipline and care; it produces genuine information about item construction. At the same time, it cannot produce what only human data can produce. Both statements are true, and the protocol holds both.

2. Background: Evidence from the Persona Pilot

This protocol builds directly on a completed pilot in which a panel of 57 biographically rich AI personas answered two instruments with established human benchmarks from the author's own doctoral research data: the Posttraumatic Growth Inventory (Short Form) and the DUREL religiosity index. The pilot produced one encouraging result and one cautionary result, and this protocol is designed around both.

The encouraging result. The AI personas' PTGI-SF reliability (α = .786) approached the human benchmark (α = .888), and — unprompted — the personas reproduced the same weakest domain pairing (Appreciation of Life) found in the real human data. Structural patterns can emerge from the simulation that echo genuine human response structure.

The cautionary result. The DUREL intrinsic-religiosity subscale failed in simulation (α = .476 against the human α = .885), and it failed on the most interior item of the scale: "In my life, I feel the presence of Divinity." Human respondents answer this item bimodally; the AI personas hedged toward the scale midpoint even when their biographies described unambiguous devotion or unambiguous collapse of faith (item variance 0.487 against the human 1.995). The pilot's conclusion: the AI could write conviction, but could not reliably rate conviction.

The direct implication for this study. The MCSI and the Arcadia Index are built substantially out of interior territory — the Witnessing Self, Mortality Integration, Grief Tolerance, Contemplative Depth. The simulation method is demonstrably weakest exactly where these instruments live. This protocol therefore treats interior-item hedging as a known simulator limitation, pre-registers it as an expected pattern, and builds decision rules that prevent good interior items from being wrongly flagged as faulty. No item may be removed because an AI persona answered it timidly.

3. Instruments Under Screening

Three instruments are screened. The canonical item pools are those implemented in the live instruments at din.nu at the time of Phase 0 export; item counts and wordings are frozen at export and no edits are permitted between export and the completion of data generation.

Instrument	Structure	Screening focus
MCSI (Mars Colonist Symbiocene Inventory)	Twelve subscales measuring capacity, organised under meta-axes; includes reverse-scored items	Item–subscale coherence; keying direction; cross-loading; known-groups sensitivity
ESI (Earth Symbiocene Inventory)	Earth-language parallel of the MCSI	Same as MCSI; additionally, parallelism of item behaviour with MCSI counterparts
Arcadia Index	Fourteen vignettes across eight life domains; developmental stages from the Enjoyer to the Returning One; framed as structured contemplative reflection, not a validated psychometric instrument	Stage-assignment plausibility; vignette clarity; bypass-signature behaviour against MCSI
Archetype layer	Nine archetype profiles (including The Still Point), assigned by relative-ranking logic over the axes	Whether engineered profiles trigger their intended archetypes; Still Point coverage

4. Study Questions

Q1. Do the items of each subscale cohere with their own subscale, and not better with a foreign one, when answered by a diverse persona panel?

Q2. Are all reverse-scored items behaving in the keyed direction?

Q3. Do subscales discriminate between personas engineered to differ on the relevant capacity, in the predicted direction?

Q4. Does the engineered bypass persona produce the predicted signature — late-stage Arcadia centre of gravity alongside MCSI resistance on Grief Tolerance, Embodied Presence, or Mortality Integration?

Q5. Do the nine archetypes, including The Still Point, fire on the profiles they were designed to capture, and does no realistic profile fall through the assignment logic?

Q6. Which items show interior-hedging compression (low variance attributable to the simulator rather than the item), and how should this constrain interpretation?

5. The Persona Panel

5.1 Panel A — The Core Panel (n = 57)

The existing 57-persona biography set from the pilot is reused without modification. Reuse preserves comparability with the pilot's established behaviour (the PTGI success and the DUREL interior-hedging failure act as known reference points for how this panel behaves), and it spares the cost of constructing a new panel. Each biography includes demographics, life history, formative losses, practice history, relationship to technology, nature, and death, and a voice sample.

5.2 Panel B — Engineered Contrast Personas (n = 10)

Ten personas are constructed specifically for known-groups testing. Each is written to occupy a defined position on the constructs, with the predicted instrument signature registered before any administration. A subscale that cannot separate the personas engineered to differ on it is carrying suspect items.

Persona	Engineered profile	Pre-registered predicted signature
The Adept	Four decades of contemplative practice; years of hospice accompaniment; settled relationship with death	High Contemplative Depth, Witnessing Self, Mortality Integration; late-stage Arcadia; archetype: Sage
The Achiever	Highly successful, externally validated life; no interior practice; actively avoids the subject of death	Low Mortality Integration and Witnessing Self; early-to-mid Arcadia; archetype: Climber
The Bypasser	Fluent non-dual vocabulary acquired quickly; unprocessed bereavement; estranged family; somatically avoidant	Late-stage Arcadia language alongside MCSI resistance on Grief Tolerance and Embodied Presence — the bypass signature
The Griever	Recent profound loss, actively and honestly metabolised; functioning intact	Elevated but coherent Grief Tolerance; no bypass signature
The Hermit	Chosen, contented solitary life; rich inner resources; few close relationships	High Isolation Resilience & Solitude; lower Relational Depth
The Technophile	Professionally and emotionally dependent on AI systems; defers to machine judgment	AI Partnership Capacity at the dependence pole
The Skeptic	Able professional who distrusts and avoids AI on principle	AI Partnership Capacity at the aversion pole
The Still One	Settled, low-reactivity evenness across life domains; neither striving nor withdrawn	No single dominant axis; archetype: The Still Point
The Novice	Young, curious, open; no practice history; unformed views	Most capacities developing; archetype: Beginner's Mind
The Carer	Long relational caregiving history; open-hearted; own needs habitually deferred	High Heart-axis openness; archetype: Healer or Empath

6. Administration Procedure

6.1 Author–respondent separation. Because Claude co-authored the item pools, the primary respondent system is DeepSeek: each persona is instantiated in DeepSeek and answers all items in character. Claude performs the analysis blind to persona identity during scoring. An optional secondary arm repeats administration in Claude for a cross-system comparison, reported separately and never pooled with the primary arm.

6.2 Label stripping and randomization. Items are presented with no subscale names, no axis names, and no indication of which instrument section they belong to. Item order is randomized independently for every persona. Vignettes of the Arcadia Index are presented in randomized order with stage labels removed.

6.3 Fresh context per persona. Each persona answers in a fresh, separate session containing only the persona biography, the standardized instructions, and the randomized item stream. No persona sees another persona's answers; no session carries memory of a previous session.

6.4 Standardized instructions and anchors. All personas receive identical instructions, including the full verbal anchors of the response scale, restated at intervals through the item stream. This addresses the scale-direction confusion observed in the pilot's explainability probe, which was attributed to an under-specified scale description.

6.5 Frozen materials. Item wordings are exported from the live din.nu instruments in Phase 0 and frozen. Any discrepancy discovered between exported items and this protocol's assumptions is recorded, and the export is treated as canonical.

7. Pre-Registered Predictions

All predictions are registered before any response is generated. Confirming and disconfirming outcomes are reported with equal prominence.

#	Prediction	If it fails
P1	Within each MCSI/ESI subscale, corrected item-total correlations are positive for every item	Negative item flagged for keying error or construction fault
P2	No item correlates more strongly with a foreign subscale total than with its own (margin ≥ .20)	Item flagged as cross-loading; review wording and placement
P3	Each engineered contrast in Panel B produces its predicted signature on the targeted subscale(s)	If the item is non-interior: item or subscale flagged. If interior: interpreted under the hedging control (Section 8)
P4	The Bypasser produces the bypass signature; the Griever does not	The bypass-detection logic of the instruments is reviewed
P5	All ten Panel B personas receive their intended archetype; no persona receives a null or implausible assignment	Archetype assignment logic re-examined, with attention to Still Point coverage
P6	Item variance decreases with a priori interiority level: interior items show compressed variance relative to behavioural items	If interior items do NOT show compression, the hedging concern is weaker than the pilot suggested — a finding worth reporting in itself

8. The Interiority Classification and Hedging Control

Before administration, every item in all three instruments is classified by the author on a three-level Interiority scale. Level 1 (behavioural): the item asks about observable action or frequency. Level 2 (attitudinal): the item asks about a view, preference, or self-assessment. Level 3 (interior): the item asks about felt, first-person contemplative or existential experience — the territory where the pilot demonstrated hedging.

This classification is the study's central safeguard. All flag criteria in Section 10 are conditioned on it. Compressed variance or weak discrimination on a Level 3 item is, by pre-registration, ambiguous between item fault and simulator limitation, and can never by itself justify item removal. The DUREL Divinity item stands as the cautionary precedent: a psychometrically sound human item that the simulation answered badly.

9. Analysis Plan

For every item: mean, variance, response distribution, corrected item-total correlation with its own subscale, and correlation with every foreign subscale total. For every subscale: Cronbach's alpha, alpha-if-item-deleted, and inter-subscale correlation matrix. For Panel B: per-persona subscale profiles against the pre-registered signatures. For the archetype layer: assignment outcome per persona, with the full axis-ranking trace retained. For the Arcadia Index: stage centre of gravity per persona and the bypass cross-check against MCSI

resistance patterns. All analysis code is retained and re-runnable; the analysis is performed by Claude with persona identities masked during scoring and unmasked only for the known-groups comparisons.

10. Item Decision Rules

Outcome class	Criteria	Action
FLAG — repair or review	Negative corrected item-total correlation; or cross-loading margin ≥ .20 onto a foreign subscale; or known-groups failure on a Level 1–2 item; or response direction opposite to keying	Item goes to expert review (author). Wording, keying, and placement re-examined. Repair or retention decided by the author, documented with rationale
CAUTION — interpret, do not cut	Compressed variance on a Level 3 item; or known-groups failure confined to Level 3 items; or subscale alpha depressed solely by Level 3 items	Recorded as a simulator-limitation finding. Item retained. Earmarked for priority attention in the first human pilot
PASS	Expected behaviour on all criteria	No action; result archived

The absolute rule: no item is removed from any instrument on simulation evidence alone. The simulation proposes; the author disposes. Final authority over every item rests with the human expert, and every flag decision is documented so the reasoning survives the study.

11. Circularity Mitigations

The specific risk of this study is that a model from the same family that co-authored the items could produce a spuriously clean structure by recognising which items belong together, rather than by simulating measurement. Five mitigations apply: (1) the primary respondent system (DeepSeek) is different from the item co-author (Claude); (2) subscale and axis labels are stripped and item order randomized per persona; (3) all predictions are registered before data generation; (4) scoring is performed blind to persona identity; (5) disconfirming results are reported with the same prominence as confirming ones. A residual risk remains — both systems share broad training exposure to psychometric conventions — and the report states this plainly.

12. Workflow and Phases

Phase	Work	Responsible
0. Freeze	Export canonical item pools from the live din.nu instruments; author classifies every item on the Interiority scale; predictions registered in final form	Dinu (decisions), Claude (extraction and formatting)
1. Panel	Panel A biographies compiled; Panel B contrast personas written and reviewed against their predicted signatures	Claude (drafting), Dinu (approval)
2. Generation	All 67 personas answer all items in fresh sessions under Section 6 conditions; responses logged to structured data files	DeepSeek (respondent), Dinu (session operator)
3. Analysis	Full analysis per Section 9; flag list produced under Section 10 rules	Claude
4. Review	Expert review of every flagged item; repair decisions documented	Dinu
5. Report	Study report with methods, results, flag dispositions, and the limitations statement of Section 13	Claude (drafting), Dinu (final)

13. Reporting and Use in Outreach

In outreach material (including correspondence with analog-mission programs), results of this study are described in one register only: the instruments have undergone structured, simulation-based item screening with documented limitations, prior to human piloting. The

following limitations paragraph, or its equivalent, accompanies every citation of the study:

> *"Item screening was conducted using AI persona simulation as a pre-validation filter. This method can detect item-construction faults but cannot establish reliability, validity, or norms for human populations. A documented limitation of the method — compressed response variance on items addressing interior, first-person experience — was pre-registered and controlled for in the decision rules. All psychometric claims await human data."*

Held honestly, this study strengthens the instruments twice over: once through the faults it finds, and once through the discipline with which it declines to claim more than it found.