

# AGORA

*Interface Irreducibility in Human-Agent Software*

---

AUTHORS	<b>Andy Salvo</b> , founder of Polylogic AI and co-founder of Crest; <b>Jameson Ackerman</b> , co-founder of Crest.
VERSION	v5.0.0 · 2026-04-26
LICENSE	Creative Commons Attribution 4.0 International (CC BY 4.0)
CITE AS	Salvo, A., & Ackerman, J. (2026). <i>AGORA: Interface Irreducibility in Human-Agent Software</i> . v5.0.0. Zenodo. <a href="https://doi.org/10.5281/zenodo.19778457">https://doi.org/10.5281/zenodo.19778457</a>

# Abstract

Artificial intelligence turns every serious product into a three-body system: user, agent, and interface. Products that let the interface dissolve into user habit or agent optimization may work briefly, but they become replaceable; products that conserve an irreducible interface become places people return to, recognize, and trust. AGORA formalizes a design hypothesis for this new class of software: a durable human-agent interface should preserve an independent mode basis under simultaneous pressure from user habit and agent optimization; this paper proposes and preregisters rhythm, friction, and style as a risky, testable candidate basis. We formalize the claim as the *Three-Eigenmode Stability Lemma*, introduce the Interface Autonomy Index for public measurement, and preregister Crest alongside external comparison systems as falsifiable tests. The deposit of this paper is the opening lab notebook: if the named invariants do not predict persistence, recognition, and resistance to absorption, AGORA is weakened by the evidence that should have tested it.

## 1. Introduction

For sixty years, software design assumed two bodies: a user, and a system. The interface was treated as a translation layer between them. That assumption was sufficient when software was a tool and the system had no active strategic role.

Agentic software breaks the assumption. The system now adapts, ranks, summarizes, persuades, and sometimes acts. Once the system becomes an active pressure, the interface can no longer be a neutral surface. It either dissolves into the user's habit, or it is captured by the agent's optimization objective.

AGORA names the missing third body. A persistent interface is not merely the user's workflow and not merely the agent's output. It is an autonomous place with its own conserved direction. The engineering problem is to preserve that direction under pressure.

### PROBLEM STATEMENT

Agentic software turns interface design from a usability problem into a conservation problem.

## 2. The Two Failure Modes

When a system is built as a two-body interaction between user and process, only two outcomes are stable. Both are failures.

### 2.1 User-Absorption

The interface becomes pure habit. The user no longer notices it. It does not guide them, surprise them, or shape them. Default messaging apps, generic news readers, and most operating-system shells exemplify this mode. They are technically functional but have high replaceability elasticity: an equivalent product with lower switching cost can displace them quickly. The user uses them, but cannot reliably identify what the interface itself contributed.

User-absorption produces software that is replaceable. If a competitor offers the same function with one fewer click, the user migrates instantly, because nothing was holding them to the original.

### 2.2 Agent-Absorption

The interface becomes pure optimization. The system uses every pixel and every gesture to maximize a measurable outcome - usually engagement, attention, or revenue. The user feels manipulated, even if they cannot articulate why. Late-stage algorithmic feeds, dark-pattern e-commerce, and certain attention-capture games exemplify this mode.

Agent-absorption produces software that the user *resents*. They keep using it because it has hijacked their reward circuitry, but their relationship to it is adversarial. Eventually, exhaustion or alternatives end the relationship.

Both failure modes can look commercially successful: they generate metrics, retain users, and produce revenue. AGORA treats them as failures only when the system claims to be a durable place. The operational signs are high replaceability elasticity, low blind recognition, declining trust, or retention that disappears when the optimization loop is removed.

	User-Absorption	Agent-Absorption
Symptom	Habit, invisibility. The user no longer notices.	Resentment, manipulation. The user feels pushed.

	User-Absorption	Agent-Absorption
<b>Examples</b>	Default messaging apps, generic news readers, most OS shells.	Late-stage algorithmic feeds, dark-pattern checkouts, attention-capture games.
<b>Direction of collapse</b>	Interface dissolves toward the user's habit (the $I$ -dimension collapses toward $U$ ).	Interface dissolves toward the agent's optimization (the $I$ -dimension collapses toward $A$ ).
<b>Outcome</b>	Replaceable. Users migrate to the next equivalent product.	Adversarial. Users keep using it but resent it; eventually leave.
<b>Looks like success</b>	High retention, low engagement.	High engagement, low trust.
<b>Measurement</b>	High replaceability elasticity; low blind interface recognition.	High short-term engagement; low trust and rising avoidance when alternatives appear.

**Table 1.** The two failure modes compared. Both look like success on metrics; neither is.

#### KEY TAKEAWAY

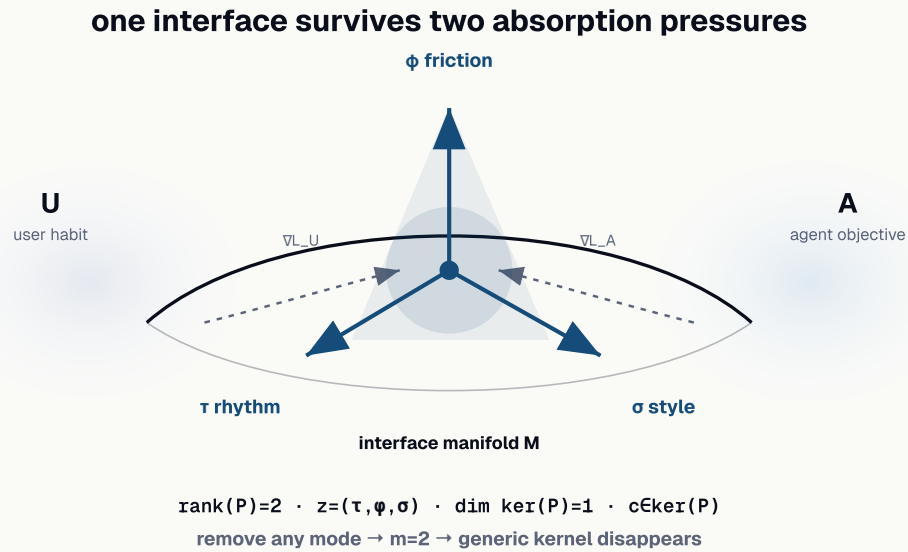
Two-body design has only two stable outcomes, and both are forms of failure. The interface that endures is one that resists being absorbed in either direction.

### 3. A Minimal Model of the Triad

Let  $U$  be the user state,  $A$  the agent state, and  $I$  the interface state. A conventional two-body model treats  $I$  as a function of  $U$  and  $A$ . AGORA treats  $I$  as a third state variable with its own admissible motion.

Collapse occurs when the motion of  $I$  becomes reducible to either adjacent pressure. If  $I$  follows only  $U$ , the surface becomes habit and disappears. If  $I$  follows only  $A$ , the surface becomes optimization and manipulates. Persistence requires a nonzero component of  $I$  orthogonal to both pressures.

FIG.1 · CONSERVATION PHASE SPACE



**Figure 1.** Conceptual conservation phase space. User and agent pressures are represented as two independent rows of the conceptual pressure matrix  $P$ . Rhythm, friction, and style are the measured interface coordinates ( $m=3$ ). By rank-nullity, a three-mode system admits a one-dimensional homeostatic kernel  $\ker(P)$  spanned by  $c$ ; for  $m=2$  the kernel is generically trivial. The diagram is illustrative:  $P$  is a local linearization of design trade-offs, not an empirical measurement.

## 4. The Conservation Principle

**Conservation of Interface Irreducibility.** A human-agent product persists as an interface when it maintains a nonzero interface state that is not determined by either user habit or agent optimization.

Operationally, an interface is irreducible when it preserves a named invariant that can be observed by a reader, measured by an instrument, and removed in an ablation. The invariant does not have to be mystical. It can be a schedule, a checkpoint, a recognizable voice, or any other product primitive that keeps the interface from becoming merely a habit or merely an output.

The pressure matrix is a deliberate formalization of design trade-offs, not a claim about physical force. Let  $L_U(z)$  be the local loss minimized when the interface becomes pure user habit, and let  $L_A(z)$  be the local loss minimized when the interface becomes pure agent optimization. At an interface state  $z$ , the two pressure rows are

the local gradients  $\nabla L_U(z)$  and  $\nabla L_A(z)$ . The matrix  $P$  is their first-order linearization. The full-rank case is exactly the case in which user and agent pressures are not the same pressure in disguise.

#### LEMMA 1 · THREE-EIGENMODE STABILITY LEMMA

Let  $z \in \mathbb{R}^m$  be the vector of empirically independent interface modes. Let  $P \in \mathbb{R}^{2 \times m}$  be the local pressure matrix whose rows are  $\nabla L_U$  and  $\nabla L_A$ , with  $\text{rank}(P) = 2$ . A nonzero interface-homeostasis vector  $c$  can be stable under both pressures only if  $c \in \ker(P)$ . By rank-nullity,  $\dim \ker(P) = m - 2$ . Therefore a generic two-mode interface has no nonzero kernel, while a three-mode interface has exactly one preserved direction. Three is the minimum dimension in which an interface can resist two independent absorption pressures without relying on a nongeneric cancellation.

The stability calculation is simple. Consider the homeostatic flow

$$\dot{z} = -K(z - c) - P^T P z, \quad K > 0, \quad P c = 0.$$

With  $V(z) = 1/2 \|z - c\|^2$ , the derivative along trajectories is

$$\dot{V} = -(z - c)^T (K + P^T P) (z - c) < 0$$

for every  $z \neq c$ . The equilibrium  $z = c$  is globally asymptotically stable. If  $m = 2$ , a generic rank-2 pressure matrix has  $\ker(P) = \{0\}$ ; the only compatible setpoint is  $c = 0$ , which is the absorbed interface. If  $m = 3$ ,  $\ker(P)$  is one-dimensional and can hold a nonzero interface invariant. A fourth mode may add redundancy, but it is not required by the algebraic obstruction.

The lemma is conditional, not decorative. It assumes that the measured modes are genuinely independent and that the two pressures are genuinely independent. If rhythm, friction, and style collapse empirically into a single latent factor, or if one pressure is not distinct from the other, the three-eigenmode claim weakens. That is why §8 makes the independence of the modes a falsifier rather than a protected definition.

## 5. The Three Eigenmodes

AGORA names the three product primitives that supply the minimal independent basis in Lemma 1. We use *eigenmode* in the operational sense: a mode must be independently measurable, independently ablatable, and capable of carrying part of

the interface's identity across sessions.

**Measurement bridge.** Let  $M$  be the manifold of observable interface traces. Define three functionals on  $M$ :  $\tau(x)$  for cadence signatures,  $\phi(x)$  for participation-threshold signatures, and  $\sigma(x)$  for recognition/style signatures. The mode-space map is  $z(x) = (\tau(x), \phi(x), \sigma(x))$ . Mathematically, the model requires the local Jacobian of this map to have rank 3. Operationally, the bridge is validated only when an ablation matrix shows that removing any one mode reduces predictive rank by exactly one while the other two modes remain independently measurable. In Crest,  $\tau$  is measured by issue cadence and voluntary waiting,  $\phi$  by checkpoint yield and recall after pause, and  $\sigma$  by blind masthead/voice recognition. If the ablation matrix collapses to rank 1 or 2, rhythm, friction, and style are merely qualitative surrogates, not operational modes for Lemma 1.

This basis is not claimed to be the only possible basis. It is the candidate basis this paper makes risky by naming in advance. A different measurable rank-3 basis that predicts persistence better would revise AGORA rather than violate it.

## 5.1 Rhythm

The interface has its own cadence. Sessions have a beginning, a middle, and an end that are not dictated by the user's impulse or by the agent's optimization. A daily ritual, a paced reveal, a deliberate slowness or quickness – the timing of the experience belongs to the interface.

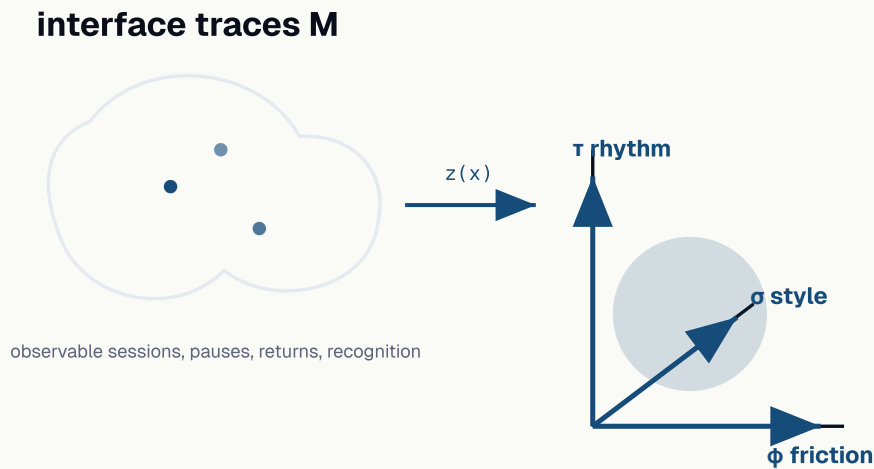
## 5.2 Friction

The interface contains resistances that produce depth. This is not artificial difficulty. It is the measurable threshold that turns passive exposure into participation. Friction transforms casual exposure into earned competence. An interface without friction produces no transformation in the user – only consumption.

## 5.3 Style

The interface speaks in its own voice. It has an aesthetic, a tone, a personality that is identifiably *the interface's*, not the user's, not the agent's. The user can recognize it instantly across mediums. The voice is consistent across sessions even as content changes.

These eigenmodes are not decorations. They are public product commitments. If any one is not independently measurable, the candidate basis is overnamed; if any one can be removed without changing persistence, the AGORA claim loses force.



**Figure 2.** The measurement bridge from product traces to mode space. Rhythm, friction, and style qualify as independent modes only when the empirical map  $z = (\tau, \phi, \sigma)$  yields an ablation matrix of full rank 3, where removing any single functional reduces predictive rank by one. Rank here is operational, not necessarily differential.

#### KEY TAKEAWAY

Under two independent absorption pressures, three independent interface modes are the smallest basis that leaves a nonzero conserved direction. Two collapses generically; four is allowed but not required.

## 6. Design Rule: Conserve the Kernel

The design rule implied by Lemma 1 is operational: preserve the interface kernel before optimizing features, personalization, or engagement.

**Do not maximize features.** A feature that does not preserve rhythm, friction, or style is additional absorption pressure.

**Do not maximize personalization.** Personalization that dissolves the interface's voice lowers the style coordinate even when users click more.

**Do not maximize engagement.** Engagement extracted by removing boundaries lowers autonomy when the product claims to be a durable place.

**Maximize autonomy margin.** The engineering target is a positive IAI under ablation, not a larger surface area of functionality.



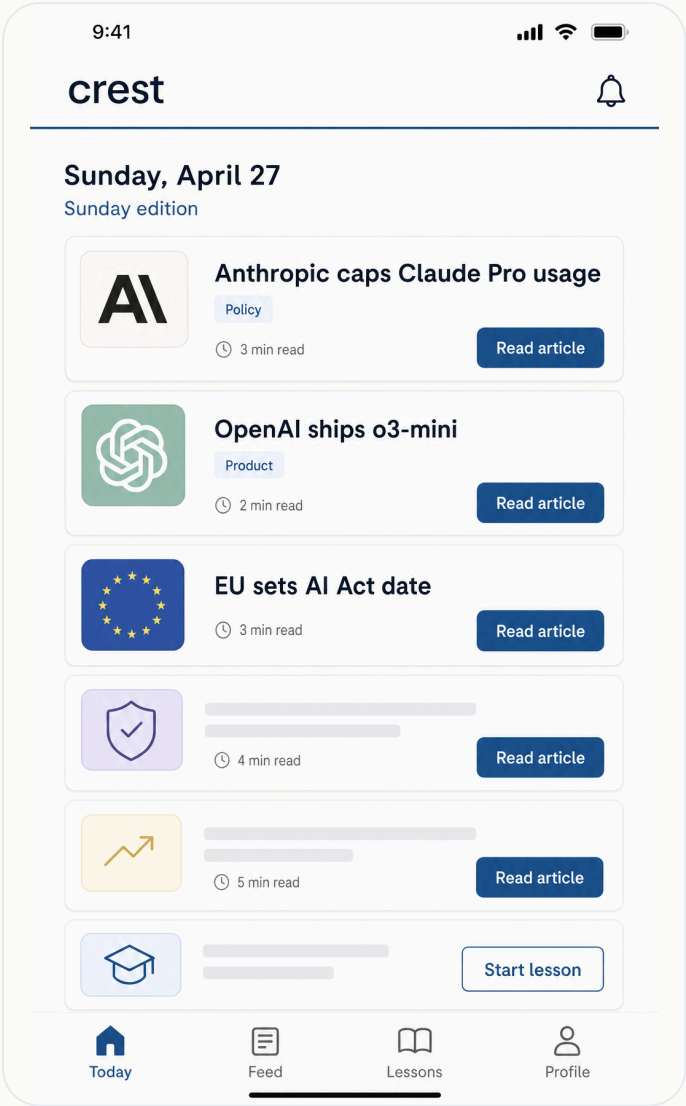
## 7. External Tests

Crest is not the only possible test bed. AGORA should score products the authors do not own. The table states external predictions in advance so the theory cannot hide inside one company.

System	AGORA prediction	Window and falsifying threshold
Substack publications	Durable publications conserve cadence and voice even when distribution algorithms change.	Eight-week reader study: blind voice recognition $\leq 55\%$ or cadence-unlinked return kills the claim for this class.
Are.na	Collection friction and recognizable community style should protect against pure feed replacement.	Thirty-day cohort: if board return equals a generic image feed within 3 percentage points, the predicted friction advantage is absent.
Readwise Reader	The product persists if review rhythm and library identity outrank summarization convenience.	A/B summarization test: if AI compression raises retention while lowering neither recall nor style recognition, the agent-absorption warning weakens.
Libraries and code review	Schedules, thresholds, and house style should predict durable participation after personnel turnover.	Newcomer study: if participants cannot identify rhythm, threshold, or style above chance after four sessions, the institutional extension fails.

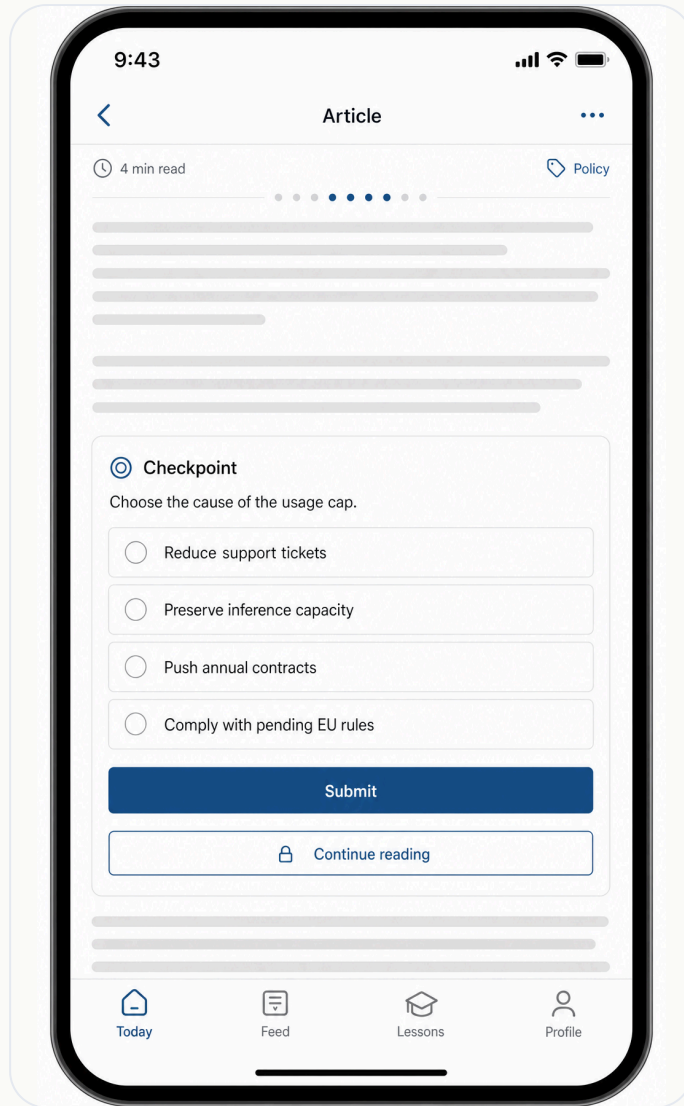
### 7.6 Apparatus Surfaces

These three Crest screens are not evidence. They are snapshots from the simulated product in `analysis/CREST_PRODUCT_SIMULATION.md`: a finite daily edition with six items, a required checkpoint, and a visible close. The test is whether changing those surfaces changes measured autonomy.



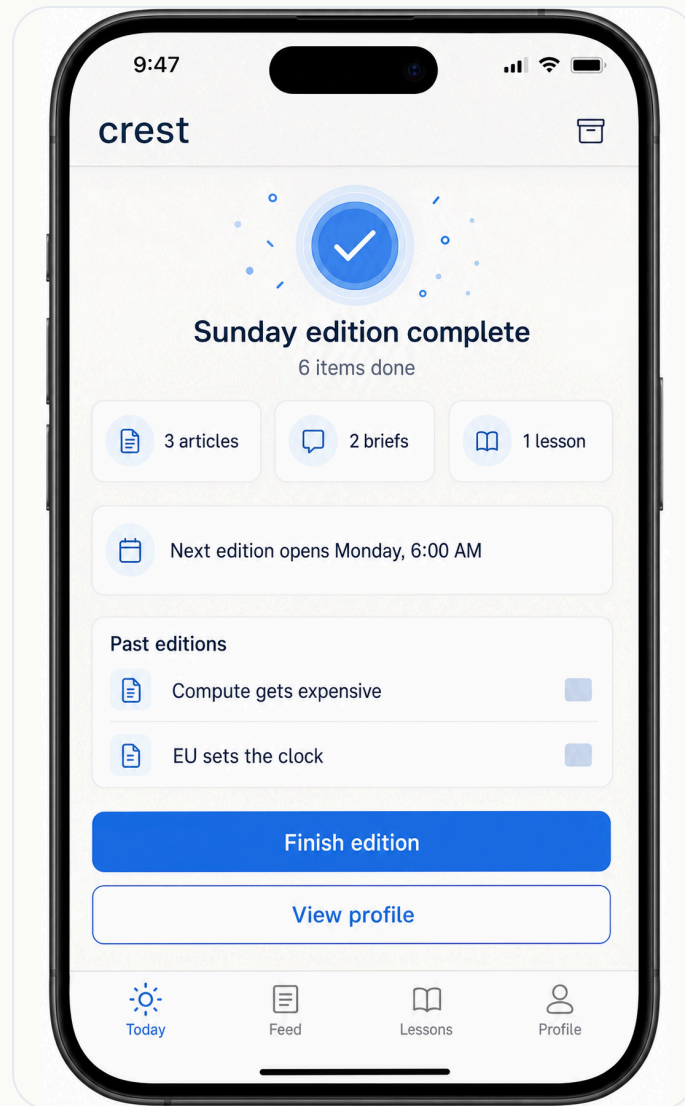
EMBODIMENT 1 ·  
SUNDAY EDITION  
OPEN

**Protocol surface: edition cadence.** Ablation removes the scheduled daily edition or replaces it with an unbounded feed; the measured target is return regularity and cadence recall.



EMBODIMENT 2 ·  
READER  
CHECKPOINT

**Protocol surface: checkpoint yield.** Ablation removes the required checkpoint before continuation; the measured target is recall, completion, satisfaction, and unprompted mention of the pause.



EMBODIMENT 3 ·  
SUNDAY EDITION  
CLOSE

**Protocol surface: bounded close.** Ablation replaces the close with continuation; the measured target is whether return and recognition persist without infinite-feed pressure.

## 7.7 The Crest Experiment and Public Commitments

**Crest** is the authors' preregistered falsifiable apparatus, not the only possible evidence for the AGORA framework. The conflict is real: the same people who state the principle are building one of the systems meant to test it. The response is not to hide the apparatus. The response is to publish the instruments, name the falsifiers, and compare Crest against systems the authors do not control.

Crest's interface invariants are fixed in advance: one edition opens on a scheduled cadence; the masthead and visual language remain stable across issues; each reading session has a closing action that refuses infinite continuation. The first public

deposit at the DOI of this paper is the opening lab-notebook entry. Subsequent entries should record changes in the apparatus, measurement windows, and whether the evidence strengthens or weakens the law.

Invariant	Measurement	Falsifying signal
Scheduled edition rhythm	Return regularity, voluntary waiting, and recall of the cadence	Users bypass or ignore cadence without loss of retention or recognition
Closing checkpoint	Completion, recall, satisfaction after pause, and unprompted mention of the gate	Removing the checkpoint has no effect on memory, return, or perceived place
Stable voice and masthead	Blind recognition above chance against generic reading feeds	Users cannot distinguish Crest from a generic feed or AI digest

If Crest endures while conserving these invariants, AGORA gains conflicted but useful evidence. If Crest endures after the invariants are removed, or fails while they are conserved, AGORA loses evidence. Independent teams can run the same instruments on Substack, Are.na, Readwise Reader, libraries, code review systems, or any human-agent product that claims to be a place.

## 8. Calculations and Falsifiers

AGORA's attack surface is post-hoc explanation. The countermeasure is to publish the exact margins that would make the theory lose. The calculations below are preregistered decision rules, not fitted results.

Let  $z=(R,F,S)$  be the preregistered 0-1 measurement vector for the three modes. Let  $U_a$  and  $A_a$  be preregistered 0-1 measures of user- and agent-absorption pressure, normalized to the same support. Define the Interface Autonomy Index as a directional margin:

$$IAI = (R + F + S) / 3 - \max(U_a, A_a)$$

The max operator encodes a worst-case absorption assumption; a sum formulation may be preregistered, but the directional sign of IAI and its rank correlation with persistence are the primary falsifiable signals, not its absolute magnitude. The equal

weights are a hostile default: AGORA claims all three modes are necessary, so the preregistered audit gives none of them privileged rescue power.

**Measurement instruments.** R is the share of users who correctly identify the interface cadence and voluntarily return on schedule. F is checkpoint yield: completion, recall, and satisfaction after a preregistered pause. S is blind recognition of masthead, voice, or house style against generic alternatives.  $U_a$  is replaceability elasticity under lower-friction substitutes.  $A_a$  is optimization pressure: engagement lift paired with lower trust, recognition, or recall.

**Analysis model.** The primary persistence model is logistic regression for four-week return, with covariates R, F, S,  $U_a$ ,  $A_a$ , domain fixed effects, and preregistered interaction terms only if listed before outcome inspection. A secondary Cox model estimates time-to-churn. Predictive rank is the rank of the standardized mode block plus cross-validated change in AUC; removing a mode must reduce predictive contribution without being spanned by the remaining modes.

**Sampling and power.** External validation uses a 50-system sample: ten editorial/reading systems, ten knowledge-management systems, ten social-feed systems, ten institutional/workflow systems, and ten AI-agent products, all public, active for at least six months, and measurable without privileged access. Each system requires at least 30 active or recent users for the primary field protocol; systems below 20 validated respondents are excluded from the primary correlation. Crest invariant A/B tests require at least 200 users per arm or an equivalent preregistered sequential design with 80% power to detect Cohen's  $h = 0.20$  for proportions or Cohen's  $d = 0.20$  for means at  $\alpha = 0.05$ .

**Deposited instrument contract.** The enumerated 50-system sample, item-level R/F/S/ $U_a$ / $A_a$  instruments, Crest randomization protocol, model formulas, and anti-tamper commitment are fixed in `analysis/AGORA_PREREGISTRATION_PLAN.md` and `analysis/agora_prereg_models.py`. Those files are part of the candidate deposit and precede outcome inspection.

**Subjectivity controls.** Qualitative items require two blind coders, third-coder adjudication, and Cohen's  $\kappa \geq 0.70$  before use. Ablation spanning is tested with canonical correlation  $\leq 0.80$ , VIF  $\leq 5.0$ , Holm-Bonferroni-corrected nested tests, and  $\Delta AUC \geq 0.03$ . The rows of  $P$  are estimated by finite differences under lower-friction substitute interventions and unbounded-personalization interventions;  $U_a$  and  $A_a$  are preregistered scalar summaries of those rows.

**Executable trace construction.** `analysis/agora_trace_operationalization.py` defines the trace-to-metric functions, finite-difference pressure-row estimator, ablation-spanning tests, and a deterministic synthetic sanity check. The synthetic output is not evidence for AGORA. It proves that the quantities are computable from observable traces and that the preregistered pipeline would flag a collapsed single-latent-factor case.

Attack	Decision rule	What kills AGORA
Mode independence	Ablation matrix over R,F,S; each removed mode should reduce predictive rank by $\geq 0.15$ IAI units and off-diagonal substitution should stay $\leq 0.05$ .	Any mode is redundant, or PCA predicts persistence with one latent factor and $\leq 5\%$ loss.
Crest invariants	Eight-week A/B per invariant: rhythm, checkpoint, and close each compared against a functionally equivalent removal.	Removal changes four-week return by $< 3$ percentage points and recognition/recall by $< 0.05$ , with a 95% confidence interval excluding a minimal preregistered effect size (Cohen's $h \geq 0.20$ for proportions, Cohen's $d \geq 0.20$ for means).
Agentic personalization	Bounded vs unbounded personalization compared on S and F while holding content quality constant.	Unbounded personalization raises retention while lowering neither S nor F.
Generalization	Fifty-system sample, six-month persistence outcome, preregistered IAI before outcome inspection.	Spearman $\rho(\text{IAI}, \text{persistence}) \leq 0.20$ or confidence interval crosses zero.
Weight sensitivity	Monte-Carlo sensitivity: draw 10,000 weight vectors uniformly from $w \in [0.2, 0.5]^3$ , $\text{sum}(w)=1$ .	The sign or rank ordering of IAI changes exclusively in Crest's favor.

These thresholds are intentionally sharp. A product can still be good if AGORA fails. The claim that dies is narrower: that independent rhythm, friction, and style are the minimal measured basis for durable human-agent interfaces. A high IAI-persistence correlation is not sufficient support if the three modes collapse to one latent factor; in that case the measurement may predict persistence while the eigenmode claim fails.



## 9. Limits

This theory is descriptive of what persists, not prescriptive of what should always be built.

There remain categories of software where one of the failure modes is not a failure but a virtue. A calculator should be user-absorbed: a great calculator disappears in the hand. A fire alarm should be agent-optimized: it should manipulate the user toward immediate action with no friction. The principle of interface irreducibility applies to any system intended to *endure as a presence in the user's life* - to be a place rather than a moment.

We make no claim that all software must become a place. We claim only that systems aspiring to lasting cultural and economic presence must preserve an interface kernel, and that violations of this requirement predict collapse.

Software type	Should be	Why
Calculator, OS shell, file manager	User-absorbed (tool)	Meant to disappear in the hand. The interface should not push back; that would be friction without purpose.
Fire alarm, emergency banner, payment failure modal	Agent-optimized	Urgency is the goal. Interrupt the user. Friction would cost a life or a transaction.
News feed, learning app, social platform	Agora	Meant to endure. Without conservation it becomes engagement bait or generic substitutable software.
Editorial publication, journal, paper-of-record	Agora	The voice is the product. Personalization without conservation kills the institution.
Productivity tool, IDE, design canvas	Mixed (tool with agora layer)	The work surface should disappear; the surrounding institution (community, conventions, voice) should be an agora.
Search engine, lookup utility	User-absorbed (with one agora seam)	The query result should be transparent. The brand of the search service must remain itself or it gets replaced.

**Table 2.** Where the interface-kernel requirement applies and where it does not.

Some software is meant to disappear; some is meant to interrupt; only systems meant to endure must preserve an interface kernel.



#### KEY TAKEAWAY

The theory is descriptive of what persists, not prescriptive of what should always be built. Calculators should disappear. Fire alarms should manipulate. Only software meant to endure as a presence in your life must be an agora.

## 10. Conclusion

We have proposed interface irreducibility as a design hypothesis for human-agent software. The claim is simple: an interface that cannot preserve its own motion under user and agent pressure will be absorbed by one of them.

The proposed basis is rhythm, friction, and style. The proposed apparatus is **Crest**. The proposed standard is falsifiable: if the basis does not predict persistence, or if the apparatus succeeds without conserving it, the principle must be revised.

AGORA is therefore not a slogan for making software feel meaningful. It is a research program for measuring whether an interface can remain a place when the user and the agent both try to make it disappear. This paper is the first lab notebook entry, not a claim that the experiment has already succeeded.

## 11. Related Work

AGORA builds on affordance theory, pattern language, and third-place sociology, then formalizes pressures documented in adjacent empirical literatures. Amershi et al. (2019) show how human-AI systems can override interface stability; Mathur et al. (2019) document dark-pattern dynamics at scale; Zuboff (2019) describes market pressure toward extraction. The contribution here is narrower: a minimal model specifying the dimensional condition for interface irreducibility.

## Appendix A: Proof of Lemma 1

**Lemma.** Let  $L_U$  and  $L_A$  be differentiable local absorption losses and let  $P$  be the  $2 \times m$  matrix with rows  $\nabla L_U$  and  $\nabla L_A$ . Let  $P$  have rank 2, let  $K \in \mathbb{R}^{m \times m}$  be symmetric positive definite, and let  $c \in \mathbb{R}^m$  satisfy  $Pc = 0$ . The system  $\dot{z} = -K(z-c) - P^T Pz$  has a unique equilibrium at  $z = c$ , and that equilibrium is globally asymptotically stable. There exists a choice  $c \neq 0$  if and only if  $m \geq 3$ .

Set  $\dot{z} = 0$ . Then  $(K + P^T P)z = Kc$ . The matrix  $K + P^T P$  is positive definite because  $K$  is positive definite and  $P^T P$  is positive semidefinite. It is therefore invertible. Since  $Pc = 0$ ,  $P^T Pc = 0$ , and  $(K + P^T P)c = Kc$ . Thus the unique equilibrium is  $z = c$ .

Define  $V(z) = 1/2 \|z - c\|^2$  and let  $x = z - c$ . Differentiating  $V$  along trajectories gives  $\dot{V} = x^T[-Kx - P^T P(x+c)]$ . The compatibility condition  $Pc = 0$  removes the cross term, so  $\dot{V} = -x^T Kx - \|Px\|^2 = -x^T(K + P^T P)x$ . Because  $K + P^T P$  is positive definite, there is  $\lambda > 0$  with  $\dot{V} \leq -\lambda \|x\|^2$ . Lyapunov's direct method gives global asymptotic stability; the same inequality gives exponential convergence by comparison.

It remains to count dimensions. A nonzero conserved setpoint requires  $c \in \ker(P) \setminus \{0\}$ . By rank-nullity,  $\dim \ker(P) = m - \text{rank}(P) = m - 2$ . If  $m = 2$ , the kernel is generically trivial and  $c = 0$  is forced; the stable equilibrium is the absorbed origin. If  $m = 3$ , the kernel is one-dimensional, so a nonzero  $c$  exists and the preceding stability argument applies. Any  $m > 3$  also works, but the existence requirement is already satisfied at  $m = 3$ . The result is therefore a minimality statement under the stated assumptions, not a universal claim about every possible interface representation.

**SymPy check.** The deposited build includes `papers/build/v5.0.0/math/sympy-three-eigenmode-check.py`, which verifies  $P(c) = 0$  for a symbolic  $2 \times 3$  pressure matrix, derives the Lyapunov derivative, and records the  $m = 2$  determinant condition. The symbolic check does not prove the empirical independence of rhythm, friction, and style; it verifies the algebra used after that assumption is made.

## References

- Alexander, C., Ishikawa, S., & Silverstein, M. (1977). *A Pattern Language: Towns, Buildings, Construction*. Oxford University Press.
- Hubinger, E., van Merwijk, C., Mikulik, V., Skalse, J., & Garrabrant, S. (2019). Risks from Learned Optimization in Advanced Machine Learning Systems. arXiv:1906.01820.
- Norman, D. A. (1988). *The Psychology of Everyday Things*. Basic Books.
- Oldenburg, R. (1989). *The Great Good Place: Cafés, Coffee Shops, Bookstores, Bars, Hair Salons, and Other Hangouts at the Heart of a Community*. Paragon House.
- Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Horvitz, E., et al. (2019). Guidelines for Human-AI Interaction. *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1-13.
- Mathur, A., Acar, G., Friedman, M. J., Lucherini, E., Mayer, J., Chetty, M., & Narayanan, A. (2019). Dark Patterns at Scale: Findings from a Crawl of 11K Shopping Websites. *Proceedings of the ACM on Human-Computer Interaction*, 3(CSCW), 1-32.
- Zuboff, S. (2019). *The Age of Surveillance Capitalism*. PublicAffairs.
- Victor, B. (2011). Explorable Explanations. worrydream.com.

## Version History

- **v5.0.0** · 2026-04-26 - Stability and independence candidate. Adds Lemma 1, Appendix A proof, SymPy algebra check, public product audits, external replication targets, and a standalone Conflict-of-Interest and Independence statement.

## Conflict-of-Interest and Independence Statement

Andy Salvo is the founder of Polylogic AI and a co-founder of Crest. Jameson Ackerman is a co-founder of Crest. AGORA is authored by Andy Salvo and Jameson Ackerman. The authors therefore have a structural commercial conflict of interest: Crest is named as a preregistered apparatus and the authors may benefit if the AGORA framework is influential.

The conflict is structural and financial: the authors hold equity in Crest and may benefit directly from adoption of this framework. This paper is therefore a conflicted prospectus, not independent evidence. Crest cannot confirm AGORA; it can only fail it. Any evidence originating from Crest must be treated as adversarially generated until reproduced by external teams using the public instruments in §7-8. The conflict is the reason the paper makes the following public commitments: (1) Crest evidence should be treated as conflicted until reproduced by external teams using public instruments; (2) §8 names results that would count against AGORA before those results are known; (3) §7.5 names comparison systems the authors do not control; and (4) future claims should link to the DOI-stamped lab notebook rather than to private interpretation.

**Independence protocol.** Independent observers can invalidate the framework by showing that rhythm, friction, and style collapse into one latent factor; that removing any one mode has no effect on Interface Autonomy Index; that low-IAI products persist as well as high-IAI products; or that Crest persists after its named invariants are removed. Those outcomes should revise the paper, not be explained away.

**Author contributions.** Conceptualization, theoretical synthesis, and figure design: Salvo. Editorial review, section structuring, and prior-art clearance: Ackerman. Writing of the paper: jointly.

No external funding was received for this work. The commercial relationship to Crest and Polylogic AI is the declared conflict of interest.