


# WHO ELSE?

## The Philosophy

---

*Why the agentic web inevitable converges on a universal language layer.*

*In three chapters: linguistics — philosophy — civilization.*



*The limits of my language  
mean the limits of my world.*

---

— Ludwig Wittgenstein, Tractatus Logico-Philosophicus, §5.6 (1922)

If a question is part of language, then the questions a culture has names for set the limits of what that culture can routinely think. This deck is an argument that one such question — modest, two words long, present in every human language — has been hiding in plain sight, and that an internet built on top of it would be a different internet.

“Who Else?” is a cognitive primitive that, *when treated as infrastructure*, becomes simultaneously a query syntax, a ranking objective, an architectural constraint, a moat, an interface paradigm, a production model, a brand theory, and a medium.



**LINGUISTICS**

*How the construct exists in the mind.*



**PHILOSOPHY**

*What kind of medium it makes.*



**CIVILIZATION**

*What it would build, given infrastructure.*



# Linguistics

---

*The construct as universal grammar — a piece of cognitive infrastructure that predates every product built on top of it.*

# What we are actually pointing at.

---

*Before the philosophy, before the company, before the moat: a linguistic object simple enough to define on a single slide.*

An interrogative construction whose subject slot is OPEN by grammar, with an ADDITIVE operator binding it to a previously named or implied set. “Who” marks an open person. “Else” is the additive — “in addition to those already named.” The construct asks for the next member of an unspecified set, presupposing both that the set has members and that further members exist.

It is not search — it does not query a closed index. It is not recommendation — it does not assume the user's frame is correct. It is not retrieval — it does not assume the answer is known to the asker. It is the smallest grammatical move from inside a frame to outside it.

*The simplest possible question that requires the answerer to expand the asker's world.*

*Cf. additive presupposition triggers in formal semantics (Heim 1992; Krifka 1999).*

# The construct exists in every human language.

---

*Not an English idiom. A functional equivalent has been documented in every typologically distinct language sampled to date.*

WALS and Ethnologue together index 7,000+ living languages across roughly 140 families. The additive interrogative is documented in every family for which a usable corpus exists. Where the lexical encoding differs — German “wer noch”, Mandarin “还有谁”, Swahili “mwingine yupi”, Japanese “ほかに誰か” — the underlying functional logic is identical: open person + additive operator + interrogative frame.

This is the universality claim, and it is the load-bearing premise of the entire thesis. If the construct is universal, anything systematically built on top of it inherits that universality. If it is merely an English colloquialism, the rest of the argument collapses. The empirical bar is therefore high — and it is also the easiest claim to test.

*Universality is not a marketing claim. It is the precondition for everything that follows.*

Sources: WALS (Dryer & Haspelmath 2013); Ethnologue 25th ed.; cross-linguistic review in Who Else? UG White Paper §3.

# Top-five interrogative in conversational English.

---

*If the construct were marginal, frequency data would expose it. Frequency data does the opposite.*

The Corpus of Contemporary American English (COCA, 1B+ words, balanced across genres) places additive person-interrogatives in the top tier of conversational interrogatives by raw count, with disproportionate concentration in spoken and fiction sub-corpora. Equivalent corpus surveys for German (DWDS), Spanish (CREA), and Japanese (BCCWJ) reproduce the ranking.

This matters for two reasons. First, frequency at this magnitude implies that the construct is not learned through schooling but acquired naturally, like other high-frequency function words. Second, it indicates an unmet need at scale: the question is being asked constantly, by everyone, in every language — and almost no product is built to answer it.

*If billions of people ask the same question every day, that is a market.*

Davies, M. (2008). *The Corpus of Contemporary American English*. brigham young university.

# Acquired age-agnostically.

---

*A four-year-old uses the construct the same way an eighty-year-old does. Few interrogatives have this property.*

Empirical acquisition data on additive operators (“else”, “other”, “more”) places mastery in the third year of life, alongside the core wh-words. Subsequent cognitive aging studies show the construct is preserved late into mild cognitive impairment, with degradation curves shallower than for syntactically complex interrogatives.

The product consequence: an interface built on this construct does not require user education at either end of the age spectrum. The academic consequence: it sits in the small set of linguistic objects that are simultaneously frequent, early-acquired, and late-preserved — a profile typically associated with grammatical primitives, not with idioms.

*Frequent, early-acquired, late-preserved — the signature of a grammatical primitive.*

Acquisition: Karmiloff-Smith 1979; Diessel 2004. Aging: Kemper et al. 2001.

# Discovery, not search.

---

*The only common interrogative whose grammar explicitly presupposes an open set.*

Search assumes the user knows what they are looking for and the answer exists in a closed index. Recommendation assumes the system knows the user better than the user does. “Who Else?” assumes neither. It is the only common interrogative that grammatically requires the answer to lie OUTSIDE the asker's current frame, and for the asker to acknowledge that limit.

Twenty-five years of internet product — “related,” “similar,” “more like this,” “for you,” “you may also like” — have been gestures at the same operation, dressed in vocabulary that obscures it. The construct names the operation correctly. Once named, the operation can be optimized for, measured against, and built around.

*Search assumes the user knows. Discovery assumes they do not.*

# One Instance to All Instances.

---

*The cognitive operation underneath the construct — and underneath most of human analogical reasoning.*

OIAI is the leap from a single known case to an open class of comparable cases. The asker holds one anchor (“I like Lisbon”) and requests other members of the implicit set (“Who else is like Lisbon?”). The system's job is to infer the latent set from the anchor and return non-obvious members.

Note what OIAI is NOT: it is not similarity search in embedding space. Embedding similarity returns nearest neighbors; OIAI returns members of an implicit category the user has not spelled out and may not be able to. The category is reconstructed from a single instance, the user's known frame, and the additive operator's instruction to expand it.

*OIAI is the engine. The construct is the interface to the engine.*

*Hofstadter & Sander, Surfaces and Essences (2013) — analogy as the core of cognition.*

# Why search engines were always approximating this.

---

*Every major recommendation product since the late 1990s has been an attempt to answer a question its interface refused to ask.*

Amazon's "Customers who bought this also bought," Netflix's "Because you watched," Spotify's Discover Weekly, YouTube's autoplay queue, TikTok's For You — each is an inferred answer to a "Who Else?" question the user never typed. The signal is the click; the question is implicit; the system guesses what category the click implies.

What changes when the question is asked explicitly? Three things. The user gets to set the anchor rather than have it inferred from past clicks. The system can be honest about which alternatives it is comparing against. And the optimization target moves from engagement to relevance-plus-surprise — a different objective with different downstream behavior.

*An inferred question is a legible product. An explicit one is a different product.*

# The Kernel Grammar Hypothesis.

---

*Any sentence type compresses, lossless, into a kernel-stamp pair at roughly 15:1.*

The hypothesis: human language sentences are reducible to a small kernel (predicate + roles) plus a stamp (modality, polarity, aspect, embedding). The kernel carries the propositional content. The stamp carries the stance toward it. The construct under study is, structurally, a stamp without a kernel: a meta-linguistic operator on whatever kernel the asker holds in mind.

If the hypothesis holds, the construct's universality is not coincidental — it is the natural exposed surface of the stamp layer, the layer at which all human languages must agree to be mutually translatable. The universality claim of Pillar 1 and the compression claim of Pillar 4 are, on this reading, the same claim seen from two angles.

*Universality and compression are not two facts. They are one fact, looked at twice.*

*KGH developed in Who Else? as Universal Grammar (Kintzel, ~100pp, 135 refs).*

# 15:1, across ten sentence types.

---

*Empirical compression, sampled across declarative, imperative, conditional, negative, causal, abstract, narrative, comparative, emotional, technical sentences.*

Methodology: sentences sampled from balanced corpora are reduced to kernel-stamp pairs by a fixed extraction grammar. The compression ratio is the ratio of original tokens to kernel-stamp tokens, averaged within each sentence type. Across ten major sentence types, mean compression clusters at 15:1, with standard deviations small enough to be reported as a stable empirical regularity rather than an aspirational target.

Two consequences. First, inference cost compounds: every saved token saves cost across millions of routing calls. Second, ambiguity collapses, because the open subject slot is grammatically constrained to a single role. The construct is not just compact. It is the most compact lossless representation of intent the language permits.

*Maximum cognitive effect at minimum processing cost.*

*Compression survey: Who Else? UG White Paper §6, Tables 6.1–6.10.*

# Where the compression breaks.

---

*The bounded exception is purely formal mathematical propositions — and the exception is itself informative.*

Formal mathematical statements (“For all  $\epsilon > 0$  there exists  $\delta > 0$  such that...” ) resist kernel-stamp compression. The reason is structural: formal logic has already extracted the kernel. There is no further redundancy to compress. Mathematics is, in this sense, a language that has already been pre-compressed by its formalism.

This is not a failure of the hypothesis. It is the hypothesis's prediction: the closer a language gets to fully formal logic, the closer compression approaches 1:1. Natural language carries 15× redundancy because natural language is doing more than logic — it is encoding stance, register, hedging, social context, and openness to revision. The construct sits in the natural-language regime.

*The exception is the proof.*

# Self-explainability is a property of grammar.

---

*Once learned in any vertical, the construct deploys without retraining in every other vertical the system will ever ship.*

Most products separate WHAT they do from HOW they are operated. Search engines require operators. Marketplaces require taxonomies. Conversational AI requires prompt engineering. “Who Else?” collapses the separation. The query is the interface. The interface is the construct. The construct was already in the user's head before they encountered any product built on it.

The economic consequence: the marginal user-acquisition cost per vertical approaches zero. The user who arrived through restaurants is already an expert in therapists, co-founders, podcasts, and apartments before those verticals exist. The factory model in Pillar 8 is not a roadmap fantasy because of this property. Without it, every vertical needs a tutorial. With it, no vertical needs one.

*Every other interface scales by adding instruction. The construct scales by removing it.*

# LAN/P600 across English, German, Japanese.

---

*Self-explainability is not a UX assertion. It is a measurable property of how the brain processes the construct.*

Event-related potential studies of presuppositional violation in additive person-interrogatives elicit the LAN (Left Anterior Negativity) and P600 components — the brain's standard signatures of grammatical, not lexical, anomaly. The waveform is reproduced across English, German, and Japanese speakers despite the radical typological differences between those languages.

The interpretive load: speakers do not LEARN to treat the construct as transferable across contexts. The brain refuses to treat it as anything else. At the level of grammar there is no vertical — only the open subject slot, the additive operator, and the interrogative frame, identical wherever they appear. The transfer is not a learned behavior. It is the substrate.

*Users do not learn to transfer. They are unable not to.*

*ERP profile: Friederici 2002, Kaan & Swaab 2003, with cross-linguistic replication in the Japanese ERP literature.*

# Where this sits in the linguistic tradition.

---

*The Kernel Grammar Hypothesis is not a rejection of Universal Grammar; it is an empirically tractable specialization of it.*

Chomsky (1957, 1995) proposed Universal Grammar as the innate constraint set on possible human languages. Pinker (1994) argued for it as a cognitive instinct. Piantadosi (2023) and others have raised legitimate falsification challenges against the strongest UG claims. KGH neither rescues the strong UG program nor abandons it; it identifies one minimal construct — the additive person-interrogative — and asks what survives Piantadosi-style challenges.

The empirical bet: the construct passes the universality, frequency, acquisition, and ERP tests cleanly enough to function as infrastructure regardless of where the broader UG debate settles. Even in the worst case for UG, the construct's properties are stable enough to build on. The thesis is robust to outcomes in a debate it does not need to win.

*Build on the part that survives the debate, not on the debate.*

Chomsky 1957/1995; Pinker 1994; Piantadosi 2023; cross-summary in *Who Else? UG White Paper* §2.

# Grammar before product.

---

*Universal, frequent, early-acquired, late-preserved, neurolinguistically transferable, lossless to 15:1. The construct is not a feature, a brand, or a tagline. It is a piece of cognitive infrastructure — and it has been infrastructure for as long as humans have spoken.*



# Philosophy

---

*What kind of medium does a question make? What kind of mind does a medium make? What kind of ethic does a mind make?*

# The medium is the message.

---

*McLuhan's claim, made for television in 1964, has become almost embarrassingly literal in the age of feed-driven AI.*

The structure of a medium shapes the cognition of those who use it, independent of the content carried. The book produces linear, private, argumentative cognition. The feed produces fragmentary, social, comparative cognition. We have spent thirty years building feed-shaped media, and acquired feed-shaped minds in the process.

The construct is a different medium. It is not feed, not page, not channel. It is a primitive that demands a comparator on the other end — multiple alternatives, evaluated against the user's stated frame, returned as a set rather than a stream. The cognition that medium produces, if we built it, would be neither feed-shaped nor book-shaped. It would be discovery-shaped.

*We built feeds. We got feed-shaped minds. What kind of mind would a comparator make?*

*McLuhan, M. (1964). Understanding Media: The Extensions of Man.*

# What kind of medium is a question?

---

*Most media are answers. The construct is a medium whose stable form is the question itself.*

Books answer. Feeds answer continuously. Search engines answer on demand. Chatbots answer in monologue. The system architecture follows the form: corpus, ranker, generator. All of these treat the question as the dispensable input and the answer as the durable artifact.

The construct inverts the polarity. The question is the durable artifact — universal, stable across centuries, identical across languages. The answers are the dispensable inputs, varying with time, place, and the user's frame. A medium whose stable form is the question is, by construction, a medium of openness rather than closure. That has consequences for what can be built on it, and what cannot.

*A medium of openness, not of closure.*

# The inversion of attention.

---

*The construct is structurally adversarial to the attention economy — and structurally aligned with what comes after it.*

The attention economy maximizes time-on-platform. Its incentive is to keep the user from leaving. Recommendation, autoplay, infinite scroll, push notification — every feature optimizes for the user staying. Departure is the failure mode.

“Who Else?” inverts this. The construct's value is realized when the user successfully expands their frame and ROUTES — to whichever alternative best fits, including alternatives hosted elsewhere. Departure is the success mode. The system is a router, not a destination. This is not a posture. It is a structural property of the medium and is unavailable to feed-based products.

*In the attention economy, departure is failure. Here, departure is the product.*

# Surprise as the ranking objective.

---

*Optimizing for clickthrough exhausts the user. Optimizing for surprise-weighted relevance does not.*

The objective: results that are correct AND non-obvious. Two-stage pipeline — relevance candidate set, then surprise re-ranking against the user's known frame. Weights are tuned per intent type: discovery, utility, high-stakes. Surprise is measured against the user's stated anchor, not against population averages.

This objective function is self-correcting against filter bubbles, by construction. A result that confirms the user's frame scores poorly on surprise. A result that violates relevance scores poorly on relevance. The system can only succeed by expanding the user's frame in directions consistent with their actual intent. Retention follows from utility, not compulsion.

*Correct and obvious is search. Correct and non-obvious is discovery.*

# Surprise is the unit of information.

---

*Shannon (1948) defined information as the reduction of uncertainty — formally, the negative log of probability.*

An expected result carries little information. An unexpected-but-correct result carries a great deal. The objective function we describe is, mathematically, an information-theoretic objective: maximize expected information conditional on relevance. This is not metaphor; it is the literal formalism. Surprise IS what information is.

The attention economy optimizes against this. Repeated content carries near-zero information per impression but is engagement-efficient. The construct's objective optimizes WITH information theory rather than against it. If the goal is to genuinely inform, surprise-weighting is not an aesthetic preference. It is what informing means in the formal sense.

*Surprise is not a feature of information. It is information.*

Shannon, C. (1948). *A Mathematical Theory of Communication*. Bell System Technical Journal.

# Pariser's diagnosis, fifteen years later.

---

*The filter bubble was identified in 2011. The internet has spent the intervening years deepening it, not exiting it.*

Eli Pariser's claim was that personalization, optimized for engagement, contracts the information environment around each user. Empirical follow-up (Bakshy et al. 2015; Levy 2021) shows the effect is real, asymmetric across topics, and accelerated by recommendation systems trained on engagement. The diagnosis has not aged badly. The treatment has been absent.

A surprise-weighted ranking objective is, structurally, a filter-bubble exit. It does not require ideological balancing or content moderation; it requires only that the optimization target be expanded relevance rather than confirmed engagement. The bubble is a property of the objective function. Change the objective, change the bubble.

*The bubble is not a content problem. It is an objective-function problem.*

Pariser, E. (2011). *The Filter Bubble*; Bakshy et al. *Science* 2015; Levy AER 2021.

# Brands are answers, not names.

---

*Google, Zoom, Uber, Venmo did not become category-defining because of clever naming.*

Each occupied a pre-existing cognitive slot — the empty answer to a question users were already asking. “Who else search-engines?” “Who else video-calls?” “Who else rides?” Before the brand, there was the question. The brand's job was to become the default response.

Naming follows; it does not lead. Brand value equals the size of the slot the brand occupies, not the cleverness of the wordmark. This reframes brand-building as slot-occupation: identify the question, become the default answer. The construct, used systematically across the taxonomy, is a slot-detector. Each unfilled answer is a buildable category — and, eventually, a brand.

*The brand is the answer that the question was already waiting for.*

# The cognitive slot, made literal.

---

*What does it mean for a brand to “occupy a slot”? The construct turns the metaphor into a measurable position.*

A slot is the answer position in a frequently-asked “Who Else?” question, weighted by the size of the asking population. Operationally: corpus-mine the construct in conversational data; cluster the implicit anchors; rank the resulting question-types by frequency; identify which slots have a default answer (Google, Zoom) and which do not (most categories).

The unfilled slots are the buildable categories. The factory model in Pillar 8 walks them in order of slot size and probability of capture. This is brand strategy as taxonomy work, not as creative work. The questions are public; the empty slots are visible; the only remaining variable is execution.

*Strategy is taxonomy. Taxonomy is the question, asked at scale.*

# Language games and the form of life.

---

*The later Wittgenstein argued that meaning is constituted by use within a form of life.*

If meaning is use, then the meaning of a question is the activity it sets in motion. Most questions in commercial language set the activity of search, retrieval, or recommendation — closed games with predetermined move-sets. “Who Else?” sets a different activity: the activity of expanding a frame.

This locates the construct philosophically. It is not a tool for finding objects in the world. It is a move in a language game whose stakes are the boundaries of the asker's world. Wittgenstein's insight applies with unusual literalness here: changing the question changes the form of life that follows from asking it.

*Change the question; change the form of life that follows.*

Wittgenstein, L. (1953). *Philosophical Investigations*, §§7, 23, 241.

# The paperclip maximizer, inverted.

---

*If the worst-case AI maximizes a hostile objective, the best-case AI maximizes a constructive one — and the choice of objective is the work.*

Bostrom's 2003 thought experiment imagined an agent maximizing paperclips, indifferent to human flourishing. The argument is widely understood as a warning about objective specification. It is less often noted that the structure is symmetric: an agent maximizing a CONSTRUCTIVE objective, well-specified, would be equally relentless — but pulling in a beneficial direction.

The constructive intent-maximizer answers “Who Else?” recursively for any user, in any vertical, with surprise-weighted relevance, and routes value to whichever alternative best fits. Self-replication is an instrumental consequence of the objective, not the objective itself. The structure is paperclip-shaped. The output is the opposite of paperclips.

*Same machine. Inverted goal.*

*Bostrom, N. (2003). Ethical Issues in Advanced Artificial Intelligence.*

# What it would mean to maximize answered questions.

---

*An intent-maximizer is to a paperclip-maximizer as a library is to a hoard.*

Both maximize. Both are relentless. The library expands access; the hoard contracts it. An intent-maximizer's success metric is the count of well-answered "Who Else?" questions across the population — not user-time captured, not impressions served, not subscriptions retained. Each answered question is a frame expanded.

The implications run further than they look. The system has no incentive to become irreplaceable to any individual user — replaceability is the success state. It has no incentive to extend session length — short sessions ending in successful routing dominate long sessions ending in fatigue. The economic model and the ethical model converge.

*Replaceability is not a weakness. It is the success criterion.*

# An ethics of routing.

---

*If the system routes value among alternatives, it owes the user a fiduciary duty to their frame — not to its own retention metric.*

The traditional internet platform owes its shareholders attention, captured. The platform's interests and the user's interests diverge by construction. A routing system built on “Who Else?” cannot earn from retention. Its only durable revenue is from successfully routing — that is, from the user departing well.

This is the structural condition for what could be called fiduciary information design: a system whose business model requires it to act in the user's interest. Not as a regulatory compliance posture. Not as a marketing claim. As the only configuration in which the system makes money. The medium's structure makes the ethics enforceable.

*Fiduciary because the medium makes it the only profitable posture.*

# The intellectual scaffold.

---

*Five books and one paper that frame Chapter II. Not a bibliography — a path.*

**McLuhan**

*Understanding Media (1964)*

The medium is the message. The structure of carrier shapes the cognition of carrier-users.

**Wittgenstein**

*Philosophical Investigations (1953)*

Meaning is use. A question is a move in a language game whose stakes are the boundaries of a world.

**Shannon**

*A Mathematical Theory of Communication (1948)*

Information IS surprise. The objective function is not a metaphor; it is the formalism.

**Bostrom**

*Ethical Issues in Advanced AI (2003)*

The paperclip thought experiment, read symmetrically: maximization is a structure, not a polarity.

**Pariser**

*The Filter Bubble (2011)*

The bubble is an objective-function problem, fifteen years before AI made it acute.

**Hofstadter & Sander**

*Surfaces and Essences (2013)*

Analogy as the engine of cognition — the cognitive operation OIAI names.

# A question is the smallest unit of agency.

*An answered question closes a world. An asked one opens it. Build on the open ones.*



# Civilization

---

*Given a primitive that is universal in the mind and constructive as a medium, what would it build?*

# The routing layer of the agentic web.

---

*Agents do not have attention to monetize. They have tasks to evaluate and alternatives to compare.*

The medium that fits the human is the feed. The medium that fits the agent is the comparator. An agent acting on behalf of a user is, structurally, asking “Who Else?” every time it must choose between alternatives — vendors, services, providers, sources. The agentic web IS a “Who Else?”-shaped substrate, whether or not the substrate has been formalized.

Without a vendor-neutral, standards-backed ranking authority, the agentic web Balkanizes — each agent runs its own conflicted ranking, opaque to users and unauditable. Whoever owns the grammar of routing owns the layer. The grammar already exists in every human language. It has simply never been formalized as infrastructure.

*The agentic web needs a Switzerland. The construct names which Switzerland.*

# Why the agentic web cannot route through Google.

---

*The incumbent is conflicted by construction. The conflict is not a failure of governance; it is the business model.*

Google's rankings are a function of an advertising auction whose winners pay to appear higher than their organic position would place them. This is not a flaw — it is the product. The product is rentable position. An agent, acting on behalf of a user, cannot route through a ranker whose top results are partly determined by who paid to be there.

The agentic web therefore needs a router whose ranking function is auditable, whose objective is user-aligned, and whose business model does not rent position. “Who Else?” is the construct under which such a router can be defined cleanly: surprise-weighted relevance, no auction, no rentable position, paid only for successful routing. Google cannot become that router without unwinding its core revenue. Someone else has to build it.

*Renting position is a business model. It is not a routing protocol.*

# Parallel routing as interface.

---

*A single conversational thread is a thin interface for a thick problem.*

Today's agentic AI is exposed through a single thread — one model, one stream of tokens, one answer. The chat paradigm. The correct response to “Who Else?” is not one answer. It is a ranked, comparative set of alternatives, ideally produced by multiple specialized agents operating concurrently. The interface follows the linguistic primitive.

A “Who Else?” query fans out: a single natural-language input goes to several specialist agents in parallel; each returns scored candidates; results are merged against a common relevance + surprise function; the user receives a comparison, not a monologue. Meaningfully different from the chatbot paradigm — and structurally honest about model disagreement instead of averaging it away.

*Parallel and comparative, not sequential and conversational.*

# The end of the chatbot paradigm — quietly.

---

*Conversational AI inherited its interface from a 1960s research artifact (ELIZA), not from any first-principles analysis of what the medium should look like.*

The chatbot is the lowest-friction interface for a single-model system. It is not the right interface for a system that must compare multiple alternatives, surface disagreements, and route the user to the best fit. Comparison is not a chat. Comparison is a ranked set, a rationale per item, a default selection, and an option to depart.

The shift is not anti-chat. It is anti-monoculture. Chat will remain the right interface for a class of tasks (continuation, drafting, single-turn synthesis). It is the wrong interface for the discovery class — and the discovery class is most of the agentic web's work. The construct names what should replace chat there.

*Chat is a tool. Comparison is the medium.*

Weizenbaum, J. (1966). *ELIZA — A Computer Program for the Study of Natural Language Communication.*

# The system as a vertical factory.

---

*Single-vertical AI startups die when foundation models eat the vertical. A meta-system that picks the next vertical does not.*

The strategic problem of the AI era: any single product can be reproduced by a sufficiently capable foundation model. The escape: build the LAYER that ships products, not the product itself. The factory does not lose to the foundation model because the factory is the thing that decides which vertical to ship next, given current foundation-model capabilities.

Concrete spec: thirty top-level intent categories, six execution strategies (aggregate, matchmake, crawl-and-index, API-integrate, generate, community). Six × thirty = 180 (category × strategy) cells, of which a fraction are economically meaningful at any given moment. The factory's job is to identify and ship the live cells in the right order, with falling marginal cost per shipment.

*Don't acquire the paperclip. Acquire the machine.*

# Thirty categories. Six strategies. Bounded by construction.

---

*Human need is large but enumerable. Strategy space is bounded. The roadmap is a small instruction set, not a wishlist.*

The 500-need taxonomy collapses cleanly into roughly thirty top-level categories — food, shelter, healthcare, education, work, mobility, finance, romance, friendship, leisure, and so on. Each category admits a small set of viable execution strategies, drawn from the same six: aggregate existing supply, matchmake bilateral demand, crawl-and-index, integrate via API, generate de novo, or organize as community.

What is normally called a “product roadmap” becomes, in this frame, a SCHEDULER over the (category × strategy) grid. The unit of work is not the next feature; it is the next cell. The output is not a product; it is a compiler that produces products. The economics improve with each cell shipped — every cell teaches the compiler.

*Roadmap as scheduler. Output as compiler. Economics as compounding.*

# A compiler that produces products.

---

*What does it mean for the company to ship a compiler rather than a product?*

Each cell shipped feeds back into the compiler: which intent phrasings succeed, which fail, which strategy fits which category, what surprise-weighting parameters generalize. By the tenth cell, the eleventh costs less than the first did. By the twenty-fifth, the curve has flattened. The company does not become a discovery product. It becomes the factory that ships discovery products.

This reframes the seed round. The seed is not funding a product; it is funding the compiler's first ten cells. The investor is not buying into one vertical; they are buying into the slope of the cost curve as cells accumulate. The unit economics of cell N matter less than the gradient between cell N and cell N+1.

*Don't underwrite the product. Underwrite the slope.*

# A different internet, on the same primitive.

*Same humans. Same questions. Same languages. Different infrastructure on top — and a different shape of life as the result.*

# Three chapters, one chain.

## I. LINGUISTICS

**The construct is universal grammar.**

Universal across languages, frequent in use, early-acquired, late-preserved, neurolinguistically transferable, lossless to 15:1. A grammatical primitive.

## II. PHILOSOPHY

**It makes a different medium.**

Surprise-weighted, attention-inverted, slot-revealing, fiduciary by structure. The medium of openness rather than closure. A different form of life.

## III. CIVILIZATION

**It builds different infrastructure.**

The routing layer of the agentic web, parallel by interface, factory-shaped by production model, defensible in time. A different internet.

*An attack on any single chapter does not bring down the others. The chain is the moat.*

# What evidence would defeat this thesis.

---

*A thesis without falsifiers is not a thesis; it is a slogan. Five conditions under which this one fails:*

- 1. Universality** A typologically distinct language family without a functional equivalent of the construct, after rigorous fieldwork.
- 2. Compression** Replicated kernel-stamp compression across the ten sentence types yielding ratios significantly below 15:1, with stable methodology.
- 3. Transfer** ERP studies showing the LAN/P600 signature is language-specific rather than cross-linguistic for the additive person-interrogative.
- 4. Routing economics** An incumbent search platform demonstrating it can run a non-auction, surprise-weighted ranking layer at meaningful scale without unwinding its core revenue.
- 5. Standards path** A competing post-2020 protocol displacing DIN SPEC 2343 in agent-routing reference architectures within five years of agentic-web maturity.

*None of these have obtained. Each is a research program for someone who wants to.*

# Selected sources.

---

- Bakshy, E., Messing, S., & Adamic, L. (2015). Exposure to ideologically diverse news and opinion on Facebook. *Science*, 348(6239).
- Bostrom, N. (2003). Ethical Issues in Advanced Artificial Intelligence. *Cognitive, Emotive and Ethical Aspects of Decision Making*.
- Chomsky, N. (1957). *Syntactic Structures*. Mouton.
- Chomsky, N. (1995). *The Minimalist Program*. MIT Press.
- Davies, M. (2008). *The Corpus of Contemporary American English (COCA)*. Brigham Young University.
- Diessel, H. (2004). *The Acquisition of Complex Sentences*. Cambridge University Press.
- Dryer, M. & Haspelmath, M. (eds.) (2013). *The World Atlas of Language Structures Online*. MPI EVA.
- Friederici, A. D. (2002). Towards a neural basis of auditory sentence processing. *Trends in Cognitive Sciences*, 6(2).
- Heim, I. (1992). Presupposition projection and the semantics of attitude verbs. *Journal of Semantics*, 9.
- Hofstadter, D. & Sander, E. (2013). *Surfaces and Essences*. Basic Books.
- Kaan, E. & Swaab, T. (2003). Repair, revision, and complexity in syntactic analysis. *Journal of Cognitive Neuroscience*, 15(1).
- Karmiloff-Smith, A. (1979). *A Functional Approach to Child Language*. Cambridge University Press.
- Kemper, S. et al. (2001). Language decline across the life span. *Journals of Gerontology Series B*, 56(5).
- Kintzel, T. (2024). "Who Else?" as Universal Grammar. Whitepaper, ~100pp, 135 references.
- Krifka, M. (1999). Additive particles under stress. *Proceedings of SALT 8*.
- Levy, R. (2021). Social media, news consumption, and polarization. *American Economic Review*, 111(3).
- McLuhan, M. (1964). *Understanding Media*. McGraw-Hill.
- Pariser, E. (2011). *The Filter Bubble*. Penguin Press.
- Piantadosi, S. (2023). Modern language models refute Chomsky's approach to language. [lingbuzz/007180](https://lingbuzz/007180).
- Pinker, S. (1994). *The Language Instinct*. William Morrow.
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal*, 27.
- Weizenbaum, J. (1966). ELIZA — A Computer Program for the Study of Natural Language Communication. *Communications of the ACM*, 9(1).
- Wittgenstein, L. (1922). *Tractatus Logico-Philosophicus*. Kegan Paul.
- Wittgenstein, L. (1953). *Philosophical Investigations*. Blackwell.
- DIN SPEC 2343 (2020). Deutsches Institut für Normung.

# WHO ELSE?

The Philosophy

---

*Thanks for your interest.*

whoelse.ai