

Theory of Communication Under Uncertainty

Michael Ghilissen
Public Health
ULiège

I Introduction and Premise

Communication rarely occurs under conditions of complete knowledge. The sender does not fully understand the world being described; the receiver does not perfectly reconstruct the sender's intended meaning; the transmission environment introduces noise; and the consequences of interpretation may be nonlinear.

Communication is therefore not merely a transmission problem. It is also an uncertainty problem.

Classical communication models often ask whether information has been successfully transmitted from sender to receiver. A theory of communication under uncertainty asks a more demanding question:

How well does consequential meaning survive the disturbances to which real communication is exposed?

The central objective is therefore not simply to maximize information transfer, but to:

Preserve consequential signal while minimizing fragility under uncertainty and perturbation.

Communication quality cannot be assessed solely by whether a receiver correctly understands a message at a particular moment. It also depends on whether the message's consequential core remains intelligible, useful, and defensible when the message is shortened, challenged, repeated, removed from context, transmitted across audiences, or exposed to changing evidence over time.

This framework is particularly relevant to public health. Public-health communication operates under persistent uncertainty: evidence evolves; risks are probabilistic; audiences differ in prior beliefs, incentives, and health literacy; messages are repeatedly compressed by media and social networks; and seemingly minor misunderstandings can sometimes produce substantial behavioral or population-level consequences.

Public health therefore provides a demanding empirical domain for the theory. The central question is not simply whether information was received or recalled, but whether consequential meaning survives perturbation without generating disproportionate harm.

This claim should not be overstated. Communication under uncertainty is not unique to public health, nor does every communication failure arise from fragility. Public health's value lies in

providing a consequential setting in which the proposed constructs can be defined, challenged, and tested.

2 Communication Fragility

The central construct of the theory is communication fragility.

Communication fragility describes the sensitivity of a message's consequential interpretation or effects to perturbation.

Let:

M = a public-health message

p = a perturbation to the message or its environment

Y(M,p) = the consequential interpretation or behavior associated with message M under perturbation p

Conceptually:

$$F(M) \propto |\Delta Y| / |\Delta p|$$

A message is fragile when a relatively small perturbation produces a disproportionately large change in consequential interpretation or behavior.

A message is robust when its consequential core remains sufficiently stable despite relevant perturbations.

An antifragile communication system, by contrast, does more than resist disturbance: it uses bounded perturbations to improve subsequent communication by revealing ambiguity, eliminating weak claims, sharpening definitions, or improving the communicator's model of the audience.

The distinction matters. Fragility is not synonymous with misunderstanding. A message may be misunderstood because it is poorly written, because the audience lacks necessary knowledge, because the receiver is motivated to reject it, or because the environment is deliberately adversarial. The theoretical claim is narrower: fragility concerns the sensitivity of consequential meaning to perturbation.

Consider a vaccination message:

"The vaccine substantially reduces the risk of severe disease, although breakthrough infections remain possible."

Suppose social media compresses it to:

"Vaccinated people can still become infected."

The extracted statement may remain factually correct, yet its consequential meaning can shift dramatically. Some receivers may infer that vaccination provides no meaningful benefit. A

relatively modest perturbation—extracting one component from its explanatory context—can therefore produce a large change in interpretation.

This is contextual fragility.

At least six forms of fragility warrant distinction:

<i>Fragility</i>	<i>Core issue</i>	<i>Public-health question</i>	<i>Example</i>
Compression	Does essential meaning survive shortening?	Does the recommendation survive compression?	Vaccine guidance reduced to a headline
Contextual	Does meaning survive removal from its original context?	What happens when a statement circulates independently?	“Vaccinated people can still become infected”
Challenge	Does the claim survive informed scrutiny?	Which claims fail under serious criticism?	Questions about vaccine effectiveness or adverse events
Trajectory	Does the communication strategy remain viable when repeated?	Can authorities communicate this way repeatedly without degrading trust or effectiveness?	Repeated changes in epidemic guidance
Audience	How does consequential meaning vary across receivers?	Which elements remain stable across populations, and which require adaptation?	The same alcohol-risk message interpreted by adolescents, clinicians, and dependent users
Tail	Can a small component generate disproportionate adverse consequences?	Could one phrase dominate the effects of the entire message?	Poor wording around a rare adverse event

These dimensions move communication fragility from metaphor toward a potentially measurable construct.

They also expose a limitation of conventional evaluation. Public-health communication is often assessed through average comprehension, recall, reach, engagement, or intention. A campaign may perform well on all five while remaining highly fragile. Average performance and robustness are not equivalent.

3 Consequential Signal

The theory distinguishes information transmitted from consequential signal.

More information does not necessarily produce better communication. Additional detail may increase understanding, but it can also increase cognitive load, ambiguity, distraction, internal contradiction, and opportunities for misinterpretation.

A useful conceptual representation is:

$$S^c = R / (L + A + D)$$

where:

S^c = consequential signal

R = decision-relevant information

L = cognitive load

A = ambiguity

D = distraction or irrelevant information

This equation should not be interpreted as an established functional relationship. It is a conceptual representation of a testable proposition: communication quality may not increase monotonically with information volume.

This leads to the principle of communication by subtraction:

Remove before adding.

A five-slide presentation may carry more consequential signal than a thirty-slide presentation. Three paragraphs may communicate a strategy more effectively than a thirty-page report when the additional material contributes more cognitive burden than decision-relevant understanding.

Subtraction, however, is not inherently beneficial. Excessive simplification can remove qualifications, uncertainty, contraindications, or distributional differences that are themselves consequential. The relevant objective is therefore not brevity. It is signal density without consequential distortion.

The distinction becomes especially important in vaccine communication.

A technically comprehensive communication may contain efficacy estimates, confidence intervals, relative and absolute risks, contraindications, adverse-event frequencies, immunological mechanisms, surveillance methodology, and subgroup analyses. All may be scientifically relevant.

They are not necessarily equally relevant to every decision-maker.

For a member of the public, the consequential core might instead be:

Who should be vaccinated → why → expected benefit → important risks → what to do → where to obtain it

Communication by subtraction therefore means removing information whose marginal contribution to decision-relevant understanding is smaller than the cognitive or interpretive burden it creates.

This generates an empirically testable—and potentially uncomfortable—hypothesis:

Adding technically correct information can sometimes reduce effective public-health communication.

The converse must also be tested: removing technically complex information can improve comprehension while simultaneously degrading informed decision-making. A theory of

consequential signal must therefore evaluate not only what audiences remember, but what information they actually require to make an appropriately informed decision.

4 Semantic Survival Under Compression

Compression provides a direct stress test of communication.

Consider successive representations of the same argument:

5,000 words → 500 words → 100 words → 1 sentence

The relevant question is not how many words disappear. It is how much essential consequential meaning survives.

Conceptually:

$$SS(c) = \text{essential consequential meaning retained after compression} / \text{essential consequential meaning in the original message}$$

A strong message should exhibit relatively high semantic survival across a reasonable range of compression.

This suggests a useful diagnostic proposition:

Difficulty compressing a message may reveal conceptual weakness, not merely summarization difficulty.

The qualification matters. Some ideas are irreducibly complex. Clinical recommendations, probabilistic evidence, ethical trade-offs, and heterogeneous treatment effects may legitimately resist aggressive compression. Failure to reduce a thirty-page analysis to a slogan is not, by itself, evidence of poor thinking.

The stronger claim is narrower: when communicators cannot identify what must survive compression—and what can safely disappear—the underlying hierarchy of ideas may be insufficiently clear.

Compression therefore serves two functions. It tests the communication, and it can expose weaknesses in the architecture of the underlying argument.

Communicable-disease communication provides a useful example.

A full epidemiological message might state:

Current evidence suggests that transmission occurs primarily through respiratory exposure, with risk increasing with duration, proximity, ventilation conditions, and prevalence. Recommendations may change as evidence develops.

At approximately twenty words:

Infection risk rises with close, prolonged exposure, particularly indoors; reduce exposure and follow current public-health guidance.

As a headline:

Avoid close indoor contact during periods of high transmission.

The empirical question becomes:

At what degree of compression does consequential meaning begin to collapse?

Researchers could expose participants to progressively compressed versions and measure whether they retain the direction of risk, relevant uncertainty, recommended action, important exceptions, and appropriate confidence.

A message whose consequential interpretation remains stable under substantial compression exhibits greater semantic survival than one whose meaning changes sharply after modest shortening.

5 Mental Health and Substance Use: Tail Fragility

Mental-health and substance-use communication illustrates the importance of tail consequences.

Average outcomes can conceal concentrated harm. One sentence, image, label, or extracted phrase may dominate the consequences of an otherwise effective campaign.

Suppose a mental-health campaign produces the following distribution:

90% interpret the message appropriately.

8% misunderstand it without meaningful consequences.

2% interpret one phrase in a harmful or stigmatizing way.

An average comprehension score of 90% may suggest success. Yet that conclusion ignores the severity of outcomes within the remaining distribution.

A more relevant conceptual quantity is:

$$E[\text{Consequence}] = \sum P_i \times C_i$$

where:

P_i = probability of interpretation i

C_i = consequence associated with interpretation i

If a rare interpretation carries sufficiently severe consequences, optimizing average comprehension is inadequate.

This matters in communication concerning addiction, overdose, suicide prevention, severe mental illness, treatment seeking, and other contexts in which low-frequency misunderstandings may carry high costs.

The appropriate question becomes:

Which component of this message could generate disproportionate harm if misunderstood, decontextualized, or circulated independently?

Tail fragility therefore shifts evaluation away from average message performance toward the distribution and severity of plausible consequences.

6 Lifestyle Behavior and Repeated Communication

Lifestyle communication illustrates trajectory fragility.

Consider messages about exercise, diet, alcohol, smoking, sleep, or weight.

A message such as:

“Sitting is killing you.”

may perform well as a one-off intervention. It attracts attention and may generate high recall or immediate intention to change.

But single-exposure performance answers only one question.

A theory of communication under uncertainty asks another:

What happens when this becomes the habitual communication strategy?

Repeated exposure to alarming messages may produce habituation, skepticism, anxiety, reactance, fatalism, or declining trust. Alternatively, repetition may reinforce norms and sustain beneficial behavior. The direction cannot be assumed; it must be measured.

Thus:

Immediate communication payoff $\neq$ time-path communication payoff

A useful conceptual representation is:

Trajectory value = cumulative behavioral benefit – cumulative communication harm

Again, this is not a validated mathematical identity. It identifies the quantities that a dynamic theory should consider.

This distinction is particularly important in public health because institutions do not communicate once. They communicate repeatedly, often over decades.

Trust should therefore be treated not merely as a response to an individual message but partly as a stock shaped by accumulated communication experience.

7 Audience Fragility and Non-Communicable Disease

Non-communicable diseases demonstrate why a message cannot be considered robust simply because its average recipient understands it.

Consider:

“Regular physical activity reduces cardiovascular risk.”

A healthy 25-year-old, a sedentary 55-year-old, a person with obesity, someone recovering from myocardial infarction, a physician, and someone with limited health literacy may interpret the same statement differently.

The relevant quantity is therefore not simply:

$P(\text{correct interpretation})$

but variation in consequential interpretation across audiences:

$F_a \approx \text{variation in } Y(M,a) \text{ across audiences } a$

Yet audience variation is not inherently a communication failure. Different populations may legitimately require different interpretations or actions.

The objective is therefore not uniformity. It is to preserve the core consequential meaning where stability is warranted while allowing justified adaptation where it is not.

For example:

Core: physical activity is generally beneficial.

Audience-dependent layer: appropriate intensity, duration, contraindications, and need for clinical supervision.

This motivates a layered public-health architecture:

Core health message $\rightarrow$ Recommended action $\rightarrow$ Evidence $\rightarrow$ Population-specific qualification $\rightarrow$ Uncertainty

Such layering may reduce audience fragility without assuming that one formulation can or should produce identical effects across heterogeneous populations.

8 Credibility Through Exposure and Accountability

Traditional communication models can be represented as:

Sender → Message → Receiver

Under uncertainty, receivers may also evaluate the institutional environment surrounding the message:

Sender → Accountability/Exposure → Message → Receiver

A receiver may reasonably ask:

What happens if this claim proves wrong?

Compare:

“This strategy should increase revenue.”

with:

“I recommend investing €2 million in this strategy, and I will be accountable for the outcome.”

The propositional content may be similar, but the signals surrounding the claims differ.

A conceptual model of credibility might therefore include:

$$C = f(\text{Evidence, Track Record, Accountability/Exposure})$$

Exposure does not make a proposition true, and lack of exposure does not make it false. Nor should personal risk substitute for expertise or evidence. A confidently exposed communicator can still be wrong.

Exposure is therefore best treated as a potential credibility modifier, not a truth criterion.

In public health, the more defensible construct may be institutional accountability rather than personal exposure.

Compare:

“We are confident that this intervention is safe.”

with:

“We recommend this intervention based on current evidence. Adverse events are continuously monitored, the data will be published, and the recommendation will be revised if the evidence changes.”

The second statement makes the institution's accountability architecture visible. It identifies mechanisms through which the claim can subsequently be tested, challenged, corrected, or withdrawn.

Credibility may therefore depend partly on visible commitments to surveillance, transparency, correction, and updating.

This distinction is particularly important when evidence evolves. Pretending to possess certainty may create trajectory fragility: when recommendations later change, legitimate scientific updating can be interpreted as inconsistency or deception.

Calibrated uncertainty may therefore protect credibility better than artificial certainty.

9 Revealed Conviction: When Behavior Contradicts Messaging

Communication occurs through more than words.

Represent a communication act as:

$$M = (C, E, X)$$

where:

C = propositional content
E = expressed salience or conviction
X = observable action or commitment

Consequential meaning becomes:

$$Y = f(C, E, X)$$

Suppose a CEO declares:

“This project is our highest priority.”

Yet the CEO allocates no resources, spends no time on it, accepts no responsibility for it, and never mentions it again.

The verbal channel communicates priority. Observable behavior communicates something else.

Conversely, sustained investment of time, resources, reputation, or institutional capacity can communicate conviction even when verbal emphasis is limited.

This produces the revealed-conviction principle:

When verbal claims and costly observable behavior diverge, receivers may update their beliefs toward the behavioral signal.

This principle is related to, but distinct from, exposure. Exposure concerns vulnerability to consequences. Revealed conviction concerns information conveyed through observable allocation and behavior.

In public health, an institution may declare:

“Mental health is one of our highest priorities.”

while providing limited access to psychological services, long waiting times, and minimal funding.

The verbal message communicates high priority. Resource allocation may communicate low revealed priority.

Similarly, an institution may repeatedly emphasize prevention while directing the overwhelming majority of resources toward treatment.

Budgets, staffing, service provision, and institutional behavior can therefore become part of the communication environment.

However, behavioral signals require cautious interpretation. Resource allocation reflects constraints, competing obligations, political processes, and implementation capacity—not conviction alone. Revealed behavior is informative, but it is not a transparent measure of intention.

10 Tail Consequences

Communication consequences need not be symmetrically or normally distributed.

Organizations commonly evaluate communication through averages: engagement, sentiment, recall, click-through, comprehension, or audience response.

Yet a small number of communication events may account for a disproportionate share of consequences.

One sentence can dominate an interview. One extracted quotation can overwhelm a forty-page report. One phrase can become the enduring representation of an institution or individual.

Communication analysis should therefore examine not only expected or average effects but also tail exposure.

A practical test follows:

Which component of this message could dominate its consequences if extracted and circulated independently?

This changes the unit of analysis. The relevant object is no longer merely the average quality of the document, campaign, or speech, but the distribution of potential consequences across its components.

A message can therefore perform well on average while remaining extremely fragile in its tail.

The empirical challenge is substantial. Tail outcomes are, by definition, uncommon and difficult to estimate. The concept is theoretically useful only if researchers specify what constitutes an extreme consequence, over what time horizon, and for which population.

II Communication Trajectories

Communication should also be evaluated dynamically.

A strategy that succeeds once may deteriorate when repeatedly used.

Suppose provocative communication reliably generates attention but occasionally produces severe reputational or behavioral harm. Evaluating many communicators using the strategy once differs fundamentally from evaluating one institution using the same strategy repeatedly over decades.

This creates a distinction between cross-sectional communication payoff and time-path communication payoff.

The practical rule is:

A viable communication strategy must survive repetition.

The relevant question is therefore not merely:

“Did this message work?”

but:

“If we communicate this way repeatedly, what happens to trust, credibility, interpretation, behavior, and exposure to severe downside?”

Communication strategies may prove more durable when they rely on consistent evidentiary standards, calibrated claims, transparent correction, and visible updating rather than tactics that generate recurrent short-term gains at the cost of occasional severe damage.

This remains a hypothesis, not a universal prescription. Repetition can strengthen as well as weaken communication. The theoretical contribution lies in making the time path itself an object of evaluation.

I2 Antifragile Communication Systems

Fragility, robustness, and antifragility are most defensible when applied to communication systems and processes, rather than isolated sentences.

A fragile communication process deteriorates under stress.

A robust process withstands stress.

An antifragile process uses bounded stress to improve future communication.

Conceptually:

$$Q(t+1) > Q(t)$$

when stress applied at time t reveals information that the system successfully incorporates into communication at $t+1$.

A practical pipeline might be:

Draft → Challenge → Revision → Adversarial review → Revision → Release →
Feedback → Update

Scientific seminars provide a familiar example. Questioning can expose weak assumptions, ambiguous definitions, unsupported claims, or unanticipated interpretations. The argument that emerges may be stronger than the original.

The essential mechanism, however, is not criticism.

It is learning from criticism.

Stress without detection, feedback, institutional memory, or adaptation merely damages a system. Stress becomes productive only when the system can identify failure and incorporate what it learns.

Antifragile Public-Health Communication

Public-health recommendations routinely encounter new evidence, criticism, misinformation, unexpected behavioral responses, and heterogeneous interpretations.

A fragile institution may experience correction as threat:

Criticism → Defense → Escalation → Loss of credibility

A robust institution withstands criticism.

An antifragile institution learns from it.

For a vaccination campaign:

Draft → Community testing → Misinterpretation detection → Revision → Release →
Monitoring → Feedback → Revised message

Suppose audiences repeatedly misunderstand:

“90% effective”

as:

“10% of vaccinated people will become infected.”

An antifragile system does not treat this merely as evidence of public ignorance. It treats recurrent misunderstanding as information about the communication architecture.

The next message changes.

The system has learned from perturbation.

This is materially different from resilience. Resilience restores function after stress. Antifragility, as defined here, requires demonstrable improvement because of information revealed by bounded stress.

The burden of proof is therefore higher: not every adaptive communication system is meaningfully antifragile.

I3 Layered Communication

Robust communication should be architecturally layered.

A message can distinguish:

Core claims → Supporting evidence → Examples → Interpretation → Speculation

These layers need not have identical survival requirements.

Core claims should exhibit high semantic survival, strong evidential support, and low contextual fragility. Supporting details may change without destroying the central argument. Interpretive claims should be distinguishable from observations. Speculative elements should be explicitly identifiable as such.

This modular architecture offers two advantages.

First, failure in a peripheral component need not invalidate the entire communication structure.

Second, criticism becomes diagnostically useful because it identifies which layer requires revision.

A communication system can therefore evolve without pretending that every previous formulation was equally central, equally certain, or equally durable.

In public health, layering also offers a way to reconcile simplicity with informed choice. The objective is not to hide complexity but to organize it according to consequence and audience need.

I4 Time as a Filter

Messages also differ in how they survive over time.

Ask what remains consequential after:

10 minutes → 1 day → 1 month → 1 year → 5 years

Not every message should survive for years. Some communication is deliberately ephemeral, and some recommendations should change rapidly when evidence changes.

For communication intended to endure, however, time provides another stress test.

Define temporal survival as the extent to which consequential meaning remains useful as immediate context, novelty, rhetorical salience, and surrounding conditions decay or change.

This idea resembles Lindy-type reasoning but should not be confused with a literal application of the Lindy effect.

Survival through time is evidence to interpret, not proof of quality. False beliefs can persist; obsolete messages can remain memorable; and good scientific communication may appropriately become outdated when knowledge advances.

Temporal survival therefore matters only relative to the intended lifespan and epistemic status of the message.

I5 Integrated Model

The framework can now be stated more parsimoniously.

Communication produces consequential meaning through interactions among:

Signal × Credibility × Semantic Survival × Revealed Conviction

while being exposed to:

Noise + Fragility + Tail Exposure

These expressions should remain conceptual rather than mathematical until their functional relationships are empirically established.

The central theoretical proposition is simpler:

High-quality communication preserves consequential meaning across relevant perturbations while avoiding disproportionate adverse consequences.

Public health provides a useful test domain because the same theoretical constructs generate different observable vulnerabilities:

<i>Domain</i>	<i>Typical uncertainty</i>	<i>Potentially dominant fragility</i>
Vaccines	Effectiveness, safety, changing evidence	Contextual + challenge + tail
Communicable disease	Transmission and evolving epidemic conditions	Compression + trajectory
Non-communicable disease	Long horizons and heterogeneous individual risk	Audience + compression
Mental health & substance use	Stigma, heterogeneity, potentially severe consequences	Audience + tail
Lifestyle behaviors	Small effects accumulated through repeated behavior	Trajectory + audience

These classifications should be treated as hypotheses rather than fixed assignments. The dominant fragility may vary by message, audience, institution, medium, and historical context.

The underlying theoretical claim nevertheless remains consistent:

High-quality public-health communication preserves consequential meaning across relevant perturbations, adapts appropriately when evidence changes, and limits disproportionate adverse consequences.

I 6 Testable Propositions

The value of the framework ultimately depends on whether its constructs generate falsifiable predictions rather than merely persuasive descriptions.

The following propositions provide an initial empirical agenda.

P1 — Subtraction

Beyond a threshold of informational sufficiency, removing information with low decision relevance will improve comprehension and decision quality, provided that consequential qualifications are preserved.

P2 — Compression survival

Messages with greater semantic survival under controlled compression will exhibit greater stability of consequential interpretation across receivers.

P3 — Accountability and credibility

Holding propositional content and evidence constant, visible mechanisms of accountability, monitoring, and correction will affect perceived credibility.

P4 — Revealed conviction

When verbal claims conflict with costly observable behavior, receivers will increasingly use behavioral evidence to infer institutional priorities as the discrepancy becomes more salient.

P5 — Tail fragility

Communication consequences will sometimes be substantially more concentrated than

average-performance measures imply, with a small number of message components accounting for a disproportionate share of extreme outcomes.

P6 — Trajectory fragility

Some communication strategies that perform well under single-exposure evaluation will perform less well under repeated exposure because their effects on trust, interpretation, or behavior accumulate nonlinearly over time.

P7 — Adversarial improvement

Messages subjected to bounded adversarial review followed by structured revision will exhibit lower subsequent challenge fragility than comparable messages receiving conventional editorial review alone.

P8 — Layering

Messages that explicitly distinguish core claims, evidence, interpretation, and speculation will exhibit greater semantic survival under compression and decontextualization.

P9 — Error correction

Communication systems with visible mechanisms for acknowledging and correcting errors will recover credibility differently following demonstrated error than systems that prioritize consistency with previous claims.

P10 — Conceptual coherence

After controlling for writing quality and irreducible subject complexity, high compression fragility will be associated with lower underlying conceptual coherence.

These propositions deliberately avoid assuming that the theory is correct. Several could fail empirically. That is a strength rather than a weakness: a framework becomes theoretically useful when evidence can discriminate between its claims and plausible alternatives.

I 7 Practical Protocol

Before releasing consequential communication, subject it to five stress tests.

Compression

Can the consequential core survive at 50%, 10%, and one-sentence length?

Extraction

What happens if the most consequential sentence circulates independently?

Challenge

Which claims fail under an informed adversarial reader?

Audience shift

Which elements change meaning when interpreted by receivers with different knowledge, incentives, priors, or health literacy?

Repetition

What happens if this becomes the communicator's habitual strategy?

After release, add a sixth test:

Learning

What did actual misunderstanding, criticism, behavioral response, or unexpected use reveal that should change the next message?

The objective is not to create communication that cannot be attacked. Such a goal is unrealistic and may encourage empty, defensive, or excessively cautious language.

Nor is the objective to make every message invariant to context. Good public-health communication must sometimes change because evidence, risks, and recommended actions change.

The objective is more demanding:

Preserve the consequential core when it should remain stable; make uncertainty visible when it matters; change the message when evidence justifies change; prevent small distortions from producing disproportionate harm; and build communication systems that learn systematically from their own failures.

Under this formulation, communication under uncertainty is not a theory of perfect messages.

It is a theory of how consequential meaning survives, fails, adapts, and improves in an uncertain world.

[Note: Intellectual note: This section draws on ideas from Nassim Nicholas Taleb's Incerto, focusing on fragility, antifragility, exposure, tail risk, and time. Any errors in extending these concepts to communication are mine.]