# From Stream to Space

## Planning for Context in Persistent Social-Media Campaign Information

## 1. Introduction

Social media does not contain only recent information. A user looking at a company profile, following a link, using a search function, or reading a discussion can encounter posts that were published days, months, or years earlier. An older product announcement can still be read. A statement made during an earlier campaign can still be found. A post can be quoted or referenced long after its original publication. In each case, information created at an earlier point in time remains available in the present. (Treem & Leonardi, 2013)

What may no longer be as readily available is the situation in which that information was originally published. A product may have changed. A campaign may have ended. A comparison may have referred to an earlier product generation. A statement may have responded to circumstances that were obvious at the time but are no longer obvious to someone encountering it today. The information can still be understood as something that was published, while parts of the context that originally helped to interpret it may no longer accompany it. (Brandtzaeg & Lüders, 2018)

For social-media campaign planning, this creates a simple but consequential difference. Campaigns are planned and executed within periods of active communication. They have current objectives, publication schedules, distribution periods, and communication contexts. The information produced during those periods, however, can remain available after active campaign distribution has ended. (Fill & Turnbull, 2019)

**The temporal boundary of active campaign distribution is not necessarily the temporal boundary of the information produced by that campaign.**

This difference does not mean that persistent digital information is a new phenomenon. Social-media content can remain accessible, searchable, referenced, or otherwise available after publication. Campaign research also recognizes processes and effects that continue beyond active distribution. The issue addressed here is therefore not simply that content can remain online or that campaign effects can continue after a campaign has ended. (Treem & Leonardi, 2013; Kumkale & Albarracín, 2004)

The issue becomes more specific when information and the context required to interpret it do not remain available in the same way. At the time of publication, a campaign communication may be easy to understand because its surrounding conditions are apparent. A comparison refers to a known product. A statement relates to a current event. A claim is accompanied by a

qualification or source. An image is understood together with its caption. A message belongs to a campaign whose purpose is recognizable.

If the information remains available after that immediate communication situation has passed, some of those conditions may no longer be apparent. This does not necessarily change the information itself. It changes what a person or system encountering that information may need in order to interpret it without making assumptions that were unnecessary when it was originally published.

This distinction is important because the continued availability of campaign information is not the same as the continuation of the campaign. Nor is it necessarily a delayed campaign effect. A campaign can cease active distribution while information produced during the campaign remains available. That information can later be encountered without continuous exposure, continued campaign activity, or any demonstrated connection to an earlier communication effect.

The planning problem therefore begins neither with persistence alone nor with the possibility of a later encounter. It begins where information that remains available depends on context whose absence could materially change how that information is understood.

This problem is particularly relevant to social media because the chronological flow of a feed is only one way in which published information can be reached. Older information can also be found through profiles, links, references, search, quotations, recommendations, or other forms of retrieval. Its availability is therefore not necessarily tied to the period in which it was actively distributed.

This distinction provides the basis for the title *From Stream to Space*. The **Stream** describes the familiar condition in which social-media information is organized through ongoing publication and distribution. Attention is directed toward what is appearing and circulating now. **Space** describes a different condition: information remains available and can become relevant without its relevance depending on its original position in that temporal flow.

Stream and Space are not successive stages, and Space is not a predicted future form of social media. Both conditions can coexist.

Developments in retrieval and recommendation make this distinction increasingly significant. Information can be selected according to criteria other than publication time alone. AI-mediated retrieval adds further possibilities for selecting, combining, summarizing, and transforming existing information in response to a current task. These developments do not create the underlying problem. Information could outlive its original communication context before such systems existed. They do, however, increase the practical importance of asking what happens when older information participates in new contexts. (Campos et al., 2014; Lewis et al., 2020)

Campaign planning cannot determine whether persistent information will later be encountered, how it will be reached, what other information will accompany it, or what purpose it might serve. What can be examined before publication are dependencies already present in the communication, such as source, qualification, comparison basis, temporal scope, or relations between multimodal components. The planning problem therefore concerns whether loss of such context would materially affect later interpretation and, where it would, whether the relevant distinctions can remain sufficiently recoverable without compromising the communication's immediate purpose.

The research problem addressed by this paper is therefore that **campaign-produced information may remain available beyond the active distribution period, while the original campaign context may not remain available in the same way. A later encounter can consequently occur under conditions that differ from those under which the information was originally planned, distributed, and interpreted.**

This leads to the controlling research question:

**How should social-media campaign planning account for context dependencies in information that may remain available beyond its period of active campaign distribution?**

The paper treats campaign planning as the point at which this information problem can still be considered before publication. It does not propose new theories of persistence, context, representation, retrieval, or claim validity. Instead, it examines what follows for campaign planning when information created for a temporally bounded communication process can remain available beyond that process.

The resulting framework begins with the immediate purpose of the communication and remains deliberately selective. Context dependency alone does not create a requirement to preserve campaign content or make every communication independent of context. Recoverability becomes a planning concern only where loss of context would materially affect later interpretation.

The scope of planning remains correspondingly limited. Before publication, planners can influence the communication itself and the relevant conditions attached to it. After publication, they cannot control whether it will later be retrieved, quoted, combined with other information, transformed by a system, or interpreted under changed circumstances.

**Campaign planning determines initial conditions, not future relational outcomes.**

The task is therefore not to predict what will happen to campaign-produced information in the future, but to identify where materially important context should remain recoverable before the original communication situation passes out of planning control.

## 2. Social Media Campaigns as Temporal Distribution

Social-media distribution gives information a temporal position. A campaign asset enters an ongoing flow of other publications, reactions, events, and competing communication. Its initial visibility and relevance are therefore shaped not only by what the asset contains, but also by when it appears and what is happening around it at that time. (Fill & Turnbull, 2019)

This temporal position performs an important communicative function. A product announcement published during a launch is encountered alongside other launch communication. A response to a public event appears while that event is current. A campaign statement can rely on references, discussions, images, or assumptions that are readily recognizable because they belong to the same period of communication. Time is therefore not merely a scheduling variable. It contributes to the conditions under which campaign information is initially encountered and interpreted.

Social-media feeds make this temporal organization particularly visible. New information continually enters the flow, while earlier information loses its initial position within it. Recency, current interaction, ongoing distribution, and surrounding communication influence what becomes visible and what receives attention. This is the condition described here as the

**Stream**: information is encountered within an ongoing process of temporal distribution. (Campos et al., 2014)

The Stream should not be reduced to chronological ordering. Social-media feeds are selected and ranked rather than simply displayed by publication date. Nevertheless, the communication encountered through them remains situated within a changing present. What is current, what is being discussed, what has recently happened, and what other communication is presently available all contribute to the conditions under which an asset is received.

This also explains why campaign context can initially remain largely implicit. During active communication, some interpretive conditions are supplied by proximity. The campaign itself connects individual assets. Current events supply references. Product pages, accompanying posts, media coverage, comments, and other communication can make relationships apparent without requiring every individual asset to restate them.

Such contextual support is efficient because campaign communication is not normally designed as a collection of isolated information objects. Individual assets participate in a larger communication situation. A short post can therefore work precisely because it does not reproduce everything required to understand the campaign from first principles.

The temporal position of an asset, however, is not permanent. As new communication enters the Stream, the relations that were obvious through temporal proximity can become less immediate. The launch moves into the past. Public discussion changes. Other campaign assets become harder to find, and linked information may be updated.

This is where temporal distribution becomes relevant to the problem developed in this paper. The Stream can supply context while information occupies its original communication situation, but it does not guarantee that the same context will accompany the information outside that situation.

The distinction is not a criticism of temporal campaign planning. Contextual compression is often necessary for effective communication. Requiring every campaign asset to carry a complete and permanently self-sufficient account of its context would undermine many legitimate forms of social-media communication. The relevant question is therefore not how to eliminate dependence on the Stream, but which dependencies become significant when information can remain available without the temporal surroundings that originally supported its interpretation.

## 3. Beyond the Distribution Period

Once information loses its original position in active distribution, its continued existence cannot be described adequately by attention alone. A post may receive little visible activity while remaining accessible. A video may no longer circulate prominently while its page continues to exist. A statement may disappear from current communication while remaining reachable through a profile, link, search, reference, or other path. The decline of observable presence therefore does not necessarily imply the disappearance of the information. (Treem & Leonardi, 2013)

This distinction separates **presence** from **persistence**. Presence concerns whether information is currently observable within a communication situation. Persistence concerns whether something of the information remains available beyond that situation. The two can coincide, but they need not.

A campaign publication can consequently pass through periods in which it has little or no observable presence without ceasing to exist as available information. **Dormancy** provides a useful analogy for this condition. A dormant seed is not actively growing, yet its lack of visible activity does not establish its absence. Similarly, campaign information can remain available without currently participating in observable distribution or attention.

The analogy is deliberately limited. Dormancy is not proposed as a lifecycle stage through which social-media information necessarily passes. Information does not have to move from active distribution into dormancy and later into another defined state. Some publications continue to receive attention, some disappear, some remain accessible without further observable activity, and some may later be encountered again. Dormancy describes only the possibility that continued informational existence and current observable presence can diverge.

This is also why dormancy should not be confused with a delayed communication effect. A period without observable activity followed by a later encounter does not demonstrate that an earlier campaign effect remained latent and subsequently re-emerged. Nor does it establish continuity of attention. Without additional evidence, it establishes only that information available from an earlier communication can be present again at a later point. (Kumkale & Albarracín, 2004)

The same boundary separates persistence from the **Long Tail**. Long-tail patterns describe some form of continuing attention, use, traffic, or effect beyond an earlier concentration of activity. Persistent information does not require such continuity. Information can remain available even where no continuing pattern of attention can be demonstrated. A later encounter with persistent campaign information is therefore not necessarily the tail of the original campaign. (Anderson, 2006)

These distinctions reveal that several temporal conditions must be kept separate. Information can be published at one point, become observable or accessible again at another, and produce a response or other consequence at yet another. Continued availability does not establish continued presence, renewed presence does not establish how information was retrieved, and retrieval does not establish that an effect followed.

**Temporal presence is not equivalent to informational existence.**

Persistence itself also requires further differentiation. An original publication can continue to exist in substantially the form in which it was published. But continued existence of the original object is only one way in which information can persist.

A publication can disappear while information derived from it remains. A quotation can survive after the original post has been removed. A screenshot can retain a representation after the source page is unavailable. A summary can preserve a statement without preserving the publication from which it was derived. A later reference can establish that something was communicated even when the original object can no longer be reached.

**Object persistence is therefore not equivalent to information persistence.**

This distinction has an important consequence for source recoverability. Information about an earlier publication can remain available even when its original source is no longer directly accessible. In other cases, the source may still exist but its status may be uncertain, changed, or difficult to establish. Source conditions should therefore not be reduced to a simple distinction between persistence and disappearance. For later interpretation, a source may be available, unavailable, or unresolved.

Information persistence consequently does not guarantee source recoverability. Conversely, loss of the original source does not necessarily terminate the availability of information derived from it.

The distinction becomes more complex because a social-media publication is rarely a single undifferentiated unit. Text, image, video, caption, link, quoted material, and other elements can contribute differently to what is communicated. These components do not necessarily persist together.

A post may remain while a linked page changes. An image may survive after its caption is separated from it. A quotation may retain the wording of a claim while losing a qualification that originally accompanied it. A screenshot may preserve visible composition while no longer providing access to linked evidence. Persistence of one component therefore does not establish persistence of the relations through which the original communication was interpreted. (Klein et al., 2014)

This exposes three different questions that should not be collapsed into one.

The first is whether the relevant object or representation continues to exist. This is a question of **physical persistence**.

The second is whether the distinctions required to reconstruct what the communication originally meant remain available. This is a question of **semantic persistence**.

The third is whether the conditions required to determine what the information can still validly support remain identifiable. This is a question of **validity persistence**.

These forms of persistence can diverge. A post can remain physically accessible while a changed link or missing qualification makes its original meaning harder to reconstruct. Its original meaning can remain understandable while changed external conditions mean that a historical claim should no longer be treated as a statement about the present. Conversely, an original object can disappear while sufficient source information and context survive elsewhere to establish what was communicated.

Persistence should therefore not be treated as a single state. The continued existence of an object, the continued recoverability of its meaning, and the continued ability to determine the bounds of valid use are related but distinct conditions.

This distinction also clarifies the role of **addressability**. Information can persist without being practically reachable through a usable path. Conversely, information that remains addressable can potentially be referenced or accessed even when it is no longer part of active distribution. Addressability therefore concerns whether a recoverable path to information exists; it does not establish that the information is currently being retrieved. (Rights of Robots, 2026b)

An addressable publication may never be selected again. Information may also become addressable indirectly through a quotation, reference, derivative, or other relation after the original publication itself has become unavailable. Continued addressability of information therefore does not necessarily require continued addressability of the original object.

This separation is important because actual retrieval is a further event. Information may exist and be addressable without being retrieved. Retrieval establishes that information has entered a

particular information process; it does not by itself establish why it was selected or what consequence followed.

These distinctions do not turn persistence into a planning obligation. Many campaign assets can remain available without creating any material problem. Some may remain understandable without their original surroundings. Others may have no significant informational role beyond their immediate communication purpose.

The importance of persistence lies instead in what it makes possible: information and the conditions that originally supported its interpretation can cease to follow the same path. The publication may remain while a source disappears. A claim may remain while its qualification becomes difficult to recover. A visual may remain while its relation to a caption or supporting information is lost. A derivative may continue while the status of its source becomes unresolved.

The transition beyond active distribution is therefore not a transition into an archive, a dormant stage, or a new lifecycle phase. It is the point at which the continued availability of information can no longer be assumed to include the continued availability of the conditions under which that information was originally understood.

Persistence establishes that this separation is possible. It does not establish whether the separation changes what the information can reasonably mean or support.

## 4. From Persistence to Changed Context

Persistence becomes relevant to interpretation when the information that remains available and the conditions required to understand it do not persist in the same way. This does not mean that older information necessarily loses its meaning or validity. The more precise question is which distinctions remain recoverable when information is encountered outside the communication situation in which it was originally produced. (Brandtzaeg & Lüders, 2018; Marwick & boyd, 2011)

Campaign communication routinely relies on context. A product comparison may depend on a particular model, baseline, configuration, or measurement condition. A commitment may depend on a stated time horizon and scope. A quotation may depend on its source. An image may acquire part of its communicative significance through a caption or accompanying statement. A response to an event may rely on circumstances that did not need to be explained when the event was current.

During active communication, many of these dependencies can remain implicit because the surrounding campaign supplies them. This is not an informational defect. Communication depends on shared references and does not need to reproduce its entire interpretive basis within every individual asset.

The problem changes when the asset remains available while those supporting conditions become less accessible. A comparison can remain while its baseline becomes difficult to identify. A statement can remain while the situation to which it responded has changed. A link can lead to an updated page rather than the information that accompanied the original publication. A visual can remain while the relation that once connected it to a qualification or source is no longer apparent.

The continued existence of the asset therefore cannot by itself establish that its original meaning remains equally recoverable. This is the distinction between physical and semantic persistence introduced in the previous section. A representation can survive while some of the distinctions required to interpret it within its original bounds become harder to reconstruct.

Semantic persistence does not require every aspect of the original situation to remain available. The relevant requirement is narrower. The distinctions necessary to understand what the communication established must remain sufficiently recoverable for the particular use being made of it.

This is why context should not be treated as a single surrounding container that is either present or absent. Different parts of context perform different functions. Some identify what a statement refers to. Some establish when it applied. Some qualify its scope. Some identify its source or evidence. Others contribute tone, cultural significance, or relations between components of a multimodal asset. Their loss can therefore have very different consequences.

The same principle applies to validity. A statement can remain understandable as a historical statement without remaining valid as a description of the present. An announcement that accurately records what an organization intended at one time can remain valid evidence of that announcement even if the intention was later revised. A product claim can remain an accurate record of what was stated about an earlier product version without establishing the same property for a later version.

This requires a distinction between **historical validity** and **current contextual validity**. Historical validity concerns what the information can support about the communication and conditions to which it originally referred. Current contextual validity concerns what that information can support when it is used in relation to a present situation.

The age of information alone determines neither. Older information does not become invalid merely because time has passed. Equally, continued accessibility does not make an earlier statement valid for current conditions. **Content age is not validity decay, and present relevance is not the same as current contextual validity.**

This distinction is particularly important for claims. A claim has a bounded reach: it concerns particular properties, entities, relations, conditions, or circumstances. The evidence supporting that claim is bounded as well. Later availability does not enlarge either boundary.

If a campaign established a claim about one product version under specified conditions, retrieval of that statement in another context does not turn it into a claim about every later version. If evidence supported a comparison under one measurement basis, continued availability of the numerical result does not broaden the conditions under which that evidence applies.

The relation between **Claim Reach** and **Validity Reach** is therefore inherited rather than recreated by later retrieval. What the claim says and what its supporting basis can justify remain bounded by the original conditions. A later context can make the claim relevant to a new question, but relevance does not itself extend what the claim established. (Rights of Robots, 2026a)

This also separates **Addressability** from Claim Reach and Validity Reach. Addressability concerns whether information can be reached or referenced. Claim Reach concerns what a claim asserts. Validity Reach concerns what its basis supports. An increase in the ways through which information can be found does not enlarge either the claim or its evidentiary support.

Transformation makes this boundary especially important. A later quotation, summary, screenshot, repost, or machine-generated representation may carry information from an original

communication into a different setting. The derivative remains related to its source, but it is not identical to the original communication.

A quotation can preserve wording while dropping a qualification. A summary can preserve a proposition while generalizing its scope. A screenshot can preserve visible content while losing access to supporting evidence. In such cases, the original claim has not expanded. Instead, a later representation may have produced a differently bounded claim that must be evaluated on its own terms.

This distinguishes the propagation of an existing claim from the derivation of another claim. Repeating or carrying an identifiable claim forward does not itself change its original reach. A transformation that alters its scope, conditions, referent, or qualification can create a derived claim whose relation to the original must remain distinguishable. (Rights of Robots, 2026a)

**Derivation preserves a relation to the source; it does not preserve identity with the source.**

This is also why provenance matters. If later information is used merely to establish that a particular statement was made, the source, statement, and relevant time may provide sufficient context. If the same information is used to establish a current product property, additional evidence and current conditions may be required. The context required for interpretation therefore depends partly on what the later use is trying to establish. (W3C, 2013)

Recoverability must accommodate this variation. Material context can be **embedded** in the surviving information, **linked** through a reliable relation, **reconstructed** from identifiable supporting information, or in some cases **supplied when the information is retrieved**. No single mechanism is required.

These alternatives matter because recoverability is not the same as self-containment. A social-media post does not need to reproduce every source, qualification, historical condition, or related asset within itself. What matters is whether the distinctions required for a materially consequential interpretation can still be established when needed.

Nor does recoverability require certainty where the available information cannot provide it. If a source has disappeared, attribution has become uncertain, or the original context cannot be reconstructed sufficiently, the appropriate result may be to constrain what can be concluded. Missing context should not be replaced by invented context merely to produce a complete interpretation.

This principle also applies to intent. The fact that an older communication can be interpreted meaningfully in a later context does not establish that its author intended that later interpretation. Later interpretability is not evidence of earlier intent. New contexts can reveal previously unnoticed relevance, but they can also generate new interpretations that did not form part of the original communication situation.

Meaning and intent must therefore remain distinguishable. A later information process may rediscover a relation that was already present, or it may construct a new relation between persistent information and a current question. Without evidence, the existence of the later relation should not be projected backward as an original intention or historical relation.

The same caution applies to multimodal communication. A post can function as one publication while containing several semantically distinct components. Text, image, caption, link, and quoted material can contribute different claims, references, or qualifications. **Compositional unity does not imply semantic unity.**

A visually persuasive representation also does not establish the evidentiary basis of a claim merely because the two appear together. Where a material claim depends on evidence, representational plausibility must not silently substitute for that evidence. The relevant relations between representation, claim, source, and supporting basis need to remain distinguishable.

Context change therefore presents a problem of bounded interpretation rather than a general problem of aging information. Persistent information can remain historically meaningful, become newly relevant, lose some contextual support, or participate in later representations without any of these conditions automatically determining the others.

The planning significance arises only where the loss of a distinction would materially alter what the information appears to establish. Context Dependency identifies that vulnerability. **Materiality determines whether the vulnerability warrants planning action.**

Where the dependency is material, the objective is semantic durability rather than contextual permanence. Semantic durability concerns whether the distinctions required for bounded later interpretation remain sufficiently recoverable across time and recontextualization. It does not require contexts to remain unchanged.

A persistent information environment can therefore accommodate changing contexts. What it must not do is silently collapse past and present conditions into one undifferentiated state. Historical and current contexts can differ while both remain interpretable, provided their relevant distinctions remain recoverable.

## 5. AI-Mediated Retrieval and the Changing Information Environment

Persistent information becomes operationally relevant when there are paths through which it can enter a present information process. This is where the distinction between **persistence**, **addressability**, and **retrieval** established above becomes important. Information can persist without remaining practically addressable, remain addressable without ever being retrieved again, and be retrieved without thereby becoming correct, current, authoritative, or consequential in the situation in which it appears.

Search already demonstrates this separation from original distribution. A user looking for information about a product, organization, event, or topic can encounter material that was published outside the current campaign period. The publication becomes relevant because of its relation to the present inquiry rather than because it still occupies a prominent position in its original distribution sequence.

Recommendation systems create another path. Their selection of information does not need to reproduce publication order. Recency can remain an important signal while operating alongside other indications of relevance. Older information can therefore become selectable without implying that chronological organization has disappeared or that recency no longer matters. (Campos et al., 2014)

AI-mediated retrieval extends this condition. Information can be selected in response to a question or task, combined with information from other sources, summarized, or represented within a generated response. A person can consequently receive information derived from campaign communication without necessarily encountering the original campaign asset itself. (Lewis et al., 2020)

This possibility does not make AI the cause of the problem examined in this paper. The difference between active distribution and continued information availability predates contemporary AI-mediated retrieval. So do search, references, links, quotations, recommendation, and other paths through which older information can become accessible.

The significance of AI lies instead in the changing conditions of selection and representation. Information originally produced for one communication situation can become relevant to another information task because of its relation to that task rather than its original temporal position. Machine-mediated retrieval can therefore increase the practical importance of distinctions that were previously supplied by the campaign environment.

A numerical claim, for example, can be relevant to a current question even when the campaign in which it appeared has long ended. Retrieval can make the claim present to the current task. It does not make the claim current in the sense of changing its historical scope or evidentiary basis. The distinction between **present relevance** and **current contextual validity** remains necessary.

The same applies when several pieces of information are selected together. Their co-retrieval establishes that they have been brought into the same present information process. It does not establish that they were historically related, that one responded to another, that they share the same scope, or that a causal relation exists between them.

Semantic similarity likewise establishes neither equivalence nor historical relation. Two publications can concern the same organization while referring to different products, periods, markets, or conditions. Two claims can address the same subject while relying on different evidence. A later statement can revise an earlier one. Their relevance to the same question does not erase these differences.

AI-mediated systems can also create representations rather than merely reproduce retrieved objects. A generated answer may summarize a source, extract a proposition, combine several sources, or express an inferred relation between them. This creates an additional boundary that must remain explicit: a system's representation of a relation is not evidence that the represented relation exists in the world or existed historically.

A system may identify two statements as related because they concern similar subjects. That establishes something about the system's information process. Whether the statements are actually historically connected, whether one supports the other, or whether they describe the same state requires separate evidence.

The distinction is epistemically important. **Existence, observation, representation, and validation are not interchangeable conditions.** Something may exist without having been observed. A system may represent something that has not been validated. Validation of a representation does not itself imply that the underlying world has changed. Retrieval systems operate across these distinctions but do not eliminate them.

Transformation creates a related problem. When information appears through a generated summary or synthesis, the resulting representation may retain some distinctions from its sources while omitting others. Source, date, qualification, comparison basis, or uncertainty may become less visible even where the substantive proposition appears to have been preserved.

This does not make transformation inherently unreliable. Summarization, extraction, comparison, and synthesis are legitimate information operations. The relevant question is

whether the distinctions necessary for the particular interpretation remain available or recoverable.

A further distinction concerns the status of relations themselves. Retrieval can create new present arrangements of persistent information. Two publications that were never connected during their original distribution can become relevant to the same current task. That present arrangement can be useful without being projected backward as an original relation between the publications.

Relations therefore have their own provenance and temporal conditions. A relation established by an original campaign, a relation documented by a later source, and a relation generated through present retrieval are not interchangeable merely because they connect the same information objects. (W3C, 2013)

This becomes especially important where information appears contradictory. An earlier and a later statement can both be correctly retrieved. Their difference may reflect changed circumstances, correction, disagreement, different scope, or different evidence. Retrieval alone does not resolve which interpretation applies. It creates the informational situation in which those distinctions have to be evaluated.

The same boundary applies to effect. Retrieval establishes that information has entered a particular process. It does not establish that a person has accepted it, acted on it, changed a belief, or produced any other externally observable consequence. **Retrieval creates an opportunity for effect; it is not itself evidence of effect.**

Presence, retrieval, and effect therefore remain distinct. A consequence can also occur later than the encounter from which it developed.

For campaign planning, none of this creates a requirement to optimize information for particular AI systems. Ranking criteria, retrieval architectures, interfaces, and machine-processing practices can change. Designing a planning framework around a particular technical mechanism would make the framework dependent on conditions outside the planner's control.

The more stable planning concern lies upstream. If information materially depends on knowing its source, date, referent, scope, qualification, comparison basis, or relation to supporting information, that dependency exists before any later retrieval takes place. The planning question concerns whether such distinctions remain sufficiently recoverable, not whether a particular system will later choose to use them.

AI-mediated retrieval therefore changes the practical environment without changing the underlying Research Core. It increases the range of circumstances in which persistent and addressable information can be selected, transformed, and combined outside its original temporal distribution setting. It does not establish the future selection of any particular campaign asset, and it does not determine the validity or effect of information once selected.

## 6. From Stream to Space

The preceding sections establish two different ways in which campaign information can acquire relevance. During active distribution, the **Stream** describes an information condition in which temporal distribution is a dominant organizing principle: temporal position contributes to visibility and supplies part of the context through which information is interpreted. The **Space** describes a different information condition. Persistent information can remain addressable through structures and relations other than its original temporal position, allowing it to become

relevant to a present task without depending on its continuing position in the distribution sequence in which it originally appeared.

Stream and Space are not different classes of content. The same publication can participate in both conditions. Nor does information necessarily move from one to the other through a defined lifecycle. Some information disappears. Some remains available but is never encountered again. Some continues to attract attention. Some becomes relevant after a period with little observable presence.
Space therefore does not describe a later stage of Stream. It describes a condition under which original temporal position is no longer the only meaningful path through which persistent information can be addressed.

This makes **addressability** central to Space. Persistent information need not be currently retrieved, but it must remain connected to a recoverable path through which it can potentially participate in another information process.

Space consequently cannot be reduced to storage. An archive can preserve an object because it belongs to the past. Space concerns the possibility that persistent information can participate in a present information situation through relations that make it relevant now. (ISO, 2025)

This does not make all persistent information equally addressable. An original publication can remain directly accessible. Another may be reachable only through a reference. Information from a removed publication may survive in a quotation, screenshot, summary, or other derivative. In still other cases, the source or relation may remain unresolved. Space therefore accommodates different conditions of access without treating them as equivalent.

Nor does addressability establish importance. Something can be addressable without being retrieved, retrieved without being useful, and useful without being valid for every purpose for which it might be selected. Space increases the possible relations through which information can participate in a present process; it does not predetermine the outcome of those relations.

This distinction also prevents **Addressability**, **Claim Reach**, and **Validity Reach** from collapsing into one concept. Information may become addressable through additional paths without the underlying claim saying anything more than it originally said. Increased addressability likewise does not expand the evidence that supports the claim. A broader relational environment is therefore not a broader claim and not broader validity.

Space also does not remove time from information. In important respects, the opposite is true. When temporal proximity no longer supplies context automatically, the historical position of information can become more important.

An announcement made in one year and a later statement concerning the same subject may both be addressable in relation to a current question. Their joint availability does not place them in the same state. One may describe an earlier intention and the other a later outcome. Preserving the temporal distinction can be necessary to understand both.

The same applies to apparently contradictory information. Different statements can represent different states, perspectives, evidence, or circumstances. Space should make their relation inspectable without requiring those differences to disappear.

This is especially important in social media because persistent information does not consist solely of settled factual propositions. It includes situated claims, perspectives, experiences,

reactions, announcements, interpretations, and assertions made by different sources under different conditions.

A useful information space cannot therefore depend on resolving these different contributions into a single undifferentiated account. Its value can instead lie in preserving enough distinction to determine who or what a statement is attributable to, what it concerns, when it was made, under what relevant conditions it was expressed, and how its status may subsequently have changed.

The objective is **continuity without collapse**. Persistent information can become connected to new questions and other information while the distinctions required to interpret those connections remain available.

Connection itself must therefore remain bounded. Two objects can be related because they concern the same product without making equivalent claims about that product. Two statements can address the same event while representing different perspectives. An earlier and later claim can concern the same referent while describing different states.

**Connection does not establish equivalence.**

The same principle applies to aggregation. Bringing several information objects together can expose relationships, differences, or unresolved questions. It does not necessarily resolve them. A collection of related statements is not automatically a unified conclusion.

**Aggregation is not necessarily resolution.**

This matters because relations are themselves information. A relation can have a source, a time, and a basis. A relation explicitly established by the original publication differs from one documented by a later source, inferred from shared characteristics, or created by present retrieval.

Space therefore requires attention not only to the provenance of objects but also to the provenance of relations between them. Without that distinction, a present arrangement of persistent information could be mistaken for a relation that existed when the information was originally produced.

The relational environment of an information object can consequently grow while the historical identity of the original object remains bounded. Later references, responses, quotations, summaries, classifications, and retrieval contexts can accumulate around it. Those additions become part of its evolving relational environment; they do not retroactively become part of the original publication.

This protects another important boundary: **relational growth must not imply object expansion.**

A derivative illustrates the distinction. A screenshot can be related to an original post without becoming that post. A summary can represent information from a source without inheriting the complete identity or context of the source. A later response can concern an earlier statement without becoming part of the earlier statement's original communication situation.

The same boundedness applies to meaning. A later relation can make a previously unnoticed aspect of information relevant. It can also produce an interpretation that did not exist in the

original situation. Later interpretability therefore cannot, by itself, establish earlier intended meaning.

Space must accommodate both possibilities: rediscovery of distinctions already supported by the information and formation of new present relations around persistent information. What matters is that the provenance and status of those relations remain distinguishable.

This is also why ambiguity and plurality cannot simply be treated as defects to be eliminated. Two perspectives can remain different without one being an incomplete version of the other. An unresolved relation can remain unresolved. A source can remain unavailable or uncertain. Preserving these conditions can be more accurate than forcing them into a unified representation.

Space is therefore not a promise of complete semantic resolution. It is a condition in which persistent information can remain addressable while relevant differences in source, time, context, claim, representation, validity, and relation can remain recoverable.

Multimodal information reinforces this requirement. A publication can appear as one interface object while its meaning depends on several components and their relations. A caption may qualify an image. A link may provide the evidentiary basis for a numerical statement. Quoted material may belong to another source. Visual containment alone does not establish that all components have the same informational identity or evidentiary status.

Space must therefore preserve distinctions not only between publications but also where necessary within them. Compositional unity does not imply semantic unity.

The transition from Stream to Space should consequently not be understood as a claim that social media is becoming timeless or that chronological feeds are disappearing. Stream remains a powerful and continuing form of information organization. Space identifies the additional condition created when persistent information remains addressable beyond dependence on that temporal organization.

For campaign planning, this creates a bounded consequence. Campaign assets are produced for communication in the Stream. Some of the information they contain may remain addressable under the conditions described as Space. Planning does not need to predict which information will later be retrieved or how its relational environment will develop.

## 7. A Planning Model for Persistent Campaign Information

The implications of persistent campaign information do not justify treating every social-media asset as a future information resource. Campaign communication is produced for a current purpose, and that purpose remains the primary planning condition. Persistence becomes a planning concern only where information may remain relevant beyond that purpose and where loss of context could materially change its later interpretation.

The planning model therefore begins with **Current Purpose**. The planner first determines what the communication must achieve within the campaign situation for which it is being created. A promotional asset may need to communicate an immediate offer. A product communication may need to establish a particular comparison. An announcement may need to communicate a commitment. The possibility of later information use does not override these immediate functions.

The second consideration is **Potentially Persistent Information**. The question is not whether the complete asset will remain technically available indefinitely, nor whether it will actually be retrieved later. The planner identifies information within the communication that could plausibly retain informational relevance beyond the immediate campaign purpose.

Where no such relevance is apparent, the analysis stops. Persistence alone does not create a requirement for additional treatment.

A temporary promotion illustrates this boundary. Its price, availability, or call to action may be explicitly tied to a limited period. The communication can be highly dependent on its campaign context while having little reason to function as an enduring information resource. Designing it for hypothetical future interpretation could add unnecessary information and weaken its immediate function.

Where potentially persistent information is present, the third consideration is **Context Dependency**. The planner asks what a later user or system would have to infer if the surrounding campaign were no longer available.

For a quantified product claim, this may include the comparison baseline, measurement conditions, product version, market, or supporting evidence. For an announcement of a future commitment, it may include the date on which the statement was made, the entity to which it applied, the intended time horizon, and the scope of the commitment. For a multimodal asset, relevant information may be distributed across image, text, caption, overlay, link, or quoted material.

Context Dependency identifies vulnerability, not an automatic requirement for intervention. Social communication routinely depends on context, and much of that dependence is harmless outside the immediate communication situation.

The fourth consideration is therefore **Materiality**. The planner asks whether loss of the identified context would materially change what the persistent information appears to establish.

This is not a probability estimate. Materiality does not ask how likely future retrieval is, how many people might encounter the information, or whether an AI system will select it. It concerns the significance of the missing distinction if a later encounter occurs.

A quantified claim can have high materiality because loss of its comparison basis can change what the number appears to mean. A future commitment can have high materiality because removal of its temporal scope can make a historical statement appear to describe a current obligation or state. By contrast, an expressive campaign asset can depend heavily on cultural or emotional context while the loss of that context creates no materially misleading factual interpretation.

**Context dependency identifies vulnerability; materiality determines whether that vulnerability warrants planning action.**

Only where the dependency is material does the fifth consideration, **Recoverability**, require action. The objective is not to preserve the complete campaign environment. It is to reduce the amount of materially important context that a later interpretation would otherwise have to guess.

Recoverability can take the forms established above: material context may be embedded, linked, reconstructable, or supplied when the information is later retrieved. The appropriate form depends on the communication and on the distinction that needs to remain recoverable.

A product claim such as “40% lighter” demonstrates the issue. The wording alone does not identify what the product is lighter than, under which configuration the comparison was made, or where the basis for the comparison can be established. If those conditions are material to the claim, some recoverable relation to them is required.

The same principle applies to multimodal communication. A numerical statement may appear as an image overlay while its qualification appears in a caption and its evidence behind a link. The publication can function as one campaign asset while the information required for bounded interpretation is distributed across several components. If one component persists without the others, the apparent claim can change.

The planner therefore does not ask whether every component should be preserved. The question is whether the material relation between them remains sufficiently recoverable. Visual coherence or representational plausibility cannot substitute for the evidence required to support a material claim.

Future-oriented statements require a different treatment. A commitment published at one point in time can later remain useful as historical evidence that the commitment was made. It does not need to remain an accurate description of the organization's current state to retain that historical function.

Durability therefore does not require timelessness. A statement can remain durably interpretable precisely because its original temporal conditions remain identifiable.

Derivatives create another case. The original publication may disappear while a quotation, screenshot, summary, or other representation persists. Planning cannot control whether such derivatives will be created or how they will later circulate. It can, however, influence whether materially important distinctions such as source, referent, qualification, or temporal scope are recoverable from the information initially published.

This does not make every possible derivative the planner's responsibility. Derivation preserves a relation to an original source but does not preserve identity with it. A later transformation that changes the claim, removes qualifications, or introduces new assertions must be evaluated as a later representation rather than treated as an expanded version of the original campaign asset.

Expressive and culturally situated communication provides the opposite stress case. Such content may rely strongly on timing, shared references, humor, emotion, or cultural circumstances. Attempting to make every such dependency permanently explicit could destroy the communicative quality that makes the asset effective in the first place.

High Context Dependency therefore does not imply high Materiality, and neither implies universal preservation. The planning model is deliberately selective.

The model progressively reduces the set of assets requiring intervention. Current Purpose protects the immediate communication objective. Potentially Persistent Information excludes material with no plausible continuing informational relevance. Context Dependency identifies what later interpretation could lose. Materiality excludes dependencies whose loss does not materially alter interpretation. Recoverability applies only to what remains.

This boundedness is essential. The model does not assign persistence scores, predict future retrieval, estimate AI visibility, require archival treatment, or introduce a new campaign-performance metric. It does not ask planners to anticipate every possible future interpretation.

Its function is narrower: to identify, before publication, cases in which potentially persistent information depends on context whose later absence would materially change what that information appears to establish, and to determine whether the necessary distinctions can remain recoverable without compromising the communication's current purpose.

## 8. Operational Integration into Campaign Planning

The planning model does not require a separate campaign process. Its function is to introduce a bounded diagnostic perspective at the point where the actual information structure of an asset is sufficiently concrete to evaluate.

At the briefing stage, campaign objectives, audiences, messages, channels, and other strategic conditions are still being established. Introducing a mandatory persistence assessment at this point would require planners to evaluate information dependencies before the relevant content and representations necessarily exist. No separate *From Stream to Space* gate is therefore required in the briefing. (Fill & Turnbull, 2019)

The same applies to concept and strategy development. Awareness that some campaign information may remain available beyond active distribution can inform decisions, particularly where a campaign contains substantial claims, commitments, comparisons, or source-dependent information. Awareness alone, however, does not justify an additional approval stage.

The primary diagnostic point lies in **late asset development or pre-production**. At this stage, the actual communication can be examined: wording, claims, images, overlays, captions, links, sources, qualifications, temporal references, and relations between components are sufficiently defined to determine whether a material context dependency exists.

The assessment begins with the asset's current communication purpose, which remains controlling. Potentially persistent information is considered only where it has plausible informational relevance beyond that purpose. If no such information is present, no additional treatment is required.

Where potentially persistent information exists, Context Dependencies can be examined against the actual asset. Questions about referent, source, temporal scope, comparison basis, qualification, evidence, or relations between multimodal components can then be evaluated concretely. Materiality determines whether an identified dependency requires action.

Where a material dependency exists, Recoverability is addressed through the least intrusive treatment sufficient for the communication. The objective is not to maximize recoverability: additional context can impose costs on clarity, attention, visual composition, and immediate communication function. The intervention therefore remains bounded by the material dependency actually identified.

Existing review and approval processes provide the appropriate verification point. Review does not repeat the full analysis and does not introduce a separate approval logic. It verifies, where a material issue was identified, that the relevant distinction has not been lost during production and that the implemented treatment still supports the intended communication.

This placement preserves the distinction between **planning control** and later information behavior. Before publication, planners can influence the representation being released: its claims, qualifications, temporal markers, sources, references, component relations, and other relevant communication conditions.

Publication marks the **Control Boundary**. After publication, the planner cannot determine whether the asset will remain available, whether a platform will continue to expose it, whether another party will quote or capture it, whether a search or AI-mediated system will retrieve it, whether a derivative will be created, or which other information may later be presented alongside it.
The planner likewise cannot determine the future social context, knowledge state, relevance, or interpretation surrounding the information. These are possible later relational outcomes rather than controllable properties of the original campaign plan.

**Campaign planning determines initial conditions, not future relational outcomes.**

This boundary prevents Recoverability from becoming an open-ended responsibility for future information use. A campaign team can make a source, date, qualification, comparison basis, or other material distinction identifiable in the original communication, but it cannot guarantee that a later screenshot, quotation, derivative, or generated summary will preserve it. Such later transformations can be evaluated in relation to the original information, but their existence does not retrospectively expand what the campaign planner controlled at publication.

No additional campaign KPI follows from the model. Impressions, engagement, conversions, reach, and other performance measures concern campaign performance under their respective measurement conditions. The possibility that information remains persistent or addressable does not create a meaningful performance metric by itself.

Nor does the model introduce a persistence score, semantic-durability score, future-retrieval probability, or similar proxy. Such measures would imply predictive capabilities that the framework does not establish and would risk converting a bounded interpretive problem into an unsupported optimization objective.

Campaign end also does not create a new operational stage. The framework does not require teams to archive every asset, preserve every surrounding campaign element, maintain every link indefinitely, or conduct a general post-campaign preservation exercise.

Where an organization already has legal, records-management, product-information, evidence-retention, or archival obligations, those requirements remain governed by their respective purposes. The planning model does not replace or expand them merely because the same information originated in social-media communication. (Upward, 1996; ISO, 2025)

A related but distinct pre-publication requirement exists in U.S. advertising substantiation. For objective advertising claims, the Federal Trade Commission requires advertisers to possess a reasonable basis before dissemination. That requirement concerns whether the claim is adequately substantiated when made; the present model asks a different question: whether contextual distinctions whose later loss would materially alter interpretation should remain recoverable if campaign-produced information persists beyond active distribution. The two questions can overlap where evidentiary context matters, but FSTS Materiality is not the FTC test of materiality to consumer decision-making. (Federal Trade Commission, 1984)

The operational intervention is therefore deliberately limited: potentially material context dependencies are assessed when the communication is concrete enough to evaluate;

necessary distinctions are addressed before publication where this can be done within the planner's control; existing review verifies the resulting treatment; and publication marks the boundary beyond which future relational outcomes are not treated as campaign-planning decisions.

The possible informational life of campaign content can extend beyond active distribution without requiring the campaign-planning process itself to extend across that later informational life.

## 9. Discussion

The framework developed here addresses a specific mismatch between campaign distribution and the later informational conditions of campaign-produced content. Social-media campaigns are planned within bounded periods of activity, but the information they produce does not necessarily share the same temporal boundary. The resulting issue is not persistence alone. It is the possibility that information remains available while the conditions required for its bounded interpretation change, disappear, or become differently accessible.

This distinction separates the framework from accounts concerned primarily with continuing attention or delayed communication effects. A campaign can produce effects after active distribution has ended, and attention to campaign content can decline gradually rather than stop at a fixed point. Neither condition is equivalent to a later informational encounter with persistent content. Such an encounter can occur without demonstrating continuous attention, prior exposure, or a delayed effect of the original campaign. (Kumkale & Albarracín, 2004; Anderson, 2006)

The end of active distribution therefore remains a meaningful campaign boundary. What it does not necessarily provide is an informational boundary. Treating these boundaries separately avoids extending campaign management indefinitely while still accounting for information that can remain available beyond the period for which its distribution was planned.

Persistence alone is insufficient to determine the significance of this condition. Physical, Semantic, and Validity Persistence can diverge, just as historical validity can remain intact without establishing current contextual validity.

The same boundary applies to claims and transformations. Increased Addressability or renewed relevance does not enlarge Claim Reach or Validity Reach. Quotations, screenshots, summaries, reposts, and machine-generated representations can propagate information or produce derived claims, but later relational growth does not expand the historical identity of the original communication. **Derivation preserves relation, not identity.**

Source recoverability remains purpose-dependent. Information derived from an unavailable original can still support some historical determinations without necessarily supporting present substantive claims. The relevant provenance therefore depends on what a later use attempts to establish.

This principle extends to identity. Content identity, source identity, account identity, author identity, and attribution should not be assumed to coincide merely because they once appeared together in the same communication environment. Identifier continuity can assist identification, but it does not by itself establish uninterrupted identity across every relevant dimension.

The same caution applies to interface presentation. Several informational components can appear within one post or one retrieved representation without becoming one information

object. Text, image, quotation, source link, caption, and derivative material can have different origins and evidentiary functions. Interface containment does not establish information-object identity, just as compositional unity does not establish semantic unity.

These distinctions become increasingly important under the conditions described as Space. Persistent information can become addressable through relations other than its original temporal position. This expands the relational environment in which information can participate without expanding the historical object itself.

Space should therefore not be understood as an archive. An archive can preserve information from the past without making the relational conditions described here its defining characteristic. Space concerns the ability of persistent information to participate in present information processes through addressable relations while relevant differences in time, source, context, claim, validity, and representation remain distinguishable. (ISO, 2025)

This produces an important consequence for social information. Social-media environments contain claims, perspectives, experiences, announcements, reactions, and interpretations produced by different sources under different conditions. Persistence does not turn these into a single account of reality.

An earlier statement and a later correction can both remain informative. Different participants can describe the same event differently. A claim can remain historically important even where its proposition is later disputed or shown to be incorrect. In such cases, the durable informational fact may include that a particular source expressed a particular claim under particular conditions at a particular time.

The value of Space therefore does not depend on resolving multiple persistent statements into one truth. It can instead depend on preserving the distinctions required to determine what each statement is, where it came from, what it referred to, when it applied, and what later information changed its status.

This is not an argument for treating all perspectives as equally valid. Preserving distinctions is what makes evaluation possible. Evidence can support one claim more strongly than another. A statement can be corrected or contradicted. A source can be unreliable. The requirement is that aggregation should not silently perform that evaluation by collapsing materially different information into an undifferentiated representation.

**Aggregation is not necessarily resolution.**

Co-retrieval illustrates why relational provenance matters. Bringing persistent objects into the same present information process establishes a current arrangement, not necessarily a historical relation between them. A relation established within the original communication therefore differs from one documented later, inferred from shared properties, or generated through present retrieval.

AI-mediated retrieval intensifies the practical importance of this distinction without creating it. Retrieval, representation, and validation remain separate: a system can represent a proposition or relation without thereby establishing its truth, current validity, source authority, historical existence, or effect. Nor does later interpretation establish earlier intended meaning.

This prevents semantic recoverability from becoming reconstruction of presumed intent. Recoverability concerns the distinctions necessary to bound interpretation, not the production of a definitive account of what an author privately intended.

These informational conditions do not expand campaign control. The planner's intervention remains concentrated before publication and only where loss of a contextual distinction would materially change what potentially persistent information appears to establish. Materiality prevents Context Dependency from becoming a universal preservation requirement, while Recoverability limits intervention to the distinctions required for bounded interpretation.

The contribution of *From Stream to Space* is consequently not the discovery that digital content can persist, that communication effects can be delayed, that information can be retrieved semantically, that algorithms can select older content, or that provenance and context matter. These phenomena belong to broader existing research. (Treem & Leonardi, 2013; Brandtzaeg & Lüders, 2018; Kumkale & Albarracín, 2004; Campos et al., 2014; Lewis et al., 2020; W3C, 2013)

The contribution lies in relating them to a prospective campaign-planning problem: information created under temporally bounded distribution conditions can remain available and addressable beyond those conditions, while the context required for bounded later interpretation may persist differently. This creates a planning problem before publication without turning future information behavior into an extension of campaign control.

## 10. Limitations and Research Agenda

The framework developed in this paper is conceptual. It establishes distinctions and a bounded planning logic; it does not establish how frequently campaign information remains available beyond active distribution, how likely later encounter is, or how often contextual dependencies produce materially different interpretations. Technical persistence alone cannot provide such estimates because platform availability, indexing, searchability, account status, access conditions, external references, and other factors can change independently.

The same limitation applies to Addressability. The framework distinguishes persistence from the existence of a recoverable path to information, but it does not provide a measure of how addressable a particular information object is. Nor does Addressability predict retrieval. An object can remain addressable without being selected for any later information process.

Empirical research should therefore distinguish continued existence, observable presence, addressability, actual retrieval, and subsequent effect rather than using one as a proxy for another. Longitudinal studies could examine how these conditions diverge across platforms and over time.

Platform differences require particular attention. Social-media systems differ in search functions, recommendation mechanisms, account structures, deletion practices, visibility controls, link behavior, interfaces, and access available to external systems. The framework does not assume that information persists or remains addressable in the same way across these environments.

Research on temporal behavior could examine how campaign information moves between periods of observable activity and low or absent presence without treating such patterns as evidence of a fixed lifecycle. This would also allow the bounded use of Dormancy to be tested empirically without converting it into a required stage between Stream and Space.

A second research area concerns **context loss**. The framework assumes that different contextual components can persist differently, but it does not establish which losses most frequently change interpretation. Controlled studies could separate loss of temporal markers,

source information, referents, qualifications, comparison baselines, supporting evidence, and relations between multimodal components.

Such research would also help distinguish Physical, Semantic, and Validity Persistence empirically. An object can remain technically available while participants can no longer reconstruct its original meaning or determine the conditions under which a claim was supported. Measuring these differences would provide a stronger empirical basis than treating persistence as a binary property.

The distinction between historical and current contextual validity also requires empirical examination. Studies could test how users and systems interpret older statements when dates, versions, changing circumstances, corrections, or subsequent states are more or less recoverable. This would help determine under which conditions historical information is incorrectly interpreted as current information.

**Materiality** requires separate validation. The framework deliberately does not define a score or universal threshold. Its function is to distinguish contextual vulnerabilities that warrant planning action from those that do not. Research should therefore examine whether different evaluators can identify materially consequential context loss with sufficient consistency and whether judgments vary systematically by content type or use case.

Recoverability presents another empirical question. Context can be embedded, linked, reconstructed from related information, or supplied when information is retrieved. The framework does not establish that these mechanisms are equally reliable. Future studies could compare how effectively they preserve bounded interpretation under different forms of persistence and transformation.

This question is particularly important for multimodal information. Text, image, caption, overlay, quotation, and linked evidence can persist independently. Research could examine which component relations are most vulnerable to separation and when such separation changes the apparent scope, source, or evidentiary status of a claim.

Transformation requires similar attention. Quotations, screenshots, summaries, reposts, and generated representations can preserve some properties of an original communication while changing others. Empirical work could examine when Claim Propagation remains distinguishable from Claim Derivation and which provenance signals are sufficient to prevent derivatives from being mistaken for their sources.

Source recoverability should not be assumed to be binary. The original source may remain available, become unavailable, or become unresolved. Future research could investigate how these different conditions affect the interpretation and validation of surviving derivatives and how much provenance is required for different kinds of historical or current claims.

Relations between persistent objects also require study. Search and AI-mediated retrieval can bring information together that did not share an original communication context. Experiments could examine when users or systems infer unsupported equivalence, historical connection, contradiction, or causality from co-retrieval and which provenance or temporal distinctions reduce such inference.

This is particularly relevant to social information containing multiple situated claims and perspectives. The framework argues that aggregation should not automatically collapse these distinctions. Further work is needed to determine how information systems can preserve

differences in source, time, claim status, evidence, and perspective while still enabling comparison and synthesis.

AI-mediated retrieval represents an important but deliberately bounded research area. This paper does not establish how frequently AI systems retrieve older campaign information, which retrieval architectures do so, or how particular models select, transform, or cite such information. Those are empirical questions and can change with system design.

Research should therefore avoid treating AI retrieval as a uniform mechanism. Relevant studies could compare whether source, date, qualification, referent, and evidentiary relations survive retrieval and transformation under different systems and tasks. They could also examine when generated representations introduce relations that were not present in the retrieved sources.

Such work must retain the distinction between retrieval and validity. Successful retrieval does not establish that information is current, authoritative, correctly interpreted, or supported for the use being made of it. Likewise, the representation of a relation by a system does not establish that the relation exists independently of that representation.

A further limitation concerns the economics of persistent information. Space can contain more addressable information than any particular information process can practically retrieve, contextualize, validate, and synthesize. **Token economics may influence which addressable information is economically selected and how much contextual reconstruction can be performed, but this paper does not model those allocation mechanisms.** Addressability therefore remains distinct from the resource constraints governing whether and how an available informational path is operationally used.

**Future research should examine whether machine-processing costs create an economic selection layer between theoretical addressability and operational retrieval, and whether the structural recoverability of persistent content affects the cost of its later interpretation.**

Finally, the framework has a deliberate boundary of responsibility. It identifies decisions available before publication but does not assign responsibility for every later transformation, retrieval, relation, or interpretation to the original campaign planner. Empirical and normative research may establish additional responsibilities in particular organizational, regulatory, or technical settings. Those questions require separate treatment rather than expansion of the planning model beyond its present control boundary.

## 11. Conclusion

Social-media campaigns are temporally bounded communication activities, but the information produced through them does not necessarily share the same temporal boundary. A publication can remain available after active distribution has ended, information derived from it can survive after the original object has disappeared, and persistent information can later become relevant through relations that were not part of its original distribution setting.

The resulting planning problem is not persistence alone. It arises where information remains available while the context required for its bounded interpretation persists differently. A statement can survive while its temporal scope becomes less apparent. A numerical claim can remain while its comparison basis becomes difficult to recover. A derivative can preserve information while obscuring its source. An earlier statement can become relevant to a current question without becoming a statement about current conditions.

The distinction between **Stream** and **Space** captures these different informational conditions. Stream is organized strongly through temporal distribution and its accompanying contextual proximity. Space describes the additional condition in which persistent information remains addressable through structures and relations other than its original temporal position. The two coexist; Space is neither a later lifecycle stage nor an archive.

Addressability does not itself establish retrieval, effect, broader Claim Reach, or broader Validity Reach. Nor does persistence guarantee that meaning, validity, source, or context remain equally recoverable. Historical information can remain interpretable without becoming a statement about current conditions, and later representations can remain related to an original communication without becoming identical to it.

Space therefore depends on preservation of relevant distinctions rather than their collapse into a unified representation. Persistent claims, perspectives, corrections, states, and derivatives can become jointly addressable without becoming equivalent or historically related. AI-mediated retrieval increases the practical relevance of this condition but does not cause it. Retrieval and machine-generated representation do not themselves establish validity, authority, historical relation, or effect.

For campaign planning, the consequence is deliberately bounded. Current Purpose remains primary. Potentially Persistent Information warrants intervention only where loss of a Context Dependency would materially change what the information appears to establish. Materiality determines whether action is warranted; Recoverability concerns whether the distinctions required for bounded later interpretation can remain available without compromising the communication's immediate function.

This does not imply universal preservation, self-contained assets, prediction of future retrieval, new performance metrics, additional approval processes, or post-campaign archival obligations.

The appropriate point of intervention remains before publication, when the communication and its material dependencies are still within planning control. Publication constitutes the Control Boundary. The planner can influence the claims, representations, qualifications, temporal markers, sources, references, and relevant relations being released, but cannot determine their future ranking, retrieval, transformation, combination, interpretation, or effect.

**Campaign planning determines initial conditions, not future relational outcomes.**

The framework therefore answers the Research Question by locating responsibility neither in indefinite management of persistent content nor in prediction of future information use. Social-media campaign planning should account for context dependencies where potentially persistent information contains distinctions whose later loss would materially alter bounded interpretation, and it should address their recoverability within the conditions controllable before publication.

The contribution is correspondingly limited. *From Stream to Space* does not propose a new theory of digital persistence, delayed communication effects, semantic retrieval, provenance, claim validity, or machine-mediated information selection. It connects these established problem areas to the temporal boundary of campaign planning and identifies the planning consequence created when active distribution can end before the informational availability of campaign-produced content ends.

The framework remains conceptual. Empirical research is required to determine how often these conditions occur, which context losses become material, and how different retrieval and transformation environments preserve or remove relevant distinctions.

These questions do not alter the central boundary established here. The informational life of campaign-produced content can extend beyond the period of active campaign distribution, while the planner's controllable intervention remains bounded to the conditions that can be established before publication.

## References

Anderson, C. (2006). The Long Tail: Why the Future of Business Is Selling Less of More. Hyperion.

Brandtzaeg, P. B., & Lüders, M. (2018). Time collapse in social media: Extending the context collapse. Social Media + Society, 4(1). https://doi.org/10.1177/2056305118763349

Campos, P. G., Díez, F., & Cantador, I. (2014). Time-aware recommender systems: A comprehensive survey and analysis of existing evaluation protocols. User Modeling and User-Adapted Interaction, 24(1–2), 67–119. https://doi.org/10.1007/s11257-012-9136-x

Federal Trade Commission. (1984). FTC Policy Statement Regarding Advertising Substantiation. 104 F.T.C. 839. https://www.ftc.gov/legal-library/browse/ftc-policy-statement-regarding-advertising-substantiation

Fill, C., & Turnbull, S. (2019). Marketing Communications: Touchpoints, Sharing and Disruption (8th ed.). Pearson.

International Organization for Standardization. (2025). ISO 14721:2025 Space Data System Practices — Reference model for an open archival information system (OAIS) (3rd ed.). ISO.

Klein, M., Van de Sompel, H., Sanderson, R., Shankar, H., Balakireva, L., Zhou, K., & Tobin, R. (2014). Scholarly context not found: One in five articles suffers from reference rot. PLOS ONE, 9(12), e115253. https://doi.org/10.1371/journal.pone.0115253

Rights of Robots. (2026a). Claim Validity Across Transformation (Version 1.3). Research paper, rightsofrobots.com.

Rights of Robots. (2026b). HCS-08 — Context Infrastructure Interface Specification (Working Draft 0.1). Human Context System Specification Series.

Kumkale, G. T., & Albarracín, D. (2004). The sleeper effect in persuasion: A meta-analytic review. Psychological Bulletin, 130(1), 143–172. https://doi.org/10.1037/0033-2909.130.1.143

Lewis, P., Perez, E., Piktus, A., Petroni, F., Karpukhin, V., Goyal, N., Küttler, H., Lewis, M., Yih, W.-t., Rocktäschel, T., Riedel, S., & Kiela, D. (2020). Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems, 33.

Marwick, A. E., & boyd, d. (2011). I tweet honestly, I tweet passionately: Twitter users, context collapse, and the imagined audience. New Media & Society, 13(1), 114–133. https://doi.org/10.1177/1461444810365313

Treem, J. W., & Leonardi, P. M. (2013). Social media use in organizations: Exploring the affordances of visibility, editability, persistence, and association. Communication Yearbook, 36(1), 143–189. https://doi.org/10.1080/23808985.2013.11679130

Upward, F. (1996). Structuring the records continuum part one: Post-custodial principles and properties. Archives & Manuscripts, 24(2), 268–285.

World Wide Web Consortium. (2013). PROV-O: The PROV Ontology. W3C Recommendation. https://www.w3.org/TR/prov-o/