

Diffusion Before Adoption: A Model of Compositional Latent Diffusion

Abstract

Diffusion is often observed through adoption, but the unit that moves through a population need not be the unit eventually adopted. We introduce latent diffusion as social propagation of adoption-relevant elements through actors who need not adopt the focal object those elements enable, and compositional latent diffusion as the case in which adoption depends on which distinct elements coexist. An agent-based model shows that transmission can build a population of carriers before adoption expands and that actors not currently adopting can carry much of the element traffic. Finite retention governs whether circulating elements can coexist in sufficient configurations. In paired simulations with identical retained histories, increased transmission has a consistently smaller return when readiness depends on distinct element identities than when it depends on input count or source diversity. Diffusion can therefore be well under way before adoption becomes visible, while conversion depends on what elements actors can hold together.

Introduction

Diffusion is a population-level process in which socially transmissible material, practices, or states propagate among actors. Adoption is an actor-level state in which a focal object is taken up. We use focal object for the practice, behavior, innovation, or other larger unit whose adoption is being tracked. Researchers commonly use adoption events as observable evidence that diffusion is occurring (Rogers 2003; Strang and Soule 1998). That inference is straightforward when the transmitted and adopted units coincide. But what moves socially can instead be something that makes adoption possible, and it can travel through a population well before the focal object is adopted.

Some practices and innovations depend on multiple adoption-relevant elements. We use element for a substantively distinguishable unit that can be acquired, retained, and retransmitted separately from the focal object and that contributes to its adoption. An element may be a constituent part of the focal object, such as one technique within a multipart practice, or a socially transmissible enabling condition that helps make adoption possible. In either case, elements are not arbitrary subdivisions of an otherwise unitary exposure. Adopting a complex practice such as algebra, for example, requires command of smaller elements such as numbers and arithmetic operations, which can be learned, used, and taught without algebra itself.

Clinical care bundles provide a concrete analogue. Central-line infection-prevention bundles, for example, combine distinguishable practices such as hand hygiene, maximal-barrier precautions, chlorhexidine skin antisepsis, and appropriate catheter-site selection. Park et al. (2017) implemented such a bundle in a surgical intensive care unit using peer tutoring in which nurses prepared and taught portions of the prevention guidelines to one another. Because the practices can be taught and learned separately, knowledge of one can be acquired and transmitted without the larger bundle being implemented.

We use configuration for the set of adoption-relevant elements that coexist in an actor at a given time. When adoption requires a sufficient configuration, substantial social transmission can occur even while the focal object itself remains rare or absent.

We call this process latent diffusion: the social propagation of adoption-relevant elements through actors who need not adopt the focal object those elements enable. Stephan (2023) uses the same term for a related measurement problem in which diffusion status is unobserved and adoption is observed. Our use instead names a substantive process in which adoption-relevant elements themselves propagate through nonadopters. Compositional latent diffusion is the case in which adoption depends on which distinct

elements coexist in a configuration rather than only on how much input has accumulated. Latent diffusion need not be compositional. A specific argument for joining a new party, for example, could circulate through people who have not joined and contribute to joining without requiring any particular combination of distinct arguments. By contrast, implementing a central-line bundle requires distinct practices to be held together. Element identity is also separate from source diversity (Centola and Macy 2007; Wang et al. 2015): source diversity concerns who supplied an input, whereas element identity concerns what the input contains.

These distinctions have been developed separately in prior research. Our contribution is to bring them together in a model in which adoption-relevant elements can circulate through nonadopters while adoption depends on which of them coexist in a sufficient configuration. When the transmitted unit and the adopted unit differ in this way, circulation and adoption need not move together.

We use an agent-based model to derive consequences of this possibility rather than to estimate how common latent diffusion is empirically. In the baseline simulations, transmission builds a substantial population holding relevant elements without adopting the corresponding focal object, and nonadopters account for roughly half or more of successful relevant-element transfers. Conversion from circulation to adoption depends on whether enough relevant elements coexist to form a sufficient configuration, a state we call readiness. To isolate the specifically compositional claim, paired simulations hold complete retained histories fixed while evaluating those same histories according to input count, source diversity, or element identity. Across the tested capacities and loss conditions, increased transmission has a smaller return when sufficiency depends on distinct element identities than under the other two rules. Without spontaneous loss, distinct-element readiness is the only rule with a clearly negative return at both capacities. Finally, targeted transmission after adoption can produce additional adoption along the trajectory without a corresponding endpoint gain when supporting configurations later deteriorate.

Transmission, Composition, and Adoption

The review below takes up four distinctions from prior research: transmission from participation, acquisition from implementation, redundant from nonredundant social inputs, and unitary objects from objects composed of multiple features. None alone implies latent diffusion.

One line of prior work separates transmission from participation and acquisition from adoption. In their study of microfinance diffusion, Banerjee et al. (2013) estimated that participants were about seven times as likely to transmit information as informed nonparticipants. Yet information transmitted by nonparticipants still accounted for roughly one-third of eventual participation. Attewell (1992) showed that organizations could know about a technology while still lacking the technical know-how required for implementation, often acquiring that know-how through mediating institutions. Fichman and Kemerer (1999) similarly distinguished acquisition from deployment, while Lesthaeghe and Vanderhoeft (2001) modeled behavioral transitions as depending on several conditions being satisfied together. Together, these studies establish two possibilities: actors can contribute to onward transmission without participating in the focal behavior, and adoption-relevant states can accumulate before adoption. Latent diffusion connects them by allowing what is acquired before adoption to remain socially transmissible.

Threshold and complex-contagion research addresses a different form of nonredundancy. Granovetter (1978) modeled adoption as depending on how many others have already acted. Centola and Macy (2007) and Centola (2010) show how reinforcement from multiple independent sources changes the conditions under which diffusion spreads, a line of work subsequently developed across a range of settings (Guilbeault, Becker, and Centola 2018). Wang et al. (2015) explicitly represent memory for nonredundant sources, so another input from the same person need not be equivalent to one from somebody new. Compositional latent diffusion introduces a separate dimension: the identity of the transmitted content.

Inputs from several people can carry the same element, while several distinct elements can come from one person. Source diversity and element identity are therefore independently variable dimensions of an exposure history and can operate together in the same diffusion process.

Other work motivates representing diffusion content as distributed, divisible, and subject to capacity constraints. Carley (2001) represents knowledge as distributed across people and relations rather than reducible to a single observed state. Weng et al. (2012) show how competing information changes diffusion when attention is limited. Axelrod (1997) represents culture as multiple features that actors can share in part without sharing as a whole. Research on knowledge recombination likewise treats innovations as configurations whose constituent elements can acquire new functions when combined (Xiao, Makhija, and Karim 2022). Recombination asks how elements become innovations. Our question is what happens when such elements themselves diffuse before the larger object is assembled and adopted. Once elements circulate separately, widespread circulation does not ensure that the elements required for adoption will coexist within the same actor. Retention therefore becomes consequential for conversion from circulation into adoption.

Related diffusion models show why the level at which propagation is represented matters. Rossman (2014) looks upward from a particular innovation to the broader category whose legitimacy can affect its diffusion. Our move looks downward from the focal object to adoption-relevant elements that can circulate before it is adopted. Goldberg and Stein (2018) make a different change in transmitted content by modeling the diffusion of perceptions about which practices belong together rather than transmission of the practices themselves. Compositional latent diffusion instead treats the transmitted units as states that contribute directly to adoption of a focal object without being identical to that object. Their role in diffusion follows from their ability, as distinguishable adoption-relevant states, to be acquired, retained, and retransmitted separately while contributing to adoption.

Figure 1 represents these unit relationships schematically. Direct transmission takes the form $A \rightarrow A$, in which the transmitted and adopted units coincide. Count-based accumulation takes the form $x+x+x \rightarrow A$, in which adoption depends on accumulated input without requiring different input identities. Compositional latent diffusion takes the form $a+b+c \rightarrow A$, in which readiness depends on which distinct elements coexist. The shorthand distinguishes relationships between transmitted content and the focal object rather than defining mutually exclusive classes of diffusion models.




Figure 1. Conceptual architecture of compositional latent diffusion. The socially transmitted unit need not be the focal object that is eventually adopted. Adoption-relevant elements can circulate separately, including through actors who have not adopted the focal object. Direct transmission is represented as $A \rightarrow A$, count-based accumulation as $x+x+x \rightarrow A$, and compositional latent diffusion as $a+b+c \rightarrow A$. In the compositional case, element identities determine whether a sufficient configuration is

present. Finite retention can prevent required elements from coexisting when new arrivals displace material already held. The schematic isolates content identity from input amount. Source diversity is a separate dimension that can operate alongside either.

These foundations imply a layered relationship between propagation and adoption. Latent diffusion concerns propagation: adoption-relevant elements can move socially through actors who have not adopted the focal object. Compositional latent diffusion adds a condition on conversion: readiness depends on which elements coexist in a configuration. Finite retention is not part of either definition, but it governs whether sufficient configurations can form and persist once elements circulate separately. Activation is a separate condition that determines whether a sufficient configuration becomes manifest, the model state corresponding to adoption. An enabling condition counts as an element only if it can itself be acquired, retained, and retransmitted. Fixed resources, private dispositions, and conditions that merely trigger manifestation therefore lie outside the element-level diffusion process.

The model addresses four questions that follow from this architecture. First, when elements can circulate independently of adoption, how much realized transmission passes through nonadopters, and how extensively do relevant elements accumulate before the focal object becomes manifest? Second, how does finite retention affect the conversion of circulating elements into adoption by determining whether earlier and later elements can coexist? Third, holding retained histories fixed, does increased transmission have different consequences when readiness depends on input count, source diversity, or element identity? Finally, when actors who already manifest the focal object increase transmission of its supporting elements, does the persistence of assembled configurations shape what that transmission accomplishes? We also examine whether two candidate signatures, abrupt takeoff and simultaneous arrival of complementary elements, accompany the process.

Model

The model represents the theoretical process in deliberately abstract form. Literal constituent parts and socially transmissible enabling elements are represented identically as members of an object's element set, so the model does not distinguish their dynamics. Its purpose is to separate element-level circulation from focal-object adoption and then vary the conditions under which circulating elements can form and maintain a sufficient configuration.

Population and Network

The model contains 100 agents connected by a Watts-Strogatz small-world network with mean degree six and rewiring probability 0.1. This structure provides clustered local neighborhoods while retaining some longer-range connections. At each tick, agents act in random order. Each agent contacts one randomly selected network neighbor.

We distinguish contact from transmission. A contact creates an opportunity for an element to move. Conditional on contact, a transmission attempt occurs with a specified probability, set to 0.1 in the baseline specification. Experimental treatments vary this transmission probability while leaving the number of contacts unchanged. The manipulation therefore changes how often contacts generate transmission attempts rather than directly setting the number of successful acquisitions. Realized throughput remains an outcome of the diffusion process.

Elements and Focal Objects

Eight element types circulate through the population, and two focal objects can be adopted. Each focal object has a four-element requirement set, meaning that four element types are relevant to that object. The two requirement sets share one element, so seven element types are relevant to at least one focal object. The eighth is relevant to neither and serves as background material. Elements that do not contribute to a particular focal object, including an element relevant only to the other object, still circulate and compete for retention capacity.

At baseline initialization, each agent is independently seeded with each element with probability 0.15. Relevant material is therefore already present before transmission begins. Because this initial distribution affects carrier prevalence, we report carrier prevalence at initialization alongside its later values. This allows carriers built through subsequent transmission to be distinguished from those present at the start.

A focal object's four-element requirement set is distinct from the configuration sufficient for readiness. In the baseline specification, an actor becomes ready when it retains any three distinct elements from the relevant four-element set. The four elements are therefore associated with the focal object, but no individual element is required in the baseline condition. Because manifestation is reversible, the distinction between the requirement set and the readiness threshold also matters after adoption. An actor holding all four relevant elements can lose one and remain above the readiness threshold with three, whereas an actor holding exactly three falls below readiness if one is lost. Other analyses vary the number of distinct elements sufficient for readiness while holding the four-element requirement set fixed.

Transmission and Retention

When a transmission attempt occurs, one element currently held by the sender is selected. Anyone holding an element can transmit it, regardless of whether that actor has adopted a focal object the element serves. This rule makes diffusion through nonadopters possible. Whether nonadopters actually become important to realized transmission is a property of the resulting trajectories rather than something imposed by the transmission rule.

Agents retain at most one copy of each element type. If a recipient already holds the selected element, no acquisition occurs and the transmission opportunity is spent. Otherwise, the element enters the recipient's holdings.

Agents also have finite retention capacity, defined as the maximum number of distinct elements they can hold simultaneously. Baseline retention capacity is six. If a new element arrives when an actor is already at capacity, one currently held element is selected at random and displaced to admit the arrival. Elements that do not contribute to a focal object compete for the same retention capacity as elements that do. Additional transmission can therefore supply a missing element while also increasing turnover among material already retained.

Some experiments additionally introduce spontaneous loss, in which retained elements can disappear independently of new arrivals. Spontaneous loss and displacement are separate mechanisms. Displacement occurs when an incoming element replaces material in a full retention set. Spontaneous loss removes retained material without requiring an incoming element.

Readiness and Activation

After transmission and any applicable loss, the model recomputes each actor's readiness. An actor is ready for a focal object when it currently retains a sufficient configuration of elements. Under the baseline rule, this means at least three distinct elements from that object's four-element requirement set.

Readiness is not itself adoption. The model includes an additional activation condition that determines whether a ready actor manifests the focal object. We use manifestation for the model state corresponding to adoption. A nonadopter is accordingly an actor not currently manifesting the focal object. Because manifestation is reversible, nonadoption is a state at the time of observation rather than a lifetime category. An actor can therefore possess a sufficient configuration without yet manifesting the focal object.

The baseline specification uses scheduled activation, analogous to a focal object becoming useful or permitted because of an exogenous event. Because the schedule is externally imposed, it separates activation timing from element circulation and allows accumulation and retention to be examined under controlled timing. We repeat the central accumulation and capacity analyses under locally triggered activation, which has no external schedule: a ready actor manifests a focal object once at least half of its network neighbors are ready for, or already manifest, that same object. Comparing these regimes shows which trajectory features persist when the activation rule changes. Immediate manifestation upon readiness, in which ready actors manifest at once, provides a further robustness regime (online supplement, S3).

Manifestation is reversible. If spontaneous loss or displacement later reduces an actor's holdings below the readiness threshold, the focal object ceases to be manifest. Readiness, activation, and manifestation therefore describe distinct parts of the process.

Unless otherwise stated, reported adoption is the mean object-specific manifest share: the number of manifest actor-object pairs divided by the number of actors times the number of focal objects. With 100 actors and two focal objects, the denominator is 200 actor-object pairs. This differs from the share of actors manifesting at least one focal object, which is used only where explicitly identified in the carrier diagnostics (online supplement, S2.5).

Unless otherwise stated, baseline simulations use retention capacity six, initial element density 0.15, transmission probability 0.1 conditional on contact, and 200 replications per experimental cell with predetermined random seeds. Full parameter values, update order, initialization procedures, and robustness analyses appear in the online supplement.

Verification and Mechanism Discrimination

The simulations serve two inferential purposes: verifying that the code implements the stated model and discriminating among explanations for patterns the model generates. The second matters because similar aggregate diffusion patterns can arise from different processes (Palloni 2001; Hébert-Dufresne, Scarpino, and Young 2020), and robustness across parameter values does not by itself identify which feature of the model produces a result.

For verification, automated mechanism assertions test whether specified events produce the expected state changes, structural invariants are checked during runs, fixed seeds establish reproducibility, and reported quantities are audited against stored replication-level output. These checks establish correspondence among the formal specification, its implementation, and the reported results.

Mechanism discrimination relies on comparisons that hold the other model rules fixed while changing the focal feature. Some alter the activation rule while leaving the transmission rule unchanged. Because these comparisons do not preserve the same stochastic history, they establish robustness to the activation regime rather than a same-history contrast. Other comparisons use paired runs in which retained histories, delivered material, or delivery opportunities are held fixed while the theoretically relevant rule changes. These stronger controls can distinguish candidate explanations that would otherwise remain compatible with the same aggregate pattern, including by showing when a candidate mechanism fails to produce the contrast it predicts.

Results

Elements Before Adoption

Scheduled activation remains the baseline, but we begin the pre-adoption results with locally triggered activation because it removes the fixed waiting period that could mechanically create a carrier interval. Immediate manifestation upon readiness provides the stricter robustness boundary (online supplement, S3). Because any actor holding an element can retransmit it regardless of whether that actor has adopted the focal object it serves, the question for the simulations is whether relevant elements merely accumulate among nonadopters or also move onward through them.

We define a carrier with respect to a focal object as an actor who holds at least one element relevant to that object but does not currently manifest it. The broad carrier measure is calculated separately for each focal object and then averaged across objects. A stricter measure requires at least two relevant elements without manifestation. Both measures are proportions of actors. We use latent substrate for the share of adoption-relevant actor-element cells currently occupied across the population.


Relevant elements accumulate before adoption expands (Figure 2, Panel A). Under locally triggered activation, no external schedule determines when manifestation begins. At each replication's own adoption onset, mean latent substrate has already reached 0.46. Because onset varies across replications with a standard deviation of about nine ticks, this quantity is calculated within each run and then averaged rather than read from the mean trajectory at the mean onset. Substantial adoption-relevant material is therefore already present in the population when manifestation begins to rise.

Transmission also builds a carrier population beyond what initialization supplies (Figure 2, Panel B). Under locally triggered activation, the broad carrier measure begins at 0.474 because independent seeding already gives many actors at least one relevant element. It then rises to a peak of 0.852 before declining as manifestation expands. The stricter measure more clearly isolates accumulation produced after initialization. Only 0.109 of actors satisfy the two-element carrier criterion initially, while subsequent transmission raises that share to 0.562. The resulting carrier population therefore cannot be attributed simply to initial seeding.

Accumulation alone would not establish latent diffusion if actors holding relevant elements did not pass them onward. We therefore record the sender's manifestation state after every successful transfer of an element relevant to at least one focal object (Figure 2, Panel C). Under locally triggered activation, 0.49 of successful relevant-element transfers (SD 0.04) originate with senders who manifest none of the focal objects served by the transmitted element. Under scheduled activation, the corresponding share is 0.60 (SD 0.04). For an element shared by both focal objects, a sender is classified as manifest if either associated object is currently manifest. These counters record completed transfers without altering model states or transmission probabilities.

The difference between 0.49 and 0.60 shows how activation timing affects the share of transmission occurring through nonmanifest actors. Scheduled activation keeps actors nonmanifest for longer by construction, so more relevant traffic occurs before manifestation. Locally triggered activation shortens that interval, yet nonmanifest senders still account for roughly half of successful relevant-element transfers. Substantial pre-adoption circulation therefore persists when the external activation schedule is removed, even though the timing and shape of manifestation change. Under immediate manifestation upon readiness, the strictest robustness regime, nonmanifest senders still carry a substantial share of relevant transmission and transmission still builds carriers well beyond their seeded level (online supplement, S3). The supplement also shows that the stepwise adoption pattern under scheduled activation is produced by that activation rule.

Together, the carrier trajectories and transfer records establish the central pre-adoption pattern: relevant elements accumulate among nonmanifest actors, and those actors participate extensively in onward transmission before manifestation expands.



The two shares differ by activation regime; they are not invariant model constants.

Figure 2. Element circulation before adoption. Panel A compares latent substrate with mean object-specific manifestation under locally triggered activation. Latent substrate is the share of adoption-relevant actor-element cells currently occupied. The annotation at adoption onset, 0.46, is calculated within each replication at that run's own onset and then averaged because onset timing varies across runs. Panel B

shows two object-specific carrier measures under the same activation regime: the broad measure requires at least one relevant element without manifestation, while the stricter measure requires at least two. Both are calculated separately by focal object and then averaged. Dotted reference levels indicate their prevalence at initialization, allowing transmission-built accumulation to be distinguished from seeded carriage. Curves end at the last tick for which all 200 replications remain active, and diamonds report mean run-end values. Panel C classifies successful relevant-element transfers according to whether the sender currently manifests any focal object served by the transmitted element. Nonmanifest senders account for 0.49 of such transfers under locally triggered activation and 0.60 under scheduled activation.

Retention and Adoption

Once adoption-relevant elements are circulating, conversion depends on whether enough of them can coexist within the same actor to form a sufficient configuration. Figure 3, Panel A varies both retention capacity and the number of distinct elements sufficient for readiness. When two elements are sufficient, mean final adoption crosses 0.50 at a retention capacity of about 3.7. Requiring three elements shifts that crossing to about 5.6, and requiring all four shifts it to about 7.2. As more distinct elements must coexist, greater retention capacity is required to reach comparable adoption levels. Circulation alone therefore does not determine adoption. Conversion also depends on whether actors can retain the required configuration amid the other material they hold.




Figure 3. Retention capacity limits adoption. Panel A shows mean final adoption as the readiness threshold and retention capacity vary. White markers indicate the interpolated capacity at which mean adoption reaches 0.50, using the two neighboring capacities actually simulated. The half-adoption points are approximately 3.7, 5.6, and 7.2 when two, three, and four distinct elements are sufficient for readiness, respectively. Panel B holds retention capacity at six and varies the density of elements present at initialization. Final adoption changes little despite the large increase in starting material.

The capacity dependence is not explained by scarcity of starting material. Holding retention capacity at six while increasing initial element density from 0.05 to 0.60 changes final adoption only from 0.67 to 0.65 (Figure 3, Panel B). Actors begin with substantially more material, but their ability to retain elements simultaneously is unchanged. Increasing what is initially available therefore does not remove the constraint imposed by limited coexistence.

The same relationship persists when activation is locally triggered rather than scheduled. Under locally triggered activation, final adoption rises from 0.13 at retention capacity four to 0.78 at capacity six and is effectively saturated by capacity seven. Under scheduled activation, the corresponding values are 0.13, 0.65, and 0.99. Changing the activation rule changes the level and timing of manifestation, but greater

retention capacity continues to make sufficient configurations easier to form. The dependence also persists under immediate manifestation upon readiness (online supplement, S3).

These results identify retention as a condition governing conversion from circulating elements into adoption. As the number of distinct elements required for readiness increases, the retention capacity needed for substantial adoption rises systematically, and supplying much more material at initialization does not remove that dependence. This relationship does not by itself establish a specifically compositional effect, because a count-based readiness rule could also become harder to satisfy under limited capacity. Isolating the additional consequence of element identity therefore requires holding retained histories fixed while varying whether readiness depends on input count, source diversity, or element identity. The token-memory comparison in the next section does so.

Element Identity

The specifically compositional claim is that readiness depends on which retained elements are present, not only on how many relevant inputs an actor holds or how many different people supplied them. We therefore hold complete retained histories fixed while changing only the criterion by which those histories become sufficient for readiness.

For this comparison, we use a purpose-built token-memory variant of the model. A token is a retained input that records both the element it contains and the sender who supplied it. Unlike the main model, which permits at most one retained copy of each element type, this variant allows duplicate element tokens. Those duplicates make it possible to vary input count independently of element identity and therefore to construct an equal-count comparison. The variant otherwise retains the same network structure, contact and transmission process, retention capacity, eviction rule, and spontaneous-loss mechanism.

The same complete retained token history is evaluated under three readiness rules. Under equal-count readiness, any three relevant retained tokens are sufficient, including duplicates. Under distinct-sender readiness, three relevant tokens must come from three different senders, regardless of their element identities. Under distinct-element readiness, three different elements from the focal object's requirement set must be retained, regardless of who supplied them. Paired histories are identical at every tick in token holdings, admission order, element labels, sender labels, losses, and evictions. Only the rule determining whether those holdings constitute readiness differs.

The design therefore separates input quantity, source diversity, and content identity. Three copies of one element supplied by three people satisfy distinct-sender readiness but not distinct-element readiness. Three different elements supplied by one person do the reverse. In the clinical example, learning the same prevention practice from three colleagues supplies source diversity without supplying three different practices, while one trainer who teaches three different practices supplies the required content diversity from a single source. The distinct-sender arm thus tests whether source nonredundancy alone reproduces the response associated with element identity.

The central result is a stable ordering in the return to increased transmission (Figure 4). Across both tested capacities and both loss conditions, raising transmission probability from 0.05 to 0.20 has a smaller return under distinct-element readiness than under either distinct-sender or equal-count readiness. At retention capacity six without spontaneous loss, the return is -0.058 (SE 0.006) under distinct-element readiness, $+0.011$ (SE 0.008) under distinct-sender readiness, and -0.009 (SE 0.009) under equal-count readiness. At capacity five, the corresponding returns are -0.032 (SE 0.005), $+0.033$ (SE 0.009), and $+0.017$ (SE

0.010). Distinct-element readiness is therefore the only rule with a clearly negative response at both capacities when spontaneous loss is absent.

Introducing spontaneous loss at rate 0.01 changes the signs but not the ordering. At capacity six, the returns are +0.059 under distinct-element readiness, +0.185 under distinct-sender readiness, and +0.215 under equal-count readiness. At capacity five, they are +0.018, +0.140, and +0.173, respectively. Spontaneous loss gives additional transmission a replenishment role because lost material can be supplied again, so increased transmission helps under all three rules. What persists across loss regimes is a consistently smaller return when readiness depends on assembling particular element identities.




Figure 4. Element identity changes the return to transmission. The schematic illustrates the independence of source diversity and element diversity. Three copies of one element from three senders satisfy distinct-sender readiness but not distinct-element readiness, while three different elements from one sender do the reverse. Paired runs have identical token holdings and retention histories at every tick and differ only in how those holdings are evaluated for readiness. Panels A and B show final adoption under lower and higher transmission without spontaneous loss. Panel C reports the paired high-minus-low contrast at both tested capacities under no-loss and loss conditions. Across all four conditions, the return to increased transmission is smaller under distinct-element readiness than under the other two rules. Without spontaneous loss, distinct-element readiness is the only rule with a clearly negative return at both capacities. With spontaneous loss of 0.01, all three returns are positive. Intervals show ± 1.96 times the paired standard error of the within-seed difference across 200 replications.

The paired histories show why finite retention alone cannot explain this ordering. At capacity six, all three readiness rules experience about 45 evictions per actor under lower transmission and about 195 under higher transmission. Higher transmission therefore produces the same additional turnover across paired arms. What differs is what the retained material can accomplish. Under equal-count readiness, another

relevant token can contribute even when it duplicates an element already held. Under distinct-sender readiness, another token can contribute when it introduces a new source even if its content is redundant. Under distinct-element readiness, neither additional quantity nor a new source substitutes for a missing element. The required identities must coexist in the retained configuration.

Differences in absolute adoption levels also do not reproduce the ordering. The three rules differ in how easily they can be satisfied, so their adoption levels are not themselves evidence for compositionality. At capacity five without spontaneous loss, distinct-element readiness responds negatively while both comparison rules respond positively. Distinct-sender adoption lies between the other two arms in level, but its response to increased transmission is not intermediate. With spontaneous loss, distinct-element adoption remains the lowest while its transmission response becomes positive. A calibrated-count diagnostic in the online supplement provides an additional check where that comparison is feasible: its adoption level falls below the distinct-element arm while its response to increased transmission remains positive. The ordering in transmission returns therefore cannot be recovered from absolute adoption levels alone.

The paired comparison supports a specific inference: identical retained histories have systematically different consequences when readiness depends on element identity rather than input count or source diversity. Because traffic, turnover, and source histories are held constant across readiness rules, those features cannot account for the difference in response. The design does not identify every micro-event through which increased transmission changes a configuration, such as a redundant arrival displacing a useful token from a nearly sufficient set. Nor does the distinct-sender arm test compositional latent diffusion against complex contagion as a general theory. Its narrower role is as a separability control for source nonredundancy within this design: it tests whether source nonredundancy reproduces the response associated with content identity. Because readiness is evaluated over identical histories, the comparison is unchanged when manifestation follows readiness immediately (online supplement, S3).

A separate paired comparison tests whether the compositional result depends on complementary elements arriving together. Each recipient receives one copy of each of four elements from the focal object's requirement set across four fixed delivery opportunities. In the separate-delivery condition, one element arrives at each opportunity. In the joint-delivery condition, all four arrive together at one of those opportunities. Paired runs hold the recipient, delivery opportunities, total delivered content, background traffic, network, and random seed fixed, and all deliveries pass through the ordinary retention process. Flow in this experiment is varied by the frequency of delivery blocks, not by transmission probability.


Joint arrival produces no consistent advantage. At capacity six, the joint-minus-separate difference is -0.001 (SE 0.003) at the lower block frequency and -0.002 (SE 0.004) at the higher block frequency. At capacity five, the corresponding differences are $+0.008$ (SE 0.003) and -0.008 (SE 0.002). All four contrasts are smaller than 0.01 in magnitude, and their signs are inconsistent. Under the conditions tested here, simultaneous arrival therefore does not account for the compositional result. What matters is whether the required element identities coexist in retention, not whether they arrive at the same moment.

Retention After Adoption

The preceding results concern whether circulating elements can coexist long enough for a focal object to become adoptable. But a sufficient configuration can also become unstable after manifestation. Because manifestation is reversible, subsequent loss or displacement of supporting elements can push an actor below the readiness threshold. The final experiment asks how the persistence of these configurations changes the consequences of targeted post-manifest transmission.

We first vary targeted post-manifest transmission, in which the transmission probability of an element is multiplied when its sender manifests a focal object the element serves (Figure 5, Panel B, without spontaneous loss or post-manifest protection). Tripling targeted transmission changes final adoption little, from 0.65 to 0.64, even though the gap between peak and final adoption increases from 0.14 to 0.18. Higher targeted transmission therefore produces additional manifestation along the trajectory without a corresponding endpoint gain. Some configurations that become sufficient under higher transmission subsequently lose the elements needed to remain manifest.

Panel A adds spontaneous element loss and varies post-manifest protection, which reduces spontaneous loss for elements supporting a focal object after that object becomes manifest. Protection does not prevent displacement when new elements enter a full retention set. With no protection from spontaneous loss, tripling targeted transmission raises final adoption from 0.28 to 0.43. With full protection, endpoint adoption is much higher overall, rising from 0.88 to 1.00 across the same transmission comparison. The gain from tripling targeted transmission grows with protection, from 0.15 without protection to about 0.30 at protection 0.75, and then narrows to 0.12 under full protection as adoption approaches its ceiling. Targeted transmission can therefore help produce sufficient configurations, but its endpoint consequence depends partly on whether those configurations subsequently persist.



Protection reduces spontaneous loss after manifestation only. It does not prevent displacement when a new element enters a full retention set.

Figure 5. Post-manifest retention changes the consequences of targeted post-manifest transmission. Panel A crosses targeted transmission level with protection against spontaneous loss for elements supporting a manifested focal object. Protection reduces spontaneous loss after manifestation but does not prevent displacement through finite retention. Cell labels report mean final adoption. Panel B shows the targeted transmission comparison without spontaneous loss or post-manifest protection. Higher targeted transmission raises manifestation along the trajectory without increasing final adoption, producing a larger gap between peak and endpoint manifestation. Together, the panels show that targeted transmission can generate adoption that is subsequently lost and that endpoint adoption depends on the stability of supporting configurations after manifestation.

These results separate formation of a sufficient configuration from its persistence after manifestation. Targeted post-manifest transmission can produce additional manifestation that later disappears when supporting elements are lost or displaced. Protection against spontaneous loss increases the persistence of manifested configurations, but because it does not prevent displacement, the experiment does not establish the consequences of stabilizing retention more generally.

Discussion

The model separates circulation of socially transmissible material from manifestation of the focal object, two processes treated as one when diffusion is represented only through adoption. The results place nonmanifest actors inside the diffusion process as conduits for adoption-relevant material rather than merely as potential future adopters. Classic diffusion research documents a related temporal separation. Iowa farmers commonly heard of hybrid seed corn, often first from salesmen, well before planting it, while neighbors mattered most at the point of adoption (Ryan and Gross 1943). The present model adds a different possibility: what travels before adoption can itself remain socially transmissible through actors who have not adopted the focal object.

Compositional latent diffusion adds a second distinction. The consequences of exposure can depend on content identity rather than only on input amount or source diversity. In the paired token-memory experiment, identical retained histories produce different readiness and adoption outcomes under different sufficiency rules. Across the tested conditions, increased transmission has its smallest return under distinct-element readiness. Without spontaneous loss, that return is clearly negative at both capacities. With loss, it becomes positive but remains smaller than under the comparison rules. The result therefore supports a specific claim: the same exposure and retention history can have different consequences when sufficiency depends on what the retained inputs contain. It does not imply that increasing transmission generally reduces adoption. Distinct-sender readiness does not reproduce the distinct-element response, so the comparison separates content identity from source nonredundancy within the design, while allowing that source diversity and compositionality can operate together in empirical diffusion.

Finite retention constrains whether circulating elements can form and maintain sufficient configurations because the required elements must coexist within an actor. One empirical analogue is the difficulty of spreading multipart practices: individual components may circulate widely even when few actors can sustain the combination required for implementation. Increasing the readiness requirement systematically shifts the retention capacity needed for substantial adoption, while greatly increasing initial element density does not remove that dependence. Retention also matters after manifestation. Targeted post-manifest transmission can produce adoption along a trajectory without producing a durable endpoint gain when supporting configurations subsequently deteriorate. The element-level process therefore has three analytically distinct layers: propagation determines what material moves, compositional requirements determine what must coexist for readiness, and retention affects whether those configurations can form and persist.

These distinctions imply specific empirical requirements. Where latent diffusion operates, adoption histories can understate how far adoption-relevant material has already traveled. Establishing latent diffusion requires showing that distinguishable adoption-relevant elements move onward through actors who have not adopted the focal object, not merely that such actors possess knowledge, prerequisites, or readiness. Establishing compositional latent diffusion additionally requires showing that element identities or configurations matter beyond accumulated input and source diversity. Empirical designs must therefore observe some portion of the content actors acquire or transmit, not only contact and eventual adoption. Table 1 summarizes these requirements.

Table 1. Empirical diagnostics for latent and compositional latent diffusion

Claim	Evidence that would support it	What is insufficient by itself
Latent diffusion	Nonadopters acquire socially transmissible adoption-relevant elements and pass those elements onward to others.	Finding knowledge, enabling conditions, or readiness among nonadopters without showing onward social transmission.
Compositional latent diffusion	Element identities or configurations matter for readiness or adoption beyond the amount of input received and the number of distinct sources.	Repeated exposure, accumulated input count, or source diversity without evidence that content identity matters.
Retention-constrained conversion	Conversion from circulating elements to adoption depends on whether a sufficient configuration can coexist within available retention, and persistence depends on whether supporting elements remain available.	Limited retention or element loss in isolation. These conditions shape the consequences of latent diffusion but do not establish the process itself.

Note. Abrupt takeoff, a particular numerical gap between latent substrate and adoption (online supplement, S3), and simultaneous arrival of complementary elements are not defining diagnostics of latent diffusion.

Settings that record distinguishable elements of transmitted content together with their movement across actors and into subsequent implementation could support empirical tests of these requirements. Training or procedural records in organizations, contribution histories in collaborative technical projects, and other traceable content-level exchanges could reveal whether particular elements pass through nonadopters and later recombine in adoption. The requirement is not to observe every internal state, but to distinguish movement of adoption-relevant content from the focal adoption event.

The intervention implications concern the organization of transmission rather than a general prescription to increase or reduce exposure. Under finite retention, additional transmission can supply missing elements while also increasing turnover among material already held. In empirical settings, an analogous constraint could arise when only a limited set of adoption-relevant routines, practices, or other elements can be sustained simultaneously. The paired identity experiment shows that the consequences of additional transmission depend on what incoming material contributes to the retained configuration. The joint-versus-separate comparison provides no evidence that simultaneous arrival of complementary elements is itself advantageous under the conditions tested. Post-manifest protection against spontaneous loss addresses a different margin by changing the persistence of supporting configurations after manifestation. Interventions may therefore alter not only how much exposure occurs, but which adoption-relevant elements circulate and whether sufficient combinations can be maintained.

The theory applies only when the material preceding adoption has the properties required of an element. An element must be a substantively distinguishable adoption-relevant state or content that can be acquired, retained, and retransmitted separately from the focal object. Arbitrary subdivisions of repeated exposure do not qualify merely because they can be counted separately. Elements may be literal constituent parts or socially transmissible enabling conditions. The model represents these interpretations identically and therefore does not establish whether their empirical dynamics differ.

The model also leaves important features of the information environment unresolved. Elements are discrete and drawn from a small global pool in which seven of eight types contribute to at least one focal object. Empirical environments may contain far more irrelevant material, and the consequences of limiting the model's global pool to mostly adoption-relevant elements have not been tested here. Nor need local exposure mirror the global information pool. Clustering, homophily, and segregation can restrict actors to much narrower sets of available material. The model incorporates clustered local interaction but does not vary how network structure determines which complementary elements are locally accessible. Whether network organization facilitates or prevents the assembly of particular configurations therefore remains an open question.

The broader implication concerns the process that precedes visible adoption. Adoption-relevant elements can circulate through actors who are not manifesting the focal object, while conversion depends on whether a sufficient configuration can be retained. When element identity matters, otherwise identical retained histories can have different consequences for readiness. Visible adoption is therefore one outcome of diffusion rather than the full path by which adoption-relevant material moves through a population.

References

- Attewell, Paul. 1992. "Technology Diffusion and Organizational Learning: The Case of Business Computing." *Organization Science* 3(1):1–19.
- Axelrod, Robert. 1997. "The Dissemination of Culture: A Model with Local Convergence and Global Polarization." *Journal of Conflict Resolution* 41(2):203–226.
- Banerjee, Abhijit, Arun G. Chandrasekhar, Esther Duflo, and Matthew O. Jackson. 2013. "The Diffusion of Microfinance." *Science* 341(6144):1236498.
- Carley, Kathleen M. 2001. "Learning and Using New Ideas: A Sociocognitive Perspective." Pp. 179–207 in *Diffusion Processes and Fertility Transition: Selected Perspectives*, edited by J. B. Casterline. Washington, DC: National Academies Press.
- Centola, Damon. 2010. "The Spread of Behavior in an Online Social Network Experiment." *Science* 329(5996):1194–1197.
- Centola, Damon, and Michael Macy. 2007. "Complex Contagions and the Weakness of Long Ties." *American Journal of Sociology* 113(3):702–734.
- Fichman, Robert G., and Chris F. Kemerer. 1999. "The Illusory Diffusion of Innovation: An Examination of Assimilation Gaps." *Information Systems Research* 10(3):255–275.
- Goldberg, Amir, and Sarah K. Stein. 2018. "Beyond Social Contagion: Associative Diffusion and the Emergence of Cultural Variation." *American Sociological Review* 83(5):897–932.
- Granovetter, Mark. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83(6):1420–1443.

- Guilbeault, Douglas, Joshua Becker, and Damon Centola. 2018. "Complex Contagions: A Decade in Review." Pp. 3–25 in *Complex Spreading Phenomena in Social Systems*, edited by S. Lehmann and Y.-Y. Ahn. Cham: Springer.
- Hébert-Dufresne, Laurent, Samuel V. Scarpino, and Jean-Gabriel Young. 2020. "Macroscopic Patterns of Interacting Contagions Are Indistinguishable from Social Reinforcement." *Nature Physics* 16(4):426–431.
- Lesthaeghe, Ron, and Camille Vanderhoeft. 2001. "Ready, Willing, and Able: A Conceptualization of Transitions to New Behavioral Forms." Pp. 240–264 in *Diffusion Processes and Fertility Transition: Selected Perspectives*, edited by J. B. Casterline. Washington, DC: National Academies Press.
- Palloni, Alberto. 2001. "Diffusion in Sociological Analysis." Pp. 66–114 in *Diffusion Processes and Fertility Transition: Selected Perspectives*, edited by J. B. Casterline. Washington, DC: National Academies Press.
- Park, Sang-Won, Suhui Ko, Hye-sun An, Ji Hwan Bang, and Woo-Young Chung. 2017. "Implementation of Central Line-Associated Bloodstream Infection Prevention Bundles in a Surgical Intensive Care Unit Using Peer Tutoring." *Antimicrobial Resistance & Infection Control* 6:103.
- Rogers, Everett M. 2003. *Diffusion of Innovations*. 5th ed. New York: Free Press.
- Rossman, Gabriel. 2014. "The Diffusion of the Legitimate and the Diffusion of Legitimacy." *Sociological Science* 1:49–69.
- Ryan, Bryce, and Neal C. Gross. 1943. "The Diffusion of Hybrid Seed Corn in Two Iowa Communities." *Rural Sociology* 8(1):15–24.
- Stephan, L. S. Sanna. 2023. "A Trimming Estimator for the Latent-Diffusion-Observed-Adoption Model." arXiv preprint arXiv:2309.01471.
- Strang, David, and Sarah A. Soule. 1998. "Diffusion in Organizations and Social Movements: From Hybrid Corn to Poison Pills." *Annual Review of Sociology* 24:265–290.
- Wang, Wei, Ming Tang, Hai-Feng Zhang, and Ying-Cheng Lai. 2015. "Dynamics of Social Contagions with Memory of Nonredundant Information." *Physical Review E* 92:012820.
- Weng, Lilian, Alessandro Flammini, Alessandro Vespignani, and Filippo Menczer. 2012. "Competition among Memes in a World with Limited Attention." *Scientific Reports* 2:335.
- Xiao, Ting, Mona Makhija, and Samina Karim. 2022. "A Knowledge Recombination Perspective of Innovation: Review and New Research Directions." *Journal of Management* 48(6):1724–1777.