

Merger, Coordinate, and Legibility: Why Ensemble Casts Fail to Cohere

Flyxion
Independent Researcher

July 2026

Abstract

Ensemble casting is usually evaluated through an additive model, in which each performer contributes an independent audience and the film's total reach is treated as the sum of these contributions. This essay argues that the additive model is structurally mistaken, and that the correct object of analysis is not audience aggregation but narrative merger. The merger distinction, drawn from the categorical contrast between a coproduct and a colimit, identifies the phenomenon: it names the difference between components that have merely been placed together and components that have actually been glued into a new object. A second, more predictive distinction, between coordinate type and coordinate constraint strength, explains the phenomenon: it specifies what a colimit's gluing data must look like in a narrative in order for merger to succeed. The essay proposes that an ensemble coheres into a single cultural object only when its constituent character trajectories admit a shared coordinate strong enough to constrain their behavior, and that this structural coherence must then survive two further filters, legibility and recognition, before it becomes observable as durable cultural memory. Contrasting examples are drawn from *Heat*, *Jaws*, *Ocean's Eleven*, *Das Boot*, *12 Angry Men*, and *Alien*. The essay closes by separating compressibility from quality and by extending the framework beyond film to institutional and corporate mergers, scientific theories, and other domains in which collections of parts either do or do not become objects. A brief aside notes a structural parallel with the Lacanian quilting point in recent psychoanalytic film theory, though the two accounts approach the phenomenon from different directions and are not conflated here.

1 The additive model and its failure

The default way of talking about a large cast is arithmetic. Each recognizable performer is treated as carrying a pre-existing audience, and the studio's casting strategy is understood as an attempt to sum these audiences into something larger than any single actor could deliver alone. Chris Hemsworth is assumed to bring one audience, Mark Ruffalo another, Halle Berry a third, and so on down the credits, with the expectation that a sufficiently long and sufficiently famous cast list will produce a correspondingly large total viewership.

This model has some historical justification. In an era of limited distribution channels and comparatively undifferentiated audiences, a single dominant star could function as a near-guarantee of box office performance, and stacking several such figures onto one production plausibly did behave in something close to an additive way. But the conditions that made the additive model locally accurate have eroded. Streaming fragments attention across far more competing objects than a theatrical release calendar ever did. Social media multiplies the channels through which a film can be discussed without multiplying the coordination between those channels. Audiences that once overlapped through mass broadcast now diverge into narrower and more numerous communities, so that the pre-existing audience attached to any individual performer is both harder to estimate and less likely to combine cleanly with any other performer's audience.

The deeper problem, however, is not merely that overlap makes addition inaccurate. It is that addition is the wrong operation regardless of overlap. Even under conditions of perfect audience independence, simply placing several fan bases in proximity to one artifact does not automatically produce a new, larger audience for that artifact. It produces, at most, several audiences that happen to be aware of the same object for different reasons. Whether those audiences actually become one audience, coordinated around a shared reason to attend to the film, is a separate question that the additive model has no resources to ask, let alone answer.

2 The merger model

A more adequate framework treats the assembly of a cast as an attempted merger rather than an attempted sum. The relevant contrast is the one between a coproduct and a colimit. A coproduct is simple concatenation: several objects placed side by side without any required identification between them. A colimit is a coproduct equipped with identifications, a set of maps through which the components are glued along whatever structure they genuinely share. The value of a colimit is concentrated precisely in these identifications. Where the identifications are trivial or absent, the colimit degenerates back into a coproduct, and nothing has actually been merged even if the components have been formally combined.

Corporate and institutional mergers make this distinction easy to see because the failure mode is so familiar. Two profitable companies can merge on paper, combine their balance sheets, and consolidate their organizational charts, while their employees continue to behave as members of the original, separate organizations for years afterward. The merger exists legally without existing operationally. A university can absorb several departments into a single administrative structure without producing any corresponding intellectual

coherence between them. A software company can acquire a handful of promising startups and end up, after integration, with a worse product than any of the components had separately, because the acquisition consolidated ownership without consolidating design.

An ensemble cast is subject to the same distinction. Assembling a long list of recognizable performers is the film equivalent of combining balance sheets: it produces a nominal merger, a single credits sequence and a single marketing campaign, but it does not by itself produce a new center of gravity around which audiences reorganize their attention. Whether that center of gravity forms is an empirical question about the resulting object, not a guaranteed consequence of the assembly.

The merger distinction, however, is diagnostic rather than predictive. It tells us what to look for, a colimit rather than a coproduct, but it does not by itself tell us what makes a given attempted colimit succeed or fail. Two ensembles can both be described as failed or successful mergers without that description explaining why one failed and the other succeeded. Sections V and VI introduce a second distinction, between the type of shared coordinate an ensemble is organized around and how strongly that coordinate constrains its components, which does carry predictive weight. The relationship between the two distinctions is best stated plainly at the outset: the merger distinction identifies the phenomenon, and the coordinate distinction explains it. Everything that follows can be read as an attempt to specify, in structural terms, what the gluing data of a narrative colimit actually consists of.

3 Attraction and coherence

Two variables are frequently conflated in casual discussion of star power, and separating them clarifies a great deal. The first is attraction: the degree to which a cast generates a reason to attend to a film before that film exists as an experience in the viewer's mind. This is a matter of external signals, actors, trailers, source material, director reputation, and prerelease press, and it is naturally measured by observables such as opening weekend box office. The second is coherence: a property of the internal organization of the film once it has actually been consumed, concerning whether the experience resolves into a single, memorable object or remains a loosely associated set of performances.

These are not the same variable, and treating them as interchangeable produces systematic misdiagnosis. Attraction functions as an acquisition mechanism; it gets people into the theater. Coherence functions as a maintenance mechanism; it determines whether the resulting object persists as a stable, nameable thing in memory afterward, and whether it is capable of generating the kind of durable cultural reference that turns a film into a genre touchstone rather than a forgotten release. A film can succeed at attraction and fail at coherence, drawing a large opening audience on the strength of its cast recognition while never becoming, in retrospect, anything more than a list of names that happened to appear together. A film can also fail at attraction while eventually achieving strong coherence, opening quietly and being reassembled into a singular object only later, through critical consensus or repeated viewing.

This pattern is common enough to be recognizable without needing a single anchoring case. A film can assemble an unusually deep bench of critically respected performers, generate positive critical notice, and still underperform commercially relative to its budget,

prompting trade reviews to describe it as delivering less than the sum of its parts. That description is best read not as a claim about individual performances, which may each be strong in isolation, but as a claim about the interaction term between them: the synergy the merger was presumably intended to produce did not fully materialize. Section IV returns to a concrete instance of this pattern once the structural vocabulary needed to diagnose it has been introduced.

4 From audience to object

The natural first move is to locate the site of merger failure in the audience, treating integration as something viewers either do or fail to do as they process a film with many familiar faces. But this move has the causality backward. A colimit does not arise because observers decide, after the fact, to identify several components with one another; the identifications must already exist in the diagram itself, prior to any observer encountering it. If they are absent from the object, no amount of audience goodwill can manufacture them.

This reframes the question. The audience is not the location where integration succeeds or fails; it is, at most, a detector that reveals a structural property already present, or already absent, in the artifact. A large cast is a cardinality property, a simple count of recognizable names. Narrative integration is a structural property, a fact about the dependency relations between character trajectories, and the two properties vary independently of one another.

The claim can be stated more precisely as an inequality relating what a film's structure actually contains to what any given audience manages to perceive of it:

$$\text{Perceived coherence} \leq \text{Actual structural coherence.}$$

Audience integration cannot exceed structural integration. A viewer may fail to notice a coordinate that genuinely exists in the work, through inattention, poor pacing, or a marketing campaign that foregrounds the wrong element, and the resulting shortfall is a legibility failure of the kind discussed in Section VI. But no viewer, and no accumulation of viewers, can recover a coordinate that the work does not contain. There is no operation available to an audience that manufactures structure out of its absence. This inequality is the reason the earlier move, locating merger success or failure in the audience, has the direction of explanation backward: perception is bounded above by structure, never the reverse, and every subsequent distinction in this essay, between legibility and recognition, between compressibility and quality, is really an elaboration of how close a given audience's perception is allowed to come to that upper bound.

Films that are widely regarded as successfully integrated ensembles make the structural requirement visible. *Heat* organizes its entire cast around a single relational opposition between a professional thief and the detective pursuing him; nearly every subplot in the film is intelligible only in relation to that axis, and removing the axis would cause the surrounding structure to decompose rather than merely shrink. *Jaws* organizes its town, its mayor, its scientists, its fishermen, and its families around a single shared threat; the shark functions not as one plot element among several but as the coordinate through which every other element becomes comparable to every other. *Ocean's Eleven* organizes a comparably large ensemble around a single functional task, the execution of one heist, such that each

character’s relevance to the film is legible as an answer to the question of what role they play in that task.

The 2026 film *Crime 101*, an adaptation of Don Winslow’s novella starring Chris Hemsworth, Mark Ruffalo, Halle Berry, and Barry Keoghan, offers a useful contrast case precisely because its intended coordinate is announced in its own title. The film organizes its ensemble around a shared setting, the Highway 101 corridor along which its various criminal and investigative activities occur, and was received with positive reviews but comparatively weak commercial performance relative to its budget[1], prompting at least one trade assessment that the whole amounted to less than the sum of its parts[2]. A shared setting is a weaker gluing condition than a shared relational opposition or a shared functional task, because a setting need not force character trajectories into either convergence or opposition; it can remain a container in which otherwise independent stories merely happen to occur, rather than a relation that binds them together. Nothing in the argument depends on this single case, which serves only as one convenient illustration of a pattern that recurs across many ensemble productions; the diagnostic claim is the inequality above, not the fate of any particular film.

5 Coordinate type and constraint strength

The contrast between weak and strong coordinates introduced above suggests a typology of shared coordinates, distinguishing at minimum the relational, in which characters are defined against one another; the functional, in which characters are defined by role within a shared task; the spatial, in which characters are defined by shared location; the thematic, in which characters are linked by a shared motif; the temporal, in which characters are bound by a shared window of time; and the adversarial, in which characters are organized around a shared external threat.

It would be a mistake, however, to treat this typology as though coordinate type alone determined coherence, with relational and functional coordinates simply outranking spatial and thematic ones. A spatial coordinate can be extremely strong: the submarine in *Das Boot*, the jury room in *12 Angry Men*, and the ship in *Alien* are all, technically, shared locations, yet in each case the location is not decorative but constraining, determining what actions are physically and socially possible for every character at every moment. A thematic coordinate, conversely, can be extremely weak if it exists only as a recurring motif with no consequence for what characters can do or must confront.

The more accurate formulation therefore separates coordinate type from coordinate constraint strength, treating them as two independent dimensions rather than collapsing them into a single ranked list. Constraint strength itself admits a more precise statement than the operational description given above: it is the reduction in admissible trajectories that the coordinate induces in the characters bound to it. Let $\mathcal{T}$ denote the set of all trajectories available to a character in the absence of the coordinate, and let $\mathcal{T}_C \subseteq \mathcal{T}$ denote the subset that remains compatible with coordinate C . Stronger coordinates are those for which $\mathcal{T}_C$ constitutes a smaller fraction of $\mathcal{T}$; a coordinate that leaves $\mathcal{T}_C$ nearly equal to $\mathcal{T}$ is weak regardless of its type, and a coordinate that collapses $\mathcal{T}_C$ to a small subset of $\mathcal{T}$ is strong regardless of its type. A weak coordinate leaves most trajectories available regard-

less of whether a character engages with it; a strong coordinate removes most of them, so that only a small admissible set remains compatible with the coordinate at all. Readers interested in a psychoanalytic account of a structurally related idea, a single element that retroactively organizes an otherwise loose field into a stable whole, may find a point of contact in the Lacanian quilting point as treated in the recent film-theoretic literature[3], though that literature approaches the phenomenon through the viewer's relation to desire rather than through the admissibility structure of the work itself. The submarine hull in *Das Boot* is strong by this measure because it collapses the space of physically and socially available actions for every character aboard it. The heist in *Ocean's Eleven* is strong because nearly every meaningful action available to a character is evaluated relative to its contribution to that single task, leaving little room for trajectories the coordinate does not touch. A shared setting of the kind discussed above is comparatively weak because many character trajectories remain fully admissible without any reference to the setting itself. The relevant predictive relationship is not from coordinate type directly to coherence, but from the pair of coordinate type and constraint strength, understood in this admissibility-reducing sense, to coherence. A spatial coordinate that constrains action as tightly as a submarine hull constrains a crew can outperform a nominally relational or functional coordinate that constrains very little. What matters is not the category the coordinate falls into but the degree to which it actually limits and shapes what the characters bound to it can do, independent of that category.

6 The structure, legibility, and recognition pipeline

Even a genuinely strong coordinate does not automatically become durable cultural memory, and collapsing every failure into a single category of merger failure obscures three distinct points at which the process can break down.

The first point of failure is structural: no sufficiently constraining coordinate exists in the work, so that the components remain a coproduct dressed as a colimit regardless of how the film is edited, marketed, or received. The second point of failure is legibility: a genuinely strong coordinate exists in the work but is difficult for a viewer who watches the film to actually perceive, whether through pacing, editing, or narrative misdirection that buries the coordinate under competing material. A coordinate may exist in the script or in the production history while remaining unavailable to viewers in this sense; for the purposes of this essay, legibility concerns the coordinate as it appears in the viewed work itself, not merely as it existed during production, since a coordinate that never survives into the released cut cannot be perceived regardless of how strongly it constrained earlier drafts of the story. The third point of failure is recognition: a strong coordinate exists and is legible to those who watch the film, yet the film fails to reach a sufficient audience for that coherence to become culturally visible at all, so that the merger is real and perceptible but simply under-observed.

These three stages form a pipeline, structure preceding legibility preceding recognition, and a film can fail at any one of them independently of the other two. A film may possess excellent structural integration and excellent legibility while still disappearing commercially, its coherence never reaching enough viewers to generate durable cultural reference.

Conversely, a film may achieve enormous recognition despite weak underlying structure, becoming widely seen and widely discussed without ever resolving into a single object in the memory of those who saw it. Separating these stages prevents box office analysis from being conflated with narrative analysis, since the two are diagnosing different links in the same chain rather than the same underlying fact.

Of the distinctions introduced in this essay, the pipeline is the one least tied to film specifically, and it is worth pausing on that portability before returning to the narrower case. A scientific theory can possess strong internal structure, a genuine and tightly constraining set of relations among its components, while remaining illegible to most of the field that would need to adopt it, buried in notation or scattered across papers that never state the unifying coordinate plainly. A political movement can achieve wide recognition without ever having settled internal structure, so that its visibility outruns any coherent account of what actually binds its constituent grievances together. A research program, a book, or an online community can fail at any of the same three stages: a book can be tightly integrated and well written yet never reach readers who would recognize its integration; a community can have clear structural norms that are legible to newcomers yet never achieve enough external recognition to grow past a small founding group. Most existing accounts of why an idea, a product, or an institution succeeds or fails stop at the question of structure, asking only whether the thing is well designed. The pipeline adds two further and independent points of failure: a well-structured object can still fail because it cannot be perceived by those who encounter it, and a legible object can still fail because too few people encounter it at all.

7 Compression is not quality

A further discipline is required to keep this framework from collapsing into the assumption that whatever proves most memorable must also be best. Compressibility, the degree to which an ensemble can be reduced to a single sentence or a single organizing coordinate, is a structural and mnemonic property. It is not an aesthetic judgment, and treating it as one would produce a theory in which any film reducible to a slogan is thereby superior to any film that resists such reduction.

A slogan is highly compressible without being profound. A rich and carefully integrated narrative can be considerably harder to summarize than a simpler one, not because it is less coherent but because its coordinate structure is more intricate, involving several interacting constraints rather than one dominant axis. Coherence, in the sense developed here, should be understood as a claim about whether a shared coordinate exists and how strongly it constrains its components, not as a claim about the artistic merit of the result. A highly compressible film can be mediocre, and a richly integrated film can be difficult to compress into a single line while remaining, by any other measure, the stronger achievement.

8 Toward a predictive test

The practical value of separating coordinate type, constraint strength, and the structure, legibility, and recognition pipeline is that it suggests a test that can in principle be applied before a film is released, at the point where only the script and the announced premise are available. Given a proposed ensemble, one can ask what coordinate, if any, the character trajectories are being organized around, what type that coordinate belongs to, and how tightly it constrains the range of actions available to each character bound to it. This yields a structural forecast distinct from any prediction about box office, since box office is downstream of legibility and recognition as well as structure, and is subject to confounds, release timing, marketing spend, and competition, that have nothing to do with narrative integration.

What such an analysis could forecast is not commercial performance but the likelihood of long-term cultural coherence, the probability that a given ensemble becomes describable, years later, as a single named object rather than as a list of performers who happened to share a production. Testing this would require assembling a comparative sample of known cases, films generally regarded as successful mergers such as *Heat*, *Jaws*, *Ocean's Eleven*, *Das Boot*, and *12 Angry Men*, alongside cases of ambiguous or contested merger, and scoring each for coordinate type and constraint strength independent of any knowledge of its reception, before checking that score against the pipeline outcomes actually observed.

9 Generalization beyond film

Nothing in this framework is specific to cinema. A university merger can produce a larger institution with less intellectual coherence precisely when no shared coordinate, no common research question, pedagogical mission, or administrative constraint, binds the merged departments together more tightly than their prior separate structures did. A software acquisition can leave a company with a worse product when the acquired components are integrated administratively without being integrated around any shared functional coordinate in the product itself. A social network can add users while becoming less engaging if growth adds coproduct volume without adding colimit structure, more participants without any coordinate strong enough to bind their activity into a single coherent experience.

In each domain the same diagnostic applies. Aggregation is not integration. A merger, whether of actors, departments, companies, or users, succeeds only when the merged object becomes easier to perceive than its components, and fails whenever those components remain more salient, individually, than whatever the merger was meant to produce. The central claim of this essay can accordingly be stated in a single sentence: ensembles cohere not when many components are present, but when those components can be projected onto a sufficiently strong shared coordinate that becomes easier to perceive than the components themselves, and that coherence must then survive independent filters of legibility and recognition before it is observable as durable cultural memory.

References

- [1] *Crime 101* (2026 film). Wikipedia. Retrieved July 2026.
- [2] “Crime 101 Review: Chris Hemsworth, Halle Berry and Mark Ruffalo Headline a Lavishly Crafted Thriller That’s Less Than the Sum of Its Polished Parts.” *The Hollywood Reporter*, February 2026.
- [3] Helen Rollins. *Psychocinema. Theory Redux*. Cambridge: Polity, 2024.