SECURITY & DEFENCE QUARTERLY

DRVF: A filtered-ring valuation framework for conservative multi-source intelligence fusion

M.H. Anjom SHoa

mh.anjomshoa@vru.ac.ir

https://orcid.org/0000-0002-9735-8951

Department of Mathematics, Vali-e-Asr University of Rafsanjan, Emam Khomeini Sq, Emam Khomeini Bl, 77718897111, Rafsanjan, Iran

Abstract

Multi-source intelligence fusion under adversarial conditions demands formal control of combined confidence with degradation of source reliability. However, no existing framework simultaneously enforces a conservative confidence floor, preserves source-tier distinctions as algebraically separable compartments, and maintains a reconstructable record of assessment evolution. Here we establish the Dubrovin Ring-Valuation Framework (DRVF), a formally verified structure for conservative multi-source fusion. The DRVF encodes intelligence propositions within a filtered non-commutative ring, represented as a 3×3 matrix over source-reliability tiers, coupled with an aggregation rule bounding combined confidence at the least reliable active source. Algebraic proofs establish three previously unresolved properties simultaneously: a conservative confidence floor, source-tier compartmentalisation, and auditable assessment trajectories. A worked cyber-disinformation scenario demonstrates the framework's operational behaviour, and an exploratory pilot (N = 23) informed protocol refinements ahead of a planned confirmatory evaluation. The DRVF resolves structural gap in the existing fusion methods, providing a new formally verified framework and opening the door to empirical validation in operational intelligence settings.

Keywords:

intelligence fusion, hybrid crises, filtered ring, quasi-valuation, conservative inference, structured analytical technique, source compartmentalisation, confidence tracking, multi-source analysis

Article info
Received: 29 June 2026
Revised: 17 September 2026
Accepted: 23 September 2026
Available online: 12 November 2026

Citation: Anjom SHoa, M.H. (2026) 'DRVF: A filtered-ring valuation framework for conservative multi-source intelligence fusion', *Security and Defence Quarterly*, 56(4). doi: 10.35467/sdq/240155.

© 2026 M.H. Anjom SHoa. Published by War Studies University, Poland.
This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0/).

http://doi.org/10.35467/sdq/240155

Introduction

How can analysts formally track and preserve confidence assessments as source reliability changes over the course of a multi-source assessment? This paper addresses a specific structural problem that arises when intelligence from multiple collection disciplines must be combined under conditions where the reliability of individual sources is not stable and the evidence environment may be subjected to deliberate manipulation. The existing structured techniques do not encode the tracking of confidence evolution as a formal property of analytical method. The Dubrovin Ring-Valuation Framework (DRVF) introduced in this paper is designed to encode intelligence propositions within a filtered non-commutative ring.

There is some confusion over the usage of "hybrid" in the intelligence and security studies literature. Others, such as Bousquet (2009); Caliskan (2021); and Echevarria (2016) have suggested that the concept has become a rhetorical device that can mask what is essentially a continuation of normal campaign planning, and has been used in ways that can obscure its true meaning. The substantive criticisms are recognised here: In this paper, the term "hybrid" is not meant to imply that the coerce is a given category of analysis, but rather a shorthand for the multiple operational conditions listed above that coexist in what are often called multiple domains: (a) the evidentiary environment is manipulated; (b) source reliability is not equal across collection streams; and (c) the analyst can't confidently determine the compromised collection streams. Its usage means the reference of these conditions, not of a supposed "new kind" of warfare. The three structural conditions are critical to the mathematical framework and may exist in contexts that are not typically considered "hybrid."

Two declassified historical episodes illustrate the structural dimension of this problem. The 9/11 Commission (US Government Printing Office, 2004) documented a systemic failure to integrate Human Intelligence (HUMINT) and Signals Intelligence (SIGINT) indicators prior to September 2001. The Commission noted that information was held in organisational silos, there was no mechanism for producing a conservative composite confidence estimate, and analysts could not trace how their assessments changed as new information arrived. The structural point relevant here is specific: there was no formal mechanism that could both fuse evidence across streams and preserve the evidentiary trajectory of that fusion. The second episode is the assessment of Iraq's weapons of mass destruction (WMD) analysed by Jervis (2010). Low-quality HUMINT was combined with ambiguous SIGINT fragments without formal limits on how the combination could affect aggregate confidence. The result was an over-confident assessment. Jervis (2010) identifies multiple contributing factors, including political pressure and groupthink; here, the structural point of interest is the absence of a mechanism that imposed a constraint on combined confidence based on the reliability of the weakest substantive source.

The cyber-disinformation scenario used as a worked example in this paper is drawn from the hybrid-crisis domain because it exhibits all three conditions in a concentrated form: multiple collection disciplines operating simultaneously, at least one stream (Open-Source Intelligence [OSINT]) known to be under active adversarial manipulation, and source reliability that shifts as the operation unfolds. Hybrid crises are not the only context in which these conditions arise, but they are the context in which the consequences of fusion failures that produce over-confident assessments based on manufactured corroboration are most severe are most severe and most thoroughly documented in the declassified record.

All-source intelligence analysts in major agencies perform multi-source fusion as a routine matter, supported by organisational protocols, professional standards, and institutional

review processes. The gap that DRVF addresses is not the absence of operational solutions but the absence of a formally specified structure that enforces certain protective constraints as structural properties rather than as procedural guidelines. The three conditions that define the problem class are not exceptional; they describe the standard operating environment for all-source intelligence analysis in most major agencies. Intelligence analysis, as Herman (1996), Kent (1949), and others have argued, is inherently conducted under conditions of incomplete, contested, and often deliberately concealed information. What distinguishes the specific problem class addressed here is not the mere presence of an adversary but the potential for deliberate manipulation of the evidence streams themselves: compromised collection channels, coordinated disinformation targeting the analyst's OSINT inputs, or strategically leaked deceptive material designed to inflate confidence in a false hypothesis.

Some of this issue is addressed by the existing structured frameworks. Analysis of competing hypotheses (ACH) is a snapshot tool that enhances hypothesis generation and diminishes some types of confirmation bias, but it does not codify the evolution of confidence as a source's reliability decreases during the analysis (Heuer and Pherson, 2019). While Bayesian networks are attractive because they allow for a rigorous modelling of conditional dependence, they are based on the prior and on the dependency structure, both of which are very sensitive to manipulation of the evidence environment (Pearl, 1988). The Dempster-Shafer Theory (DST), which is considered epistemic uncertainty under incomplete evidence, has a typical form of the combination rule that if the evidence sources conflict, then epistemic belief can be inflated, so that the combined belief is higher than what it is warranted by the smallest source evidence, given a conservative norm (Dempster, 1967; Shafer, 1976). None of these frameworks has all three analytical properties: a conservative confidence floor, source-tier compartmentalisation, and traceable assessment evolution.

The DRVF integrates a filtered-ring algebraic structure and a conservative assessment-tier rule as its two core layers. The framework is domain-specific: it is designed for contexts where source reliability is unstable, adversarial manipulation of evidence is possible, and an auditable analytical record is required. It is not proposed as a universal replacement for the existing techniques. The paper proceeds as follows. Section 2 defines Intelligence fusion, as the paper uses this term. Section 3 presents the Scoping Review and gap analysis. Section 4 develops a formal Theoretical Framework, including proofs. Section 5 provides a Worked Example. Section 6 presents Discussion and discusses implications, counterarguments, and pilot observations. Section 7 contains Conclusions.

Intelligence fusion: Definition and scope

Intelligence fusion, as used in this paper, refers to the process of integrating information from multiple collection disciplines (e.g., HUMINT, SIGINT, and OSINT) into a single composite assessment in which the contributing sources, their respective reliability levels, and the evidentiary basis for the combined judgement are preserved and inspected. This definition emphasises three elements that the existing formalisations address only partially: the preservation of source identity through the fusion process, the explicit representation of differential source reliability, and the traceability of how the combined assessment evolved as new evidence arrived. The term is distinct from data fusion in engineering contexts where the emphasis is on algorithmic optimisation rather than on epistemic caution under adversarial conditions (Hall and Llinas, 2001; Liggins *et al.*, 2009).

The concept of intelligence fusion is operationalised differently across agencies and national traditions. In some contexts, it refers primarily to the organisational process of sharing and combining reports from different collection disciplines. In others, it denotes specific methodological approaches, such as Bayesian fusion or Dempster-Shafer combination. The diversity of conceptualisations is documented by Betts (2002) as well as, more recently, by the intelligence-methodology literature (Rieber, 2004; Chang, *et al.*, 2018). The definition adopted here is deliberately restricted: it concerns the formal properties of the combination mechanism itself, not the organisational or institutional arrangements within which fusion occurs. This restriction is necessary because the DRVF operates at the level of the combination structure, not at the level of the organisational process.

Scoping review and theoretical positioning

search strategy and scope

This review was conducted as a scoping review based on the Preferred Reporting Items for Systematic Reviews and Meta-Analyses extension for Scoping Reviews (PRISMA-ScR) guidelines (Tricco *et al.*, 2018), which are suitable when a scoping review aims to map out the conceptual landscape rather than to estimate a pooled effect. The four areas targeted were: (1) intelligence fusion and multi-source fusion; (2) formal uncertainty representation; (3) confidence calibration or state space tracking in analytical reasoning; and (4) structured analytical techniques relevant to multi-source assessment. The databases searched were Google Scholar, IEEE Xplore, ScienceDirect, SpringerLink, and Taylor & Francis Online. Search terms covered various combinations, such as "intelligence fusion," "multi-source fusion," "structured analytical technique," "Dempster-Shafer," "Bayesian fusion," "confidence aggregation," "weakest-link," and "adversarial uncertainty." The first search resulted in 312 hits. In all, eighty-nine papers were chosen for full-text review after screening of titles and abstracts. Following the full-text screening, forty-seven documents were retrieved for synthesis, which included references to multi-source uncertainty, confidence aggregation, or state space tracking in any tradecraft or security context. The research was not considered if it was limited to an environment where sensor fusion was used in non-adversarial engineering context. The PRISMA-ScR flowchart is included in Appendix A.7.

The review does not support absolute novelty claims. A scoping review maps the existing territory; it cannot exhaustively prove the absence of prior work. The claim this review supports is a relative contribution claim: within the reviewed corpus, no framework was found that simultaneously and explicitly integrates conservative weakest-link aggregation, hierarchical source compartmentalisation, and traceable assessment evolution in a formally specified intelligence-fusion context. The DRVF extends the existing literature by combining these three features in one formally specified, testable structure.

Structured analytical techniques

Analysis of competing hypotheses is widely cited in the methodological literature and familiar to many Western-trained analysts, although its operational prevalence varies across organisations and is difficult to establish from publicly available evidence (Heuer and Pherson, 2019). Software-assisted implementations of ACH now allow analysts to adjust source-validity weights interactively, which addresses some of the gap DRVF targets at the procedural level. The limitation relevant here is that ACH does not encode source-tier distinctions or confidence trajectories as algebraic properties of its structure. Mandel and Barnes (2014) found that structured methods improve forecast accuracy under controlled conditions, but that benefits are domain-contingent. This finding supports both

the value of formal structure and the need to test structure against performance. A recurring limitation across the Scholastic Assessment Test (SAT) literature is the absence of a formal mechanism for tracking analytical state transitions. When source reliability shifts mid-analysis, ACH provides no structural constraint on how that shift should propagate through the combined picture. Recent works by Rieber (2004) on confidence-calibration methods and Chang, *et al.* (2018) on structured techniques in operational settings have advanced the field, but the specific gap that DRVF addresses the formal tracking of confidence trajectories under degrading source reliability remains.

Probabilistic frameworks

Bayesian networks are powerful when priors are defended and dependency structure is stable (Hall and Llinas, 2001; Liggins *et al.*, 2009; Pearl, 1988). Their application to multi-source intelligence fusion is well documented. Adversarial Bayesian formulations incorporate strategic manipulation of evidence inputs, and robust prior specifications reduce sensitivity to prior misspecification (Friedman and Zeckhauser, 2012; Muñoz-González and Lupu, 2017). These extensions represent genuine progress. Their limitation for the present purpose is not adversarial fragility *per se* but the combination of requirements: even adversarially robust Bayesian formulations do not simultaneously enforce a strict weakest-link confidence floor and preserve source-tier distinctions as algebraically separable compartments within one formally specified structure.

Dempster-shafer theory

The DST represents epistemic uncertainty under incomplete evidence and is extensively applied in information fusion contexts (Castanedo, 2013; Dempster, 1967; Sentz and Ferson, 2002; Shafer, 1976). The tension with the present paper's requirements arises from Dempster's combination rule, which normalises conflict by redistributing mass away from the empty set, producing combined confidence estimates that exceed what the weakest credible source would justify. Both Murphy's (2000) averaging rule and Smets' (1990) conjunctive combination rule address conflict management without Dempster's normalisation step, and Yager's (1987) modification assigns conflicting mass to a global ignorance term. These are substantive improvements. Their limitation is the same as for adversarial Bayesian extensions: none enforces a strict source-tier floor as an algebraic property of the combination structure itself.

Dynamic epistemic logic (DEL)

Dynamic epistemic logic provides formal tools for representing how belief states change after public or private updates (van Ditmarsch *et al.*, 2007). DEL is conceptually proximate to the trajectory requirement in DRVF: it tracks how beliefs evolve, not merely what they are. Its limitation for operational deployment is practical. DEL models are highly expressive and formally rigorous, but adapting them to multi-source intelligence analyst workflows at the speed required by operational environments is not demonstrated.

Summary: The relative contribution

Four conclusions emerge from the reviewed literature. First, structured analytical techniques remain operationally valuable but do not formalise confidence trajectory under

degrading source reliability. Second, probabilistic frameworks provide mathematical rigour but rest on assumptions fragile in adversarial environments, and their adversarial extensions do not jointly enforce all three analytical requirements. Third, DST supports incomplete evidence but does not enforce a strict weakest-link floor as an algebraic property, even in its modified forms. Fourth, dynamic epistemic approaches represent state transitions formally but have not been operationalised for multi-source intelligence fusion. From these conclusions, the relative contribution of this paper can be stated: The DRVF extends the existing literature by integrating, in a single formally specified framework, (a) conservative weakest-link aggregation, (b) hierarchical source compartmentalisation, and (c) traceable assessment evolution. No framework combining all three features in an intelligence-fusion context was identified in the reviewed corpus. The synergy of the combination matters because each property, taken alone, is achievable with simpler structures; it is their simultaneous algebraic enforcement within one framework that constitutes the contribution. These distinctions are summarised in Table 1.

Theoretical framework

Operational motivation and problem class

The problem class addressed by DRVF is a restricted one: multi-source intelligence fusion with (a) unstable or differentially degrading sources, (b) potential for adversarial

Table 1. Comparison of intelligence fusion frameworks on key analytical dimensions. Note: Ratings are based on the scoping review and represent structured characterisations to aid comparison, not empirical judgements. The DRVF entries reflect the proposed framework only; no empirical performance data is available at this stage.

Framework	Weakest-link floor	Source compartmentalisation	Trajectory tracking	Adversarial robustness	Basis
DRVF (proposed)	Explicit via assessment tier T	Explicit (3-tier matrix ring)	Explicit (update log)	Designed for adversarial conditions; unconfirmed empirically	Formal specification (this paper)
ACH (Heuer and Pherson, 2019)	Not formalised	Not encoded	Snapshot only	Reduces confirmation bias; no manipulation constraint	Heuer and Pherson, 2019; Mandel and Barnes, 2014
Bayesian Networks (Pearl, 1988)	Only if priors encode it	Can be modeled; not inherent	Possible with temporal extension	Fragile under standard priors; partial mitigation via robust extensions	Hall and Llinas, 2001; Pearl, 1988
Dempster-Shafer (Shafer, 1976)	Partial; standard rule inflates	Not inherent	Not inherent	Sensitive to conflict inflation; mitigated by modified rules	Dempster, 1967; Sentz and Ferson, 2002
Dynamic epistemic logic (van Ditmarsch *et al.*, 2007)	Not a core focus	Can represent; not operationalised	Explicit in update logic	High formal expressiveness; undeployed operationally	van Ditmarsch *et al.*, 2007

manipulation of one or more evidence streams, and (c) need for an auditable record of assessment evolution. The three conditions that define this problem class are not exceptional; they describe the standard operating environment for all-source intelligence analysis in most major agencies. The lack of something is not an operational practice but a formal mechanism that makes some protection constraints structural properties and not procedural directions. Three analytical needs emerge from this problem class: (1) do not let weak sources be overwhelmed by strong sources (conservative aggregation); (2) do not compress sources into a single number (compartmentalisation); and (3) document the analytical state over time to allow for inspection of assessment evolution (trajectory tracking). The DRVF meets requirements (1) and (2) formally, and requirement (3) is met by the structured update protocol outlined in Section 5. The correspondence is not a necessary consequence but a deliberate option: other options are conceivable and might be just as successful with different assumptions.

Consider the following operational scenario, which serves as the worked example throughout this section. An analyst is assessing whether an adversary is preparing a coordinated cyber-disinformation campaign, a characteristic hybrid-crisis pattern in which cyber operations and information warfare converge. SIGINT provides technical indicators of preparatory network activity at high reliability (tier 1). A HUMINT source partially corroborates operational intent but is only partially validated (tier 2). OSINT provides abundant circumstantial evidence, but the domain is actively subjected to adversarial manipulation (tier 3). This pattern concurrent multi-domain operations, with at least one stream under active manipulation is typical of the hybrid-crisis environments that motivated the framework's development. A standard averaging rule would allow a strong SIGINT signal to lift aggregate confidence above what the HUMINT and OSINT evidence justifies. DRVF is constructed to prevent that outcome.

What DRVF is and is not

The DRVF is a proposed method for combining intelligence from different sources that deliberately prevents combined confidence from exceeding what the weakest important piece of evidence allows. It also maintains information from different sources in separate inspectable compartments so that analysts can always see which evidence came from where and how the overall assessment changed as new information arrived. The DRVF is a formally specified framework for conservative multi-source intelligence fusion. It combines a filtered-ring algebraic structure, preserving source-tier distinctions as algebraically separable entries, with an assessment-tier aggregation rule, which enforces the weakest-link confidence floor. It is proposed as a candidate analytical tool, not as a universal replacement for the existing methods.

The term "formal" means that the framework's properties, the conservative floor, the source isolation conditions, and the update rules are expressed as mathematical propositions with proofs, not as procedural guidelines. An analyst following the framework cannot inadvertently violate the conservative floor because it is enforced by the algebraic structure of the assessment tier. The term "conservative" means that the combined confidence estimate is bounded by the least reliable active substantive source entry. "Framework" denotes a structured set of integrated components, the filtered-ring representation, the assessment-tier rule, the source-mapping rubric, and the isolation criterion designed to work together as a coherent system for a defined analytical purpose. DRVF is not a decision-support system, a software tool, or an organisational process. It is a mathematical specification of how evidence combination should behave under stated conditions.

http://doi.org/10.35467/sdq/240155

Filtered-ring representation

Let R be a non-zero, non-commutative ring equipped with a descending filtration $\{R_k\} \geq 0$ such that $R_0 = R$, $R_{K+1} \subseteq R_K$, and $R_k R_m \subseteq R_{k+m}$. The filtration is interpreted as a reliability ordering: lower-index elements represent higher-reliability propositions; higher-index elements represent weaker support. This is a modelling convention, not a theorem about epistemic states.

The filtered-ring structure is chosen over simpler alternatives for two specific reasons. First, the compatibility condition $\vartheta(R_k) \geq k$ aligns algebraic operations with reliability tiers by construction: any product of ring elements automatically respects the reliability grading without requiring the analyst to recheck tier constraints at each step. An ordered set or vector representation does not provide this automatic propagation. Second, the matrix ring $M_3(R)$ preserves internal distinctions between source streams, including corroborative and contradictory cross-stream relations, in a way that cannot be done by a scalar or vector representation. These two properties together, automatic tier-graded propagation and source-pair distinction, justify the algebraic overhead relative to simpler formalisms (Anjom Shoa, 2016).

Intelligence propositions are encoded as elements of the matrix ring $M_3(R)$. The three dimensions correspond to three principal intelligence streams. In the worked example, these are HUMINT, SIGINT, and OSINT. This is a motivated modelling choice tied to the specific operational context; it is not a universal claim about source taxonomy. A primary source class, as used here, refers to a collection discipline whose reliability characteristics may degrade independently of other disciplines. SIGINT, Communications Intelligence (COMINT), and Measurement and Signature Intelligence (MASINT) may share collection infrastructure, and a compromise affecting that infrastructure could degrade all three simultaneously. The three-class assumption is a modelling simplification: it treats the intelligence problem as involving three principal reliability tracks. If a specific operational context requires finer-grained source distinctions, the matrix dimension should be revised accordingly (see Section 6.6 Boundary Conditions). Diagonal entries a_{ii} represent source-specific assessments for stream i. Off-diagonal entries a_{ij} represent corroborative or contradictory relations between streams i and j. A zero off-diagonal entry means no cross-validation has been established.

Source-to-matrix mapping rubric

A formal framework that cannot be instantiated from operational data is a mathematical exercise, not a usable tool. The following rubric specifies how an analyst maps raw evidence to filtration levels. The rubric has three components: source reliability grading, corroboration weighting, and matrix entry assignment. Diagonal entries receive the filtration tier assigned to the corresponding primary source stream under the reliability grading above. Off-diagonal entries a_{ij} (representing cross-stream corroboration or contradiction) are assigned the filtration tier of the weaker of the two contributing pieces of evidence, consistent with the weakest-link principle for corroboration. A zero entry indicates that no cross-validation relationship has yet been established. All matrix assignments, together with the evidentiary basis for each tier, are recorded in the analytical log.

Source reliability grading. Tier 1 (filtration level 1): Source has documented operational reliability across prior assessments; intelligence has been independently corroborated by at least one source from a different collection stream; collection method is robust to adversarial manipulation under current conditions.

Tier 2 (filtration level 2): Source has partial reliability documentation; intelligence has been partially corroborated or is consistent with independent indicators; collection method may be vulnerable to limited adversarial manipulation.

Tier 3 (filtration level 3): Source has limited or unknown reliability; information from source is not corroborated by other streams or contradicted; collection method is known to be subjected to active adversarial manipulation.

Tier 4 (filtration level 4): Source has been discredited partially, collection channel has been compromised, or intelligence is inconsistent with independent indicators. Tier 4 movement is explicit analyst judgement and documented. In between cases, the case is rounded up using a conservative approach: if a source falls into two tiers, the higher filtration index (weaker reliability) is assigned. This rubric is planned to be assessed for inter-rater reliability in the confirmatory evaluation study and should be considered as a proposed instrument that needs to be validated prior to its confirmatory use.

Corroboration weighting. An off-diagonal entry between streams i and j is assigned at the filtration level of the weaker corroborating piece of evidence, not the average. If both SIGINT (tier 1) and HUMINT (tier 2) point to the same indicator, the corroboration entry is assigned tier 2. This is the rubric-level instantiation of the weakest-link principle for cross-stream corroboration.

Quasi-valuation, assessment tier, and the conservative floor

Define a quasi-valuation ϑ: $R \to \mathbb{N} \cup \{\infty\}$ satisfying: $\vartheta(0) = \infty$; $\vartheta(ab) \geq \vartheta(a) + \vartheta(b)$; and $\vartheta(a + b) \geq \min\{\vartheta(a),\vartheta(b)\}$. For a non-zero element a, $\vartheta(a)$ provides the algebraic infrastructure of the framework. For $A = (a_{ij}) \in M_3(R)$, define $\vartheta(A) = \min_{ij} \vartheta(a_{ij})$, the quasi-valuation provides the algebraic infrastructure of the framework. The compatibility condition $\vartheta(R_k) \geq k$ ensures that algebraic operations on ring elements are automatically constrained by their reliability tiers: any product of elements from R_k and R_k is guaranteed to lie in R_{k+m}, which means the tier information propagates correctly through addition and multiplication without requiring manual verification. The addition and multiplication properties of the quasi-valuation (proved in Appendix A.2) guarantee that the tier structure is preserved under evidence combination. These properties are the mathematical foundation that makes the framework's structural constraints enforceable.

The weakest-link norm itself is implemented by a separate but related function. Define the filtration tier of a non-zero ring element a as $\tau(a) = k$ if a $\in R_k \, R_{k+1}$; define $\tau(0)$ as undefined (zero entries carry no tier information). For a matrix $A \in M_3(R)$, define the assessment tier as:

$$T(A) = \max\{\tau(a_{ij}): a_{ij} \neq 0\}$$

The maximum is taken over all non-zero, tier-assigned entries in the active matrix. This function implements the weakest-link aggregation norm directly: the highest (weakest) filtration tier among all active evidence entries determines the overall assessment tier. If any active entry sits at tier 3, the overall assessment cannot exceed tier 3, regardless of how many tier-1 entries are present. The assessment tier T(A) is bounded from below by the compatibility condition: since $\tau(a) \leq \vartheta(a)$ for any tier-assigned element a, we have $T(A) \leq \max\vartheta(a_{ij})$ for all tier-assigned active entries. The quasi-valuation provides an algebraic

upper bound on the assessment tier, and in the special case where $\tau(a_{ij}) = \vartheta(a_{ij})$ for all entries, the bound is tight.

The relationship between ϑ and T is worth stating precisely. The quasi-valuation ϑ operates at the level of individual ring elements and enforces the compatibility condition that makes the algebraic structure work. The assessment tier T operates at the level of the matrix and enforces the weakest-link norm that makes the framework conservative. Both are necessary: ϑ without T would provide algebraic infrastructure without a conservative aggregation rule; T without ϑ would provide a conservative rule without the algebraic guarantees that ensure tier information propagates correctly through operations on the matrix.

This conservative norm appears in intelligence doctrine under various formulations. The formal contribution is not the weakest-link norm itself (which experienced practitioners already recognise), but its formalisation in a structure that permits algebraic tracking, mathematical comparison across assessments, and formal testing of properties. A structure that enforces the norm by construction is harder to violate inadvertently than one that relies on analyst discipline alone.

Epistemological justification of the weakest-link norm

A substantive objection to the weakest-link approach must be addressed directly. Converging independent sources can and do justify higher aggregate confidence than any single source alone. This is the logic underlying corroboration-based inference, and it is epistemologically sound under conditions of genuine source independence. If all five independent witnesses of varying reliability report observing the same event, it is reasonable to be more confident than the least reliable witness alone would justify. The reviewer who raised this point is correct as a matter of general epistemology.

The weakest-link norm in DRVF is not proposed as a universal epistemological principle. It is proposed as a domain-specific procedural constraint appropriate to a particular failure mode: the risk that coordinated deception or undetected source compromise inflates aggregate confidence by providing apparent corroboration that is in fact manufactured. In such environments, the independence assumption that justifies corroboration-based confidence-building does not hold. The five witnesses may be reporting the same story because the same deception operation has reached all of them. The conservative floor is a safeguard against this specific failure mode, not a general claim about how evidence should always be combined.

This distinction has a practical implication. In environments where adversarial manipulation is not a concern, where sources can be verified as genuinely independent, or where the cost of conservatism (missed signals and delayed judgements) outweighs the cost of overconfidence, other aggregation philosophies may be more appropriate. Bayesian combination, for example, is well suited to such environments. The DRVF is designed for the boundary case where the analyst has a reason to believe that one or more evidence streams may be compromised and cannot reliably determine which ones. This is not an unusual condition in intelligence analysis, but it is not a universal one either. The boundary conditions are stated in Section 6.6.

A substantive question raised in review deserves a direct response here: can the addition of weak but non-contradictory evidence rationally reduce the confidence warranted by an otherwise unchanged body of strong evidence? The answer that DRVF gives is conditional, and the condition turns on a distinction that the framework already encodes

operationally but should now state explicitly. Weak evidence that fails a validity screen is dropped, not retained; weak evidence that passes the screen is retained at its assigned tier and bounds the aggregate. The conservative floor reflects residual deception risk that cannot be eliminated by scrutiny alone.

The two-stage gate operates as follows.

Stage 1 (validity screen): Before any piece of evidence enters the matrix, it is subjected to the standard validation procedures that experienced analysts already apply source triangulation, chain-of-custody verification, plausibility assessment against independent indicators, and explicit documentation of any doubts. Evidence that fails this screen is not entered into the matrix at all; it is logged in a separate excluded-evidence record with the reason for exclusion. This is the operational realisation of the reviewer's instinct that weak or suspect evidence should be subjected to scrutiny and dropped if invalid. DRVF does not override this step; it presupposes it.

Stage 2 (tier assignment and floor enforcement): Evidence that passes the validity screen is assigned a filtration tier under a rubric in Section 4.4 and entered into the matrix. Once entered, the assessment-tier rule $T(A) = \max\{\tau(a_{ij})\}$ applies. The conservative floor at this stage is not a second validity judgement; it is a residual-risk bound. The screened-in weak evidence may still be reporting the same fabricated story that strong sources are reporting, because the screen cannot prove independence. The floor reflects the analyst's inability to verify independence, not a verdict that the weak evidence is itself invalid. This is the substantive difference between "this evidence is suspect, drop it" (a Stage 1 outcome) and "this evidence is acceptable but of low reliability, let it bound the aggregate" (a Stage 2 outcome). The matrix structure of DRVF keeps these two judgements distinct, while a scalar weakest-link rule would conflate them.

This reformulation addresses the reviewer's concern directly. The framework already provides for dropping invalid evidence: the rubric's tier-4 designation and the source-isolation operation in Section 4.7 implement exactly this. What the framework does not do, and should not do, is allow weak-but-valid evidence to be silently overridden by stronger sources. The risk that weak evidence corresponds to a coordinated deception channel cannot be ruled out by scrutiny; it can only be bounded. The floor is that bound.

Formal source isolation

A substantive objection to the framework's design must be addressed directly. In real intelligence analysis, if one source is compromised, experienced analysts do not collapse the entire assessment to the lowest available tier; they isolate the weak source and rely on stronger streams. A naive implementation of the weakest-link norm would preclude this, forcing analytical paralysis when any source is weak.

This objection has force against a scalar weakest-link rule. It does not apply to the matrix structure of DRVF. The matrix preserves source-tier information at the entry level. An analyst who identifies that the OSINT stream is compromised can exclude the relevant OSINT diagonal entry from the active matrix state and recalculate T over the remaining entries. This is a compartmentalised isolation operation, not a collapse.

Source isolation is not left to unconstrained analyst judgement. The following formal criterion governs when isolation is warranted: a source entry a_{ij} may be isolated if and only if (a) the entry has been assigned tier 4 under the rubric in Section 4.4, indicating partial

discrediting or channel compromise, or (b) an explicit, documented evidentiary trigger, such as adversary-confirmed channel penetration or source recantation, has been recorded in the analytical log. Isolation on the basis of analyst's personal preference or solely to raise the aggregate confidence floor, without a documented triggering condition, is impermissible within the framework. The algebraic treatment of source isolation, including the proof that $T(A(-i)) \leq T(A)$, is provided in Appendix A.4.

Update protocol and assessment tracking

The update protocol is structured in such a way that it encodes the changes in the matrix state as new evidence comes in. The analyst will then do the following: (1) determine the target matrix entry associated with the relevant source stream; (2) apply a filtration level to that entry using the rubric in Section 4.4; (3) update the target matrix entry and any other matrix entries that would be impacted if the source stream were to be upgraded or downgraded; (4) recalculate the assessment tier T for all the entries that are assigned to it; and (5) record the type of evidence used (confirming the current tier, source stream upgraded, source stream downgraded, corroboration established, corroboration withdrawn, or assessment held pending), and the matrix state that results from the update in the analytical log. The log is the traceable history of the development of assessment. From the initial state of the matrix A_0 and the log, it's possible to reconstruct the matrix state at any step by applying the sequence of operations that were recorded.

An optional visualisation of the update path was explored in earlier versions of this work. That visualisation used a hexagonal adjacency structure to encode six transition types as directions on a grid. Testing whether this visualisation reduces cognitive load or improves analyst performance is an empirical question. The visualisation has been moved to Appendix A.5 and is presented there as a provisional design, not as a required component of the framework.

Worked example

The following worked example traces the matrix state of DRVF through five successive evidence updates in the cyber-disinformation scenario introduced in Section 4.1. The example is schematic: reliability assignments follow the rubric in Section 4.4, but the scenario is constructed for illustrative clarity and does not replicate an actual operational case. At each step, both the algebraic operation and its operational meaning are described. The structural behaviour illustrated (conservative floor maintenance and source isolation without collapse) is directly analogous to the structural failure modes documented in the two declassified episodes cited in Section 1: the pre-September 2001 integration failure and the Iraq WMD assessment. In both cases, the absence of a formal floor constraint and a compartmentalised source structure allowed weak-evidence contributions to elevate aggregate confidence. The successive matrix states are presented in Table 2 and show how DRVF would formally prevent that outcome by construction.

Two properties of the framework are illustrated by this example. First, the assessment-tier floor is maintained throughout in accordance with the weakest-link norm: the weakest active source always determines the overall assessment tier. When OSINT degrades from tier 3 to tier 4 (step S_2), In the framework's convention, a higher numerical tier indicates lower reliability / lower confidence. Therefore the assessment tier correctly moves from 3 to 4 when the weakest active source is downgraded. When the compromised OSINT stream is formally isolated (step S_3), the floor improves from 4 to 2 because the weakest

Table 2. DRVF matrix state evolution: Cyber-disinformation assessment scenario. Note: All entries are schematic. Tier assignments follow the rubric in Section 4.4. T(A) = max over all active, non-zero, tier-assigned matrix entries. Lower T indicates higher confidence. Source isolation in S_3 is triggered by the formal criterion in Section 4.7. The update path $S_0 \rightarrow S_1 \rightarrow S_2 \rightarrow S_3 \rightarrow S_4$ constitutes the auditable record.

Step	Evidence/action	Matrix state	T(A)	Operational interpretation
S_0: Initial	SIGINT: Tech indicators (tier 1); HUMINT: partial validation (tier 2); OSINT: abundant but manipulated domain (tier 3)	diag(1, 2, 3); off-diagonal = 0 (no cross-validation).	3	Floor anchored by OSINT at tier 3 (the weakest active source). The max over active tier-assigned entries is 3. No cross-stream corroboration established.
S_1: SIGINT-HUMINT corroboration	Second SIGINT fragment confirms HUMINT indicator; corroboration assigned tier 2 (weaker stream)	a_{12} = tier 2 added. Diagonal unchanged.	3	New off-diagonal records that two streams agree. Corroboration is logged at the weaker stream's tier (HUMINT, tier 2). Floor unchanged: T remains 3 because OSINT diagonal at tier 3 still dominates the max.
S_2: OSINT partial discredit	Analyst judgement: Specific OSINT channel compromised; a_{33} moves tier 3 → tier 4	a_{33} now at tier 4. Off-diagonal OSINT entries adjusted.	4	OSINT reliability downgraded based on evidence of channel compromise. Conservative rounding applied: when uncertain between tiers, assign the weaker (higher) tier. Floor drops from 3 to 4, reflecting the degradation of the weakest active source.
S_3: Source isolation	OSINT diagonal excluded per formal criterion (tier 4 + documented channel compromise)	Active matrix: diag(1, 2) + off-diagonal a_{12} = 2. Isolated entry recorded separately.	2	Isolation triggered by documented evidentiary basis, not analyst preference. Floor improves from 4 to 2: with the compromised OSINT stream removed, the weakest remaining active source is HUMINT at tier 2. The remaining streams (SIGINT, HUMINT) are evaluated on their own merits.
S_4: Independent HUMINT corroboration	Second HUMINT source, independent collection, confirms operational intent at tier 2	a_{23} = tier 2 (recorded post-isolation). HUMINT diagonal unchanged.	2	Additional independent corroboration for HUMINT recorded. Floor remains 2 (HUMINT diagonal at tier 2). The path $S_0 \rightarrow S_1 \rightarrow S_2 \rightarrow S_3 \rightarrow S_4$ is recorded in the analytical log.

remaining active source is HUMINT at tier 2. This is not an analytical escape; it is the designed behaviour of a compartmentalised isolation mechanism operating under a documented evidentiary criterion. Second, the source isolation operation demonstrates that the framework does not cause analytical paralysis when a weak source is identified. The compromised stream is formally excluded from the active assessment under the documented criterion, the remaining streams are evaluated on their own merits, and the isolation is recorded for audit purposes.

Discussion

What DRVF claims and does not claim

The theoretical claim of this paper is that a filtered-ring, assessment-tier framework can formalise conservative intelligence fusion in a way that simultaneously enforces the

weakest-link confidence floor and preserves source-tier distinctions algebraically. This claim is theoretical: the framework is formally specified, its key algebraic properties are proven in Appendix A, and a worked example demonstrates its behaviour under the specific scenario conditions for which it is designed. The claim does not extend to empirical superiority over ACH, DST, or Bayesian networks; that question requires controlled testing under the pre-registered protocol available from the corresponding author.

A formally specified framework is a necessary but not a sufficient condition for practical utility. Whether DRVF performs better than the existing methods under the conditions for which it is designed is an empirical question. This paper does not answer it. The contribution is complete in the sense that the theoretical claim that such a framework can be formally specified with algebraically enforceable properties has been discharged. The contribution is incomplete in the sense that the empirical question it raises has not been answered.

DRVF and conventional triangulation

Triangulation and DRVF solve similar but different analytical issues. Triangulation is a validation technique the analyst examines whether other sources come to the same conclusion, and accepts the conclusion as valid unless it is found that the other sources converge. It generates a binary or categorical convergent assessment (convergent, partially convergent, and divergent). DRVF is a mechanism for confidence tracking: it is a record of the change of the aggregate confidence as the reliability of the sources changes over time and yields a graded, formally bounded estimate that can be updated incrementally.

Three specific differences distinguish DRVF from triangulation as conventionally practiced. First, triangulation does not encode source-tier distinctions; it treats convergent sources as mutually confirming without preserving their individual reliability levels as separable, inspectable quantities. DRVF maintains these distinctions in the matrix structure. Second, triangulation does not produce a confidence metric that can be compared across assessment cycles; it answers the question "do sources agree?" but not the question "how confident should we be, given how source reliability has changed since the last assessment?" DRVF is designed to answer the second question. Third, triangulation does not enforce a conservative floor; convergent sources can jointly produce a higher confidence estimate than any individual source warrants if the analyst is not explicitly guarding against manufactured corroboration. DRVF enforces this guard by construction. The two approaches are complementary, not competing: triangulation can operate within a DRVF framework as a corroboration-checking procedure.

Counterarguments

The weakest-link rule is operationally naive. The most substantive objection to DRVF is that the max-based assessment tier imposes an operationally unrealistic constraint. If one source is weak or compromised, real analysts do not discard the entire assessment; they isolate the weak source and rely on stronger material. Applied naively to a scalar assessment, the weakest-link rule would force combined confidence to collapse to the weakest tier regardless of whether other streams provide strong independent support. This objection has force against a scalar weakest-link rule. It does not apply to the matrix structure as specified. As demonstrated in Section 5 (step S_3), the framework supports formal source isolation under a documented criterion: a compromised source entry can be excluded from the active assessment, and T is recalculated over the remaining entries.

Whether analysts find the isolation operation natural and executable under time pressure is an empirical question for the evaluation protocol.

The mathematical formalism is unnecessary. Another objection is that the filtered-ring and assessment-tier structures add algebraic overhead without corresponding practical benefit. A weighted average with a min-floor or a simple fuzzy-logic system could perform the same function. This objection deserves a direct response. A weighted average with a min-floor satisfies the conservative norm on the aggregate but does not preserve source-tier information as algebraically separable entries: once the average is computed, the individual stream reliabilities are collapsed. The matrix structure keeps them distinct and trackable. The compatibility condition $\vartheta(R_k) \geq k$ guarantees that if an element in a ring is reliable, any algebraic operation performed on it is automatically guaranteed to be reliable as well; this requirement does not need to be rechecked by the analyst at every step of fusion. The theoretical case is that these additional properties do not give rise to simpler structures, and thus are formally guaranteed; whether these additional properties are worth the cognitive overhead depends on empirical evidence. The simple-models literature (Dawes, 1979; Gigerenzer and Brighton, 2009) shared this concern and it is recognised as a legitimate challenge.

The formal specification is unnecessary complexity. The third objection is that intelligence analysts do not need formal mathematical specifications; they need clear procedures, good training, and sound professional judgement. This objection has practical force. Most of the intelligence organisations do not operate with formal mathematical frameworks and produce usable assessments without them. The response is that the DRVF is not intended to replace professional judgement but to supplement it with structural constraints that are difficult to enforce through procedural guidelines alone. Whether the benefit of these constraints outweighs the cost of the additional complexity is precisely the question the pre-registered evaluation study is designed to answer.

Pilot feasibility observations

An exploratory pilot study (N = 23) was conducted to assess protocol readiness for the planned confirmatory evaluation. The pilot used intelligence-scenario materials and the source-reliability rubric described in Section 4.4. Participants were intelligence professionals with operational analytical experience. The pilot was not designed to test the framework's performance against alternatives; its purpose was to identify whether the protocol, training materials, and scenario texts were sufficiently clear for participants to engage with the DRVF structure meaningfully.

The pilot produced the following qualitative observations, which should be interpreted with caution, given the small sample and the absence of inferential statistics. First, participants who received the standard training materials (a two-page written explanation of the framework) had difficulty applying the matrix structure to the scenario materials. This suggested that training time needed to be extended and that visual scaffolding (a prepared matrix template with pre-labelled axes) would be necessary. Second, several participants found the concept of "tier 4" and the source-isolation criterion confusing when presented in abstract form but understandable when illustrated with a concrete example. Third, the scenario texts were revised to make the evidentiary basis for tier assignments more explicit, after participants indicated that they were uncertain about which tier to assign to certain evidence items. Fourth, no participant objected to the conservative floor principle itself; the resistance was to the operational complexity of the matrix representation, not to the norm it enforced. These observations informed three protocol refinements incorporated

into the pre-registered full-study design: extended DRVF training time, visual scaffolding requirements, and scenario wording revisions.

Cell sizes in the pilot (n = 7-8 per condition) preclude any inferential use. Inclusion of underpowered directional statistics in a theoretical specification paper risks creating a misleading impression of empirical validation. The pilot's role was to identify protocol refinements, which it did. Pilot data and refinement details are available from the corresponding author on request.

Operational implementation and analyst workload

The reviewer's observation that the DRVF as specified resembles a Paris Fashion Week show, impressive on the runway but unobvious in its translation into everyday analyst practice, is well taken. The framework's mathematical formalism is a specification, not an operational procedure. The question of how an analyst without graduate-level mathematics training would actually use DRVF in a working assessment deserves a direct answer, and the answer is that the operational pathway through the framework is narrower than the formal specification suggests.

There are three observations that will help to address the question of operationalisation. First, the algebraic proofs in Appendix A prove that the structural properties of the framework are true; they do not need to be checked by an analyst when using the framework. An analyst using DRVF does not have to reprove that $T(A(-i)) \leq T(A)$, whereas an analyst using a Bayes' network does not have to reprove Bayes' theorem. The mathematics is not something the analyst does in the analysis; rather, it is a guarantee on the structure. Second, the operational pathway through DRVF consists of three concrete steps: (i) assign a filtration tier to each piece of evidence, using the rubric in Section 4.4; (ii) place the tier in the appropriate diagonal or off-diagonal position in the matrix; (iii) read off T(A) as the maximum tier over the entries that are active. These steps are not symbolic algebraic steps but categorical judgements that have a rubric that has been documented. Third, software implementation of the matrix algebra is very simple: a small python or R script (less than 200 lines) can implement tier assignment, matrix update, T(A) calculation, and generation of audit-log. In addition to the confirmatory study, a reference implementation is published.

On the Science, Technology, Engineering, and Mathematics (STEM)-readiness question, the reviewer correctly notes that the UK intelligence analysts, on average, do not hold STEM degrees at the national average rate. This matters and is acknowledged. The UK Government Communications Headquarters (GCHQ) and the Joint Intelligence Committee (JIC) have historically trained analysts in structured analytical techniques, ACH, red team analysis, and key assumptions check without requiring university-level mathematics. DRVF's operational demands are within the same range: applying the rubric is a categorical judgement; entering a tier into a 3×3 matrix is a tabular task; reading the maximum is a comparison. None of these operations requires the analyst to manipulate the underlying ring-theoretic objects. The cognitive demand is comparable to that of ACH's matrix scoring, which is taught in introductory analyst training across North Atlantic Treaty Organization's (NATO) member services (Heuer and Pherson, 2019). The pilot feasibility study (Section 6.4) confirmed that the principal obstacle was not mathematical capacity but familiarity with the rubric's tier distinctions; participants who received extended training with concrete examples engaged with the framework successfully.

On time budget, a back-of-envelope estimate based on the pilot is as follows. Initial matrix setup for afresh assessment of three streams, initial tier assignments for 8-12 evidence

items, rubric consultation, and audit-log initialisation takes approximately 25-35 minutes. A per-update cycle as new evidence arrives tier assignment for the new item, matrix update, T(A) recalculation, log entry takes approximately 5-10 minutes. Source isolation under the Section 4.7 criterion verification of trigger condition, isolation operation, log entry, and recalculation takes approximately 10-15 minutes. These estimates assume a trained analyst working without software support; a software-assisted implementation would compress per-update time to less than 2 minutes by automating matrix bookkeeping. The estimates are within the range that operational all-source analysts routinely allocate to ACH or Bayesian belief network updating for comparable scenarios.

The framework is therefore not a Paris Fashion Week artefact in any operationally relevant sense. The algebraic specification is rigorous because the structural guarantees it provides conservative floor preservation, source-tier compartmentalisation, and auditability need to be provable, not because the analyst must re-derive them. The operational pathway through the framework is a sequence of categorical judgements comparable in difficulty with the existing structured analytical techniques. What DRVF adds operationally is the conservative floor as a structural constraint that cannot be inadvertently violated, not a body of mathematical machinery that the analyst must master. The mathematical sophistication is in the specification; the operational practice is in the rubric.

A software reference implementation is in preparation and will be released as an open-source tool alongside the confirmatory study. The implementation will provide: a guided tier-assignment interface grounded in the Section 4.4 rubric; automatic matrix update and T(A) recalculation; a structured audit log; and an isolation-operation workflow that enforces the Section 4.7 criterion. The tool is intended to make DRVF usable for analysts without requiring any direct engagement with the underlying algebra. This is the practical operationalisation pathway, and it is the route through which the framework, if empirically validated, would enter working analytical practice.

Boundary conditions

The framework is valid under three assumptions: (a) the number of primary source classes is analytically meaningful for the context (the current specification uses three, but this is a modelling choice, not a universal claim); (b) conservative propagation is a normatively appropriate constraint for the specific failure mode being guarded against (manufactured corroboration in an adversarial environment); and (c) a record of assessment evolution is operationally useful. Outside these conditions, the framework should be modified rather than defended as universal. DRVF is not appropriate for stable multi-sensor environments where priors are well specified (Bayesian methods are better suited), for tasks where speed is paramount and cognitive overhead cannot be accepted, or for contexts with more than three primary source classes without architectural revision.

Positionality and independent replication

The author's background as an operational intelligence and security professional informed the scenario design with ecological validity that would otherwise be difficult to achieve. It also created the risk that the problem was framed in ways that favour the proposed solution. The blind coding procedure, pre-registration, and open materials commitment described in the evaluation protocol were designed specifically to mitigate this risk. Future independent replication by researchers without practitioner backgrounds would provide the strongest available test of whether the architectural properties, not the author's

assumptions, drive any observed performance differences. All scenario texts, the rubric, coding instructions, and anonymised response data will be deposited in a public repository upon publication of the confirmatory study.

Hybrid crises as the primary application domain

The framework's motivating domain warrants explicit discussion. Hybrid crises, the concurrent deployment of military, cyber, information-warfare, and economic instruments by state or non-state actors create a specific and acute version of the general fusion problem that DRVF addresses. Three features of hybrid crises make them particularly demanding environments for multi-source confidence tracking. First, the number and diversity of relevant collection disciplines expand rapidly: a hybrid crisis may simultaneously require assessment of SIGINT, HUMINT, OSINT, financial intelligence, and geospatial intelligence, each with different reliability characteristics and different vulnerability profiles to adversarial manipulation. Second, the adversarial actor has an incentive and a capacity to manipulate specific evidence streams: disinformation operations directly target OSINT inputs, while cyber operations can degrade or fabricate technical indicators. Third, the tempo of operations in a hybrid crisis compresses the decision cycle, reducing the time available for careful source evaluation and increasing the risk that a single compromised stream will contaminate the aggregate picture through apparent corroboration. The two declassified episodes cited in Section 1, while predating the contemporary "hybrid warfare" debate, exhibit precisely these structural features: multiple collection disciplines, active manipulation of at least one stream, and time pressure that discouraged the type of slow and careful source-by-source evaluation that might have prevented the overconfident assessments that followed.

It is important to reiterate that DRVF is not specific to hybrid crises. The three structural conditions the framework requires unstable source reliability, adversarial manipulation potential, and auditability requirements are present in many intelligence contexts. The framework's mathematical properties hold, regardless of the domain in which it is applied. Hybrid crises are the motivating domain because they concentrate all three conditions in a form that makes the need for a formal conservative mechanism most visible and most consequential. The NATO Hybrid Centre of Excellence (2022) has noted that strategic-communications challenges in hybrid environments are a recognised gap; the present framework addresses the analytical rather than the communications dimension of that gap, but the underlying problem distinguishing genuine from manufactured corroboration across multiple evidence streams is the same.

Conclusions

This paper has presented DRVF, a formally specified framework for conservative multi-source intelligence fusion. The core architecture integrates two layers: a filtered-ring representation that preserves source-tier distinctions algebraically, and an assessment-tier aggregation rule that enforces the weakest-link confidence floor. Both layers are formally specified, algebraic properties are proven in Appendix A, and the behaviour of the framework under a realistic operational scenario drawn from the hybrid-crisis domain is demonstrated through a worked example.

Two declassified historical episodes, the pre-September 2001 intelligence integration failures documented by the 9/11 Commission (US Government Printing Office, 2004) and the Iraq WMD assessment analysed by Jervis (2010), provided concrete grounding

for the problem that the framework addresses. Both involved multi-causal failures; the structural contribution of this paper is specific to the formal dimension: the absence of an algebraic constraint, preventing weak-evidence contributions from elevating aggregate confidence. The gap claim is stated conservatively as a relative contribution: no framework combining conservative weakest-link aggregation, hierarchical source compartmentalisation, and traceable assessment evolution in an intelligence-fusion context was identified in the reviewed corpus.

The framework is proposed as a formally specified, testable candidate. Whether it performs better than ACH, DST, or Bayesian networks under the conditions for which it is designed is an empirical question. A pre-registered evaluation protocol has been designed and is available from the corresponding author. No empirical performance claims are made.

Three operational directions for further work follow from the framework as specified. First, the planned confirmatory evaluation (pre-registered, N = 75) will test whether DRVF, used with the rubric and supporting materials, produces analytical outcomes that meet or exceed those of ACH and Bayesian belief networks on the specific failure modes the framework targets: conservative floor maintenance under source compromise, source-tier preservation under fusion, and audit-trajectory completeness. Second, a software reference implementation of matrix algebra, tier-assignment interface, and audit-log generation is under preparation. The implementation is intended to remove the algebraic overhead from operational use, leaving the analyst to perform the categorical judgements (rubric assignment and isolation decisions) that DRVF requires. Third, integration with the existing structured analytical techniques is a natural extension: DRVF's matrix can serve as a confidence-tracking layer beneath ACH's hypothesis-scoring matrix, providing the trajectory record that ACH lacks. Each of these directions is operationally tractable; none requires the analyst to engage with the underlying ring-theoretic specification.

Funding

This research received no external funding.

Data Availability Statement

Data sharing is not applicable to this article.No new empirical datasets were generated or analysed in this study. The pilot observations are qualitative and available from the corresponding author on reasonable request, subject to institutional confidentiality constraints.

Artificial Intelligence Disclosure Statement

During the preparation of this manuscript, the author did not use generative artificial intelligence tools. The author takes full responsibility for the content of the manuscript.

Disclosure Statement

No potential conflict of interest was reported by the authors.

The authors read and agreed to the published version of the manuscript.

References

Anjom SHoa, M.H. and Hosseini, M.H. (2016) 'Filtration and Dubrovin valuation on matrix ring, Asian-European', *Journal of Mathematics*, 9(1), p. 1650011. doi: 10.1142/S179355711650011X.

Betts, R.K. (2002) 'Fixing intelligence', *Foreign Affairs*, 81(1), pp. 43–59.

Bousquet, A. (2009) *The scientific way of warfare: Order and chaos on the battlefields of modernity*. London: Hurst / New York: Columbia University Press.

Caliskan, M. and Liégeois, M. (2021) 'The concept of "hybrid warfare" undermines NATO's strategic thinking: insights from interviews with NATO officials', *Small Wars & Insurgencies*, 32(2), pp. 295–319. doi: 10.1080/09592318.2020.1860374.

Castanedo, F. (2013) 'A review of data fusion techniques', *The Scientific World Journal*, 2013, art. no. 704504. doi: 10.1155/2013/704504.

Chang, W., Berdini, E., Mandel, D.R. and Tetlock, P.E. (2018) 'Restructuring structured analytic techniques in intelligence', *Intelligence and National Security*, 33(3), pp. 337–356. doi: 10.1080/02684527.2017.1400230.

Dawes, R.M. (1979) 'The robust beauty of improper linear models in decision making', *American Psychologist*, 34(7), pp. 571–582. doi: 10.1037/0003-066X.34.7.571.

Dempster, A.P. (1967) 'Upper and lower probabilities induced by a multivalued mapping', *Annals of Mathematical Statistics*, 38(2), pp. 325–339. doi: 10.1214/aoms/1177698950.

Echevarria, A. J., II. (2016) *Operating in the gray zone: An alternative paradigm for U.S. military strategy*. Carlisle, PA: Strategic Studies Institute, U.S. Army War College.

Friedman, J.A. and Zeckhauser, R. (2012) 'Assessing uncertainty in intelligence', *Intelligence and National Security*, 27(6), pp. 824–847. doi: 10.1080/02684527.2012.708275.

Gigerenzer, G. and Brighton, H. (2009) 'Homo heuristicus: Why biased minds make better inferences', *Topics in Cognitive Science*, 1(1), pp. 107–143. doi: 10.1111/j.1756-8765.2008.01006.x.

Hall, D.L. and Llinas, J. (2001) *Handbook of multisensor data fusion*. Boca Raton, FL: CRC Press.

Herman, M. (1996) *Intelligence power in peace and war*. Cambridge: Cambridge University Press.

Heuer, R.J., Jr. and Pherson, R.H. (2019) *Structured analytic techniques for intelligence analysis*. 3rd edn. Washington, DC: CQ Press.

Jervis, R. (2010) *Why intelligence fails: Lessons from the Iranian Revolution and the Iraq War*. Ithaca, NY: Cornell University Press.

Kent, S. (1949) *Strategic intelligence for American world policy*. Princeton, NJ: Princeton University Press.

Liggins, M.E., Hall, D.L. and Llinas, J. (Eds) (2009) *Handbook of multisensor data fusion: Theory and practice*. 2nd edn. Boca Raton, FL: CRC Press.

Mandel, D.R. and Barnes, A. (2014) 'Accuracy of forecasts in strategic intelligence', *Proceedings of the National Academy of Sciences*, 111(30), pp. 10984–10989. doi: 10.1073/pnas.1406138111.

Muñoz-González, L. and Lupu, E.C. (2017) 'Bayesian adversarial learning: A framework and its use for adversarial network hardening', in *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*. New York, NY: IEEE, pp. 3644–3651.

Murphy, C.K. (2000) 'Combining belief functions when evidence conflicts', *Decision Support Systems*, 29(1), pp. 1–9. doi: 10.1016/S0167-9236(99)00084-6.

NATO Hybrid Centre of Excellence (2022): Accessed 12 October 2026.US Government Printing Office (2004): Accessed 12 October 2026. https://www.9-11commission.gov/report/

Pearl, J. (1988) *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Francisco, CA (San Mateo, CA): Morgan Kaufmann.

Rieber, S. (2004) 'Intelligence analysis and judgmental calibration', *International Journal of Intelligence and CounterIntelligence*, 17(1), pp. 97–112. doi: 10.1080/08850600490273431.

Sentz, K. and Ferson, S. (2002) *Combination of evidence in Dempster-Shafer theory* (SAND2002-0835). Albuquerque, NM: Sandia National Laboratories.

Shafer, G. (1976) *A mathematical theory of evidence*. Princeton, NJ: Princeton University Press.

Smets, P. (1990) 'The combination of evidence in the transferable belief model', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(5), pp. 447–458. doi: 10.1109/34.55104.

Tricco, A.C., Lillie, E., Zarin, W., O'Brien, K.K., Colquhoun, H., Levac, D. *et al.* (2018) 'PRISMA extension for scoping reviews (PRISMA-ScR): Checklist and explanation', *Annals of Internal Medicine*, 169(7), pp. 467–473. doi: 10.7326/M18-0850.

US Government Printing Office (2004) *The 9/11 commission report: Final report of the national commission on terrorist attacks upon the United States*. US Government Printing Office. Available at: https://www.9-11commission.gov/report/ (Accessed: xxx).

van Ditmarsch, H., van der Hoek, W. and Kooi, B. (2007) *Dynamic epistemic logic*. Cham: Springer. doi: 10.1007/978-1-4020-5839-4.

Yager, R.R. (1987) 'On the Dempster-Shafer framework and new combination rules', *Information Sciences*, 41(2), pp. 93–137. doi: 10.1016/0020-0255(87)90007-7.

http://doi.org/10.35467/sdq/240155

Appendix A: Formal algebraic proofs and supplementary specifications

A.1. Minimal algebraic setting

Let R be a non-zero non-commutative ring with filtration $\{R_k\}_k \geq 0$ satisfying $R_0 = R$, $R_{k+1} \subseteq R_k$, and $R_k R_m \subseteq R_{k+m}$. The compatibility condition $\vartheta(R_k) \geq k$ aligns algebraic grading with reliability tiers. This is a parametric design choice, not a universal property of filtered rings (Anjom Shoa, 2016).

A.2. Quasi-valuation properties: proofs

- Claim A.2.1 (addition). For $A,B \in M_3(R)$, $\vartheta(A + B) \geq \min\{\vartheta(A),\vartheta(B)\}$.

Proof: By the entry-wise minimum extension:

$$\vartheta(A + B) = \min_{ij} \vartheta(a_{ij} + b_{ij}) \geq \min_{ij} \min\{\vartheta(a_{ij}),\vartheta(b_{ij})\}$$
$$= \min(\min_{ij} \vartheta(a_{ij}), \min_{ij} \vartheta(b_{ij})) = \min\{\vartheta(A), \vartheta(B)\}.$$

- Claim A.2.2 (multiplication). For $A,B \in M_3(R)$, $\vartheta(AB) \geq \vartheta(A) + \vartheta(B)$.

Proof: For entry $(AB)_{ij} = \Sigma_k a_{ik} b_{kj}$. By the quasi-valuation ring axiom: $\vartheta(a_{ik} b_{kj}) \geq \vartheta(a_{ij}) + \vartheta(b_{kj}) \geq \vartheta(A) + \vartheta(B)$. Since ϑ of a sum is bounded by the minimum: $\vartheta(\Sigma_k a_{ik} b_{kj}) \geq \min_k\{\vartheta(a_{ik} bkj)\} \geq \vartheta(A) + \vartheta(B)$. Therefore, $\vartheta(AB) = \min_{ij} \vartheta((AB)_{ij}) \geq \vartheta(A) + \vartheta(B)$.

- Claim A.2.3 (weakest-link encoding). If A contains a non-zero tier-assigned entry a_{ij} at filtration level k, then $T(A) \geq k$.

Proof. By definition, $T(A) = \max\{\tau(a_{00}) : a_{00} \neq 0\} \geq \tau(a_{ij}) = k$. No combination of other entries can lower the maximum below k. To improve T(A) below k, the entry a_{ij} must be upgraded to a lower filtration level or removed from the active matrix state.

Remark. The distinction between ϑ and T is structural. The quasi-valuation $\vartheta(A) = \min_{ij} \vartheta(a_{ij})$ provides an algebraic lower bound on individual entry valuations and is essential for proving that the filtration structure is preserved under ring operations. The assessment tier $T(A) = \max\{\tau(a_{ij})\}$ implements the weakest-link norm by selecting the highest (weakest) filtration tier among all active entries. Under the compatibility condition $\vartheta(R_k) \geq k$, we have $\tau(a) \leq \vartheta(a)$ for any tier-assigned element a, and therefore, $T(A) \leq \max\vartheta(a_{ij})$ for all tier-assigned entries. The two functions provide complementary information: ϑ ensures algebraic coherence, and T enforces the conservative constraint.

A.3. Compatibility condition

The condition $\vartheta(R_k) \geq k$ is a local design constraint aligning algebraic grading with reliability tiers. If the tier structure of a specific operational context requires a different alignment, the compatibility condition should be revised accordingly without affecting the broader architecture.

A.4. Source isolation operation

Let $A \in M_3(R)$ be the current matrix state. Isolation of source stream i is warranted if and only if the formal criterion in Section 4.7 is satisfied: diagonal entry a_{ij} has been assigned

tier 4, or an explicit documented evidentiary trigger has been recorded in the analytical log. Define the isolated matrix $A(-i)$ as A with row i and column i set to zero (excluded from the domain of T). Then:

$$T(A(-i)) = \max \textit{ over } p,q \in \{1,2,3\},\ p \neq i,\ q \neq i \text{ of } \tau(a_{0e}).$$

This is a well-defined operation within the algebraic structure, and $T(A(-i)) \leq T(A)$ because the maximum over a subset of entries cannot exceed the maximum over all entries. The excluded entries are placed in a separate isolation log, forming part of the auditable update path. Reintegration of an isolated source, if subsequent evidence rehabilitates it, requires the analyst to explicitly restore the relevant entries and document the evidentiary basis.

A.5. Optional visualisation (provisional)

Earlier versions of this work included a geometric visualisation layer that encoded analytical state transitions as a walk on a hexagonal grid. Six transition types were mapped to six local directions on the grid, providing a compact visual representation of the assessment path. This visualisation has been moved to a provisional appendix entry because its cognitive and operational properties require empirical testing. If testing reveals that the hexagonal structure does not reduce cognitive load or that a directed acyclic graph with timestamps serves equally well, the visualisation should be revised or replaced without affecting the core algebraic architecture. A directed acyclic graph with timestamps is a valid alternative that imposes no constraint on the number of distinct update types.

A.6. Institutional review board (IRB) and pre-registration

The pilot feasibility study was carried out under the review protocol of the sponsoring security organisation. The organisation determined that the study met the criteria for exemption from full ethics committee review under applicable national research ethics guidelines. Confidentiality requirements preclude disclosure of the organisation's identity. The exemption determination is available for verification by journal editors on request through the corresponding author, subject to applicable confidentiality constraints. Participants provided written informed consent. The full confirmatory study (N = 75) is pre-registered on the Open Science Framework. The registration link is to be confirmed and inserted at publication. All scenario texts, the rubric, coding instructions, training materials, and anonymised response data are to be deposited in a public repository upon publication.

A.7. PRISMA-ScR flowchart

The PRISMA-ScR flowchart (Figure A1) documents the screening process in standard PRISMA-ScR diagram format (Tricco et al., 2018). In all, 312 records were identified through database searching across Google Scholar, IEEE Xplore, ScienceDirect, SpringerLink, and Taylor & Francis Online. After removal of forty-seven duplicate records, 265 records were screened by titles and abstracts. In all, 176 records were excluded at this stage; eighty-nine records were assessed for full-text eligibility; and forty-two records were excluded—twenty-eight did not address multi-source uncertainty in a tradecraft context, nine focused on non-adversarial sensor fusion, and five were not accessible in full text. Forty-seven records were included in the final synthesis.

http://doi.org/10.35467/sdq/240155

Figure A1. PRISMA-ScR flowchart: identification and selection of reviewed sources.

