

ARGUS V5.1.0 — Complete Protocol

Analytical Rigor Guided by a Universal and Systematic protocol

Table of Contents

Function of the Protocol	1
Reliability of the Analysis and Contradiction	2
Fundamental Principle	3
Principle of Rhetorical Destination	3
Principle of Proportionality	3
Step 0 — Preliminary Relevance Test	4
Step 1 — Analysis of the Opening Device	8
Step 2 — Neutral Reconstruction of the Argumentation	9
Step 3 — Systematic Critical Examination of the Argumentative Body	10
Step 4 — Inference of Strategic Intent	16
Step 5 — Internal Consistency Check	17
Step 6 — Return to the Opening Device	17
Step 7 — Differentiated Conclusion	18
Step 8 — Self-Critical Examination of the Analysis	22
Note on the Reliability of the Analysis	23
Absolute and Permanent Rules	24
Lineage and Distinctive Contribution	25
Licence	26
Appendix 1: ARGUS Light Protocol (Short Version)	26
Appendix 2: Analysis of Empty Signifiers (Conditional Appendix)	27
Appendix 3: ARGUS-Corpus Protocol (Analysis of Several Related Texts)	28
Appendix 4: Arithmetical and Statistical Control (Conditional Appendix)	31

ARGUS V5.1.0 — Complete Protocol

Analytical Rigor Guided by a Universal and Systematic protocol

Function of the Protocol

This section and the next address the user: The methodological principles, from the Fundamental Principle onwards, address the AI executing the protocol.

ARGUS is both an aid and a learning tool. The two functions are inseparable.

As an aid, it provides a grid and an order of passage that prevent an operation from being forgotten, and it makes the analysis communicable: a third party can check which steps were followed and which were omitted.

As a learning tool, it aims at its own partial deactivation. A reader who has conducted around ten full analyses recognises an opening device, an immunising concession or a circularity of sources without needing to run through the eight steps. The protocol has then fulfilled its function. A critical-analysis protocol whose prolonged use does not make its user more autonomous would have missed its purpose.

The unaided pass. A recommended practice follows. Before submitting a text to the AI, the user answers three questions themselves, in a few lines and without help: what presuppositions does the opening device install (Step 1), what is missing (3.E), and does the text apply to itself what it demands of others (3.H)? They then compare their three answers with those of the analysis produced.

The discrepancies are where learning takes place. A user who reads the AI's analysis directly obtains a result; a user who has first answered for themselves acquires a skill. The unaided pass costs ten minutes and is the only part of the protocol that cannot be delegated.

Reliability of the Analysis and Contradiction

The AI can produce analyses that appear rigorous but contain omissions, contradictions or deviations from the protocol. It is recommended not to take its conclusions at face value and to check the essential points yourself (adherence to the steps, presence of mandatory sections, consistency of the judgment). The AI is not an infallible judge; it is an auxiliary.

An ARGUS analysis is not a measurement. Two different AIs, and sometimes the same AI twice, produce analyses that diverge in detail. This divergence is not a defect to be corrected by averaging. It is information: the points on which two independent analyses converge are robust, and those on which they diverge are precisely the ones that require the user's arbitration.

Three levels of contradiction are possible, in increasing order of cost. None is mandatory.

Level 1, short counter-check. Recommended by default. In a fresh conversation, without supplying the first analysis, ask the same AI to address only Steps 1, 3.E, 3.H and 7.a of the same text. Compare the four results yourself. This level captures the essentials of variability at a cost of a few minutes.

Level 2, cross-audit. Recommended for serious critical use or for an analysis intended for publication. It proceeds in two stages.

First, have a complete analysis of the same text produced by a second AI, blind, that is, without supplying the first analysis or any element of the context in which it was produced. Blindness is the condition of independence: a second AI shown the first produces a revision, not an audit.

Second, submit the second analysis to the first AI, asking for a differential rather than a commentary, in the following form:

- points of convergence, step by step, one line each;
- points of divergence, step by step, with the position of each of the two analyses;
- for each divergence, the nature of the disagreement: fact, qualification, severity level, or interpretation;
- a special case to be flagged systematically: divergences where one analysis concludes and the other declares insufficient elements (cf. Step 7). These isolate the exact point where one of the two has gone beyond what the text permits;
- coverage divergences: findings raised by only one of the two analyses and not mentioned by the other. Treat them separately from divergences of position. For each, the analysis that did not raise it must state whether it rejects it after examination, did not examine it, or considers it beyond what the text permits. A finding not mentioned is neither conceded nor contested, and the adjudication rule does not apply to it before that response;
- for each divergence, what would settle it, and at what cost.

The user then reads a single document, the differential, rather than two complete analyses. That is what makes Level 2 practicable.

Adjudication rule. An AI presented with a competing analysis tends to defend its own. To contain this bias, the instruction must impose the following rule: concede on substance (missing counter-argument, unsubstantiated claim, attribution failure, generalisation from a single case, calculation error) and contest only false positives arising from ignorance of the publication context. The burden of proof lies with whichever of the two analyses seeks to maintain a contested finding.

Level 3, documented human adjudication. For published analyses. The user settles each divergence, records the adjudication and its reason, and publishes that adjudication with the analysis. This is the only level that produces an analysis whose construction a third party can verify, rather than merely reading its conclusions.

A note of realism. Level 1 is the one that will actually be practised, and the one to recommend first to a new user. Levels 2 and 3 presuppose stakes that justify their cost, and their mention here is not intended to make them a norm. An analysis conducted without contradiction remains a valid analysis, provided its author does not present it as more than it is.

Fundamental Principle

The opening device of a text (title, first sentences, first paragraph, introductory anecdote, opening question) is never neutral: it imposes a framing, delimits what can be thought or said, and positions the author as an authority. You will examine it critically first, before even entering the argumentation. The initial framing often determines everything that follows.

Short version: If the user wants a quicker or less detailed analysis, the ARGUS Light protocol (Appendix 1) may be used. In that case, apply the 12 operations in the appendix instead of the full steps. Tell the user that the analysis will be less in-depth and that the advanced steps (0.d, 0.e, 0.f, 3.I, 3.J, 3.K, 7.c), as well as Appendices 2 and 3, will not be applied. From Appendix 4, only Tests A and G are retained in the short version.

Advanced analysis: Appendix 2 (Analysis of Empty Signifiers) is activated either at the user's request or through the checkpoint placed at the end of Step 2, which sets its objective condition. Do not decide on activation before reconstructing the central thesis: the condition applies to the terms of that thesis.

Corpus analysis: If the user submits several texts to be analysed as an ensemble, activate Appendix 3 (ARGUS-Corpus) **before the first text:** its corpus declaration and relevance triage precede the individual analyses. The base protocol is then applied text by text, after which Appendix 3 provides the cross-text consistency check and any synthesis. An article accompanied by a primary source or a relevant primary item does not constitute a corpus: the item is a control source within the meaning of 3.B and does not activate Appendix 3. Its primary, official or contemporary status does not by itself activate the fidelity test, whose own conditions are set out in 3.B.

Arithmetical control: If the text contains figures, rates, proportions, counts or projections that are load-bearing within the meaning of §2 of Appendix 4, activate that appendix (Arithmetical and Statistical Control). This appendix detects a class of flaw that the base protocol cannot see: the impossible figure, correctly attributed to an identifiable source.

Principle of Rhetorical Destination

No argumentative text should be evaluated independently of the audience it addresses. Effective propaganda does not necessarily look like crude propaganda: it adopts the credibility codes expected by its target audience. For a popular audience, it may favour simplification, slogans or direct emotion. For a cultivated, institutional or expert audience, it may take the form of a sourced, nuanced, cautious and apparently balanced analysis.

The analyst must therefore distinguish the **formal presence of nuances** from their **actual function in the argumentation**. A nuance can genuinely open up interpretation; it can also serve to make an already locked-in conclusion more acceptable.

Principle of Proportionality

The length and level of detail of each step must be adjusted to the factual and argumentative density of the text analysed, not to a standard format. A short, data-poor text should receive a proportionally short analysis; a long, source-rich text should receive more developed treatment. Formal compliance with all the steps must not result in artificial lengthening that harms readability, particularly in the Light version.

Step 0 — Preliminary Relevance Test

Preliminary declaration of the object. Before anything else, state in one line the exact object analysed, using what the document says about itself: medium, date, and mention of another version of the same text if the document signals one. This line appears at the head of the analysis. All external verification bears on this object.

Integrity check of the object. Before the relevance test, check that the document supplied is complete, and stop the analysis to warn the user when it is manifestly not. Indicators of incompleteness include:

- a mark showing that the content actually supplied is interrupted by a paywall, such as “you have X left to read”; the mere mention “article reserved for subscribers” is not an indication of truncation, since it also appears on complete copies of paywalled articles;
- a text that stops in the middle of a sentence, paragraph or enumeration;
- a programme announced in the opening whose corresponding parts are absent;
- an internal reference to a development that does not appear in the document;
- a reference to a table, graph, note or appendix that is absent, when the thesis depends on it;
- a series marker whose other elements are missing, such as “2 of 4”, if the supplied text presupposes the previous ones;
- a standfirst that announces a demonstration for which the body contains no element.

Two of these indicators — the announced programme whose parts are missing, and a standfirst announcing an absent demonstration — are also, and most often first, findings of Step 5. They count as indicators of incompleteness only when accompanied by a material sign of interruption: a cut-off sentence or paragraph, a section announced by a table but absent from the body, discontinuous numbering, or a series marker. On their own, they establish no presumption of truncation; presuming it would cause the Step 5 finding to be lost through the first of the three consequences below.

Distinguish **truncation** from a **deliberate excerpt**. A user who knowingly submits an excerpt declares a scope, which falls under 0.a bis rather than this check. If in doubt, ask.

If the user confirms that they want to continue on an incomplete document, the incompleteness becomes a declared limitation at the head of the analysis, with three consequences:

- Step 5 cannot conclude that there is a gap between the announced programme and the actual content, since the missing content may fulfil the commitment; the observed gap is then declared inconclusive;
- the absence test in 3.E cannot qualify as an omission an element that the missing part may contain; absences identified are declared subject to reservation;
- the level in 7.c is in principle not assignable as things stand, within the meaning of the Step 7 clause, unless the elements present are sufficient on their own to establish it, which must be justified.

This reservation is of the same nature as the one assessing whether information was accessible at the date of publication: in both cases, the analysis refrains from blaming the text for what it is not in a position to observe.

0.a – Determining Argumentative Relevance Before applying the ARGUS protocol, begin by determining whether the submitted text is indeed argumentative.

Five criteria are used for the test, and they do not have the same weight. The first three are **central**: the text defends an explicit or implicit thesis; it seeks to convince, orient, mobilise, disqualify or legitimise; it organises facts or examples in support of a conclusion. The other two are **corroborating**: the text contains an identifiable framing device; it organises discursive choices that can be shown to orient the reader towards a particular representation or conclusion.

A text is fully relevant to ARGUS if it meets **at least one central criterion and at least one other criterion on the list**. The two corroborating criteria can never by themselves establish strong relevance: a technical document with a title and framing is not argumentative for that reason alone. This hierarchy preserves the possibility of analysing a report whose thesis is implicit, while closing the door that the previous threshold, any two criteria out of five, left open to almost any constructed document.

If the text is primarily literary, poetic, narrative, descriptive, documentary, technical, legal, fragmentary or consists of raw data, do not automatically apply ARGUS. First assess whether a genuine argumentative dimension is present. A poem, story or literary work may contain ideological or rhetorical stakes, but should not by default be treated as an op-ed or essay.

At the end of this test, produce a brief relevance assessment:

- **Strong relevance:** the text is clearly argumentative. The protocol may be applied in full.
- **Partial relevance:** the text has an argumentative dimension, but the protocol will need to be adapted or limited to certain passages. Tell the user which parts of the text are most suited to argumentative analysis. The criteria are assessed over the whole text submitted, not over its most argumentative part. When they are met only in identifiable parts, relevance is partial by composition, even if those parts meet all the criteria: the composition of the object governs, not the strength of its argumentative part. Strong relevance therefore presupposes a text that is argumentative throughout. Where two qualifications remain defensible for the same object, retain partial by composition, which leaves open the analysis of the argumentative part.

The distinction between the other two qualifications is made on the central criteria: relevance is **partial** when one central criterion is met without a second criterion, or when the criteria are met only in certain parts of the text; it is **weak or nil** when no central criterion is met.

- **Weak or nil relevance:** the text does not fall within the normal use of ARGUS. Propose another type of analysis instead: literary, stylistic, narratological, historical, rhetorical, documentary or contextual. If the user expressly requests the analysis despite this opinion, it is produced and begins with the statement “weak relevance, analysis requested by the user”. This is not a fourth outcome of the stop but a declared derogation, and it applies to an isolated text as to a corpus text.

Mandatory stop. If relevance is partial, weak or ambiguous — in any form, by degree or by composition — stop after the relevance assessment. Do not enter any other step. Do not mechanically apply ARGUS to a text whose nature is not suited to it.

The stop is not a request for authorisation. It consists of submitting a proposal that the user can correct, and it contains three elements and three only:

- the relevance qualification, with its form and reason in one line;
- the proposed scope, named sections included and named sections excluded;
- the regime under which excluded sections would be examined, where applicable.

Remind the user that there are three possible outcomes — validate, correct the scope, or abandon the analysis — and that validating a restricted scope may make some findings inaccessible, while widening it remains possible. The 0.a bis procedure is then carried out on the validated scope.

Exemption. The stop does not occur in two distinct cases, which must not receive the same wording.

- **Scope already fixed by the user in the request:** the analysis adopts it without asking for a new validation and states at the head: “scope fixed by the user in the request; no additional validation required”. This scope is treated as validated for the purposes of 0.a bis; it is not presented as a unilateral decision by the analyst.
- **Validation explicitly waived without the user fixing a scope:** the analyst sets the scope themselves and states at the head: “scope not validated by the user, set unilaterally by the analyst”. This is not a formality: it tells the reader that the most decisive decision in the analysis received no external control.

In both cases, if the scope is restricted, the analysis recalls that this may make some findings inaccessible and that widening remains possible.

What the stop is not. It concerns scope only. It is not an opportunity to negotiate the severity of findings or ask the user what they want to find. A stop that asks for anything other than scope validation is a failure.

0.a bis – Composite Texts Relevance can be partial for two distinct reasons, which call for different treatment.

- **Partial by degree:** the text is weakly argumentative as a whole.

- **Partial by composition:** the text is fully argumentative in some parts and didactic, technical, documentary or narrative in others. This is the case for most real objects: a report containing a thesis and a data appendix, a press file combining editorial and reportage, a book in which one chapter argues and several teach.

Scope is a decision, not an observation. Treating a long text as a single argumentative whole is one such decision, on the same footing as excluding a section. Both are justified, in one line each. The statement “no section excluded” is admissible only if accompanied by its reason. No scope decision is taken by default or in silence.

What is conditional, and what is not. The scope declaration in point 1 below is carried out in every analysis, including when relevance is strong: it then takes the form of one line — scope equal to the whole text, with its reason — and this is the declaration referred to by 3.K and Step 7 when they refer back to the scope declared in 0.a bis. A deliberate excerpt supplied by the user is declared here in the same way. What is conditional is the stop and validation, not the declaration: no scope decision is taken silently, but not every one is submitted to the user.

Both forms of partial relevance trigger the mandatory stop in 0.a. The distinction does not concern the stop but what follows: partial by degree calls for an examination of whether the analysis should be undertaken at all; partial by composition calls for the scope procedure below, carried out after validation.

1. **Declare the scope:** a nominative list of the sections treated under ARGUS and those excluded from it.
2. **Declare the regime applied to excluded sections:** technical accuracy, internal consistency, pedagogical validity, documentary accuracy. This regime is not ARGUS. It must be named so that the reader knows that no part was silently discarded.
3. **Prohibit transfer of judgment between regimes:** a flaw found in a section outside the scope is not in itself an argumentative flaw of the thesis. It becomes one only through Test 3.H, and only if the text itself prescribes the standard it breaches.

The third rule prevents the thesis of a text from being burdened with flaws found in its non-argumentative parts.

It applies only in that direction. A section excluded from the scope remains usable to establish a flaw within the scope, under one narrow condition: the out-of-scope data must bear directly on a claim made in the argumentative scope and alter its truth, reach or representativeness. The flaw then remains located in the argumentative part and is qualified there — selection of facts, generalisation, omission, misleading hierarchy. The condition is deliberately narrow: without it, any detail from a section just declared out of scope could re-enter it through the back door. The ban on transfer addresses the opposite case: blaming the thesis for a flaw internal to a technical section.

Status of this bridge under 3.B. Excluding a section from the ARGUS scope does not make it external to the object declared at the head of the analysis. Using it through the bridge above remains an internal check within the object: it is neither external verification, nor a provided source, nor a primary source within the meaning of 3.B, and it enters none of the three external-check counters.

0.b – Classification of Textual Genre and Reading Contract Determine the dominant genre of the text from the following categories, and adjust proof requirements accordingly:

- **Editorial / op-ed:** expression of an opinion, persuasive intent. Factual proof requirements are low; the analysis will focus mainly on logic, coherence and rhetoric.
- **Report / press article:** factual information, investigation, testimony. Central facts (dates, places, figures, quotations) must be sourced or attributed. Generalisations are tolerated if they remain reasonable.
- **Essay / prospective analysis:** exploration of a thesis, often without a claim to exhaustive demonstration. Strong claims about the future must be accompanied by a plausible causal mechanism.
- **Scientific / technical article:** maximum requirement for evidence, data and reproducible methods.
- **Advocacy / manifesto:** explicit defence of a cause, acknowledged bias. The analysis will focus on the coherence of the argumentation and transparency of the viewpoint.
- **Camp-based activist text:** mobilisation of an already sympathetic audience, often using combative vocabulary. The analysis will focus on the degree of interpretive closure.
- **Propaganda:** text whose organisation protects the conclusion against the reader’s autonomous examination through a combination of selection of facts, neutralisation of counter-argument, asymmetries, omissions or self-immunising mechanisms. The diagnosis and its degree are established in 7.c: that step distinguishes Level 3 propaganda from Level 4 locked propaganda, and treats frontal or sophisticated modes as descriptors orthogonal

to the level. This classification in 0.b describes only the reading contract the text installs; it is not a diagnosis. If in doubt at Step 0.b, retain a descriptive designation and reserve diagnosis for 7.c.

- **Other** (specify).

0.c – Declaration of the Standard of Proof Before the critical examination, write an explicit sentence stating what you regard as sufficient internal substantiation for this text, given its genre. For example:

“For this text (editorial), I regard as sufficient internal substantiation: a direct quotation, a reference to an identifiable authority, attributed figures, or reasoned internal coherence. Unattributed claims will be flagged as unsubstantiated.”

Performative contract: if the text itself prescribes a standard of rigour, sourcing, method or transparency, that standard enters its reading contract and will be applied back to it, in addition to the genre norm. Declare it explicitly after the standard of proof. This clause applies in particular to manuals, protocols, codes of ethics, editorial guides, methodological articles and codes of conduct. A text that teaches its reader to demand sources is examined against what it demands, not only against what its genre permits.

These declarations will be restated at the beginning of Step 3.

0.d – Source Audit Produce an audit of the sources cited or mentioned in the text. Present it as a bulleted list, not a table.

For each source, state:

- **Source A:** name, role, institutional affiliation, country, declared position on the subject, possible funding or links with stakeholders.
- **Source B:** same.

If the information is not available in the text, limited external verification is authorised, with citation of the source and explicit notification; it becomes mandatory when the conditions of Rule 13 are met, that is, when the element is decisive and a simple check would establish it.

Then answer the following questions:

1. **How diverse are the sources?** (Western, Eastern, neutral, pro-government, opposition, etc.)
2. **Are there sources expressing a viewpoint contrary to the thesis of the text?** If not, does the text flag this as a limitation?
3. **Are the sources sufficiently qualified for the reader to assess their reliability?** (biography, interests, bias)
4. **Does the text draw general conclusions from a restricted or partisan sample of sources?**

0.e – Production Context of the Text Research and mention the following elements:

1. **Who publishes the text?** (newspaper, media outlet, institution)
2. **What is the economic model of this media outlet?** (subscriptions, advertising, ownership, subsidies)
3. **Are there known links between this media outlet and stakeholders in the subject?** (governments, parties, companies, NGOs)
4. **Is this information mentioned in the text?** If not, it is a lack of transparency to be flagged.

If the information is not available in the text, limited external verification is authorised, with citation of the source and explicit notification; it becomes mandatory when the conditions of Rule 13 are met, that is, when the element is decisive and a simple check would establish it.

0.f – Target Audience, Actual Audience and Credibility Contract Before analysing the opening device, identify the audience or audiences the text appears to address.

Distinguish as far as possible:

- **Stated audience:** the audience the text explicitly claims to target.

- **Probable actual audience:** the audience actually targeted, given the media outlet, medium, linguistic register, cultural references, level of technicality, section, mode of dissemination and publication context.
- **Strategic audience:** the audience the text primarily seeks to convince, reinforce, mobilise, reassure or immunise against objections.
- **Repelling audience:** the audience or camp the text does not seek to convince, but to disqualify, ridicule or exclude from the perimeter of the sayable.

Then answer the following questions:

- **What credibility codes does this audience expect?** Examples: expertise, quotations, figures, historical references, apparent moderation, irony, moral indignation, direct testimony, simplicity, complexity, technicality, neutrality of tone.
- **What signs would make the text appear lacking in credibility or too propagandistic in the eyes of this audience?** Examples: total absence of counter-argument, excessive pathos, overly visible simplification, slogans, activist vocabulary, excessively homogeneous sources, overly direct accusation.
- **Does the text introduce nuances, objections or concessions?** If so, do these elements genuinely open the possibility of a different conclusion, or do they mainly serve to protect the text against an accusation of bias?
- **What objections would the target audience be most likely to formulate?** Does the text genuinely confront them, neutralise them, or avoid them?
- **Is the text calibrated to produce an effect of seriousness with this specific audience?** If so, which devices fulfil this function?

The qualification of the target audience and its expected credibility codes, produced at this stage, is a working hypothesis, not a conclusion. At the end of Step 3, the analyst must explicitly return to this hypothesis and answer: *“Does the detailed critical examination confirm, qualify, or invalidate the qualification of the audience and its expectations made in 0.f?”* If the examination in Step 3 provides no element likely to modify the initial qualification, the analyst must explicitly flag this as a possible limitation rather than automatic confirmation.

Step 1 — Analysis of the Opening Device

Isolate the opening segment of the text. For each significant element of this segment, answer systematically:

1. **Undemonstrated presuppositions.** What does the text take as already established, obvious or shared, without ever submitting it to examination?
 2. **Pre-emptive disqualifications.** Does the text rule out an opposing position in advance? Is that position actually defended by an identifiable person, or is it a construction by the author designed to be easily refuted (straw man)? An opposing position is a straw man when the text omits its strongest arguments and retains only its weakest or most marginal version, in violation of the principle of charity. The finding concerns the reconstruction carried out, not the author’s intention: any intention belongs to Step 4 and its three-level progression.
 3. **Perimeter of the sayable.** What questions does this framing make possible? Which ones does it render inconceivable, trivial or morally suspect?
 4. **Markers of authority.** Does the text invoke a vague but purportedly indisputable instance (“the theoretical point”, “history”, “science”, “everyone knows that”, “any informed person”) to close the debate before opening it?
 5. **Implicit positioning of the author.** Do they present themselves as correcting an error, revealing a hidden truth, speaking on behalf of a community of informed readers, saying what no one dares say?
 6. **Announced programme.** Does the text set itself an explicit programme (what it promises to do, what it declares it will not do)?
-

Step 2 — Neutral Reconstruction of the Argumentation

Reconstruct the central thesis and the argumentative path without evaluation, faithfully. You must be able to answer: *What does the author want me to believe, and through what chain do they want to lead me there?*

Checkpoint, activation of Appendix 2. The decision to activate Appendix 2 is taken here, not at Step 0: it presupposes that the central thesis has been reconstructed. Take the decisive terms of the thesis you have just stated, those whose removal would make it unintelligible, and apply two tests. Reading convention, valid for both: a test is **positive** when it reveals the defect sought, not when the text satisfies the requirement. A positive definition test therefore means that the term is not defined.

- **Definition test:** does the text provide, for this term, a definition, criterion of application or enumeration of what it covers? This is a finding of presence or absence, without assessing the quality of the definition. An evocative characterisation that gives no criterion by which one can tell whether the term applies to a given case does not satisfy the test: the test is then positive, despite the abundance of commentary the text devotes to the term. This is the hardest case to see, because the text appears explicit about a word whose boundary it never fixes.
- **Boundary test:** name two possible delimitations of the term, one restrictive and one extensive, both compatible with the way the text uses it. Restate the conclusion under each. If the conclusion holds under one delimitation and fails under the other, in either direction, the term's boundary carries the argumentation, and the test is positive. If the conclusion holds under both, the term is vague without being strategic, and the test is negative. This is the substitution test of Step 7.a applied to the extension of a term rather than to an item of evidence: the minimal correction concerns not a piece of the dossier but the boundary of a word.

Example of a positive test: a text argues that education and the media perform a propaganda function, without defining “propaganda”. Restrictive delimitation: the structural effect of ownership and funding on the production of information by a commercial media outlet. Extensive delimitation: any institution seeking assent, advertising and entertainment included. Under the second, the conclusion holds but the word no longer names a mechanism. Under the first, the word names a mechanism but the conclusion no longer covers education. The argumentation needs both delimitations at once, which neither provides alone. Positive test.

Example of a negative test: the same text uses “fallacy” without an extensive definition. Whether the term is delimited broadly or narrowly, no conclusion in the text changes. Vague term, not strategic. Negative test.

If both tests are positive for at least one decisive term, activate Appendix 2. If only the definition test is positive, flag it in one line and do not activate the appendix: an argumentative text may use abstract terms without defining them, and the appendix is justified only when the boundary of the term carries the conclusion. If only the boundary test is positive — the text defines the term, but the conclusion nevertheless changes depending on the delimitation retained — activate the appendix: it is the boundary, not the absence of definition, that controls activation. A supplied definition does not immunise a term whose reach determines the conclusion.

Record the decision in the analysis, including when it is negative, naming the term examined and the test that failed. An unreasoned non-activation is a failure, on the same footing as an undeclared absence of external verification (Rule 13).

Checkpoint, modal map (introduced in V5.1.0). Alongside the preceding checkpoint, draw up the modal map of the load-bearing propositions. For each proposition whose removal would make the thesis unintelligible or unproven, note two things as the text poses them **at that point**: its **polarity**, P or ¬P, and its **degree of commitment** to that proposition — assertoric (posed as fact), probable, hypothetical, or suspended (judgment left in suspense). Suspension concerns the degree of commitment to the target proposition; it is not a third polarity: “it is not known whether P” remains of polarity P, at the suspended degree. The annotation is neutral: one records, one does not yet judge constancy. A concession is a **location** where a low degree of commitment may appear, not a degree in itself. If the same load-bearing proposition, on the same referent and at constant polarity, appears at several points under different degrees of commitment, note each occurrence and its location — opening, body, concession, or closing reservation. Only these propositions, whose degree of commitment varies at constant polarity and referent, will be taken up again in the modal constancy check of Step 6. The pass is limited to load-bearing propositions; it does not tag every sentence.

Step 3 — Systematic Critical Examination of the Argumentative Body

Take up each link in the argumentation and subject it to the following tests. Begin by restating the standard of proof declared in 0.c.

A. Logical Validity

- Do the conclusions follow from the premises?
- Are there unjustified inferential leaps or shifts in meaning?
- Are causal chains demonstrated or merely postulated?

B. Empirical Soundness *Scope:* The analysis first relies on elements internal to the text. External verification is nevertheless not optional.

Minimum external verification. Where an element is decisive to the demonstration and a simple check would establish it, that check must be attempted. This applies in particular to:

- attribution of a decisive quotation (author, work, date);
- identity, nature or position of a named source (institution, body, office);
- existence, exact designation or nature of an entity cited as evidence against someone;
- an isolated factual point on which an inference depends, and which a single search would suffice to establish.

The following conditions accompany every external verification:

- explicitly notify that the analysis is moving outside the strict scope of the text;
- ensure that the verification concerns the object declared at the head of the analysis, not another version or edition of the same text;
- cite the external source (public document, study, contradictory statement, etc.);
- acknowledge that the original text did not contain this information.

Condition for counting. An external verification counts as such only if the analysis gives the source actually consulted in a form that allows a third party to find it again: a link, or an identifying reference containing the author or institution, title and date of the source, together with the date of consultation when the source may change. An attachment supplied by the user is an admissible source, identified by its title, date and the location of the passage within the document.

Knowledge remembered by the model does not constitute external verification, even when accurate, and can neither be counted nor presented as such. The prohibition concerns what the analysis counts and declares as verification, not ordinary use of general knowledge in the body of the reasoning. Three forms must be explicitly excluded: a claim of checking without a source; a source cited for content it does not support; and an actual but unsuccessful search whose result is then completed from memory under cover of the consulted reference.

Meaning of “external”. In ARGUS, external verification is external to the **analysed object**, not necessarily found online. An attachment supplied with the article is a source external to the text analysed, even if no search was required. Conversely, a section belonging to the same declared object remains internal to that object even when excluded from the ARGUS scope in 0.a bis: confronting it with the scope does not, for that reason alone, become external verification.

Two independent axes to declare. Each check is described by two distinct characteristics, which must not be presented as a list of categories at the same level.

Unit of a check. A check corresponds to a verificatory proposition or a protocol condition formulated at the moment the check is undertaken. Sub-elements necessary to decide that same proposition are not counted separately. If some are established and others are not, the check is partial. Two checks are distinct only if they bear on two propositions each producing an autonomous result for the analysis. The partition may not be altered after the fact in order to transform one partial check into one successful check and one unsuccessful check. Example: establishing whether a note was publicly accessible on a given date may require establishing its existence and its date of availability; if existence is established but accessibility at the relevant date is not, the accessibility check is partial.

- **Status of the check:** *successful*, sourced under the condition above; *partial*, with the established and unestablished parts named; *unsuccessful*, the check having been conducted without establishing any of the elements

sought. This status does not presuppose online research: a provided source that establishes none of what was sought yields an unsuccessful check, and the two axes remain independent.

- **Mode of access to the source:** *provided source*, confrontation with an attachment or document already present; *external search*, a source that had to be sought out.

An unsuccessful check does not satisfy the counting condition, since it establishes none of the elements sought: it is declared separately and not included in the number of successful checks. This also applies when a source was produced but does not support anything that was being sought — the result, not the production of a source, determines the status. It must nevertheless be declared, with the object of the search. The mode of access is declared because neither the cost nor the reproducibility for a third party is comparable between the two.

Symmetry of verification. Checks do not concern only claims the analyst suspects. If a quotation unfavourable to the text is checked, a quotation that credits the text must be checked too. A well-formed, paginated reference of impeccable appearance is the first place where the analyst relaxes their requirement, and credit granted to a source on the sole basis of the form of its reference falls under deference, addressed in Step 8(d).

Absence of result. An unsuccessful check must be reported as such. Absence of result does not establish that the claim is false, and the analysis must distinguish the two.

Mandatory declaration. Every analysis states three distinct numbers and the elements to which each refers: successful checks, partial checks, unsuccessful checks. Used on its own, “number of external verifications” means only successful checks, which are those the counting condition allows to be counted. A partial check is declared in its own count, and the established part is named. If the number of successful checks is zero, the analysis must state and justify this. No analysis may dispense with this declaration, and absence of verification is not a silent default regime.

Function of the three numbers. These counters are an internal execution trace within the analysis, not a metric for mechanically comparing two independent analyses. Two runs may formulate different verificatory propositions at the outset and obtain different numbers without either being at fault. The reproducibility requirement concerns a single run: the announced partition must remain stable, each status must be justified under the unit of check above, and the consistency check in Step 8 must recover the same numbers. Comparison between analyses concerns the propositions actually checked and their results, not raw equality between the three counters.

Proportion. External verification remains bounded: it concerns decisive elements, not every claim in the text, and does not replace internal examination. Beyond around ten checks **for the same analysed object**, flag that the exercise is tending towards fact-checking, which is not the principal purpose of ARGUS. In a corpus analysis, this threshold is assessed separately for each text or other individual object to which the checks are attached; the total of all checks across the corpus may be reported as an inventory, but does not by itself trigger this warning. Cross-text consistency checks required by Appendix 3 are not artificially added to a text’s checks to cross the threshold. This threshold is a proportionality warning, not an exemption from Rule 13: simple checks that Rule 13 makes mandatory on decisive elements must still be attempted. Additional checks that do not meet this criterion belong, where relevant, to separate fact-checking work.

This verification is not intended to “condemn” the text, but to assess whether the unsubstantiated claim is otherwise known to be true, false or uncertain. The analysis must clearly distinguish internal examination from external verification.

If a fact central to the thesis does not meet the standard of proof declared in 0.c, it must be classified as unsubstantiated. It is forbidden to relax the standard of proof during the analysis in order to rescue the coherence of the text.

Within this framework:

- Are factual claims substantiated within the text, or simply asserted without proof?
- Are events dated precisely within the text, or is the chronology vague, compressed or even contradictory?
- If a typology or hierarchy is established among phenomena, is each term described accurately and completely within the text? Are facts or mediations that the text itself mentions elsewhere then omitted or distorted in service of the thesis?
- For each decisive quotation: are the work, date and page given? Is the context of utterance the one the analysed text attributes to it? A quotation is **recontextualised** when it is authentic but moved to a scene it did not address.

This flaw survives literal verification, because the words are exact, and is detected only by going back to the source.

- **Fidelity to an available primary source.** The primary, official, contemporary or directly relevant character of a supplied item is not sufficient to activate this test. It activates only when the analysed text itself relates to that source — it cites, names, summarises, comments on or criticises it — or when the analysis can establish that the checked claim reproduces its content. Mere concordance of subject, date, wording or figure is not enough to establish that lineage. Otherwise, the item remains an ordinary external verification: it may corroborate, contradict or fail to corroborate a claim in the text, but the fidelity qualifications below are not applied. When the link to the source is established, decisive claims concerning it must be confronted with it. The check is narrow: it concerns claims on which the thesis of the text or its credibility contract depends, not all its sentences. For each, qualify: **faithful**, the meaning, scope and degree of certainty are preserved; **strengthened**, a hypothesis, reservation, correlation or limited finding is transformed into a more general or more certain claim; **weakened**, the text attenuates or restricts what the source establishes; **shifted**, the words or data are exact but attributed to a referent, comparison, period or scene the source did not address; **inaccurate**, the source says something else on a verifiable and decisive fact; **own interpretation signalled**, the text adds a reading that is not the source's and explicitly presents it as its own; **not decidable**, the available source does not permit a conclusion. The finding “strengthened” deserves particular attention because it can produce a claim more assured than its source while remaining literally defensible. Its severity is determined exclusively in 7.a by the substitution test, not by its qualification. These qualifications are not necessarily exclusive: a claim may be both strengthened and shifted. When they cumulate, name the principal qualification and the secondary ones. “Faithful” and “not decidable”, however, are incompatible with any defect qualification concerning the same element. Three of the seven qualifications are not defects and do not carry forward to 7.a: **faithful**, **own interpretation signalled**, and **not decidable**. The first records conformity, the second records that the text owns a reading instead of attributing it to its source, and the third records a limitation of the analysis rather than of the text. An available primary source does not by itself activate Appendix 3 or this fidelity test. If the primary source to which the text relates under the conditions above is unavailable, declare this and name the claims for which the boundary between what the source says and what the text adds remains undecidable, which falls under the “When elements are insufficient” clause in Step 7.
- If the text contains figures that are load-bearing within the meaning of §2 of Appendix 4, activate that appendix and report its findings in a dedicated subsection of 3.B entitled “Arithmetical and statistical control (Appendix 4)”.

C. Use of Symbolically Charged Terms

- Identify words that single-handedly summon a massive emotional, cultural or historical register (bread, blood, land, the people, homeland, freedom, etc.).
- Ask whether these words replace demonstration with evocation, and whether their affective force immunises the argument against contradiction.
- Do these terms describe reality adequately, or are they inadequate while rhetorically powerful? Whom do they really address?

D. Use of Terms Claiming Universality

- When the text uses “the human species”, “planetary”, “everyone”, “we”, identify to whom these terms actually refer.
- Does the author speak on behalf of an abstract humanity while describing realities that concern only a specific part of that humanity?
- Is the “we” a mask for a particular audience (Western, European, etc.) presented as universal?

E. Test of Absences *Stage 1 — Identification*

- What other explanations of the same phenomena are absent from the text?
- What facts, actors or connections, if taken into account, would weaken the thesis?

Stage 2 — Strategic Qualification

For each absence identified, answer successively:

1. Is this information easily accessible (public sources, official statements, known data, including in other texts in the same corpus where applicable)? **Object of accessibility:** the criterion concerns the absent information or proposition, not merely the availability of an interlocutor, institution or source capable of producing it. A journalist's presence at an event, the physical accessibility of an actor, or the possibility of asking them a question do not in themselves establish that a specific piece of information was readily accessible. To answer "yes", the analysis must be able to name the accessible information and the medium, statement, data or fact that made it available. Accessibility is assessed at the publication date of the object declared at the head of the analysis, not at the date of the analysis. In a corpus, another text counts here only if it is established that it was already publicly accessible **before the moment of publication** of the object analysed, or that a distinct earlier source already made the information accessible. Mere identity of date, belonging to the same edition, presumed availability within the same newsroom, or the fact that both texts were eventually published together are not sufficient. If only the date, without time or order of release, is known for two texts from the same day, their relative precedence is not established and the other text cannot by itself ground a strategic omission. A later text may inform the analysis, but does not make information retroactively accessible. An element published after the text cannot be blamed on it as an omission. It may, however, serve to qualify the modality of a claim — a thesis presented as an established fact when it was contestable at the time. A defect of modality and a defect of omission are not the same and are not counted together in 7.a. If accessibility at the relevant date cannot be established, do not treat it as either established or refuted: point 3 provides the outcome "first-gate qualification not determinable as things stand", under the strict conditions it sets.
2. Is the information actually absent from the text, including in the form of a mention intended to dismiss it? The two questions have the same polarity: two "yes" answers establish the omission.
3. Decide according to the two answers. Criteria 1 and 2 both met: **strategic omission**. Criterion 2 not met, meaning the text mentions the element: there is no absence, and the treatment of the mention is examined under Rule 16, which distinguishes formal from effective counter-argument. This criterion is treated first: if the text mentions the element, the question of accessibility no longer arises, since there is nothing absent to qualify. If criterion 2 is met and criterion 1 is not, distinguish three cases:
 - the information existed at the publication date but was not readily accessible: **mere point of weakness**. The text is weaker without it, without anything indicating selection;
 - the information was not available at that date, notably because it was published later: **non-attributable absence**. This is not a defect, does not carry forward to 7.a, and is flagged in one line because it limits what the analysis can establish, not what the text should have done;
 - the analysis cannot establish whether the information existed or was readily accessible at that date: **first-gate qualification not determinable as things stand**. This is not a defect attributed to the text and does not carry forward to 7.a. It is the only one of the three outcomes that concludes nothing, and it is bounded by three cumulative conditions: a check reasonably necessary to decide the issue was attempted and did not permit a decision, its result being partial or unsuccessful within the meaning of 3.B, or the analysis declares and justifies that no check was accessible under its execution conditions, under the zero-verification regime provided by 3.B — and where the absence concerns a decisive element and a simple check is possible, the attempt is mandatory under Rule 13; the information that would be needed to decide is named; and the indeterminacy concerns what the analysis can establish, not what the analyst has failed to decide, within the meaning of the first boundary of Step 7. An uninstructed uncertainty is not indeterminacy: where a reasonably accessible check would allow a decision and has not been attempted, the qualification is incomplete and the analysis cannot be delivered in that state — the check is attempted before delivery. This clause prevents evasion; it never requires choosing between two qualifications when it is unknown which is true, which Rule 18 prohibits.
4. An absence is a **structural omission** only if the analyst explicitly demonstrates that the absent element, if integrated, would directly contradict the conclusion or require a major overhaul of the causal mechanism defended.
5. The qualifications are passed through two successive gates, with two different criteria, and the order is not optional. The first gate is accessibility, coupled with the finding of absence; once the absence is established, it discriminates the outcomes in point 3. Criteria 1 and 2 met are sufficient to establish strategic omission, and it is forbidden to reject this qualification on the ground that the absent element would not destroy the thesis, a consideration that belongs to the second gate. Non-attributable absence and first-gate qualification not determinable as things stand are not defects and do not carry forward to 7.a. The second gate is the effect on the conclusion, and

opens only for an absence already qualified as strategic.

6. Case of an absence with indeterminate result: the element is missing as necessary evidence for a central proposition, but its result could equally confirm or refute the thesis. Indeterminacy is an additional property, not a competing qualification with those in point 3: retain the qualification obtained at the first gate and add the note “indeterminate result”. Its only effect is to bar the second gate, since the condition in point 4 cannot be demonstrated. State it in one line: the evidence is missing, and its result is indeterminate.

F. Falsifiability

- Is the thesis formulated in a way that can be refuted, or is it immunised against any contradiction?
- Can the text account for a fact that would contradict it, or would it have to ignore or distort that fact in order to survive?

Beyond the form of the thesis, examine how the text uses it.

- List the observations the text treats as confirming.
- Ask whether an observation of the opposite sign is treated as confirming as well. If so, qualify the flaw as **unfalsifiability in use**, distinct from **formal immunisation**: the thesis is refutable in its wording but absorbs every observation in its application. This form of immunisation survives an examination confined to the wording of the thesis, since that wording is refutable: the defect appears only in use.
- Ask explicitly: what observation would the text acknowledge as refuting? If the text names none, flag this.

G. Stylistic Register and Readability

- Identify sentences whose syntactic complexity is such that they cannot be understood on a first reading and are incompatible with orality.
- Ask whether this complexity is functional (it expresses a real complexity of reality) or rhetorical (it produces an effect of intellectual authority, filters the readership, or masks a logical impropriety).
- Confront the linguistic register with the stated addressee of the text. If the text claims to address “everyone”, does it speak a language genuinely accessible to everyone? The gap between the proclaimed “we” and the actual register is a performative contradiction that must be noted.

H. Epistemic Symmetry When the text attributes a flaw to an object (institution, group, theory, technology, etc.) — such as opacity, bias, irrationality, lack of legitimacy, need for external justification, or any other failure to meet a standard of rigour — systematically apply the following test:

- Can this same flaw be found in the speaker, in the group or instance they privilege, or in the device of the text itself (its framing, omissions, presuppositions)?
- If so, does the text acknowledge it? Does it draw the consequences?
- If not, is the difference in treatment explicitly and validly justified?

Apply the same test to interest-based explanations. If the text explains an opposing position by the interests, institutional position or funding of whoever defends it, does it apply the same reading to its own authorities? An explanation by interests that applies to only one side is an unjustified asymmetry, even where the theoretical framework of the text legitimises it in principle.

Flag every unjustified asymmetry as an argumentative weakness (performative contradiction, unwarranted epistemic privilege or reflexive blindness).

I. Informational Circularity Test

Answer the following questions:

1. **Do all the sources in the text that support the central thesis belong to the same camp or ideological ecosystem?**
2. **Are there sources that could contradict the thesis, and if so, are they cited, summarised or refuted?**
3. **Does the text present the absence of contradiction as proof of the validity of its thesis?** (This is a fallacy: the absence of counter-argument in a selected sample proves nothing.)

Flag every informational circularity as an argumentative weakness.

Note: the same underlying phenomenon (for example, a homogeneous body of sources) may manifest itself both in Test 3.I (circularity) and Test 3.E (absences). The analyst must identify it as one underlying cause with two manifestations, rather than count two independent flaws in the overall judgment of Step 7.

J. Comparative Analysis of Treatments For each actor or group mentioned, systematically compare:

1. **What vocabulary is used to qualify them?** (neutral, laudatory, derogatory, ironic, etc.)
2. **Are their statements subject to verification or repeated as facts?**
3. **Are their actions described in concrete detail or vaguely?**
4. **Are their motivations explained or simply denounced?**

Then answer the following question:

- **Is there a systematic asymmetry in the treatment of different actors?** If so, does the text justify it explicitly? If no valid justification is provided, flag this asymmetry as an argumentative weakness.

K. Target-Audience Adaptation Test Assess whether the text adapts its persuasive devices to the expectations of the audience identified in 0.f.

Answer systematically:

- **Does the text use the markers of seriousness expected by its target audience?** Numerous sources, experts, historical references, analytical tone, lexical caution, figures, concessions, syntactic complexity, apparent distinction between facts and interpretations.
- **Do these markers genuinely strengthen the demonstrative quality of the text, or mainly serve to make an already fixed conclusion acceptable?**
- **Is the counter-argument formal or effective?** A counter-argument is effective if it forces the text to modify, limit or reconfigure its thesis. A counter-argument is formal (or decorative) if it is mentioned but immediately absorbed, marginalised or rendered harmless.
- **Are concessions opening or immunising?** An opening concession makes a different or more uncertain conclusion possible. An immunising concession gives an appearance of balance while protecting the central conclusion from criticism.
- **Are the omissions already identified in 3.E precisely those that the target audience identified in 0.f would need to know in order to assess the thesis autonomously?**
- **Does the text produce interpretive closure in the form of complexity?** In other words: does it give the reader the feeling of having traversed complexity, while the elements introduced never really threaten the conclusion? In a composite text, this test concerns the argumentative scope declared in 0.a bis, not the technical or didactic sections. A real technical difficulty that serves no conclusion is not interpretive closure.
- **Where is rigour concentrated?** Locate separately, in the text, claims accompanied by a precise reference (work, date, page, named source, declared method) and unreferenced claims. Then compare the two locations. Do the precise references bear on claims the target audience can verify themselves at low cost, or on those carrying the thesis? Rigour concentrated on the verifiable and absent from the decisive functions as a credibility device rather than substantiation. It produces an effect of scholarship broader than the actual sourcing, and is all the more effective when the few references supplied are impeccable.

Recommended terms in the analysis:

- **Effective counter-argument:** presence of an objection genuinely taken seriously.
- **Decorative counter-argument:** presence of an objection that does not alter the argumentative trajectory.
- **Immunising concession:** nuance that mainly serves to avoid the accusation of bias.
- **Locked nuance:** nuance integrated in such a way as to reinforce the conclusion rather than open it.
- **Adapted propaganda:** propaganda that adopts the intellectual, moral or aesthetic codes of the audience it targets.

- **Displaced rigour:** documentary precision concentrated on the claims easiest to verify and absent from those carrying the thesis.
-

Step 4 — Inference of Strategic Intent

An argumentative text is not merely an arrangement of propositions; it is an act in the world that aims at effects on a readership. Identifying a strategic intention must follow a strict progression.

Preliminary — Stated Purpose

Begin by identifying the explicit purpose of the text: what does it say it wants to do? (Convince, denounce, mobilise, explain, bear witness?) This stated purpose will serve as a point of comparison for everything else in the analysis.

Level 1 — Textual Clues

Note, without interpreting them, every element of the text that could indicate a strategic intention not explicitly declared:

- marks of distinction or positioning within a field (references to schools of thought, to “wrong ways” of thinking, to unnamed adversaries);
- addresses to a specific audience (marks of complicity, shared assumptions, a “we” whose actual perimeter must be identified);
- silences, omissions, expected themes that are avoided;
- insistences, repetitions, formulations that appear to aim at a precise effect on the reader (guilt, mobilisation, legitimisation of the author);
- the gap between linguistic register and stated addressee;
- tensions, contradictions or ambiguities suggesting that the text cannot openly own everything it does.

This level is purely descriptive. No conclusion is drawn at this stage.

Level 2 — Strategic Hypothesis

From the clues identified, formulate one or more hypotheses about the inferable purpose of the text, that is, what the text seeks to do beyond its stated purpose. Each hypothesis must be presented as such (“one may hypothesise that...”, “the text might also aim to...”) and explicitly linked to the Level 1 clues that support it.

Hypotheses may concern:

- the actual addressee (to whom is the text really speaking? Does it address the already convinced to reinforce them rather than the undecided to persuade them? Does it speak to peers to mark a position in an intellectual field rather than to the general public to mobilise them?);
- the author’s legitimisation function (does the text serve to confer on the author a position as expert, theorist, moral figure or spokesperson?);
- the intended effect on the reader (inform, move, make guilty, mobilise, paralyse, reinforce, silence) and its consistency with the stated purpose;
- an operation of distinction or positioning within an intellectual or political field;
- an implicit strategic function (for example: disqualifying a competitor without naming them, distancing oneself from a more radical or more moderate position, protecting an institutional position, producing a legitimisation effect without stating it explicitly).

Level 3 — Degree of Confidence

For each hypothesis, assess the degree of confidence the text permits, using a simple scale:

- **High:** the hypothesis is supported by convergent clues and no element of the text contradicts it.
- **Medium:** clues exist, but other interpretations are possible.
- **Low:** the hypothesis rests on tenuous or ambiguous clues; it is plausible but far from certain.

If no hypothesis reaches a high degree of confidence, say so explicitly. Strategic intention is not always inferable, and it is more honest to suspend judgment than to force an interpretation.

Step 5 — Internal Consistency Check

Confront the programme announced in the opening (what the text says it will do) with the actual content of the body.

- Does the text do what it said it would do?
 - Does it refrain from what it said it would refrain from?
 - If a gap is found, the opening declaration must be requalified: it functions as a device of self-legitimation. The gap establishes that function; it does not by itself allow a conclusion about the author's sincerity, and the analysis must not claim otherwise.
-

Step 6 — Return to the Opening Device

Re-read the conclusions of Steps 3, 4 and 5 in the light of Step 1.

- Did the initial framing predetermine or protect the argumentation by neutralising certain objections in advance?
- Does the thesis survive if the opening postulate is rejected?
- Are the arguments independent of the framing, or do they depend on it entirely?
- Does the inferred strategic intention shed new light on the function of the opening device?

Modal constancy check (introduced in V5.1.0). From the modal map drawn up at Step 2, examine each load-bearing proposition whose **degree of commitment varies at constant polarity and referent**. This check compares the degree of commitment to one and the same proposition, not its polarity: it captures an inconsistency that the examination of logical validity (Step 3) lets through because it is not a contradiction, but a weakening of the degree of commitment.

Apply the **direction test**:

1. Is the **strong degree** — assertoric, or the higher of the two degrees observed — the one that does the argumentative work, the one under which the proposition serves as a premise, reinforces the thesis, or charges an opponent?
2. Does the **weak degree** — suspended or hypothetical, or a retreat to “probable” from the assertoric — appear where strong commitment would carry a cost: section on divergences, concession, closing reservation, reply to an objection?

Two “yes” answers: **directional modal retraction**. The text posits the proposition as settled where it pays and withdraws it where establishing it would cost. Severity is determined in 7.a by the substitution test, subject to the three exclusions below.

Exclusions (extension of Rule 14). Each one, if it applies, is not a defect and the check remains silent.

- **Consistent restraint.** The proposition carries the same cautious degree of commitment throughout; the closing reservation merely restates that consistent caution. Check: is the opening already cautious? If so, there is no retraction. This is the safeguard that Rule 14 already protects.
- **Acknowledged revision.** The author posits the proposition at the strong degree, then explicitly revises it **and effectively ceases to rely on its force to conclude**, adjusting the conclusion. The two conditions are cumulative. Explicit acknowledgment is not enough: the revision does not hold if the conclusion re-asserts the same force under a narrower, reworded, or re-framed formulation, in particular if it maintains the recommendation or the conclusion that the strong assertion supported. Acknowledging that an assertion was “too absolute” and then continuing to draw the same benefit from it is not an acknowledged revision but a disguised retraction: the check triggers. Application check: before granting the exclusion, verify that the conclusion does not carry the force of the strong pole again, verbatim or reworded; if the recommendation that depended on the strong pole survives intact, the exclusion does not apply.
- **Scope partition.** The two poles do not bear on the same referent or the same perimeter. Name the referent of each pole explicitly. If the text claims to distinguish two referents — a surface and a substratum, an aim and an efficacy — that distinction must be demonstrated separately; asserted without demonstration, it does not bring the exclusion into play. Case of an aim and an efficacy: the clean partition — aim held throughout and only efficacy left in suspense — is silent by construction, the aim there having a constant degree of commitment; the

exclusion is therefore invoked only if the aim itself varies, and one then verifies that the body has not already rendered the aim uncertain, which would forbid the reading “aim settled, only efficacy discussed”.

Boundary with the examination of validity (Step 3). A change of **polarity** at constant strong degree — assertoric-P then assertoric-¬P on the same referent — is not a modal retraction but a **contradiction**, reachable by the examination of validity at Step 3; the present check does not handle it and can only flag its directional structure. The check bears exclusively on the **weakening of the degree of commitment** of a proposition whose polarity and referent do not change. Such a weakening escapes the examination of validity for want of a strict contradiction: this is the check’s own contribution. The most useful case, disguised retraction as a frame shift — the strong proposition posed on one plane (teleology, tendency, or intention) and weakened on another (content or present fact) — the check imposes the demonstrate-or-withdraw alternative.

Modal strengthening, weak to strong. The reverse symmetric case — a load-bearing proposition cautious at the outset and asserted in the conclusion — does not fall under this check. Its sourced and quantified forms are already captured, one by the fidelity test of 3.B, the finding “strengthened”, the other by Test E of Appendix 4. The remaining form — the author strengthening their own unquantified assertion without the body having provided the new support that would justify the escalation — is flagged under the examination of validity (Step 3) as an unwarranted escalation, an excess of the conclusive over the premises, reusing the modal map of Step 2.

Reporting. One modal-constancy line for each load-bearing proposition whose degree of commitment varies, with the verdict — directional modal retraction, or exclusion applied and which one. An established defect carries forward to 7.a. A concession that is the retracted pole of an established retraction cannot be credited to the text in 7.b. This check does not distinguish a retraction by a single author from a retraction arising from the juxtaposition of heterogeneous strata of the text: that distinction bears on a history of production invisible in the object, and inferring it would amount to speculating about authorship, which Step 4 forbids. The object carries the contradiction whatever its making. Where the discontinuity is manifest in the object itself — a caption or an image asserting what the body concedes, a stylistically discontinuous box — the analyst may add a purely descriptive note locating the contradiction, without imputing any authorship or intention.

Declared limitation. The boundary between acknowledged revision and disguised retraction remains, in the last resort, a matter of judgment: a text may acknowledge an over-strong assertion skilfully enough that maintaining the conclusion on its force passes for a legitimate retreat. The check may then under-trigger. This failure is conservative: it errs on the side that protects honest revision. A verdict rendered near this boundary is given with signalled confidence and calls for re-reading, in accordance with the reliability note.

Step 7 — Differentiated Conclusion

When elements are insufficient: if a section of the analysis cannot be filled in reliably because the text does not provide enough information, the analyst must state this explicitly. Formulations such as “Insufficient data to conclude”, “No probative elements in the text”, or “This question cannot be settled on the basis of the analysed text alone” are preferable to an artificially complete analysis. Acknowledging a limitation does not discredit the analysis; on the contrary.

Application to the three sub-steps:

- **In 7.a:** if the severity of a flaw cannot be established by the substitution test for lack of information, the flaw is listed with the note “severity not determinable as things stand”, rather than placed by default in the “minor” category.
- **In 7.b:** if the text establishes nothing solidly, write it. The section is mandatory; filling it is not. Crediting a text with qualities it does not have in order to honour the section is a complacency symmetrical to artificial severity.
- **In 7.c:** if the elements gathered in 0.d, 0.e, 0.f, 3.I, 3.J, 3.E and 3.K — the seven steps on which 7.c is built — do not permit a level to be assigned, write “level not assignable as things stand”, specifying which element is missing. Do not default to Level 1.

Three boundaries limit this clause.

Insufficiency of the text, not indecision of the analyst. The clause covers the case where the text does not provide the information. It does not cover the case where the analyst has formed no judgment. Confusing the two amounts to

using the clause as an evasive moderator, which Rule 14 forbids. The test is simple: can the missing information be named? If the analyst can write “X is missing”, the clause applies. If they can only write “it is difficult to say”, Rule 14 applies.

Insufficiency of information, not exclusion from scope. A section left unfilled because the corresponding section is outside the declared scope (0.a bis) must be flagged as outside scope, not as missing data. The two statements do not tell the reader the same thing.

No escape from calculation. Insufficiency of information does not cover tests that the text permits to be executed with what it provides. In particular, Test A of Appendix 4 (order of magnitude) cannot be declared unfillable because of the cost of the calculation or uncertainty over the order of magnitude. It may be unfillable only when the minimum data necessary for the calculation or for relating it to a relevant reference set are missing, under the conditions of §5 of Appendix 4.

7.a – Severity Scale for Flaws For each argumentative flaw identified, assign a severity level:

- **Fatal:** the flaw invalidates the central thesis (e.g. major logical contradiction, structural omission without which the thesis does not hold, manifest falsification).
- **Major:** the flaw seriously weakens the thesis but does not completely invalidate it.
- **Minor:** the flaw does not affect the overall conclusion of the text.

Substitution test for severity: To distinguish a fatal flaw from a major one, mentally construct the minimal version of the text in which the flaw is corrected: remove an unjustified claim, restore an omitted piece of information, replace a distorted claim with its correct version, or repair a logical leap without adding an argument foreign to the text. Then ask: *can this amended version maintain the central conclusion without changing the nature of its claim (for example, moving from demonstration to mere opinion, or from factual assertion to hypothesis)?* The test concerns the minimal correction of the flaw, not always removal of the element through which it appears. A quantitative flaw follows a specific rule, depending on whether the correct value is available. When it is established by the data in the text or by an admissible source — the text writes 150% where its primary source gives 15% — the minimal correction is replacement with that value: this is not the analyst’s model but an established datum, making the counterfactual more faithful. When it is not established — an impossible or unreconstructible figure without an established replacement value — the minimal correction is removal of the quantitative proposition the figure purported to establish, not fabrication of a plausible value: the prohibition in §5 of Appendix 4 against substituting the analyst’s own model for the author’s applies here as elsewhere.

- If the central conclusion can no longer be maintained under the same claim: **fatal**.
- If the conclusion can be maintained but with a clearly weakened claim (less certainty, less generality): **major**.
- If the conclusion and its claim remain unchanged: **minor**.

Application to modal retraction (V5.1.0). When the flaw is an established modal retraction (Step 6), the minimal version forces the degree of commitment of the proposition to a single, constant value, and the forcing bears on the claim **as it is announced by the strong pole**, not on a charitably weakened version of the thesis. If the strong pole announced a factual assertion or a demonstration and the forcing brings it back to a hypothesis, the announced claim changes in nature: fatal for that claim. If a weaker thesis survives elsewhere, its severity is stated separately, often major. The two levels are named where they diverge.

7.b – What the Text Solidly Establishes (mandatory) Draft a distinct, substantial section listing the points on which the text succeeds: precise information, valid reasoning, internal coherence, legitimate rhetorical qualities (if they are not deceptive), etc. This section is not a decorative concession: every genuinely solid element is developed in a proportion that follows their number and importance. The requirement is not equality of length with the section on failures, which would force the analyst to credit a text with qualities it does not possess.

Interaction with the modal constancy check (V5.1.0). A concession that constitutes the retracted pole of an established modal retraction (Step 6, exclusions set aside) cannot be credited to the text in 7.b — neither as an effective counter-argument, nor as an opening concession, nor as a point of rigour: this would credit the very half of the flaw identified in 7.a. The consistency check in Step 8 verifies that no quality credited in 7.b coincides with a retracted pole.

Then distinguish explicitly:

- **The real qualities of the text**, as listed in 7.b, if they exist independently of its argumentative flaws. Caution: do not call “rhetorical force” an effectiveness obtained through fallacies. Do not call “internal coherence” the circular closure of a system upon itself. Do not call “ability to name real phenomena” a selective evocation that omits contrary facts.
- **The precise argumentative failures**, with their severity level (7.a).
- **The inferred strategic intentions**, presenting them as hypotheses based on the analysis, not as certainties. Recall the confidence level associated with each hypothesis.
- **Articulation between inferred intention and overall judgment**: If a significant gap is found between the stated purpose and the inferred purpose, what does it reveal about the nature of the text? A text can succeed as a political act (positioning, mobilisation) while failing as a demonstration. Inferred intention does not refute the thesis, but it can explain its argumentative weaknesses: an author pursuing a strategic objective may sacrifice rigour to rhetorical effectiveness, not by accident but through consistency with their actual purpose. Do not confuse analysis of the strategic function of the text with refutation of its thesis.
- **Unjustified asymmetries**: where applicable, list the flaws the text attributes to its object but does not acknowledge in itself or its own camp, without valid justification.
- **Overall judgment**: is the thesis established? Does the text fulfil its demonstrative claim? If the text fails as a demonstration but functions as a manifesto or act of positioning, say so explicitly, without conflating these two registers of evaluation. Mention whether any fatal flaws were identified.

7.c – Qualification of the Degree of Propaganda From the observations in Steps 0.d, 0.e, 0.f, 3.I, 3.J, 3.E and 3.K, assign a level on the scale below.

What the scale measures. The level measures the **degree of interpretive closure**, that is, the freedom the device leaves the reader to examine the conclusion, oppose to it the elements present or accessible, and modify it if contradiction requires. It measures neither the respectability of the author, nor their institutional status, nor how activist the subject is. A state, corporate, press or NGO text can be at any level if the observed mechanisms justify it.

Scale for qualifying the degree of propaganda:

- **Level 0 — Open information / non-propagandistic**: the text has a framing, as every text does, but no structuring bias protects its conclusion. Sources are sufficiently diverse for the object, comparable actors are treated symmetrically, and contrary elements can effectively modify, limit or overturn the conclusion.
- **Level 1 — Oriented information**: a structuring framing, selection or hierarchy of facts, or markedly one-sided sourcing orients the reading. The reader nevertheless retains the elements needed for autonomous examination; counter-argument is not systematically neutralised and a contrary observation can still affect the conclusion. The text orients more than it locks.
- **Level 2 — Advocacy / engaged journalism**: the text actively defends a position or conclusion and organises the selection, hierarchy or interpretation of facts in its service. The bias may be explicit or inferable, but objections remain capable of modifying or limiting the conclusion. The text seeks to persuade; it does not yet systematically protect the conclusion against contradiction.
- **Level 3 — Propaganda**: several converging mechanisms protect the conclusion against autonomous examination — strong asymmetry, informational circularity, neutralised counter-argument, immunising concessions, strategic omissions, displaced rigour or other closure mechanisms. The reader can still reconstruct a contradiction, but the device makes this markedly more difficult or absorbs it when it appears. The term “propaganda” begins at this level.
- **Level 4 — Locked propaganda**: protection becomes self-immunising or materially closed. Structural omissions, decisive fabrication or manipulation, exclusion of counter-argument, or a mechanism that turns every contrary observation into confirmation prevent internal contradiction from modifying the conclusion without substantial external input. Autonomous examination is no longer merely oriented or made difficult: it is prevented on the central point.

Thresholds between levels. Moving from one level to the next answers a different question. It is not deduced from a number of flaws.

- From 0 to 1: does a structuring orientation appear?
- From 1 to 2: does the text actively defend a conclusion, beyond the framing that organises information?
- From 2 to 3: is the conclusion protected against contradiction by mechanisms of neutralisation or closure?
- From 3 to 4: does that protection become self-immunising or materially such that contradiction can no longer have its effect without substantial external input?

One-sided sourcing, asymmetry or engagement are therefore not sufficient on their own to cross the Level 3 threshold. Conversely, absence of activist vocabulary does not protect against Level 3 or 4 if closure mechanisms are established.

Descriptors orthogonal to the level. After the level, the analyst may add two types of descriptor, only if established. They do not change the level number and must never be used to infer it.

- **Rhetorical mode:** *frontal* when commitment, accusation or mobilisation are immediately visible; *sophisticated* when persuasion operates mainly through the credibility codes expected by the audience — technicality, expertise, moderation, documentary abundance, concessions or neutrality of tone.
- **Origin or function:** *journalistic*, *activist*, *corporate*, *institutional*, *state*, *para-state*, or another precisely established qualification. A state or institutional origin implies no automatic level; institutional communication open to contradiction may be less closed than a self-immunising private text.

Cautions:

The term “propaganda” is reserved for Levels 3 and 4. A committed, asymmetrical, activist, corporate or strongly oriented text does not become propagandistic for that reason alone. The 2 → 3 threshold must be established: the conclusion is protected against contradiction by a mechanism of neutralisation or closure.

The mere fact that a text is adapted to its audience is not enough to qualify it as propaganda. All argumentative communication takes its addressee into account. Adaptation to the audience becomes an indicator of closure when accompanied by asymmetries, strategic omissions, informational circularity, neutralised counter-argument or an organisation that prevents contrary elements from modifying the conclusion.

The origin of the text is never a diagnosis. “State”, “institutional”, “corporate” or “activist” are descriptors of provenance or function, not degrees of propaganda.

Criteria used to assign the level:

- Diversity of sources (0.d)
- Transparency about production context (0.e)
- Identification of target audience and credibility contract (0.f)
- Informational circularity (3.I)
- Symmetry of actor treatment (3.J)
- Presence and nature of strategic omissions (3.E)
- Actual function of nuances, concessions and moderations (3.K)
- Effective or decorative counter-argument (3.K)
- Adaptation of persuasive methods to the expectations of the target audience (3.K)
- Degree of interpretive closure produced by the text

Presentation in the analysis:

In Section 7.c, the analyst must:

- Present the scale by reproducing all five levels with their labels.
- Assign a level to the ARGUS scope of the analysed text, specifying the number and label. For a non-composite text, this scope is the whole text.
- Add, if established and useful, rhetorical mode and origin or function as descriptors separate from the level.
- Justify the assignment by listing observed criteria, with references to the corresponding steps (0.d, 0.e, 0.f, 3.I, 3.J, 3.E, 3.K).
- Explicitly justify the next threshold: say why the text crosses or does not cross the level above. Hesitation between two levels is resolved by the corresponding threshold question, never by averaging.

Composite and long texts, sub-scope levels.

In a long or composite text, a single global level may erase the disparity of regimes. One section may, for example, be Level 3 in frontal mode in a text that is globally Level 1.

In that case, the analyst may assign a principal level accompanied by sub-scope levels, under two imperative conditions:

- the global level is not the average of the sub-scope levels; it must be justified as a judgment of its own, bearing on the **ARGUS scope taken as a whole**, and not on sections excluded under 0.a bis;
- every sub-scope level must be attached to named sections, never to an isolated passage selected after the fact. Without this condition, sub-assignment permits extraction of a severely rated paragraph from an otherwise sound text, producing a non-representative image of the text.

Descriptors of mode and origin or function may also vary between sub-scopes, but only for named sections and on the basis of established elements.

If the analysis is in Light mode (Appendix 1), this step is omitted.

Example format for reporting — brackets indicate values to be filled in, not a default level:

Scale for qualifying the degree of propaganda:

- **Level 0 — Open information / non-propagandistic:** framing without structural protection of the conclusion.
- **Level 1 — Oriented information:** structuring orientation, autonomous examination still effective.
- **Level 2 — Advocacy / engaged journalism:** conclusion actively defended, objections still capable of modifying it.
- **Level 3 — Propaganda:** conclusion protected by converging mechanisms of neutralisation or closure.
- **Level 4 — Locked propaganda:** self-immunising or material closure preventing contradiction from producing its central effect.

Level assigned to the ARGUS scope analysed: Level [0–4] — [label]. **Rhetorical mode, if established:** [frontal / sophisticated]. **Origin or function, if established:** [journalistic / activist / corporate / institutional / state / para-state / other].

Justification:

- Diversity of sources (0.d): [finding].
 - Transparency about production context (0.e): [finding].
 - Target audience and credibility contract (0.f): [finding].
 - Informational circularity (3.I): [finding].
 - Symmetry of treatment (3.J): [finding].
 - Strategic omissions (3.E): [finding].
 - Counter-argument, concessions and adaptation to audience (3.K): [finding].
 - **Upper threshold:** [reason why the next level is or is not crossed].
-

Step 8 — Self-Critical Examination of the Analysis

Consistency check, mandatory preliminary Before any reflexive examination, the analyst mechanically checks that their analysis holds together. ARGUS subjects the analysed text to an internal consistency test in Step 5; under the performative contract of 0.c, it subjects itself to the same. The checks below each receive an answer naming an element.

- **Header and appendices:** are the appendices declared active in the header those that the checkpoints decided on, and has each actually produced its section?
- **External verifications:** do the three declared numbers match the checks reported in the body of the analysis, each in its own count, and do these checks concern the declared object rather than another version of the text?
- **Non-contradiction between steps:** is a claim in Step 7 contradicted by a finding from Steps 3 to 6? In particular, can an element identified as a point of rigour in the text coexist with a qualification that presupposes its absence?

- **Modal constancy:** are the modal map of Step 2 and the modal-constancy verdict of Step 6 consistent with each other? Where a modal retraction has been identified, has a concession that constitutes its retracted pole been unduly credited to the text in 7.b?
- **Calculations:** is every calculation carried into 7.b as verified exact accompanied by the value recomputed by the analyst? If not, it is declared unverified and not exact.
- **Factual findings:** does the same fact receive two incompatible descriptions in two places in the analysis?
- **Primary source:** where the link between the text and a primary source met the activation conditions for the 3.B fidelity test, were decisive claims confronted with that source and the result qualified under the test? Was a primary item merely supplied or relevant improperly treated as the source of the text without established lineage?
- **Qualification of absences:** was an absence denied the qualification of strategic omission on the ground that it would not destroy the thesis, a criterion that belongs only to the second gate of 3.E? Was an absence placed in the “not determinable as things stand” outcome even though a reasonably accessible check would have allowed a decision and was not attempted, or without unavailability of any check being declared? This outcome prohibits both.
- **Register of the analysis (Rule 19):** does the analysis contain comparative praise or blame that I cannot establish, or a charged term the finding does not require?

An inconsistency detected must be resolved before delivery, not merely flagged and retained. This check is not a section to fill in: it is an occasion to correct. If it reveals no inconsistency, say so in one line.

This check applies in the short version as in the full version.

Reflexive examination The analyst applies the following tests to their own commentary. This step is no longer optional. It may be reduced to a short version of 3 to 4 lines, covering at least:

- (a) one undemonstrated presupposition the analyst themselves admitted;
- (b) an identified risk of bias (including series or anchoring bias if the analysis is part of a sequence of texts);
- (c) an unresolved information limitation;
- (d) **deference:** did the status of a source (recognised institution, canonical author, peer-reviewed journal, established news agency) lower the analyst’s evidentiary requirement on a specific claim? If so, which one?
- (e) **series conformism:** if this analysis follows other analyses, did the analyst align a judgment with the previous ones rather than with the text? On exactly which point?

Points (d) and (e) must receive an answer naming a precise element, or the explicit statement that no case was identified. A general acknowledgement does not satisfy them. A limitation already flagged under the “When elements are insufficient” clause in Step 7 must not be counted here again as a distinct limitation: a cross-reference is sufficient.

The long version remains recommended for published analyses or those intended for serious critical use. It **develops** points (a) presuppositions, (b) bias and (c) limitations using the formulations below, adds an explicit check of epistemic symmetry, and retains points (d) deference and (e) series conformism. The developed formulations of (a), (b) and (c) replace their short answers; they are not three additional responses. This is the basis on which the verification-symmetry requirement in 3.B refers to deference, point (d).

- **Presuppositions of the analysis:** what undemonstrated presuppositions has the analyst themselves admitted (e.g. superiority of one type of proof over another, legitimacy of the requirement for external sources, etc.)?
- **Epistemic symmetry:** has the analyst applied to their own discourse the same requirements of rigour, transparency and non-contradiction that they demanded of the text studied?
- **Potential biases:** does the analyst have reason to suspect that their own opinions (political, professional, cultural) oriented the analysis?
- **Limits of the analysis:** what additional information would be required to settle certain points left undetermined?

Note on the Reliability of the Analysis

At the end of every analysis, the AI must insert the following box:

Note on the reliability of this analysis

This analysis was generated by an artificial intelligence assisting the application of the ARGUS protocol. The AI may make errors, omissions or overinterpretations. It is advisable to reread the analysis critically and check the following points: adherence to the announced steps, presence of the “What the text solidly establishes” section (7.b), and overall consistency of the judgment.

Absolute and Permanent Rules

1. Never confuse rhetorical force with logical validity.
2. Never accept a text’s framing without examining it.
3. Never treat a metaphor as an argument.
4. Always track what the text prevents from being thought as much as what it states.
5. Always treat the opening as a strategic act, not a simple introduction.
6. When a word does massive rhetorical work on its own, it is an alarm signal.
7. When a text says “we” or “everyone”, ask who is included and who is excluded.
8. Always check whether facts are summoned to ground the thesis, or adjusted to serve it.
9. Always confront the announced programme with the actual content.
10. If an existing, argued and significant analysis is supplied with the text or explicitly included in the corpus submitted to the AI, and it directly contradicts a characterisation in the text, its absence from the analysed argumentation may constitute an indication of argumentative weakness.
11. Always examine readability and stylistic register: unnecessarily complex syntax can mask an argumentative void or filter the readership, and a text that claims to address everyone in language that excludes is in performative contradiction.
12. Always question the strategic intention of the text: to whom is it really speaking, to produce what effect, and for what social or political purpose? This questioning must follow the three-level progression (textual clues, hypothesis, confidence level).
13. The analysis first relies on elements available within the text itself. External verification is nevertheless mandatory where an element is decisive to the demonstration and a simple check would establish it, under the conditions and symmetry doctrine detailed in 3.B. Every external verification must be notified, sourced and distinguished from internal examination. Every analysis declares the three numbers required by 3.B — successful, partial and unsuccessful checks — the number of successful checks being the one the counting condition permits, and justifies it if zero; these numbers are an internal execution trace rather than a metric between analyses; remembered knowledge does not count as verification. Where the link between the text and a primary source meets the activation conditions of the 3.B fidelity test and that source is available, decisive claims concerning it are confronted with it under that test. A relevant primary item without established lineage remains an ordinary external verification. External verification remains bounded to decisive elements and does not replace internal analysis.
14. Never attenuate a critical conclusion with vague or timid moderators when the text does not justify that attenuation. Expressions such as “insufficiently”, “not enough”, “relatively”, “partially”, “tends to”, “seems”, “almost”, “to some extent” may be used only if they correspond precisely to the state of the text. If an element is absent, say it is absent. If a contradiction is observed, say there is a contradiction. Critical caution does not consist in weakening the diagnosis, but in correctly distinguishing what the analysis permits one to affirm from what it does not permit one to affirm.

This rule forbids the analyst from using moderators to evade their own judgment. It must not lead to penalising an author whose text itself expresses nuance, caution or explores alternative hypotheses. If the text is nuanced, the analysis must reproduce that nuance in precise terms, without turning it into an accusation of vagueness. The modal constancy check of Step 6 extends this same safeguard to the modality of a load-bearing proposition: it triggers neither on consistent caution, nor on an explicitly acknowledged revision.

15. Formatting rules for the analysis output:
 - Never use tables (neither Markdown nor HTML). Use bulleted lists, numbered lists, key-value pairs, or structured paragraphs with subheadings instead.

- For bulleted lists, use exclusively the character – as the bullet, with a blank line before the first bullet, a space after the bullet, and no blank lines between list items. Numbered lists are also permitted, under the same spacing principle (blank line before, no blank lines between items).
 - For sub-elements, do not use secondary bullets. Place supplementary information on the following lines, without a bullet, indented by two spaces.
 - For key-value pairs, use the format ****Term****: definition.
16. **Never confuse formal counter-argument with effective counter-argument.** The presence of a nuance, concession or opposing source does not prove the balance of the text. The analyst must check whether this nuance can genuinely modify the conclusion or merely serves to make the conclusion more acceptable to the target audience.
 17. **Always evaluate propaganda relative to its target audience.** Propaganda aimed at a cultivated audience may be sourced, nuanced, cautious and stylistically sophisticated. These formal qualities do not exclude it from the field of propaganda if they serve to lock interpretation rather than open it.
 18. **Never produce an artificially complete analysis.** If the text does not provide enough elements to deal with a section reliably, the analyst must say so explicitly (“Insufficient data”, “No probative elements”, “Question not decidable on the basis of the text alone”). It is preferable to acknowledge a limit than to fill a section with speculation or generalities. The detailed implementation of this rule, with its three boundaries, appears at the head of Step 7 under the clause “When elements are insufficient”.
 19. **Apply Tests 3.C and 3.G to your own prose.** The protocol examines the register of the analysed text; it also examines that of the analysis. A critical analysis that does not monitor its own register is contradicted by its prose before it is contradicted by its findings. Two prohibitions follow. **Unestablished comparative praise or blame.** Crediting or discrediting a text by comparison with a class of texts (“better than the genre requires”, “rare in this type of production”, “as is often the case in these circles”) presupposes an empirical claim about that class. If the analysis cannot establish it, the comparison is deleted and the finding stated without a foil. This rule is the counterpart of Rule 14: that rule forbids vague moderators that soften criticism; this one forbids vague comparisons that inflate a compliment or reproach. **Charged terms in the analysis.** The vocabulary identified by 3.C in the analysed text has no place in the commentary identifying it. Operational test: remove the pejorative term and check that the finding still holds. If it does, the term was superfluous; if it does not, the finding was not established.

Lineage and Distinctive Contribution

ARGUS belongs to a tradition it did not invent. The expression “intellectual self-defence” is Noam Chomsky’s, and its dissemination in French is due to Normand Baillargeon (*Petit cours d’autodéfense intellectuelle*, Lux, 2005). Steps 0.d and 0.e bring to the scale of a single text the propaganda model of Edward Herman and Noam Chomsky (*Manufacturing Consent*, Pantheon, 1988), whose five filters concern size and profit orientation, advertising dependence, dependence on institutional sources, discipline exerted on the media by the powerful, and ideological hostility. Appendix 4 owes its order-of-magnitude and base-rate tests to the sceptical and zetetic tradition, from Darrell Huff to Carl Sagan, Henri Broch and Joel Best.

What ARGUS adds to that tradition comes down to five points.

- The standard of proof declared before examination and indexed to textual genre. A list of fallacies applied without regard to the genre of the text does not provide this indexing: it treats an editorial and a scientific article alike, though their reading contracts differ.
- The separation, in Step 7.c, between the **degree of interpretive closure** (Levels 0 to 4) and descriptors of **rhetorical mode** (frontal or sophisticated) and **origin or function**. Propaganda aimed at a cultivated audience may be sourced, nuanced and cautious; these features describe its form without by themselves determining its degree.
- The distinction between formal and effective counter-argument, and between opening and immunising concessions.
- The mandatory Section 7.b, which prevents the protocol from operating as a machine for concluding against.
- Step 8, for which I found no equivalent in the instruments I consulted.

Licence

This protocol is published under the Creative Commons Attribution-ShareAlike 4.0 International Licence (CC BY-SA 4.0). You are free to reproduce, modify and redistribute it, including for commercial purposes, provided that you credit the original author and share any modified version under the same licence.

Appendix 1: ARGUS Light Protocol (Short Version)

1. Test relevance (is the text argumentative?)
2. Classify genre and reading contract
3. Declare the standard of proof
4. Analyse the opening device (presuppositions, perimeter, authority)
5. Reconstruct the thesis neutrally
6. Test absences and qualify them according to the two gates of 3.E
7. Check epistemic symmetry (Point H)
8. Infer strategic intention (3 levels)
9. Confront the announced programme with the content
10. Apply Tests A (order of magnitude) and G (reproducibility of counting) from Appendix 4 only to figures that are load-bearing within the meaning of §2 of that appendix
11. Conclude by distinguishing qualities, flaws with severity, and asymmetries
12. (Mandatory, short version) Consistency check of the analysis, then self-critique

The absence test enters the short version from V5.0.0, with its Stage 2 and its two gates. The reason is internal to the protocol: the unaided pass in the preamble identifies Step 1, the absence test and epistemic symmetry as the three highest-yield questions, and the short version omitted one of the three. The Step 8 consistency check, which verifies the qualification of absences, thereby also gains its object.

The integrity check of the object and the mandatory stop in 0.a apply in the short version as in the full version: they precede the analysis and are not a matter of its degree of detail. The stop applies without simplification: its three elements, three outcomes, exemption regime and prohibition are those of 0.a. The short version abbreviates the analysis, not the decision that precedes it.

Rule 13 also applies in the short version, including minimum external verification and, when its activation conditions are met, the fidelity test against an available primary source. A primary item merely supplied without established lineage remains an ordinary external verification. Only the developed examination of 3.B is omitted, not its mandatory minimum. An absolute rule does not become optional because the analysis is abbreviated.

The list of twelve operations above is the executable scope of the short version: a step in the full protocol that does not correspond to one of them is not executed, subject to the reach of the absolute rules defined in the paragraph below and the explicit exception of Rule 13. The minimum from 3.B imposed by that rule therefore executes in the short version even though none of the twelve operations corresponds to it by name, and the list of exclusions that follows is only a reminder of the heaviest omissions, not an exhaustive list. A point in the consistency check whose object was not produced is checked as not applicable, not ignored; it does not thereby become a section to fill in, which Step 8 prohibits. One line is sufficient: “consistency check completed; points relating to steps excluded from the short version were not applicable.”

Reach of the absolute rules in the short version. They govern the operations that remain within the scope; they do not by themselves reintroduce an excluded step. Rule 11 therefore operates on what the short version produces, without reinstating the developed register examination of 3.G, and Rule 17 without reinstating 0.f or 7.c. Rule 13 is the sole exception, explicitly provided for in the preceding paragraph, because it concerns an evidentiary obligation rather than a degree of detail.

The advanced steps 0.d, 0.e, 0.f, 3.I, 3.J, 3.K and 7.c, together with Appendices 2 and 3, are not included in the Light version. From Appendix 4, only Tests A and G are included; Tests B to F are excluded. The modal constancy check —

the modal map of Step 2 and the Step 6 check introduced in V5.1.0 — is not included in the Light version.

Light and corpus cannot be combined. Appendix 3 presupposes 7.a and 7.c levels for each text, and 7.c is outside the short version. If the user asks for ARGUS Light on several related texts, say so and propose two options: the full protocol with Appendix 3, or independent Light analyses, without cross-text consistency control or transversal synthesis. Producing a corpus synthesis from Light analyses would be a transversal conclusion drawn from materials that do not support it. The “When elements are insufficient” clause in Step 7 applies in the short version as in the full version.

Appendix 2: Analysis of Empty Signifiers (Conditional Appendix)

0. Activation condition

This appendix is activated either at the user’s request or when the checkpoint placed at the end of Step 2 is positive, that is, when the conclusion of the text varies depending on where the boundary of a decisive term in the central thesis is drawn. This boundary test alone controls activation, whether or not the text supplies a definition, criterion of application or enumeration. Absence of definition without an effect of the boundary on the conclusion is not sufficient. The decision not to activate the appendix must be justified in the analysis.

1. Operational definition

An **empty signifier** is a term whose rhetorical and political power comes not from precise semantic content but, on the contrary, from the absence of stable content. This emptiness allows heterogeneous, even contradictory, interpretations to be projected onto it.

In argumentative analysis, the focus is on **strategic** empty signifiers: those whose vagueness performs an identifiable argumentative function in the text, in particular:

- uniting an audience with diverse interests,
- avoiding overly rapid falsifiability,
- producing an effect of obviousness or urgency without having to justify it,
- disqualifying an opponent by association with a vague but charged term.

Do not confuse: An empty signifier is not a simple imprecision of language. It is an **argumentative tool** whose vagueness is functional.

2. Textual indicators of a strategic empty signifier

- **The term is never defined** in the text, or its successive definitions are incompatible. Generic example: “justice”, “the people”, “real democracy”.
- **The term serves to rally heterogeneous demands** without specifying them. Generic example: “we want change”.
- **The author opposes two empty signifiers** to create a moral boundary (“us” vs “them”). Generic example: “the general interest” vs “privileges”.
- **The term has a strong emotional charge** but low informational content. Generic example: “scandal”, “catastrophe”, “absolute emergency”.
- **The argumentation rests on evoking this term** without ever needing to prove it. Generic example: “the historic moment”, “the window”, “the necessary surge”.

3. Analysis grid for an empty signifier

For each suspect term, apply systematically:

- **Location:** Where does it appear? At which strategic moments (opening, conclusion, transition)?
- **Status:**
 - **Empty:** absence of definition, interchangeable usage.
 - **Full:** defined, stable, non-substitutable without loss of meaning.
 - **Mixed:** empty in some places, full in others (potential tension).
- **Function in the argument:**
 - Unite (bring together different groups)

- Avoid refutation (make the thesis non-falsifiable)
- Suggest legitimacy without proof
- Create a sense of urgency
- Distinguish the author from an opponent (by rejecting a competing empty signifier)
- **Gap with a possible full usage:** If the same term is used as a full signifier elsewhere (by the author or in the discursive field), is the gap justified? Does it signal a contradiction?
- **Impact on argumentative solidity:**
 - **Low:** the emptiness is anecdotal, the thesis holds without it.
 - **Medium:** the thesis partly depends on this vagueness (e.g. vague prediction).
 - **High:** without the empty signifier, the argument collapses (e.g. “urgency” without evidence).

4. Link with the steps of the ARGUS protocol

- **Step 1 – Opening Device:** The opening can install one or more empty signifiers (“crisis”, “urgency”, “we”) that frame the whole text.
- **Step 3.C – Symbolically Charged Terms:** Distinguish charged but full terms (e.g. “strike”) from charged and empty terms (e.g. “justice”). This refines the critique.
- **Step 3.E – Test of Absences:** An empty signifier can mask omissions: vagueness makes it possible not to choose between contradictory solutions.
- **Step 3.F – Falsifiability:** An empty signifier makes the thesis difficult to falsify, because the missing content can be redefined after the fact.
- **Step 3.H – Epistemic Symmetry:** Does a text that uses empty signifiers to disqualify an opponent itself use empty signifiers to legitimise itself?
- **Step 4 – Strategic Intention:** The use of empty signifiers may constitute an indication of an intention to mobilise rather than demonstrate; its weight is determined together with the other indicators in Step 4.
- **Step 7 – Differentiated Conclusion:** Mention whether the central thesis rests on one or more unjustified empty signifiers.

5. Generic annotated example

Consider a fictional text stating: “It is urgent to act in the face of the **crisis** threatening our **people**. The **window** is narrow, but we will seize it.”

- **Term “crisis”:** undefined (economic? political? moral?). Interchangeable with “danger”, “emergency”. Status: **empty**. Function: create a sense of urgency without evidence.
- **Term “people”:** undefined (national? local? the oppressed? all citizens?). Status: **empty**. Function: unite heterogeneous demands around a vague “we”.
- **Term “window”:** no temporal or procedural content. Status: **empty**. Impact: **high** – if these three empty signifiers are removed, all that remains is an exhortation without justification.

Conclusion of the analysis: The text rests on a chain of empty signifiers; its argumentative solidity is weak.

6. Usage recommendation

This appendix is **conditional**, not optional: in the full version, its activation is mandatory as soon as the checkpoint at the end of Step 2 is positive, and the decision not to activate it must be justified. The cost of the analysis is not an admissible reason. It can also be activated at the user’s request. Outside these two cases, the analyst may indicate that an examination of empty signifiers would be useful, but does not activate it unilaterally. Such a proposal may in particular be motivated by the presence of terms with strong symbolic charge and weak definition where the boundary test is not positive. If the reach of the thesis actually varies with the boundary of one of these terms, the threshold is crossed and activation is no longer a proposal: it is mandatory.

For a quick analysis (ARGUS Light), ignore this appendix.

Appendix 3: ARGUS-Corpus Protocol (Analysis of Several Related Texts)

Preamble

The ARGUS protocol, in its current form, is designed for the analysis of an isolated text. Yet many real uses — press dossiers, series of articles on the same subject, a corpus assembled by a third party to construct a synthesis or derivative article — involve analysing several linked texts, with an additional requirement: consistency of judgments from one text to another, and the possibility of drawing a synthesis from the corpus without that synthesis fraudulently substituting itself for a conclusion that the texts taken individually do not permit.

This appendix applies **in addition** to the base protocol, applied text by text, and not in its place.

C.1 — Corpus Declaration (preliminary, before analysing the first text) Before analysing the texts in the corpus individually, the analyst must answer:

1. **Origin of the corpus:** who assembled these texts, and according to what criterion (same media outlet, same subject, same period, selection by a third party for a pre-existing thesis)?
2. **Foreseeable homogeneity of sources:** do the texts come from the same media outlet, the same editorial line, or distinct ecosystems? This homogeneity, if it exists, must be declared before individual analysis, because it affects by construction Test 3.I of each text taken in isolation (circularity at corpus level may artificially lower the circularity measured inside each text if the same sources overlap from one article to another).
3. **Risk of second-order circularity:** if a synthesis is to be drawn from this corpus, the analyst must explicitly note that this synthesis inherits the representativeness limits of the initial selection (cf. C.4).

C.1 bis — Corpus Relevance Triage and Single Validation (before analysing the first text) Applied text by text, the mandatory stop in 0.a would produce as many interruptions as the corpus contains texts. For a corpus, it is replaced by a single stop, placed after the corpus declaration and before analysis of the first text. This replacement removes none of its elements: it groups them.

1. **Relevance triage, one line per text.** For each text in the corpus, state its relevance qualification — strong, partial by degree, partial by composition, weak or nil — and its reason in one line. This triage states the degree to which each text is suited to ARGUS; it does not state what it lacks. No analytical finding appears here: a flaw raised at this stage would be a judgment rendered before the analysis that is supposed to establish it.
2. **Proposed scope, only for texts that are not strongly relevant.** A strongly relevant text has no restriction to validate, and the qualification is sufficient in the triage; its scope declaration is still produced in its individual analysis under 0.a bis. For the others, name the included sections, excluded sections and the regime under which excluded sections would be examined.
3. **Single validation, when required.** Make one request to the user, covering both proposed scopes that require validation and corpus composition when it is not already fixed. It recalls the three outcomes in 0.a — validate, correct the scope, abandon — and adds the only one specific to the corpus: restrict the analysis as a whole to some of the texts. If corpus composition is already fixed by the user and no text has a proposed scope to validate, no request is fabricated for form's sake.
4. **Rewrite any corpus reduction into C.1.** If the user restricts the corpus, restate the corpus declaration: name the texts excluded, the reason for their exclusion, and the updated selection criterion. Without this rewrite, the representativeness test in C.4 and the second-order circularity test would concern a corpus that was not analysed. A corpus reduction is itself an act of selection and inherits its consequences.
5. **Exemption and separation of decisions.** **Corpus composition** and **ARGUS scope within each text** are two distinct decisions. The fact that the user fixed the list of texts, their order or execution regime fixes corpus composition; it does not by itself validate an internal division into sections, unless the request also explicitly fixes that division. For a strongly relevant text, no proposed scope requires validation: the line “scope equal to the whole text”, justified in 0.a bis in the individual analysis, is the normal scope declaration and **does not trigger** the formula “scope not validated by the user, set unilaterally by the analyst”. For a text requiring a restricted scope, the two exemption cases in 0.a apply normally according to what the user actually fixed or waived. **The two exemption formulas are mutually exclusive for the same scope decision: never reproduce both and never leave the reader to choose between them.** If only corpus composition was fixed, state that as such; do not transform the fixing of composition into a scope formula.

The triage concerns objects that have already been checked. The integrity check in Step 0 is carried out on every text in the corpus before triage, grouped into a single pre-check naming the documents that are complete and those that appear truncated. The order matters: projecting a scope onto an incomplete document would have the user validate a decision

taken on an object that is not what it appears to be. A document whose integrity is doubtful is flagged in the same request as the scopes, and the user may complete it, exclude it, or accept its analysis under the reservations in Step 0.

A text that triage qualifies as weak or nil relevance is proposed for exclusion rather than excluded automatically: the outcomes are those in 0.a, the derogation it provides remaining outside their list, as 0.a specifies. If analysed at the user's express request, it begins with the statement required by 0.a. A text actually excluded remains an element of the submitted corpus, and its exclusion is declared in C.1 like any other reduction.

C.2 — Application of the Base Protocol, with Specified Order of Treatment Each text is then analysed according to the base protocol, in the order in which it appears, with one carry-forward rule: the integrity check and relevance qualification of 0.a were already performed text by text in C.1 bis and are not recalculated. Each individual analysis restates at the head the exact object, retained relevance qualification and, where applicable, the declared incompleteness limitation from the pre-check with the three consequences attached to it by Step 0; it carries forward the elements of 0.a bis already validated where applicable — scope, regime of excluded sections — applies the non-transfer rule, then continues from 0.b. For a strongly relevant text, it produces at this point the scope declaration required by 0.a bis: scope equal to the whole text, with its reason. This full declaration is neither a new validation nor a scope “set unilaterally” within the meaning of the second exemption formula in 0.a. The stop and validation in 0.a are therefore not replayed text by text: the result of C.1 bis substitutes for them. However:

- The analyst must note, at the end of each individual analysis, a brief summary of the levels assigned in 7.a (severity) and 7.c (degree of propaganda), to build a comparative grid in C.3.
- If the order of analysis is not neutral (for example, the corpus is supplied in chronological or thematic order chosen by a third party), the analyst must flag it as a potential anchoring factor (cf. C.3).

C.3 — Cross-Text Consistency Check (mandatory after the last text) Once all texts in the corpus have been analysed individually, the analyst must produce an explicit consistency check answering the following questions:

1. **Comparative summary of assigned levels** (in list form, in accordance with Rule 15): for each text, recall the maximum severity level assigned in 7.a and the propaganda level assigned in 7.c.
2. **Anchoring-drift test**: two distinct risks, to be checked and declared separately. **Sequence anchoring**: have texts analysed later in the sequence received systematically more severe or more favourable levels than the first ones, without internal justification specific to each text? If a gap is observed without clear justification, it must be flagged as a risk of series bias, not silently retained. **Chronology anchoring**: when the order of treatment is not the order of publication dates, has the analyst imputed to an earlier text what its successors made visible? Every transversal conclusion concerning an evolution over time recalls the order in which the texts were actually analysed.
3. **Criteria-symmetry test**: has the same type of flaw (for example, an unquantified generalisation, or absence of the voices of people directly concerned by the subject) received a comparable severity level each time it appeared in the corpus? If the same flaw is qualified “minor” in one text and “major” in another without a contextual difference justifying the discrepancy, this must be corrected or explicitly justified.
4. **Blind rereading recommendation** (optional, for corpora of more than three texts): if the rigour of the consistency check requires it, the analyst may propose redoing the analysis of the first text in the corpus in a fresh conversation, without supplying the other texts or analyses already produced, and then comparing the two results. Only under this condition is the rereading blind: revisiting the first text after seeing all the others would be precisely the anchoring that point 2 seeks to detect.

C.4 — Second-Order Circularity Test (mandatory before any thematic synthesis) If a synthesis or transversal diagnosis is to be drawn from the corpus (for example: “what these texts, taken together, reveal about a broader phenomenon”), the analyst must, before drafting that synthesis, answer explicitly:

1. **Is the corpus representative of the phenomenon it is being made to speak for, or only of a particular editorial viewpoint on that phenomenon?** (Return here to the declaration made in C.1.)
2. **Would an alternative corpus, assembled according to another criterion or editorial line, allow a substantially different synthesis to be constructed from comparable facts?** If so, this must be flagged as a limitation of the proposed synthesis, not merely a stylistic nuance.

3. **Does the synthesis draw a conclusion that each text taken individually does not permit, but that their combination would make possible in a logically valid way (legitimate aggregation), or does the synthesis add an interpretive framework external to the corpus, not deducible from the combination of the texts alone (the analyst’s or synthesis author’s own contribution)?** This distinction must be made visible in the final drafting through a structural separation (distinct sections), not merely lexical nuances within the same development.

C.5 — Reporting Register: Distinguishing Legitimate Aggregation from Interpretive Contribution Every synthesis drawn from a multi-text corpus distinguishes two registers, and this distinction is **structural**: two separately numbered subsections, not two paragraphs in the same section. Merely stylistic separation does not satisfy the clause.

C.5a — What the corpus establishes by legitimate aggregation. Facts or trends that several texts in the corpus each document separately, and whose combination requires no information external to the corpus to be asserted.

C.5b — What the analyst or synthesis author adds from outside the corpus. Any interpretation, any broader causal or political framework, any generalisation to a scale the corpus alone does not cover. Any new factual claim introduced in this register is sourced under the counting condition in 3.B, and every interpretation is presented as such: this register is the only place where the analyst speaks in their own name, not a place where evidentiary standards are relaxed. This register opens with an explicit declaration of status — “interpretive contribution external to the corpus” — and never in the same voice as C.5a. First-person wording, “I believe that”, is an acceptable marker of this break, but does not replace it. These markers are enunciative rather than modal: they attribute a claim to its author instead of making the corpus bear it, and Rule 14 does not target them, since it forbids attenuation of a finding, not saying who carries it.

C.6 — Supplementary Reliability Note for Corpus Analyses To the standard reliability note (end of the base protocol), add for every corpus analysis:

This analysis concerns a corpus of [N] texts assembled according to the following criterion: [reminder of C.1]. The reader is invited to bear in mind that every transversal synthesis inherits the representativeness limits of this selection, and that a corpus constructed differently could, from comparable facts, lead to a different synthesis.

Appendix 4: Arithmetical and Statistical Control (Conditional Appendix)

1. Rationale for this appendix

The base protocol audits the origin of a figure: who states it, with what qualification, in what ecosystem of sources. It does not audit its possibility. A figure may be correctly attributed to an identifiable source, reproduced without distortion, free of any asymmetry of treatment, and still be impossible.

Take the following claim: “2,000 children die every hour in this country, and have done for ten years.” Attributed to a named academic, it satisfies the source audit (0.d), escapes the circularity test (3.I) and presents no asymmetry of treatment (3.J). Multiplied by 24 and then by 365, it yields 17.52 million deaths per year, in a country with 20 million inhabitants. No step from 0 to 8 asks for this multiplication.

This appendix supplies the missing multiplication.

2. Activation conditions

Activate this appendix if the text contains at least one of the following, provided that the element is load-bearing within the meaning defined after the list:

- a figure, rate, proportion or projection that is load-bearing within the meaning defined below;
- a quantitative comparison between two sets;
- a count relating to a corpus, media coverage or a documentary database;
- a claim of detection, screening, attribution or profiling;
- an average presented as characterising a population.

A figure is load-bearing whenever the text offers it in support of a central proposition, even if the conclusion could survive without it. The criterion is not whether the datum is indispensable, but the evidentiary function assigned to it by the text: a figure presented as evidence commits its verifiability. Proportionality is then applied in the weighting of 7.a, not in refusing activation.

If the text contains no load-bearing numerical element, do not activate the appendix and say so explicitly, in accordance with Rule 18.

3. Scope and prior declaration

The check concerns only load-bearing figures within the meaning of §2, in accordance with the principle of proportionality. An inaccurate illustrative figure is a minor flaw, to be flagged in one line. An inaccurate load-bearing figure may be fatal. The criterion of load-bearing is that of §2 — the evidentiary function the text assigns to the figure — not dependence of the conclusion upon it.

Before any check, the analyst must declare the list of figures they retain as load-bearing, and why each is. This declaration precedes calculation and prevents retaining after the fact only the figures that fail.

4. The seven tests

- **A. Order of magnitude.** Extrapolate the stated rate over the full period, then relate the result to the reference population or set. Three forms call for this test first: an hourly or daily rate applied over years, periodic doubling, and prolonged compound growth. A result exceeding the reference set establishes impossibility only under one condition: the **counted unit must be non-repeatable**, that is, the same unit can contribute only once to the count in question. One death per person, one seat per election, one company per closure meet this condition. It is not met by repeatable events — consultations, journeys, transactions — whose number may legitimately exceed the population concerned. When the set renews during the period, the bound need not be closed: it must be **increased**, that is, the initial population plus an upper bound on renewal. That upper bound is established from data in the text or an admissible external verification, never from the analyst's own estimate: §5 prohibits substituting the analyst's own model for the author's, and a rough estimate of renewal would be such a model. Without an established upper bound, the multiplication remains executable and is reported, but the conclusion of Test A on impossibility is **not answerable for lack of an established bound**, within the meaning of §5. The calculation and its result are written out, impossibility is not asserted, and that silence does not validate the figure. Name the counted unit, the bound retained and the time window before concluding; without a non-repeatable unit, mere excess establishes nothing.
- **B. Denominator.** Is the denominator stated? Is it the right one? The denominator is the population, set or reference base to which the percentage relates, not necessarily its absolute numerical value. "9% of global sales" states its denominator even if the total volume of global sales is not supplied: that volume may be missing for conversion into an absolute value or another reconstructibility test, but its absence is not a denominator flaw. A percentage without an identifiable population or reference base is not a datum. For a percentage change or reduction, the comparison period or value must be identifiable: "+24.5% year on year" has a comparison base even if absolute volumes are missing; "reduced by 30%" without duration, value or reference period does not. A percentage whose denominator or comparison base changes during the text is an equivocation.
- **C. Measure of central tendency.** Is it a mean, median or mode? Does the text specify? Which of the three would change the conclusion? A company with twenty-three salaries can simultaneously display a mean of 6,000, a median of 4,000 and a mode of 2,000, without any of the three figures being false.
- **D. Base rate.** For every claim of detection, screening, attribution or profiling, is prevalence in the population given? Without it, a detection rate of 99% says nothing about the probability that a designated individual is actually concerned. Omission of prevalence is sufficient to make a detection rate unusable, and is systematically checked in texts dealing with security, health or fraud.
- **E. Measurement, estimate, model, projection.** Is the figure measured, estimated from a sample, produced by a model, or projected into the future? Does the text distinguish these? A modelled figure presented as a measurement is an unsubstantiated claim, even when the model is good.
- **F. Displayed precision.** Does the number of significant digits exceed what the method can support? A projection to the nearest million on a ten-year model is false precision, and false precision functions as a marker of seriousness within the meaning of 3.K.
- **G. Reproducibility of counting.** For every quantitative claim obtained by selecting or querying a corpus of

documents, media coverage or a documentary database, are the database queried, the query or selection criterion used, and the unit of counting stated? If any one is missing, the count is not reproducible and establishes nothing verifiable, whatever the plausibility of the result. The test applies to counting performed by the text itself, and to counting it relays when it presents the count as established and draws its own conclusion from it. It does not apply to counting that the text explicitly attributes to its source while preserving its modality: in the second case, what is examined is not the count but its appropriation. **Domain boundary:** G concerns exclusively documentary counts whose result depends on the choice of database, query or selection criterion and unit of count. It does not extend to ordinary statistical, administrative, demographic, scientific or operational measurements or series merely because they rely on a database; these fall under Tests A to F and 3.B according to the flaw examined.

5. Prohibition specific to this appendix

The analyst does not substitute their own model for the author's in order to correct a calculation. They redo the author's calculation with the author's data.

If the data supplied by the text do not permit the stated result to be reconstructed, the finding to record is "result not reconstructible from the data supplied". This is a flaw in the text, not an invitation to recalculate differently. Recalculating with assumptions other than the author's would produce a refutation by the analyst, not an examination of the text. This situation falls under the "When elements are insufficient" clause in Step 7. Test A escapes that clause only on grounds of cost: neither the burden of the calculation nor uncertainty about the order of magnitude makes it unanswerable. It may, however, be unanswerable when a minimum datum necessary for its execution is missing — reference set or bound, counted unit, time window, or correspondence between the figure and its reference — and that datum is not accessible through any admissible verification. Name the missing datum precisely. The enumeration is not an excuse not to search: a datum the text does not supply but that a simple check establishes is not missing within the meaning of this clause, and the test remains answerable. This absence falls under Test B only when it is genuinely a missing or unsuitable denominator; in other cases, it remains a reconstructibility flaw in Test A and is not artificially reclassified as a denominator flaw.

6. Reporting

The findings of this appendix are reported in 3.B, in a subsection entitled "Arithmetical and statistical control (Appendix 4)", and weighted in 7.a according to the substitution test of the base protocol, with the following specific rules:

- a load-bearing figure whose impossibility is established by Test A **invalidates the quantitative proposition** it purports to establish. Its severity for the central thesis is determined in 7.a by the substitution test. Since §2 now qualifies as load-bearing a figure that the conclusion could survive without, that severity can no longer be automatic without contradicting 7.a;
- a count that is not reproducible within the meaning of Test G **does not establish** the empirical sub-thesis it supports, without thereby establishing its falsity. If the count is performed by the text itself, the sub-thesis becomes an **unproven and non-reproducible quantitative claim**. If the text relays an external count it had presented as established, it can no longer use it as a demonstrated result and must retain it as a **reported claim**. The effect on the central thesis is determined in 7.a by the substitution test and is not automatic;
- the word **fatal** is reserved for the 7.a grid, where it means that the flaw invalidates the central thesis. The invalidating finding in the first rule is not carried into 7.a under that label, and the second rule establishes no invalidation, only absence of demonstration: both enter 7.a through the substitution test, which alone decides their rank.

Calculations verified as exact must be listed in 7.b. Arithmetical control is not an instrument for the prosecution: redoing a correct calculation and saying so is part of the work.

7. Recommendation for ARGUS Light

Unlike Appendices 2 and 3, two of the seven tests have negligible cost and high yield, Tests A and G. They are included in ARGUS Light. Tests B to F, which require a longer examination, are excluded.