

Trustworthiness as a Local-to-Global Property: A Sheaf-Theoretic Account of Multi-Source Information

Edward Hermann Haeusler¹, Glenda Amaral³, Giancarlo Guizzardi³, Daniel Schwabe¹,
Veronica dos Santos¹ and Sergio Lifschitz¹

¹Department of Informatics, Pontifical Catholic University of Rio de Janeiro (PUC-Rio), Brazil

³University of Twente, Netherlands

Abstract

Trustworthy Artificial Intelligence and distributed information systems require principled criteria explaining why information originating from multiple contexts should be regarded as worthy of trust, or globally coherent. Existing approaches typically model trust as an epistemic relationship, a confidence score, a reputation mechanism, or a modal property associated with propositions and agents. In this paper, we propose a different perspective. We argue that trustworthiness is fundamentally a structural property of information itself, arising from the possibility of coherently integrating locally verified information across contexts.

To formalize this idea, we employ sheaf theory and descent theory. Contexts are organized as a Grothendieck site, informational objects are represented as sections of a sheaf, and trustworthiness is characterized by effective descent: the existence of compatible locally verified sections that admit a coherent global realization. Within this framework, trustworthiness is distinguished from logical consistency. While consistency concerns the absence of contradiction, trustworthiness concerns the existence of a globally coherent informational object obtained from compatible local data.

We further show that modal stability under contextual refinement is insufficient to guarantee trustworthiness since locally persistent information may still fail to integrate globally. This failure is characterized by Čech cohomological obstructions, providing a structural account of mistrust as the impossibility of gluing locally verified information. We also relate the proposed framework to institutional semantics, where trustworthiness corresponds to the existence of globally compatible models across contextual translations.

Finally, we illustrate the approach through distributed-record and electoral-information scenarios, showing how trustworthy global outcomes emerge from compatible local informational proxies. The resulting account provides a mathematical criterion for integrating trustworthy information and suggests a sheaf-theoretic foundation for reasoning about provenance, reliability, and trust in distributed AI systems.

Keywords

Trustworthiness, Trustworthy AI, Sheaf Theory, Descent Theory, Knowledge Representation, Contextual Reasoning

1. Introduction

Trustworthy Artificial Intelligence requires principled criteria explaining why multi-source information should be considered trustworthy. Contemporary systems integrate information originating from multiple sources, institutions, databases, and agents. Although local verification mechanisms may certify information within individual contexts, a fundamental question remains:

When does locally verified information become globally trustworthy?

Most existing approaches model trust as an epistemic relation, a confidence score, a reputation mechanism, or a modal property attached to propositions and agents. In this paper, we investigate a different possibility.

Our central thesis is that trustworthiness may be understood as a structural property of the information itself. Informational objects become trustworthy when their locally verified manifestations can be coherently integrated across contexts.

To formalize this idea, we employ sheaf theory and descent theory. Contexts are organized as a Grothendieck site, information is represented by sheaves, and trustworthiness corresponds to the existence of compatible descent data, giving rise to a global section.

This perspective provides a formal answer to one possible reason for trusting multi-sourced information: global coherence emerging from local compatibility.

The main goals of this paper are to (i) distinguish between consistency and trustworthiness; (ii) to provide a structural interpretation of trustworthiness as effective descent; and (iii) to explicitly characterize mistrust as any obstruction to gluing; (iv) to show some applications to distributed records and electoral systems.

Unlike traditional epistemic accounts, which study trust between agents or institutions, this work investigates structural trustworthiness as a property of information integration in multiple contexts.

The notion of trustworthiness proposed in this paper is not intended to replace traditional epistemic accounts of trust. Rather, it introduces a complementary structural notion that addresses a different question. Epistemic trust refers to the attitudes of agents toward information, institutions, or other agents, and explains why information is accepted or relied upon. By contrast, structural trust concerns the organization of the information itself, characterizing when independently verified local information can be coherently integrated into a single global informational object. These two notions are not competing but complementary, and in practical distributed AI systems they naturally coexist: epistemic trust may guide the acceptance of local information, while structural trustworthiness provides an objective criterion for assessing whether such information admits coherent global integration.

The remainder of this paper is organized as follows. Section 2 distinguishes trustworthiness from logical consistency and argues that the former should be understood as a local-to-global property rather than merely as the absence of contradiction. Section 3 introduces the categorical and sheaf-theoretic notions required throughout the paper, including Grothendieck sites, sheaves, descent theory, modal stability, and Čech obstructions to global coherence. Section 4 establishes a connection with institutional semantics, showing how trustworthiness can be reformulated as the existence of globally compatible models across contextual translations. Section 5 develops a detailed electoral case study, illustrating how trustworthy global outcomes emerge from compatible local informational proxies and how failures of trust correspond to obstructions to descent. Section 6 discusses why effective descent provides a principled structural reason for trust and examines the implications of this perspective for contemporary AI systems, including retrieval-augmented generation, knowledge graphs, and multi-agent architectures. Section 7 situates the proposal within the literature on contextual reasoning, institutional semantics, provenance, and sheaf-theoretic approaches to distributed information. Section 8 outlines possible quantitative extensions of the framework based on global sections and cohomological invariants. Finally, Section 9 summarizes the main contributions and discusses directions for future research. For readers unfamiliar with category theory, the appendix provides a concise introduction to the categorical concepts employed throughout the paper, illustrated with examples drawn from the trustworthiness scenarios discussed in the main text.

2. Trustworthiness versus Consistency

At first sight, trustworthiness may appear to be nothing more than consistency viewed through a different mathematical lens. However, the distinction between local compatibility and global integration suggests otherwise. As observed in contextualized sheaf-theoretic studies [1], locally compatible information does not need to admit a coherent global realization. The issue is therefore not merely the absence of contradiction, but the possibility of coherent global integration.

Consistency is a logical property of a collection of propositions: it guarantees that no contradiction can be derived from them. Trustworthiness, however, addresses a different concern. It is not primarily about the absence of contradictions, but about the possibility of integrating locally validated information into a coherent global informational structure.

Consider three institutions maintaining records about the same individual. Each institution may

satisfy all of its internal consistency requirements, and pairwise reconciliation between institutions may also succeed. Nevertheless, the system as a whole may still fail to determine a unique global representation of that individual. The difficulty is not the presence of contradictions, but the failure to integrate local information into a coherent global state.

For this reason, trustworthiness should not be identified with consistency. Rather, it concerns the transition from local validation to global integration. In a multi-context information system, locally verified informational objects may remain mutually compatible while still failing to determine a coherent global informational structure.

From a sheaf-theoretic perspective, local compatibility is a necessary condition for trustworthiness, but not a sufficient one. Trustworthiness arises only when compatible local informational states can be integrated into a global section. In this sense, trustworthiness is fundamentally a local-to-global property, whereas consistency concerns only the absence of contradiction at the local level.

Philosophical accounts of trust have long emphasized that trustworthiness is fundamentally relational rather than an intrinsic property of isolated objects or propositions. Baier [2] characterizes trust as arising from relations of dependence between independent agents, O'Neill [3] argues that trustworthiness is an objective property that justifies trust rather than a merely subjective attitude, Hardin [4] distinguishes trust from the objective conditions that make it warranted, and Floridi [5] understands information as inherently relational within an informational environment.

3. Mathematical Background

Structural trustworthiness requires a formal account of information integration across contexts. Sheaf theory provides such a framework. We briefly recall the necessary categorical notions underlying our approach, focusing on the structures needed to model local-to-global information, which is used to characterize trustworthiness as a property of global coherence. We follow [6, 7].

A category $\mathcal{C}$ consists of objects, morphisms between objects, associative composition, and identity morphisms. In our setting, categories are used to represent systems of informational contexts and their refinement relations.

A *Grothendieck topology* J on a category $\mathcal{C}$ assigns to each object U a collection $J(U)$ of families of morphisms

$$\{u_i : U_i \rightarrow U\}_{i \in I},$$

called *coverings*, satisfying stability under pullback, locality, and maximality. A pair $(\mathcal{C}, J)$ is called a *site*. Intuitively, coverings represent admissible ways of reconstructing a global context from local pieces.

A *presheaf* on $\mathcal{C}$ is a functor

$$F : \mathcal{C}^{op} \rightarrow \mathbf{Set},$$

assigning to each object U a set $F(U)$ of local data and to each morphism $f : V \rightarrow U$ a restriction map

$$F(f) : F(U) \rightarrow F(V).$$

A presheaf is a *sheaf* if for every covering

$$\{u_i : U_i \rightarrow U\},$$

any family of sections;

$$s_i \in F(U_i)$$

that agree on overlaps,

$$s_i|_{U_i \times_U U_j} = s_j|_{U_i \times_U U_j},$$

admits a unique gluing

$$s \in F(U)$$

such that

$$s|_{U_i} = s_i$$

for all i .

This condition formalizes the passage from compatible local data to a global object.

In the following example: (1) contexts are financial markets; (2) sections are local currency valuations; (3) restrictions correspond to the information shared by markets trading a common currency; (4) a global section is a coherent worldwide valuation.

Example 1 (Currency Exchange). *Consider the global currency context $U = \{\text{USD}, \text{EUR}, \text{GBP}\}$, covered by the local contexts $U_1 = \{\text{USD}, \text{EUR}\}$, $U_2 = \{\text{EUR}, \text{GBP}\}$ and $U_3 = \{\text{GBP}, \text{USD}\}$. Let F be the sheaf assigning to each context V the set of currency valuations $p : V \rightarrow \mathbb{R}_{>0}$ where $p(c)$ denotes the value assigned to currency c . Exchange rates are determined by $r(c, d) = \frac{p(d)}{p(c)}$. Each local market determines a section $p_i \in F(U_i)$, representing its local valuation of the currencies it trades. Whenever two markets share a currency, their valuations are required to agree on the overlap. For example, suppose the markets report the following compatible exchange rates:*

$$1 \text{ USD} = 0.90 \text{ EUR}, \quad 1 \text{ EUR} = 0.80 \text{ GBP}, \quad 1 \text{ GBP} = \frac{1}{0.72} \text{ USD}.$$

Normalizing $p(\text{USD}) = 1$, the first market determines

$$p(\text{EUR}) = 0.90.$$

The second market therefore gives

$$p(\text{GBP}) = 0.90 \times 0.80 = 0.72.$$

Finally, the third market yields

$$p(\text{USD}) = 0.72 \times \frac{1}{0.72} = 1.$$

Hence, the three local valuations agree on every overlap and determine a unique global valuation. The sheaf condition then guarantees the existence of a unique global valuation $p \in F(U)$ whose restriction to each market coincides with its local valuation. Thus, a sheaf provides a mathematical model for integrating contextual financial information into a single coherent worldwide system of currency values.

The sheaf condition can be understood as a formal model of contextual information integration. Different contexts may contain only partial information about a common object. Sections represent the information available within each context, while restrictions describe the information shared between contexts. When the local sections agree on their overlaps, the sheaf condition states that they determine a unique coherent global object.

Descent theory studies when local data determine global objects. Given a covering $\{U_i \rightarrow U\}_{i \in I}$, a family of local sections together with compatibility on overlaps is called *descent data*, and descent is said to be *effective* when such data arise uniquely from a global section.

The previous example illustrates successful integration of compatible local information into a unique global object. Descent theory addresses the more general question of when such local data admit a coherent global realization. The following example shows that local verification alone is insufficient: even when each local context is internally correct, incompatibilities across contexts may prevent the existence of a global section.

Example 2 (Failure of Global Integration). *Suppose now that the three local markets report the following exchange rates: 1 USD = 0.90 EUR; 1 EUR = 0.80 GBP; and, 1 GBP = 1.50 USD. Assume the valuation in the first market is normalized by $p_1(\text{USD}) = 1$, which implies, $p_1(\text{EUR}) = 0.90$. Since the first and second markets overlap on EUR, compatibility requires $p_2(\text{EUR}) = 0.90$. Using the second exchange rate, we obtain $p_2(\text{GBP}) = 0.90 \times 0.80 = 0.72$. Similarly, because the second and third markets overlap on GBP, compatibility requires $p_3(\text{GBP}) = 0.72$. Applying the third exchange rate yields $p_3(\text{USD}) = 0.72 \times 1.50 = 1.08$.*

However, the first and third markets also overlap on USD. Compatibility would therefore require $p_3(\text{USD}) = p_1(\text{USD}) = 1$, which contradicts the computed value $p_3(\text{USD}) = 1.08$. Hence the local valuations are not compatible on all overlaps and therefore do not determine a global valuation over $U = \{\text{USD}, \text{EUR}, \text{GBP}\}$.

The attempt to integrate the locally verified valuations reveals that there is no global valuation. Although each market correctly verifies its own exchange rates, the family of local valuations does not satisfy the compatibility conditions required for descent. Consequently, no global section can be obtained. This illustrates the central idea of descent theory: local verification alone does not guarantee global coherence.

Failures of descent arise whenever locally verified information cannot be integrated into a coherent global object. In more sophisticated settings, such failures may be characterized by non-trivial Čech cocycles. The present currency-exchange example illustrates a failure of compatibility at the level of local descent data. The cohomological obstructions observed above concern more subtle situations in which local compatibility is preserved, but global gluing nevertheless fails.

In foreign exchange markets, triangular arbitrage removes inconsistencies among local exchange quotations until a globally coherent pricing system is restored [8, 9, 10]. From the sheaf-theoretic perspective proposed here, this globally coherent pricing system is naturally interpreted as a global section obtained by gluing the compatible local quotations. Thus, triangular arbitrage may be viewed as the dynamic process that eliminates the obstruction to the existence of a global section.

Returning to our general approach. Local validity can be interpreted through Kripke–Joyal semantics. For a context $U \in \mathcal{C}$, the forcing relation $U \Vdash \varphi$ expresses that the statement φ holds in context U .

This captures the inherently contextual nature of information: statements may be locally valid without admitting a global realization.

The site $(\mathcal{C}, J)$ represents a universe of contexts, where objects are domains of information, morphisms represent contextual relations, and Grothendieck topology J specifies which families of morphisms can be taken as admissible ways of reconstructing a context from local pieces.

The sheaf F assigns to each context its available information together with restriction maps. A covering $\{u_i : U_i \rightarrow U\}_{i \in I}$ decomposes a global context into local ones. The fiber products

$$U_{ij} = U_i \times_U U_j$$

and

$$U_{ijk} = U_i \times_U U_j \times_U U_k$$

represent pairwise and triple overlaps, respectively.

A 0-cocycle consists of a family of local sections $(s_i)_{i \in S}$ satisfying

$$s_i|_{U_{ij}} = s_j|_{U_{ij}}$$

, for all i, j . Thus, a 0-cocycle represents locally compatible information.

A 1-cochain consists of transition data g_{ij} , defined on pairwise overlaps U_{ij} , while a Čech 1-cocycle satisfies the compatibility condition

$$g_{jk} \circ g_{ij} = g_{ik}$$

, on every triple overlap U_{ijk} .

This condition guarantees that different ways of comparing local information yield the same result.

The sheaf condition ensures that every compatible 0-cocycle determines a unique global section.

In contrast, non-trivial 1-cocycles measure obstructions to gluing. Even when local compatibility holds pairwise, incompatibilities may emerge across multiple overlaps, preventing the existence of a coherent global object.

In the sequel, these notions provide the formal basis for our analysis: contexts form a site, and information is organized as a sheaf, and trustworthiness arises precisely as the existence of compatible descent data, that is, as the existence of a global section obtained by gluing locally valid information.

3.1. Modal Stability and Its Limitations

A natural approach to modeling trust is through modal operators, in particular, the necessity or $\Box$ modality. Intuitively, one may attempt to interpret trustworthiness as a form of modal stability: a statement is trusted whenever it remains valid under all refinements of the current context. Formally, given a context U , one may define

$$U \Vdash \Box \varphi \quad \text{iff} \quad \forall f : V \rightarrow U, \quad V \Vdash \varphi$$

This corresponds to a standard Kripke-style interpretation of $\Box$ in normal modal logic, in which context refinement plays the role of accessibility. However, modal stability under refinement does not guarantee global coherence. Consider again the distributed birth-record scenario involving three contexts U_h, U_r, U_m , representing a hospital, a civil registry, and a municipal database, respectively.

Let B be the sheaf of birth-record data, and suppose we have local sections: $b_h \in B(U_h)$, $b_r \in B(U_r)$, and $b_m \in B(U_m)$. Assume that each record remains valid under every refinement of its context, i.e., for $i \in h, r, m$:

$$\forall f : V \rightarrow U_i, \quad V \Vdash \text{Verified}(b_i|_V),$$

In modal terms, we have $U_h \Vdash \Box \text{Verified}(b_h)$, $U_r \Vdash \Box \text{Verified}(b_r)$, $U_m \Vdash \Box \text{Verified}(b_m)$. Thus, each record is locally stable. Nevertheless, this remains insufficient for trustworthiness.

Suppose that on overlaps the records fail to agree¹:

$$b_h|_{U_{hr}} \neq b_r|_{U_{hr}}, \quad \text{or} \quad b_r|_{U_{rm}} \neq b_m|_{U_{rm}}.$$

Such discrepancies may arise from differences in identifiers, registration conventions, naming policies, or temporal updates. In that case, the family $\{b_h, b_r, b_m\}$ fails to define compatible descent data. Consequently, there is no global section $b \in B(U_h \cup U_r \cup U_m)$, such that

$$b|_{U_h} = b_h, \quad b|_{U_r} = b_r, \quad b|_{U_m} = b_m.$$

This example shows that local necessity captures only persistence within individual contexts. Moreover, Trustworthiness requires more than persistence. It requires compatibility across contexts. Modal logic, even in strong systems such as S4-based ones, captures intra-contextual stability, but does not provide a mechanism for integrating independently verified information, suggesting that trustworthiness cannot be easily, or extensionally characterized as a modal property alone.

What is required is a framework capable of representing both local verification and global integration. Sheaf theory provides precisely such a framework. Thus, instead of focusing on persistence under refinement, it focuses on compatibility on overlaps and the existence of global sections. In this setting, trustworthiness is not a form of necessity. It is a form of descent.

3.2. Obstruction to Descent in Čech Terms

The previous subsection showed that modal stability is insufficient for trustworthiness. Even when the information remains valid under all contextual refinements, it may still fail to admit a coherent global realization. Sheaf theory provides a precise framework for characterizing this phenomenon. Let

¹If $f : A \rightarrow B$, and $D \subseteq A$, the notation $f|_D$ is the restriction of f to D

$\{U_i \rightarrow U\}_{i \in I}$ be a non-trivial covering of a context U , i.e., I has at least two elements, so to say, to contexts for U , or two points of view regarding U . Given local sections

$$s_i \in F(U_i),$$

, the descent problem asks whether there exists a global section $s \in F(U)$, such that, for all $i \in I$: $s|_{U_i} = s_i$. A necessary condition is compatibility on overlaps, i.e.:

$$s_i|_{U_{ij}} = s_j|_{U_{ij}},$$

where

$$U_{ij} = U_i \times_U U_j.$$

In the present setting, trustworthiness requires a non-trivial covering of contexts. A covering consisting only of a single context does not involve contextual integration and, therefore, cannot provide evidence of trustworthiness. Trustworthiness arises precisely from the successful integration of independently verified information across distinct contexts (multiple sources).

When multiple contexts participate in a covering, compatibility must be checked not only locally on pairwise overlaps, but also globally across the entire covering. In particular, for coverings involving three or more contexts, non-trivial Čech cocycles may arise, providing obstructions to effective descent. This observation motivates the following definition.

Local verification certifies the correctness of individual informational objects. Structural trustworthiness certifies the possibility of integrating these objects into a coherent global informational structure.

The requirement of multiple contexts is essential rather than incidental. In a single context, the notion of structural trustworthiness becomes vacuous, since there are no overlaps on which compatibility conditions can be formulated and no descent problem to solve. Structural trustworthiness only becomes meaningful when independently verified informational objects must be integrated across distinct contexts. It is precisely the relationships among these contexts—not the individual informational objects themselves—that determine whether there is a coherent global informational object. In the absence of multiple contexts, there are no compatibility conditions to satisfy, no gluing to perform, and no possibility of structural obstruction. Consequently, the only notion that remains is local verification.

Definition 1 (Trustworthiness). *Let F be a sheaf on a site $(\mathcal{C}, J)$. A section $s \in F(U)$ is trustworthy if there exists a non-trivial covering, with $\text{card}(I) \geq 2$, $\{U_i \rightarrow U\}_{i \in I}$ and a family of locally verified sections $s_i \in F(U_i)$, such that, $s|_{U_i} = s_i$ and $s_i|_{U_{ij}} = s_j|_{U_{ij}}$, for every pair i, j .*

Definition 1 should not be understood as a reformulation of the sheaf condition. The sheaf condition is a property of the sheaf itself: it requires that compatible local data admit a unique gluing for every covering in the underlying Grothendieck topology. In contrast, trustworthiness is introduced here as a property of informational objects. A section is considered trustworthy only when it arises through descent from a non-trivial covering of distinct contexts.

The novelty of the proposal, therefore, does not lie in a new gluing theorem but in the interpretation of effective descent as a criterion for trustworthiness. By excluding trivial coverings consisting of a single context, trustworthiness is tied to the successful integration of independently verified contextual information. Consequently, a global section may exist without being trustworthy in the sense of Definition 1 if it is not supported by descent data arising from multiple contexts.

From this perspective, trustworthiness is not identified with the sheaf condition itself. Rather, it is characterized as effective descent over a non-trivial topology of contexts. Trustworthiness is therefore a property of descent data rather than isolated propositions.

To make the obstruction explicit, consider transition functions g_{ij} , defined on pairwise overlaps U_{ij} . The family $(g_{ij})_{i,j}$ forms a Čech 1-cocycle when

$$g_{jk} \circ g_{ij} = g_{ik}$$

on every triple overlap

$$U_{ijk} = U_i \times_U U_j \times_U U_k.$$

This cocycle condition guarantees that all local identifications are globally coherent.

Returning to the birth-record example, suppose that $g_{hr}(b_h) = b_r$ and $g_{rm}(b_r) = b_m$, but $g_{hm}(b_h) \neq b_m$.

Equivalently:

$$g_{rm} \circ g_{hr} \neq g_{hm}.$$

The cocycle condition, therefore, fails.

Although every record remains locally verified, the compatibility data do not determine the effective descent. Consequently, there exists no global section that simultaneously recovers all local records. The failure is not merely a disagreement between the data. It is a structural obstruction to coherent identification in contexts.

Proposition 1. *A section is trustworthy if and only if it arises as the gluing of a family of compatible locally verified sections, or equivalently, if and only if the associated Čech cocycle is trivial.*

Proof. If the local sections are compatible, the sheaf condition guarantees the existence of a unique glued section. Conversely, every global section induces compatible restrictions on the covering contexts. Hence, trustworthiness holds precisely when the associated descent data define a trivial cocycle. $\square$

Trustworthiness is, therefore, a structural property captured by effective descent. Mistrust corresponds to an obstruction to descent, represented by a non-trivial cocycle preventing the existence of a coherent global section. The issue is not merely the inconsistency of data, but the impossibility of maintaining a single global informational object across all contexts.

3.3. A Concrete Obstruction to Trustworthiness

The birth certificate example can be reformulated in terms of explicit transition data. Let U_h, U_r, U_m , represent hospital, registry, and municipal contexts. Suppose that local records b_h, b_r, b_m exist in each context. Assume that transition functions g_{hr}, g_{rm}, g_{hm} relate these records on pairwise overlaps. If $g_{rm} \circ g_{hr} = g_{hm}$, the compatibility data defines a trivial cocycle, and the records glue into a global section.

Suppose instead that $g_{rm} \circ g_{hr} \neq g_{hm}$. For example, differences in date encoding can produce incompatible identifications between hospital and municipal contexts. The family then fails the cocycle condition. Although every local record remains individually verified, there is no global section.

The resulting obstruction is not epistemic. No agent trusts any source. The failure is structural. This illustrates our central claim: trustworthiness depends on the existence of coherent descent data rather than on local verification alone.

4. Institutions and Global Model Existence

This section briefly connects the previous sheaf-theoretic account with the theory of institutions. It is intended for readers who are already familiar with institutional semantics [11, 12].

Let $\mathcal{I} = (\mathbf{Sign}, \mathbf{Sen}, \mathbf{Mod}, \models)$ be an institution [11]. In a contextualized setting, different contexts U_i may be associated with different signatures Σ_i and models $M_i \in \mathbf{Mod}(\Sigma_i)$, together with signature morphisms encoding translations between contexts [12]. Local validity corresponds to the satisfaction of sentences in each context: $M_i \models_{\Sigma_i} \varphi_i$. However, as in the sheaf-theoretic setting, local satisfiability does not guarantee the existence of a global model. Compatibility between contexts requires that models agree on overlaps up to reduct along signature morphisms.

A global model corresponds to a coherent family of models that can be integrated into a single structure with respect to all translations. This is precisely a descent condition. From this perspective, trustworthiness can be reformulated in institutional terms:

A system is trustworthy if and only if locally valid models can be integrated into a global model compatible with all signature translations.

Thus, the existence of a global model plays the same role as the existence of a global section in the sheaf-theoretic setting. Trust failures occur when such a model is absent despite local satisfiability.

This reinforces the structural interpretation: trust is not a property of individual models or sentences, but of the global compatibility of contextual interpretations.

In particular, the absence of a global model can be understood as the institutional counterpart of a descent obstruction.

5. Case study: A Sheaf-Theoretic Reconstruction of Electoral Trustworthiness

The electoral process can be understood as a paradigmatic instance of distributed information whose global validity depends on the coherence of locally generated data. In contemporary electronic voting systems, the underlying voting events — the actual choices made by voters — constitute the ontological truth-makers of the election. However, once cast, these events are no longer directly accessible. The system must therefore rely on a chain of informational proxies, including local ballot reports, transmission records and centralized aggregation outputs.

This structure naturally admits a sheaf-theoretic interpretation. Let $(\mathcal{C}, J)$ be a site of electoral contexts. The objects of $\mathcal{C}$ represent the stages of the electoral process, including local polling contexts U_i , transmission contexts T_i , a central aggregation context CEN , and a global context U corresponding to the election as a whole. Morphisms encode the flow of information across these stages, while the topology J captures the fact that the global electoral outcome is determined by its local components, through coverings. $\{U_i \rightarrow U\}$. In the following paragraph we remember that the notation $\mathcal{C}^{op}$, denotes the opposite category of $\mathcal{C}$, i.e., whose obtains from $\mathcal{C}$ by reversing the direction of the arrows between objects of $\mathcal{C}$. This helps to express in a more elegant way the restriction morphisms from a Context U to any of its coverings U_i . Thus, with $I = \{1, \dots, n\}$, the category $\mathcal{C}$ is depicted in figure 1.

$$\{u_i : U_i \rightarrow U\}_{i \in I} \in J(U)$$




Figure 1: The objects and morphisms of the category $\mathcal{C}$

The Grothendieck topology J is reflected in Figure 1 through the covering family $\{u_i : U_i \rightarrow U\}_{i \in I} \in J(U)$, together with the pullback objects representing intersections of local contexts. In a general site, every pair of covering morphisms $u_i : U_i \rightarrow U$ and $u_j : U_j \rightarrow U$ gives rise to a pullback $U_{ij} = U_i \times_U U_j$, and similarly for higher-order intersections, thereby expressing the compatibility conditions required for descent.

In the electoral setting, however, each voter belongs to exactly one polling station. Consequently, whenever $i \neq j$, the local polling contexts are disjoint and $U_i \times_U U_j \cong 0$, where 0 denotes the initial object of the category, corresponding to the empty context. The same observation applies to the transmission contexts T_i . On the other hand, since each transmission context is a subcontext of the central aggregation context, the corresponding pullback satisfies $T_i \times_{CEN} CEN \cong T_i$. Thus, the pullback structure induced by J reduces, in this application, to the expected inclusion relations among the electoral contexts, making the underlying category naturally equivalent to a finite poset.

The category $\mathcal{C}^{op}$ is then the category $\mathcal{C}$ with all morphisms or arrows reversed. For example, $U_i \leftarrow T_i \leftarrow CEN \leftarrow U$ is a diagram in $\mathcal{C}^{op}$. We can see, how this reversing of direction facilitates the restriction² functions $V(CEN) \rightarrow V(T_i)$, whenever V is a functor from $\mathcal{C}^{op}$ to **Sets**.

We define an electoral sheaf $V : \mathcal{C}^{op} \rightarrow \mathbf{Set}$, where $V(X)$ is the set of valid electoral proxies in context X . In particular, $V(U_i)$ contains local ballot reports, $V(T_i)$ contains transmitted data, $V(CEN)$ contains aggregated central data, and $V(U)$ represents the final electoral result. Restriction maps formalize the requirement that information be preserved along the electoral pipeline. Within this framework, local verification corresponds to validating proxies within individual contexts. A local ballot report may be considered valid within its polling station, and transmitted data may be considered valid within its communication channel. However, as argued in the previous sections, local validity does not suffice for global trustworthiness.

In sheaf-theoretic terms, Global trustworthiness corresponds to the existence of a global section $v \in V(U)$ obtained by gluing a family of local sections $v_i \in V(U_i)$ that agree on their overlaps. Thus, the final electoral result is not an independent object, but the unique global realization of compatible local proxies.

This perspective provides a precise structural interpretation of the two main forms of rupture observed in electoral systems. The first type of rupture occurs when local proxies themselves are rejected as valid representations of voting events. In this case, no shared family of local sections is available and, therefore, no descent data can be formed. The failure occurs before any attempt at global integration.

The second type of rupture occurs when local proxies are accepted, but the transmission or totalization stages are regarded as unreliable or opaque. Here, the issue is not the absence of local data, but the failure of compatibility across contexts. In sheaf-theoretic terms, the induced descent data do not satisfy the cocycle condition, resulting in a non-trivial Čech cocycle. As a consequence, no global section can be obtained by gluing. Therefore, electoral trustworthiness can be characterized as a property of effective descent: the system is trustworthy if and only if locally validated proxies define compatible descent data; equivalently, if the associated cocycle is trivial. Conversely, mistrust corresponds to an obstruction to descent, that is, to the impossibility of integrating locally verified information into a coherent global result.

This analysis supports the central thesis of this article. Trustworthiness is not an epistemic attitude nor a modal property attached to statements or agents. It is a structural invariant of the informational geometry. In the electoral setting, trust does not arise from belief in institutions or isolated verification procedures but from the existence of a coherent global section that integrates all locally validated proxies.

5.1. The Electoral Trustworthiness Sheaf-Theoretical Analysis

Definition 2 (Electoral Site). *The electoral process is modeled as a site $(\mathcal{C}, J)$, where:*

- $\mathcal{C}$ is a category whose objects are informational contexts of the electoral process, including:

²Remember that $CEN \cap T_i = T_i$

- local contexts U_i (polling stations/voting machines),
- transmission contexts T_i ,
- a central aggregation context CEN ,
- and a global context U representing the election as a whole;
- Morphisms in $\mathcal{C}$ represent information flow and refinement along the electoral pipeline

$$U_i \rightarrow T_i \rightarrow CEN \rightarrow U;$$

- J is a Grothendieck topology such that families $\{U_i \rightarrow U\}_{i \in I}$ form coverings of the global context, reflecting that the election is determined by its local voting contexts.

Definition 3 (Electoral Sheaf). *An electoral sheaf is a functor $V : \mathcal{C}^{op} \rightarrow \mathbf{Set}$ such that:*

- For each context $X \in \mathcal{C}$, $V(X)$ is the set of valid electoral proxies in X , including:
 - $V(U_i)$: local ballot reports,
 - $V(T_i)$: transmitted data,
 - $V(CEN)$: central aggregated data,
 - $V(U)$: the final electoral result;
- Restriction maps preserve informational consistency along the electoral pipeline;
- For every cover $\{U_i \rightarrow U\}_{i \in I}$, any family of local sections $v_i \in V(U_i)$ that is compatible with overlaps, $v_i|_{U_i \cap U_j} = v_j|_{U_i \cap U_j}$, admits a unique global section $v \in V(U)$ such that $v|_{U_i} = v_i$ for all i .

Proposition 2 (Electoral Trustworthiness as Global Section). *Let $(\mathcal{C}, J)$ be an electoral site, and let $V : \mathcal{C}^{op} \rightarrow \mathbf{Set}$ be an electoral sheaf. The electoral process is trustworthy in a global context U if and only if there exists a global section $v \in V(U)$ obtained by gluing a family of compatible local proxies $v_i \in V(U_i)$, for some covering $\{U_i \rightarrow U\}$ in J , such that $v|_{U_i} = v_i$ for all i , and $v_i|_{U_i \cap U_j} = v_j|_{U_i \cap U_j}$ for all i, j .*

Proof. By definition, the electoral sheaf assigns to each local context U_i the set of valid electoral proxies available at that stage of the process. Trustworthiness requires that these local proxies do not remain merely isolated or locally validated, but instead determine a coherent global electoral result.

Suppose first that the process is trustworthy on U . Then the final electoral result must be represented by a coherent information object that integrates the locally valid proxies across the electoral pipeline. Hence, there exists a family of local sections $v_i \in V(U_i)$ covering some areas $\{U_i \rightarrow U\}$ whose informational content agrees with overlaps. By the sheaf condition, these compatible local sections are uniquely glued to a global section $v \in V(U)$ such that $v|_{U_i} = v_i$ for all i .

Conversely, suppose that there exists a global section $v \in V(U)$ obtained by gluing a compatible family (v_i) . Then the final electoral result is not an isolated declaration, but the unique global realization of locally valid and mutually compatible electoral proxies. Therefore, the electoral process admits a coherent global informational object and is thus trustworthy. Hence, electoral trustworthiness is equivalent to the existence of a global section obtained by gluing compatible local proxies. $\square$

Theorem 1 (Trustworthiness and Triviality of the Electoral Cocycle). *Let $(\mathcal{C}, J)$ be an electoral site, let*

$$V : \mathcal{C}^{op} \rightarrow \mathbf{Set}$$

be an electoral sheaf, and let $\{U_i \rightarrow U\}$ be a covering of the global electoral context. Then the electoral process is trustworthy on U if and only if the descent data determined by the local proxies define a trivial Čech cocycle.

Proof. Assume first that the electoral process is trustworthy in relation to U . By the previous proposition, there exists a global section $v \in V(U)$ whose restrictions on the covering contexts give local sections $v_i := v|_{U_i} \in V(U_i)$. Since all these local sections arise as restrictions of the same global section, they necessarily agree on pairwise overlaps: $v_i|_{U_i \cap U_j} = v_j|_{U_i \cap U_j}$. Therefore, the associated descent datum is effective, and the corresponding Čech cocycle is trivial.

Conversely, suppose that the local proxies $v_i \in V(U_i)$ define descent data whose associated Čech cocycle is trivial. This means that the compatibility data on pairwise overlaps are coherent and satisfy the cocycle condition, so that the family (v_i) is compatible in the sense required by the sheaf condition. Since V is a sheaf, there exists a unique global section $v \in V(U)$ such that $v|_{U_i} = v_i$ for all i . By the previous proposition, the existence of such a glued global section implies that the electoral process is trustworthy over U . Therefore, electoral trustworthiness is strong if and only if the associated Čech cocycle is trivial. $\square$

Corollary 1 (Electoral Ruptures as Failures of Descent). *Let (C, J) be an electoral site and let*

$$V : C^{op} \rightarrow \mathbf{Set}$$

be an electoral sheaf. Then the two main forms of electoral rupture admit the following sheaf-theoretic characterization:

1. *If the local proxy in some polling context U_i is rejected as invalid, then no shared local section*

$$v_i \in V(U_i)$$

is available and therefore no descent datum can be formed.

2. *If the local proxies are accepted, but the transmission or totalization process fails to preserve informational compatibility, then the associated descent datum defines a non-trivial Čech cocycle, and no global section representing the final electoral result can be obtained by gluing.*

Hence, electoral rupture corresponds to either the impossibility of forming descent data at the local level or to an obstruction to effective descent at the systemic level.

Proof. For (1), suppose that in some local electoral context U_i the corresponding proxy — for instance, the voting machine record or its local ballot report — is not accepted as a valid representative of the underlying voting event. Then there is no commonly accepted local section $v_i \in V(U_i)$. But descent data requires, as a first step, a family of local sections over a covering of the global context. If one or more such local sections are missing because they are not recognized as valid proxies, then no descent datum is available. Thus, the failure occurs prior to any gluing condition: the local basis for descent is not established.

For (2), suppose now that local proxies $v_i \in V(U_i)$ are accepted for all relevant local contexts. The electoral process then proceeds through the transmission and totalization stages, which induce the compatibility data that relate these local proxies across overlaps and along the electoral pipeline. If these systemic stages fail to preserve informational consistency, then the induced transition data does not satisfy the cocycle condition. In other words, the family of accepted local proxies does not define effective descent data, but rather a non-trivial Čech cocycle.

By the theorem on trustworthiness and triviality of the electoral cocycle, a non-trivial cocycle prevents the existence of a global section $v \in V(U)$ whose restrictions recover the local proxies. Hence, no coherent final electoral result can be obtained by gluing. Therefore, the two forms of rupture correspond, respectively to: first, the absence of local descent data; and second, the failure of effective descent due to a non-trivial cocycle. $\square$

6. Trustworthiness as a Reason for Trust

The previous sections characterized trustworthiness as an effective descent. A natural question remains: Why should the existence of a global section provide a reason for trust?

Our answer is structural. Consider that trustworthiness is typically obtained by integrating evidence from multiple sources, each of which should be consistent. As we have argued, it is not enough that each individual "piece of evidence" be consistent; they must be integrable into a globally coherent information. Trustworthiness does not guarantee the truth. A globally coherent information object may still be false. However, trustworthiness guarantees something different: it guarantees that independently verified contextual manifestations can be integrated into a single informational object without generating structural conflicts. This provides a formal justification for relying on the information. The justification is neither epistemic nor psychological. It does not depend on beliefs, intentions, reputation scores, or institutional authority. It depends on a verifiable structural property: the existence of compatible descent data. From this perspective, trustworthiness constitutes a reason for trust because it certifies the absence of contextual obstructions to global coherence.

In distributed information systems, trustworthy information is not only locally validated. Information survives contextual integration. Consequently, effective descent can be interpreted as a formal criterion supporting trust decisions in distributed knowledge systems. The same holds for any multi-source Information System, i.e., when information originates from more than one source, even when all data are maintained within a single organizational or computational system.

6.1. Trustworthiness in AI Systems

Contemporary AI systems increasingly operate over distributed informational environments. Large Language Models (LLMs), retrieval-augmented generation (RAG) systems, knowledge graphs, multi-agent architectures, and decision-support platforms routinely integrate information originating from heterogeneous contexts.

In such settings, local verification alone is often insufficient to justify trust. A retrieved document may be internally consistent, a knowledge graph fragment may satisfy local constraints, and an individual agent may provide a justified recommendation. Nevertheless, the resulting system may not produce a coherent global interpretation of the information.

This observation motivates a distinction between *local reliability* and *global trustworthiness*. Reliability refers to the validity of information within a given context. Trustworthiness concerns the possibility of integrating information originating from multiple contexts into a coherent informational object.

Consider a retrieval-augmented language model. Different sources retrieved may individually support a generated answer. However, if the information provided by these sources cannot be coherently reconciled, the final answer may exhibit hallucinations, inconsistencies, or conflicting justifications. The issue is not necessarily the falsity of any individual source, but rather the absence of a globally coherent integration of locally validated information.

Unlike the use of sheaf-theoretic methods to characterize contextuality in distributed empirical systems [1], the present work employs descent theory to characterize trustworthiness as a structural property of information. In our setting, the central question is not whether local observations admit a contextual explanation, but whether locally verified informational objects determine compatible descent data that admit a coherent global realization.

Documents, database entries, knowledge graph fragments, or agent output correspond to local sections. Trustworthiness requires that these local informational objects determine compatible descent data admitting a coherent global realization.

This interpretation suggests a structural criterion for Trustworthy AI. An AI system should not be regarded as trustworthy merely because its individual informational components are verified locally. Rather, trustworthiness requires that these components admit coherent integration across contexts.

In this sense, effective descent provides a formal mechanism for distinguishing trustworthy information from merely locally validated information. The criterion is independent of confidence scores,

reputation measures, or subjective assessments. It depends solely on the existence of a coherent global informational structure obtained from compatible local information.

Consequently, sheaf-theoretic trustworthiness may provide a mathematical foundation for reasoning about information integration, provenance, and reliability in distributed AI systems.

7. Related Work

The present work connects and extends three major lines of research: contextual reasoning in knowledge representation, institutional semantics, and sheaf-theoretic approaches to logic and distributed systems.

The idea that truth is context-dependent originates in McCarthy’s formalization of contexts and has been extensively developed in knowledge representation through multi-context systems and contextualized ontologies. These frameworks allow multiple, potentially heterogeneous contexts to coexist, with formal mechanisms governing their interaction. However, in most such approaches, trust is treated as an external or epistemic notion, often represented as annotations, preferences, or agent-based relations. Our approach departs from this view by internalizing trustworthiness as a structural property of information itself, arising from compatibility across contexts rather than from external endorsement.

Institutions, introduced by Goguen and Burstall, provide an abstract model-theoretic framework for reasoning across heterogeneous logical systems. In institutional semantics, contexts are represented by signatures, and morphisms capture translations between them, preserving satisfaction. Our work is compatible with this perspective, but introduces an additional geometric layer. Although institutions capture *translation* between contexts, they do not intrinsically model *coverage*, *overlap*, or *gluing*. Grothendieck topologies enrich the category of contexts with a notion of distributed coverage, enabling the reconstruction of global objects from compatible local data. Thus, our approach can be seen as a geometric refinement of institutional semantics: trustworthiness is not merely preservation under translation but effective descent with respect to a topology of contexts.

Unlike Abramsky–Brandenburger [1], where sheaves capture contextuality via empirical models, here descent is used as a criterion for global informational coherence. In particular, Kripke–Joyal semantics provides an internal interpretation of logic in sheaf toposes. Our contribution extends this line of work by identifying *trustworthiness* with effective descent. While previous approaches use sheaves to model local consistency, we emphasize the passage from local verification to global certification as the central phenomenon.

Semiring provenance models capture how data is derived from multiple sources, tracking dependencies and justifications. Data integration frameworks address consistency across heterogeneous sources [13]. Our approach complements these by providing a structural criterion for global validity: information is trustworthy precisely when it admits compatible descent. In this sense, trustworthiness can be viewed as a global coherence condition on provenance.

8. Towards Quantitative Measures of Trustworthiness

The notion of trustworthiness introduced in this paper is qualitative: a knowledge base is considered trustworthy whenever locally compatible information can be integrated into a coherent global view. An interesting direction for future research is the development of quantitative measures that can express different degrees of trustworthiness.

A first possibility is based on the number of global sections. Given a sheaf $\mathcal{F}$ over a context space X , one may consider

$$T_{\Gamma}(\mathcal{F}) = |\Gamma(X, \mathcal{F})|,$$

where $\Gamma(X, \mathcal{F})$ denotes the set of global sections. Intuitively, this quantity measures the number of globally consistent interpretations that can be assembled from the available local information. The absence of global sections indicates a complete failure of integration, whereas multiple global sections indicate the existence of several coherent global viewpoints.

A second possibility is based on cohomological obstructions. Since non-trivial elements of the first Čech cohomology group represent failures of local data to glue globally, one may define a trustworthiness index inversely related to the size of the obstruction space, for example

$$T_H(\mathcal{F}) = \frac{1}{1 + \dim H^1(\mathcal{U}, \mathcal{F})}.$$

Under this interpretation, trustworthiness decreases as the complexity of global obstructions increases.

These two approaches capture different aspects of trust. The first reflects the abundance of globally coherent interpretations, whereas the second reflects the structural difficulty of obtaining such interpretations. Future work may investigate whether a combined measure can provide a more faithful quantitative characterization of trustworthiness, integrating both the existence of coherent global views and the complexity of the obstructions that prevent their formation.

More broadly, this perspective suggests a bridge between trust assessment and sheaf-theoretic invariants, opening the possibility of a quantitative theory of trust grounded in the geometric and cohomological properties of knowledge integration.

The quantitative perspective proposed here is loosely inspired by related uses of global sections and cohomological invariants in sheaf-theoretic approaches to contextuality and consistency [1, 14]. In particular, the existence of global sections has been used to characterize non-contextual behavior, while cohomological obstructions have been employed to detect failures of global compatibility. More generally, topological invariants, such as cohomology groups, are often interpreted as measures of structural complexity. These connections suggest that quantitative notions of trustworthiness based on global sections or cohomological obstructions may provide a promising direction for future investigation.

9. Conclusion

We have proposed a structural account of information trustworthiness grounded in sheaf-theoretic semantics. Departing from traditional epistemic and modal approaches, we argued that trustworthiness is not a primitive notion attached to agents, credentials, or propositions. Instead, it is an intrinsic property of informational objects. The currency-exchange example developed throughout the paper illustrates this perspective. Although each market locally verifies its own exchange rates, a coherent worldwide valuation exists only when the local valuations satisfy the compatibility conditions required for effective descent. This exemplifies the central thesis of the paper: structural trustworthiness arises from the successful integration of locally verified information into a coherent global informational object. Unlike modal approaches, which capture persistence under refinement within a single context, the present framework captures compatibility across distinct contexts. In particular, the existence of a global section depends on coherence conditions on overlaps, which cannot be reduced to modal invariance alone. This property is not epistemic: it does not depend on belief, endorsement, or subjective evaluation. Rather, it is structural and objective, determined entirely by the organization of information across contexts. More precisely, trustworthiness *has the power to supervene* local verification and compatibility. Any variation in the global trustworthiness of an informational object must arise from a variation in the compatibility of its locally verified instances.

Category theory has been proposed to address different aspects of ontology engineering since the initial editions and throughout many editions of the FOIS conference (e.g., [15, 16, 17]). This paper adds to that body of work. Within a sheaf-theoretic setting, this insight admits a precise geometric formulation. Informational contexts form a site (C, J) , and the information itself is organized as a sheaf on this site. Local verification corresponds to forcing in the internal logic, while global trustworthiness corresponds exactly to the existence of compatible descent data. From this perspective: (a) trustworthiness is the existence of a global section obtained by gluing compatible local data; (b) mistrust is the obstruction to descent; (c) conflict is incompatibility on overlaps; (d) Coherence is governed by the topology of contexts.

From the perspective of knowledge representation, this suggests a new foundational layer: beyond contextual reasoning and multi-context systems, a *geometry of information* in which global validity is determined by compatibility and gluing.

As future work, we foresee: (a) a cohomological analysis of obstruction classes associated with mistrust; (b) applications to decentralized identity and blockchain-based verification; (c) integration with ontology engineering frameworks and institutional semantics; (d) formal connections with provenance and data integration models. Moreover, we have proposed a qualitative notion of trustworthiness grounded in descent and global consistency. An interesting direction for future research is the development of quantitative notions of trustworthiness based on sheaf-theoretic invariants, such as the number of global sections or the size of cohomological obstructions. Inspired by analogous uses of global sections and cohomology in the study of contextuality, such measures may provide a finer understanding of trustworthiness, reflecting not only whether local information can be integrated globally, but also the complexity of the mechanisms that enable or obstruct such integration.

Declaration on Generative AI

During the preparation of this work, the authors used Claude (Anthropic), ChatGPT (OpenAI) and Grammarly to support grammar checking and sentence polishing. The authors reviewed and edited all content and take full responsibility for the publication's content.

Acknowledgements

The authors from PUC-Rio would like to thank the funding agencies CNPq, FAPERJ, and CAPES for their support. We also thank the Philosophy departments of UERJ and the team of CAPES/COFECUB 88881.878969/2023-01 project. An special acknowledgment to the team of FINEP Cybersecurity project at IME-RJ and ITA. The Authors from the University of Twente acknowledge the european funding agencies supporting them.

A. A Brief Introduction to the Category-Theoretic Notions Used in This Paper

This appendix provides an informal introduction to the categorical concepts employed throughout the paper. The presentation is intentionally elementary and uses examples taken from the trustworthiness scenarios discussed in the main text.

A.1. Categories

A category consists of: (i) objects; (ii) morphisms between objects; (iii) identity morphisms for each object and; (iv) an associative composition of morphisms. In this way, if a category has three objects a , b , and c , with morphisms $f_1 : a \rightarrow b$, $f_2 : a \rightarrow b$ and $g : b \rightarrow c$, then it also has $f_1 \circ g : a \rightarrow c$, $f_2 \circ g : a \rightarrow c$ and identities id_a, id_b and id_c , such that $id_a \circ f_1 = f_1$, $1 \circ id_b = f_1$, and so on. Since the composition is associative, $a \circ (f_1 \circ g) = f_1 \circ g$. An usual category is the category *Sets*, which has sets as objects and functions between sets as morphisms. Any graph induces a category by considering paths as morphisms and concatenation of paths as composition. This construction is also called the free-category induced by a graph.

In the present work, we deal with objects representing information contexts, and morphisms represent contextual relationships or information-preserving mappings between contexts.

For example, in the electoral scenario, we may consider objects $U_1, \dots, U_n, CEN, U$ representing the polling stations, the central counting system, and the election as a whole. Morphisms $U_i \rightarrow CEN$ represent the transmission of electoral information from the polling stations to the central authority.

Similarly, in the birth-record example, objects represent institutional contexts such as hospitals, civil registries, and municipal databases, while morphisms represent information transfers between them.

A.2. Sites and Coverings

A category becomes a site when it is equipped with a Grothendieck topology.

Intuitively, a Grothendieck topology specifies which families of contexts are sufficient to reconstruct another context.

For example, in an electoral setting, the family $\{U_i \rightarrow U\}_{i \in I}$ forms a covering because the global election result is determined by the collection of local voting contexts. Similarly, in the example of the birth-record, the contexts U_h , U_r , U_m corresponding to the hospital, the registry, and the municipal records may jointly cover the informational context associated with an individual's identity. Thus, coverings formalize the idea that global information may be reconstructed from local sources.

A.3. Presheaves

A presheaf assigns information to each context. Formally,

$$F : \mathcal{C}^{op} \rightarrow \mathbf{Set}.$$

For every context (U) , $F(U)$ is the set of information objects available in (U) . For every morphism $f : V \rightarrow U$, the pre-sheaf provides a restriction map $F(f) : F(U) \rightarrow F(V)$.

In the electoral example:

- $V(U_i)$ contains local ballot reports;
- $V(CEN)$ contains aggregated electoral data;
- $V(U)$ contains global electoral results.

Restriction maps express how the information available globally can be viewed locally.

A.4. Sheaves

A sheaf is a pre-sheaf satisfying a gluing condition. Suppose a context (U) is covered by $\{U_i \rightarrow U\}_{i \in I}$. Assume that each context U_i contains a local information object $s_i \in F(U_i)$. If the local objects agree on all overlaps, $s_i|_{U_i \times_U U_j} = s_j|_{U_i \times_U U_j}$ then the sheaf condition guarantees the existence of a unique global object $s \in F(U)$ such that $s|_{U_i} = s_i$. This formalizes the transition from local information to global information.

A.5. Descent

Descent theory studies when local information determines a global object. The collection $(s_i)_{i \in I}$ of compatible local sections is called descent data. The Descent is effective whenever these local sections uniquely determine a global section.

Within the framework proposed in this paper, trustworthiness is identified with effective descent. In other words, information is trustworthy when locally verified informational objects can be integrated into a coherent global informational object.

A.6. Cohomological Obstructions

Not every family of local informational objects admits a global realization. Suppose that local records exist in the contexts U_h , U_r , U_m corresponding to hospital, registry, and municipal databases. Even if every record is locally verified, inconsistencies in identification policies may prevent the existence of

a coherent global record. Such failures are represented mathematically by non-trivial Čech cocycles. From the perspective developed in this paper, these cocycles represent obstructions to descent.

Consequently: (1) effective descent corresponds to trustworthiness; (2) obstructions to descent correspond to mistrust.

A.7. Summary

The categorical structures employed in this paper can be interpreted as follows:

Category	Collection of informational contexts
Morphisms	Contextual relationships
Coverings	Families of information sources
Presheaf	Assignment of information to contexts
Sheaf	Compatible local information can be glued
Descent	Local-to-global integration
Global Section	Globally trustworthy information
Čech Obstruction	Structural source of mistrust

This interpretation motivates the central claim of the paper: trustworthiness is fundamentally a local-to-global property of multi-source information.

References

- [1] S. Abramsky, A. Brandenburger, The sheaf-theoretic structure of non-locality and contextuality, *New Journal of Physics* 13 (2011).
- [2] A. C. Baier, Trust and antitrust, *Ethics* 96 (1986) 231–260.
- [3] O. O’Neill, *The BBC Reith Lectures 2002: A Question of Trust*, Cambridge University Press, Cambridge, 2002.
- [4] R. Hardin, *Trust and Trustworthiness*, Russell Sage Foundation, New York, 2002.
- [5] L. Floridi, *The Philosophy of Information*, Oxford University Press, Oxford, 2011.
- [6] P. Johnstone, *Sketches of an Elephant: A Topos Theory Compendium*, OUP, 2002.
- [7] R. Goldblatt, *Topoi: The Categorical Analysis of Logic*, North-Holland, 1984.
- [8] J. C. Hull, *Options, Futures, and Other Derivatives*, 10 ed., Pearson, 2018.
- [9] C. S. Eun, B. G. Resnick, *International Financial Management*, 9 ed., McGraw–Hill Education, 2021.
- [10] R. K. Lyons, *The Microstructure Approach to Exchange Rates*, MIT Press, Cambridge, MA, 2001.
- [11] J. A. Goguen, R. M. Burstall, Institutions: Abstract model theory for specification and programming, *Journal of the ACM* 39 (1992) 95–146.
- [12] A. Tarlecki, Moving between logical systems, in: S. Abramsky, D. M. Gabbay, T. S. E. Maibaum (Eds.), *Handbook of Logic in Computer Science*, 4, OUP, 1996.
- [13] T. Green, G. Karvounarakis, V. Tannen, Provenance semirings, *PODS* (2007).
- [14] S. Abramsky, S. Mansfield, R. S. Barbosa, The cohomology of non-locality and contextuality, *Electronic Proceedings in Theoretical Computer Science* 236 (2017) 1–14.
- [15] M. Johnson, C. N. G. Dampney, On category theory as a (meta) ontology for information systems research, in: *Proc. of II FOIS 2001*, 2001.
- [16] A. Zimmermann, M. Krötzsch, J. Euzenat, P. Hitzler, Formalizing ontology alignment and its operations with category theory, in: *Proc. of IV FOIS*, 2006.
- [17] N. Asher, What have we learned in formal semantics about ontology?, in: *Formal Ontology in Information Systems*, IOS Press, 2014.