forecasting

MDPI

Article

Singular Design Foresight: A Foundational Method for Auditable Anticipation and Decision Closure

Pablo Lara-Navarra [1,2,*], Antonia Ferrer-Sapena [3] and Enrique A. Sánchez-Pérez [3]

[1] Information and Communication Sciences Department, Universitat Oberta de Catalunya (UOC), Rambla del Poblenou 156, 08018 Barcelona, Spain
[2] Open Future Lab, Universidad Tecnológica ECOTEC, Km. 13.5, Samborondón EC092302, Ecuador
[3] Instituto Universitario de Matemática Pura y Aplicada, Universitat Politècnica de València (UPV), Camino de Vera s/n, 46022 Valencia, Spain; anfersa@upv.es (A.F.-S.); easancpe@mat.upv.es (E.A.S.-P.)
* Correspondence: plara@uoc.edu

Highlights

What are the main findings?

- Singular Design Foresight (SDF) formalizes Design Foresight as a replicable methodological cycle, integrating semantic projections to transform qualitative signals into traceable and analyzable indicators.
- The Stakeholder Viability Principle (SVP) defines a viability-based filtering mechanism for delimiting relevant futures, while Social Singularity identifies context-specific conditions of critical transition.
- The framework introduces a Design Science Research-aligned evaluation architecture (validity, utility, attribution) and specifies conditions for comparative and sensitivity-based assessment.

What are the implications of the main findings?

- Practitioners gain a framework for auditable decision closure under uncertainty, complementing qualitative foresight with structured and traceable representations.
- The method enables transparent linkage between evidence, interpretation, and decision, supporting more explicit and discussable strategic processes in VUCA contexts.
- The proposed architecture provides a basis for cumulative empirical research, enabling future comparative and attribution-oriented studies in Design Foresight.

Abstract

Singular Design Foresight (SDF) is proposed as a foundational methodological framework for advancing Design Foresight (DF) toward a more explicit, traceable, and evaluable scientific discipline. The framework formalizes DF as a structured cycle in which qualitative foresight inputs—such as signals, trends, and expert interpretations—are progressively transformed into analyzable representations that support decision closure under conditions of structural uncertainty. SDF combines an expert-defined conceptual universe with semantic projections to relate textual and contextual evidence to anticipatory constructs, enabling the generation of traceable indicators and structured configurations of viable futures. Within this architecture, the Stakeholder Viability Principle (SVP) functions as a filtering mechanism that delimits relevant futures according to continuity, agency, and axiological coherence, while Social Singularity captures context-specific critical transitions that shape when and why decision closure becomes necessary. The framework is organized in alignment with Design Science Research (DSR), adopting an evaluation logic centered

Check for updates

Academic Editor: Sonia Leva

Received: 6 April 2026
Revised: 20 April 2026
Accepted: 25 April 2026
Published: 2 May 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

on validity, utility, and attribution. Rather than presenting conclusive system-level validation, the article synthesizes summative evidence from previously published studies on semantic projections, singularity detection, and mixed expert–corpus foresight applications to support the plausibility, internal coherence, and operational feasibility of the proposed framework, while delimiting full integrated validation as a future research objective. SDF does not aim to provide deterministic prediction; instead, it enables auditable anticipatory representations and justified closure under uncertainty. In this sense, the framework is compatible with forecasting understood as the production of evaluable anticipations under explicit assumptions, while preserving the interpretive and situated character of strategic decision-making.

Keywords: Design Foresight; semantic projections; anticipatory representation; viability principle; decision closure; Design Science Research; Social Singularity; strategic decision-making

1. Introduction

Contemporary organizations increasingly make strategic decisions under conditions in which discontinuity, informational turbulence, and accelerated knowledge obsolescence undermine the adequacy of linear planning based on past extrapolation. In such environments, commonly described through the language of VUCA and, more recently, BANI, the central challenge is no longer simply to improve prediction, but to sustain justified action when uncertainty is structural, evidence is incomplete, and interpretation remains contested [1,2]. Under these conditions, anticipatory practice becomes less a matter of projecting stable trends forward than of organizing decision and intervention responsibly in the presence of ambiguity, volatility, and non-linearity.

Within this broader landscape, Strategic Design and Futures Studies have developed partially complementary responses. Strategic Design contributes intervention-oriented capabilities, organizational sense-making, and frameworks for linking innovation, value creation, and decision processes to design practice [3–5]. Futures Studies and strategic foresight, in turn, provide conceptual and methodological resources for examining alternative configurations of the future, including possible, plausible, probable, and preferable futures, as well as structured approaches to scenarios and long-range anticipation [6–9]. Yet the convergence of these traditions has not fully matured into a cumulative Design Foresight (DF) program with an explicit ontology, a clarified epistemology, and evaluative procedures capable of supporting explanation, comparison, and transferability beyond isolated interventions.

This limitation is related to a longstanding tension in foresight research between practical utility and disciplinary consolidation. Much of the field has developed under a strong pragmatic orientation, in which validity is often inferred from usefulness, participation, or organizational uptake rather than from explicit criteria of explanation, attribution, and comparative assessment. This generates a theoretical and methodological deficit. If theory is not merely an inventory of variables, references, or empirical observations, but a structured explanatory contribution [10], then the distinction between a theory of foresight and theorizing within foresight becomes decisive for the maturation of the field [11]. The problem is therefore not that foresight lacks practical relevance, but that its cumulative scientific architecture remains uneven: methods are often context-sensitive and productive in practice, yet difficult to compare, replicate, or evaluate in terms of methodological contribution as distinct from leadership, institutional setting, or consensus dynamics [12–14].

From the standpoint of forecasting, this deficit becomes particularly visible. Forecasting, understood broadly, requires explicit assumptions, observable outputs, and criteria by which anticipations may be evaluated, contrasted, or refined. By contrast, many foresight and design-oriented future practices generate rich narratives, exploratory scenarios, or speculative artifacts without clearly specifying how qualitative signals are transformed into analyzable variables, how such transformations can be audited, or how alternative procedures might be compared. This is not an argument for reducing foresight to deterministic prediction. Rather, it identifies a methodological requirement: if DF is to claim disciplinary status, it must articulate procedures through which signals, trends, and interpretive judgments become traceable objects of analysis, and through which decision closure under uncertainty can be justified in a disciplined manner. Recent work in hybrid trend analysis, conceptual analysis, semantic projections, and digital mixed methods suggests that such instrumentation is both feasible and increasingly necessary [15–18].

A parallel tension can be observed in design-oriented future-making. Over the last decade, speculative design, design fiction, experiential futures, and related currents have considerably expanded the repertoire through which futures may be imagined, materialized, and debated [19–23]. These approaches have been especially valuable in displacing default assumptions, exposing hidden implications, and broadening the range of discussable alternatives. Yet precisely because they expand the space of conceivable futures, they intensify a disciplinary question that remains insufficiently resolved: by what explicit and auditable criteria should some future configurations be selected, filtered, and translated into strategic commitment while others remain exploratory? This question becomes even more pressing in light of critical perspectives on defuturing, which remind us that design does not merely open futures; it also closes them, often by legitimizing some pathways while excluding others [24]. A mature Design Foresight framework therefore cannot stop at imaginative expansion; it must also account for the legitimacy, discipline, and traceability of closure.

It is at this point that DF can be reformulated as a response to a double requirement. On the one hand, it must enlarge the space of futures sufficiently to avoid default, inherited, or prematurely stabilized trajectories. On the other hand, it must enable justified closure so that organizations can commit to action under uncertainty. Recent work has described DF precisely in terms of the articulation between futurization and de-futurization, that is, between the opening of alternatives and the disciplined narrowing required for design and strategic commitment [25]. In the present article, de-futurization is not understood as the arbitrary cancellation of plurality, but as a progressive and accountable reduction in uncertainty that renders decision possible. This logic resonates with the cone-of-uncertainty tradition, according to which uncertainty narrows as inquiry, specification, and commitment advance, even if it never disappears completely [26].

This article deliberately situates itself at the disciplinary level, that is, at the level of a theory of Design Foresight rather than at the level of a single intervention protocol. The prospective-design decision episode is adopted as the principal unit of analysis because it is there that anticipation becomes consequential: signals must be interpreted, candidate trajectories filtered, and closure enacted despite incomplete and revisable knowledge. On this basis, we propose Singular Design Foresight (SDF) as a foundational methodological framework for DF. Its purpose is neither to offer yet another operational platform for strategic navigation nor to claim conclusive empirical validation of a fully integrated system. Rather, it seeks to specify the scientific conditions under which DF may become more explicit, auditable, and cumulatively assessable.

Ontologically, SDF introduces the Stakeholder Viability Principle (SVP) as a filtering principle that delimits a viability cone composed of futures compatible with the continuity,

agency, and axiological coherence of critical stakeholders. In the present formulation, critical stakeholders are not defined abstractly, but in relation to the organizational problem under consideration: they are the actors recognized as capable of influencing, validating, or enacting decisions within the episode, and their differentiated assessments help structure the viable future space. At the level of change dynamics, SDF uses social singularity not as a typology of events, but as the designation for a context-specific rupture condition in which ordinary practical reasoning becomes insufficient and decision-making must be reorganized around signals of critical transition. In this sense, the concept is closer to the logic of critical transition than to a merely metaphorical vocabulary of disruption [27–29].

Epistemologically and instrumentally, SDF proposes a hybrid architecture in which expert-constructed conceptual universes converge with computationally supported semantic projections in order to transform discursive material into analyzable relational objects. Within this architecture, semantic intensity indicators, transition detection, and shortlists of viable futures become forecasting-compatible outputs, not because DF is reduced to narrow prediction, but because its anticipatory claims become more explicit, traceable, and open to evaluation under defined assumptions [16,18,30]. Algorithmic instrumentation is incorporated here as a support for observability and auditability, not as a substitute for deliberation.

The argument of the paper is deliberately modest in one sense and ambitious in another. It is modest because it does not claim that the full SDF architecture has already been conclusively validated as an integrated forecasting system, nor that it empirically outperforms consensus-based or alternative approaches within this article. The empirical material mobilized here is illustrative and summative: it supports the plausibility, operational feasibility, and evaluability of key components, but it does not exhaust the program of contrast required for system-level validation. At the same time, the argument is ambitious because it seeks to reposition DF as an applied scientific program with explicit ontological commitments, operationalizable epistemic instruments, and a pathway toward more rigorous evaluation. In this sense, the paper aligns with the logic of Design Science Research, not as a report of completed artifact evaluation, but as a disciplined effort to specify the conceptual and methodological architecture that such evaluation would require [31–33].

Against this background, the next section states the research aim, research questions, and contributions through which the scope of the article is specified.

2. Research Aim, Research Questions and Contributions

The aim of this article is to establish the foundational conditions under which Design Foresight (DF) can be articulated as a scientifically evaluable methodological framework capable of supporting cumulative knowledge, explicit attribution, and disciplined decision closure under uncertainty [31–33]. More specifically, the paper seeks to move DF from a repertoire of situated practices—often validated through utility, participation, or contextual success—toward a structured anticipatory architecture in which futures are systematically opened through the exploration of signals and trajectories and subsequently narrowed through explicit and justifiable criteria that enable action under conditions of irreducible uncertainty [11–13]. In this sense, the contribution of the article is not to provide a finalized or empirically validated forecasting system, but to define a conceptual and methodological foundation through which DF may become more explicit, auditable, and progressively evaluable in alignment with the requirements of an applied scientific program [31,33].

The article addresses a triple gap at the intersection of design, foresight, and forecasting. First, there is a disciplinary gap: although Strategic Design and Futures Studies have evolved in parallel, their convergence has not yet resulted in a consolidated Design Foresight framework with explicit ontological commitments, clarified epistemological posi-

tioning, and cumulative mechanisms of evaluation and transfer [3–5,11]. Second, there is a methodological gap: many foresight and design-futures practices generate rich narratives, scenarios, or artifacts, yet lack explicit procedures for transforming qualitative signals into traceable and analyzable representations, and do not provide sufficient criteria for comparing alternative procedures or distinguishing methodological effects from contextual factors [10,12,14,15]. Third, there is an evaluative gap from the standpoint of forecasting: Design Foresight still lacks sufficiently explicit outputs, observable indicators, and procedures for assessing anticipatory claims, which limits its contribution to forecasting understood as a field concerned with evaluable anticipations [7,9,11].

In response to these gaps, the article is guided by five research questions:

RQ1. Under what ontological and disciplinary conditions can Design Foresight be articulated as a cumulative scientific field rather than as a loose conjunction of design and foresight practices [10,11]?

RQ2. Through what formalizable mechanism can Design Foresight move from future-space expansion (futurization) to justified decisional closure (de-futurization) under structural uncertainty [24,25]?

RQ3. How do stakeholder viability, signals, and context-specific critical transitions delimit the set of relevant futures within a prospective-design decision episode [27,28,34]?

RQ4. How can qualitative foresight inputs—such as narratives, signals, and expert judgments—be transformed into traceable, analyzable, and auditable representations through hybrid conceptual–computational instrumentation [15,16,18]?

RQ5. What kinds of outputs and evaluative criteria make Design Foresight compatible with forecasting as evaluable anticipation without reducing it to deterministic prediction [12–14]?

In response to these questions, the article proposes Singular Design Foresight (SDF) as a foundational methodological framework. Its contributions are articulated at four interrelated levels.

First, at the ontological level, the paper introduces the Stakeholder Viability Principle (SVP) as a filtering principle that defines a viability cone within the broader future space, delimiting which futures are relevant within a given decision episode and grounding this delimitation in the continuity, agency, and axiological coherence of critical stakeholders [25,26].

Second, at the conceptual–mechanistic level, the article formalizes the passage from futurization to de-futurization as the core movement through which Design Foresight proceeds from exploratory expansion to justified closure. Within this architecture, social singularity is redefined not as a typology of events, but as a context-specific rupture condition in which existing interpretive and practical frameworks become insufficient, signals of transition accumulate, and decision-making must be reorganized under heightened uncertainty [27–29].

Third, at the epistemological–instrumental level, the paper proposes a hybrid architecture in which expert-constructed conceptual universes converge with computationally supported semantic projections in order to transform discursive material into relational and analyzable objects [15,16,18]. This architecture enables the generation of semantic intensity indicators, the identification of transition-relevant signals, and the construction of shortlists of viable futures [16,18,30]. These outputs do not constitute point predictions, but they render Design Foresight forecasting-compatible by making anticipatory claims more explicit, traceable, and potentially comparable under defined assumptions [7,9,13].

Fourth, at the methodological–evaluative level, the article formalizes the SDF cycle as a structured sequence linking future-space opening, signal processing, viability filtering, and decision closure under uncertainty, while also outlining a programmatic evaluation

pathway aligned with the logic of Design Science Research. This includes the need for artifact-based evaluation, comparative designs, and sensitivity- and attribution-oriented analysis [12,31–33].

It is important to delimit the scope of these contributions. This work does not claim empirical validation of the integrated SDF framework, nor does it demonstrate superiority over consensus-based or alternative approaches within the present study. Rather, it synthesizes summative component-level evidence reported in prior studies—on semantic instrumentation, singularity detection, hybrid signal extraction, and mixed expert–corpus applications—and integrates these elements into a coherent disciplinary framework. Its contribution is therefore foundational and programmatic: it establishes the conceptual architecture and evaluative conditions under which future empirical contrast and cumulative assessment of Design Foresight may be more rigorously pursued.

3. Ontological and Epistemological Foundations of Singular Design Foresight

The articulation of Design Foresight (DF) as a scientifically evaluable framework requires making explicit its ontological commitments, that is, the assumptions about what constitutes a relevant future within a decision episode and how such relevance is delimited. In the absence of explicit criteria, the transition from exploration to decision closure risks becoming arbitrary, driven by implicit preferences, power asymmetries, or unexamined feasibility constraints. This problem is particularly acute in contexts characterized by wicked problems, where neither the problem nor its solution can be fully specified in advance and where each intervention reconfigures the problem space itself [34].

In response to this limitation, Singular Design Foresight (SDF) introduces the Stakeholder Viability Principle (SVP) as a foundational ontological construct. The SVP operates as a filtering principle that delimits a subset of futures within the broader space of possibilities, defining a viability cone [35,36] composed of futures that are compatible with the continuity, agency, and axiological coherence of critical stakeholders. Rather than assuming that all conceivable futures are equally relevant, the SVP establishes that only those futures that can be sustained, enacted, and justified by actors with decision-making capacity should be considered as candidates for decision closure.

In the present formulation, critical stakeholders are not defined abstractly but in relation to the organizational problem under consideration. They are the actors recognized as capable of influencing, validating, or enacting decisions within the given context. Operationally, these stakeholders participate in the evaluation of emerging trends and signals through structured scoring processes, typically using ordinal scales (e.g., 0–5) that capture perceived relevance, urgency, or impact.

The constitution of expert panels within the SDF process has been operationalized in prior published applications, to which the present article refers for full methodological detail [17,30]. Across these implementations, expert panel design has followed a consistent set of criteria that can be summarized along four dimensions. First, thematic coverage and active domain engagement: participants are selected on the basis of demonstrable and current involvement in the field under analysis, evidenced through active projects, publications, or institutional roles directly linked to the domain, rather than through nominal expertise alone [17]. Second, institutional and role diversity: panels have combined academic experts from multiple institutional positions—including planning directors, postgraduate deans, and vice-rectors—with operational academic management staff—coordination, area direction, and general management—so that the panels capture both strategic and procedural knowledge relevant to the decision episode [17,30]. Third, transdisciplinary and trans-generational composition, including explicit gender balance [17]: in

practice, this has meant panels spanning age ranges from early-career to senior profiles, with parity criteria applied at the point of selection, operatively targeting a maximum of 60% representation by any single gender within each panel configuration [17]. Fourth, quantitative selection thresholds ensure that only signals receiving sufficiently convergent expert evaluations—specifically, mean scores above 4 and standard deviations below 1 on a 0–5 scale, with inter-rater cosine similarity values consistently above 0.99—are retained for subsequent analytical phases [30]. These criteria are not presented as fully standardized requirements for all SDF applications; the specific configuration of each panel must be adapted to the institutional context and decision episode. However, the present article does not develop the full expert panel design protocol here, as this has been documented and empirically applied in the cited prior work, which provides the operational reference for this component of the SDF cycle. The evaluations conducted by these panels follow the structured scoring and dispersion logic described in the next paragraph.

In prior mixed expert–corpus applications, such evaluations have been combined with explicit selection rules and dispersion criteria in order to identify robust signals and trends. These assessments make it possible to detect patterns of convergence and divergence, which can then be examined through clustering procedures to distinguish groups of stakeholders according to their degree of alignment, strategic interest, or capacity for action. The resulting clusters may subsequently be refined through deliberative processes, such as facilitated workshops, in which differences are negotiated in light of institutional roles, available capabilities, and decision relevance.

This procedure enables the SVP to function not as a purely normative abstraction, but as a methodologically grounded filtering mechanism. Inclusion or exclusion of futures is not determined by abstract desirability alone, but by the convergence of stakeholder evaluations and the feasibility of coordinated action. In this sense, viability is not reduced to technical feasibility or market logic; it is defined through the interplay between continuity (the capacity to sustain organizational trajectories), agency (the capacity to mobilize resources and enact change), and axiological coherence (the alignment between envisioned futures, institutional values, and capability development). These dimensions do not yet constitute a fully standardized metric; rather, they are proposed as evaluative axes that can be progressively formalized, calibrated, and tested in future research.

Complementing this ontological layer, SDF introduces social singularity as a conceptual device for capturing moments of critical transition within a decision context [37]. In contrast to approaches that treat singularities as typologies of events or extreme scenarios, SDF defines social singularity as a context-specific rupture condition in which existing interpretive frameworks and routines of action become insufficient. Such conditions emerge when accumulated signals—derived from environmental scanning, trend analysis, corpus-based observation, and stakeholder interpretation—indicate that incremental adjustment is no longer adequate and that a qualitative reorganization of decision-making is required [27–29].

Under this perspective, singularity does not refer to a predefined class of phenomena, but to a functional condition within the decision process. It marks the point at which uncertainty becomes decisionally critical: the organization can no longer defer commitment without incurring increasing strategic risk. In this sense, singularity serves as a bridge between signal interpretation and decision closure. It does not predict a specific outcome; rather, it identifies the threshold at which the space of possible futures must be actively filtered and reduced. Importantly, prior published studies have already shown that rupture-related differentiation can be operationalized through observable variables, multidimensional models, and classification procedures, demonstrating that singularity-related transitions need not remain purely metaphorical or speculative constructs.

This interpretation can be further clarified by distinguishing between social singularity and geometric singularity within the broader logic of futures representation. Social singularity refers to the empirically grounded rupture in the problem context, while geometric singularity designates the limiting condition in which the space of alternative futures collapses into a realized trajectory. While the latter remains a conceptual abstraction, it reinforces the idea that the value of DF lies not in eliminating uncertainty, but in reducing it to a level that allows for accountable decision closure. In this sense, the SVP and singularity operate jointly: the former delimits which futures remain viable, while the latter signals when the reduction in the future space becomes decisionally unavoidable.

At the epistemological and instrumental level, SDF proposes a hybrid architecture that connects expert judgment with computational analysis. Specifically, it combines expert-constructed conceptual universes with AI-supported semantic projections in order to transform qualitative material—such as narratives, documents, and weak or emerging signals—into relational structures that can be analyzed, compared, and traced. In operational terms, semantic projections can be instantiated as bounded affinity coefficients over a defined corpus, expressing the degree to which a signal and a concept share a semantic environment. This affinity may be estimated through co-occurrence-based measures within a documentary repository, producing normalized values in the interval [0, 1] that support the construction of semantic intensity indicators and the identification of emerging patterns [16,18].

The operational instantiation of semantic projections as forecasting-compatible instruments has been developed and empirically validated in prior published work by the present research team, to which the reader is referred for full algorithmic detail Figure 1 [16,18]. Within the SDF architecture, semantic projections are formalized as bounded affinity coefficients over a defined documentary corpus: for a signal A and a concept B belonging to a conceptual universe U, the projection is computed as the ratio of co-occurring documents to total documents containing A. These individual projections are aggregated into a composite representation vector PU(A) across all elements of U, enabling the construction of semantic intensity indicators that express the degree to which a signal participates in a given anticipatory structure. The resulting values are not interpretive approximations but computable quantities: prior implementations have produced normalized trend vectors—for instance, PU(Trend) = (0.806, 0.081, 0.010, 0.586, 0.000, 0.001) across a six-concept universe—and composite indices such as a Relative Trend Index derived via McShane–Whitney Lipschitz extension [16]. These outputs constitute traceable, evaluable anticipatory representations in the sense required by forecasting: they are explicit, reproducible under the same corpus conditions, and open to comparison across decision episodes. The present article does not reproduce this technical development but builds upon it, incorporating semantic projections as the instrumental layer through which SDF transforms qualitative foresight inputs into analyzable objects.

Signal extraction, in this framework, is not treated as an ad hoc interpretive activity but as a structured hybrid process. Prior implementations have operationalized it through multi-phase pipelines that include corpus construction, information systematization, co-occurrence analysis, clustering of related entities, and expert validation. This makes it possible to move from dispersed qualitative material to traceable trend matrices and interpretable semantic structures. In this sense, SDF does not replace interpretive judgment with algorithmic determination; rather, it introduces a layer of observability and traceability that allows discursive processes to be examined under explicit criteria.

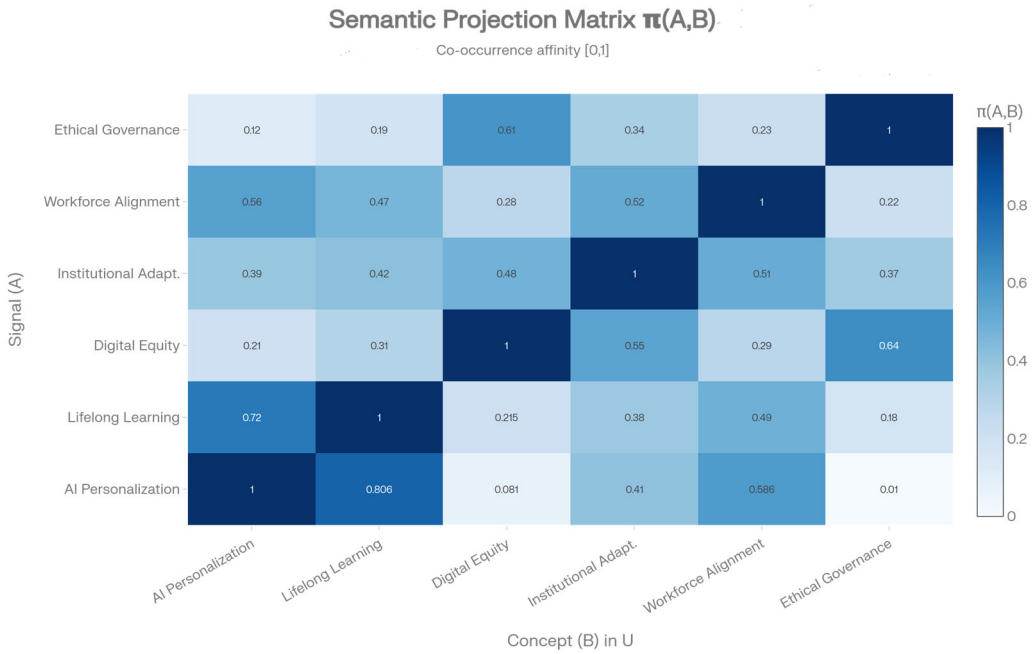


Figure 1. Semantic projection matrix π(A,B) for the higher education SDF application (Peru, 2025–2035). Each cell represents the normalized co-occurrence affinity coefficient between signal A (row) and concept B (column) within the conceptual universe U (n = 6 dimensions). Values ∈; diagonal entries represent signal self-affinity. Source: derived from Lara-Navarra et al. (2025) [30] and López Bordao et al. (2025) [18].

By integrating conceptual analysis, expert validation, and computational techniques, SDF aligns with digital mixed-methods approaches that seek to bridge qualitative interpretation and quantitative analysis without reducing one to the other [15,38]. This also helps clarify the forecasting relevance of the framework: SDF does not aim at deterministic prediction, but at the production of evaluable anticipatory representations under explicit assumptions, through which signals, transitions, and viable future configurations can be examined in a systematic and auditable way.

Taken together, the SVP, social singularity, and semantic projections define the core architecture of SDF. The SVP delimits the relevant future space, singularity identifies the conditions under which closure becomes necessary, and semantic projections provide the means to render signals traceable and analyzable. This triadic structure does not constitute a finalized model, but a conceptual and methodological scaffold that synthesizes summative evidence accumulated across previously published studies and enables DF to move from implicit practice toward explicit, auditable, and progressively evaluable knowledge production.

4. Theoretical and Disciplinary Grounding of Design Foresight

Strategic design has progressively consolidated as an approach that exceeds aesthetics and bounded problem-solving in order to operate within complex organizational and sociotechnical systems where objectives, constraints, and criteria of success remain contested. From this perspective, design is not limited to solving predefined problems, but intervenes in situations whose formulation is itself unstable. The notion of wicked problems is therefore particularly relevant, since it positions design as a discipline concerned with situations that do not admit definitive formulations or optimal solutions, but instead require situated decisions that simultaneously transform the problem space in which they intervene [34]. Within this framework, strategic design may be understood as a family of organizational practices that connect intention, meaning, and action through artifacts, narratives, and

prototypes, thereby facilitating alignment, differentiation, and value creation [3–5]. For Design Foresight (DF), this has a direct methodological consequence: if design intervenes in the problem while helping define it, then the passage from exploring possible futures to enacting decisional closure cannot remain tacit or intuitive, but requires explicit criteria of justification and evaluation.

From Futures Studies and strategic foresight, DF inherits a broad repertoire of methods for exploring plural futures, including scenarios, horizon scanning, Delphi, and backcasting, as well as conceptual distinctions among possible, plausible, probable, and preferable futures [6,7,9]. Yet the scientific consolidation of foresight continues to face a structural tension between action-orientation and the requirements of cumulative theory, explanatory mechanisms, and validity criteria. This tension is captured in the distinction between a theory of foresight and theorizing within foresight: a field matures scientifically not by merely accumulating techniques or cases, but by developing explicit epistemology, process theory, and impact theory [11]. In this sense, pragmatism alone is insufficient. If validity is equated only with utility, participation, or uptake, then explanation weakens and attribution becomes difficult: even when an intervention appears useful, it remains unclear which part of its effect derives from the method, which from the context, and which from the social dynamics of consensus or leadership [10,12,13].

This reading is reinforced by the historical development of foresight itself. What began largely as an anticipatory practice for supporting decisions increasingly aspires to function as an applied scientific field in which methods must not only orient action, but also explain, compare, and accumulate evidence under explicit criteria [7,9,39]. As a result, the demand for evaluation and value attribution intensifies: it is no longer sufficient to report usefulness retrospectively; it becomes necessary to estimate methodological contribution, conditions of success, and limits of transferability [12–14]. At the intersection with design, this requirement becomes especially pressing, because design responses to wicked problems generate outcomes that are simultaneously cognitive, social, and material. If these outcomes are to contribute to cumulative knowledge rather than remain exemplary but isolated interventions, DF requires an explicit logic of construction, observability, and evaluability consistent with the design-science paradigm [31–33].

A key disciplinary criterion therefore emerges at this interface: DF must not only expand imaginaries, but also specify under what conditions decisional closure may be considered justified. Traditions such as speculative design, design fiction, and experiential futures have significantly enlarged the future space by making alternative futures tangible through narratives, artifacts, perceptual bridges, and immersive experiences [19–23]. These approaches are highly valuable in opening the space of possibility and destabilizing default assumptions, yet they do not by themselves resolve the problem of selection: which futures should remain exploratory, which should be translated into strategic commitment, and under what assumptions such closure may be regarded as legitimate and auditable. This issue becomes sharper in light of the critique of defuturing, which reminds us that design does not merely open futures but also closes them, often by privileging pathways shaped by feasibility, institutional inertia, or market logics without sufficient epistemological reflexivity concerning the consequences of such closure [24]. The theoretical problem of DF is therefore not exhausted by opening futures; it lies equally in regulating, justifying, and documenting the passage from opening to closure.

A useful conceptual bridge for formalizing this passage lies in futures-cone representations, which visualize future space not as an undifferentiated set of options but as a structured field of constraints, trajectories, and horizons. Recent work has extended this conceptualization by arguing that futures must be understood in relation not only to a plurality of futures but also to a plurality of pasts, thereby foregrounding the causal

interconnections linking past, present, and future [26]. Whether interpreted literally or heuristically, this extension reinforces an important methodological thesis: future space is never neutral. It is conditioned by historical trajectories, causal restrictions, actor positions, and interpretive commitments. DF adopts this insight by explicitly integrating stakeholders as observers and participants who co-construct the relevant future domain, thereby clarifying who interprets, who validates, and who closes. This premise prepares the ontological role of the Stakeholder Viability Principle (SVP), introduced not as a normative add-on, but as a disciplinary criterion for delimiting relevant futures and avoiding arbitrary, purely instrumental, or tacit closure.

The evaluative axis is therefore central if DF is to be formulated as a cumulative contribution to foresight theory. Existing literature has proposed frameworks for evaluating national foresight activities and has examined how strategic foresight contributes to organizational performance and value creation [12–14]. However, even when such impacts are observed, it remains difficult to determine how much of the effect belongs to the method itself and how much to contextual contingencies unless interventions are designed with metrics, traceability, and contrast conditions that support attribution. For this reason, computational instrumentation becomes relevant not because foresight should be automated, but because observability requires forms of formalization capable of transforming narrative and interpretive material into analyzable objects. Recent hybrid approaches combining conceptual analysis, digital mixed methods, semantic projections, and expert validation indicate that narrative outputs can indeed be rendered more traceable, comparable, and, to a certain extent, reproducible without sacrificing interpretive grounding [15–18]. In this sense, DF becomes compatible with the evaluative logic of forecasting when forecasting is understood methodologically as the production of anticipations that are explicit, contrastable, and assessable under defined horizons, rather than as a promise of deterministic prediction [7,9].

To close the gap between practical intervention and cumulative knowledge, DF is aligned here with the logic of Design Science Research (DSR). Under this framing, scientific progress occurs through the construction and evaluation of artifacts—methods, models, instruments, and technological rules—that address relevant problems while allowing for the extraction of transferable principles [31–33]. Accordingly, Design Foresight is formulated at the disciplinary level not as an assemblage of tools, but as a research program defined by a minimum architecture: a clear unit of analysis, a set of operationalized constructs, a generic and replicable procedure, explicit instrumentation enabling traceability and contrast, and an evaluation design oriented toward accumulation and attribution. This formulation prevents DF from being reduced to an instrumental repertoire lacking cumulative theory and responds directly to the requirement that theory include mechanisms, conditions, and evaluable relations rather than only descriptive inventories or illustrative cases [10,11].

Within this framing, the unit of analysis adopted in this article is the prospective-design decision episode: a bounded cycle in which an organization explores future configurations, interprets signals and trends, applies viability and criticality filters, and produces a decisional closure materialized in outputs such as prioritized scenarios, roadmaps, prototypes, or decision rules under uncertainty conditions characteristic of wicked problems [34]. The method is explicitly non-deterministic. Its purpose is not to predict the future correctly in a point-forecasting sense, but to improve, in a justified and traceable manner, the capacity to anticipate, decide, and act. Scientific evaluation therefore concerns the quality of the methodological artifact and the process through which closure is achieved, rather than the impossible standard of eliminating uncertainty altogether [6,11–13].

On this basis, DF may be formalized through a set of core disciplinary constructs. The Stakeholder Viability Principle (SVP) operates as a filtering function that delimits the

set of relevant futures according to the continuity, agency, and axiological coherence of critical stakeholders. Social singularity designates a critical-transition condition in which windows of intervention become limited and decisional urgency becomes concentrated, making ordinary routines of interpretation and action insufficient [27–29]. Extended futures cones provide a topological representation of causal and temporal constraints, situating the relevant future not as an open infinity of options but as a conditioned region shaped by trajectories and observer positions [26]. Semantic projections, finally, provide a quantified representation of the affinity between signals or trends and an expert-defined conceptual universe, thereby introducing observability, inter-temporal comparability, and the possibility of contrast under explicit assumptions [15–17]. Together, these constructs connect ontology and epistemology: futures cones structure constraints, SVP bounds the relevant future domain, singularity introduces criticality, and semantic projections provide traceability.

At a generic level, the disciplinary procedure derived from Singular Design Foresight is organized as a cycle with four broad phases: configuration, futurization, de-futurization, and documentation/extraction. In the configuration phase, the decision episode is delimited in terms of scope, horizon, and constraints; critical stakeholders are identified; an initial conceptual universe is elicited through expert participation; and initial viability and criticality thresholds are defined. In the futurization phase, the future space is expanded through the combined use of semantic representation and design-oriented exploratory practices—such as scenarios, prototypes, fictions, and experiences—in order to generate candidate futures and associated sets of signals and trends [20–22]. In the de-futurization phase, the expanded future set is progressively reduced through the application of SVP filters, criticality criteria, and contextual constraints, producing a shortlist of viable futures together with an explicit record of the rules and evidence supporting closure. Finally, the documentation and extraction phase records successive versions of the conceptual universe, parameters, thresholds, evidence, filtering rules, and decisional outputs, thereby enabling the formulation of transferable rules and the accumulation of case-comparable knowledge [31–33].

This cycle is not tied to a single technological implementation. Rather, it is defined at the disciplinary level so that different artifacts may instantiate the same methodological logic under different domain conditions. What must remain invariant is not the interface or technical architecture, but the explicit relation between evidence, interpretation, filtering, closure, and traceability. In this sense, the scientific object of DF is not "the future" as a directly observable empirical entity, but the methodological artifact that connects signals, conceptualization, and decision under conditions of uncertainty [31,33]. This distinction is crucial because it clarifies that the cumulative value of DF lies in the evaluability of the artifact, not in the direct verification of future states as such.

A scientifically evaluable method also requires explicit instrumentation. In SDF, instrumentation is designed to make the transition from evidence to closure observable and auditable. The conceptual universe is treated as a versioned structure of terms, categories, and relations defined by an expert panel, with changes documented so that shifts in results can be distinguished from shifts in conceptual framing. Signals and trends are represented as semantic objects and projected onto the conceptual universe through quantified affinities, enabling the construction of semantic intensity indicators and comparison across cases and moments. Viability- and coherence-related indices are normalized and associated with explicit thresholds that support inclusion and exclusion rules. Closure is therefore not treated as tacit consensus, but as a documented decision governed by discussable criteria, traceable parameter values, and sensitivity to variation in assumptions [12,13,17,18].

Under this logic, the outputs of DF can be represented in a forecasting-compatible manner without reducing the field to point prediction. Depending on the decision episode

and the available evidence, outputs may take the form of series, such as semantic intensity indicators for key drivers; states or configurations, such as dominant future profiles or shortlisted viable futures; or events and transitions, such as signals associated with social singularity. What makes these outputs forecasting-compatible is not their determinism, but their evaluability under explicit horizons, assumptions, and procedures of comparison. Forecasting is therefore used here instrumentally and methodologically: not as the telos of DF, but as a source of criteria for explicit anticipation, traceability, and contrast [7,9,18].

Finally, the evaluative framework of SDF is organized around three dimensions: validity, utility, and attribution. Validity concerns the operational definition of constructs, the internal stability of the process under comparable inputs and parameters, and the possibility of replication across cases. Utility concerns the extent to which the method improves the grounding, clarity, and robustness of decision closure under uncertainty, without collapsing utility into validity. Attribution concerns the ability to estimate what part of the resulting shortlist or closure is due to methodological choices—such as the use of SVP, semantic projections, thresholds, or conceptual-universe configurations—rather than solely to context or group dynamics [11–14,31]. In the present article, this evaluative architecture is proposed programmatically rather than reported as fully executed. Its purpose is to establish the conditions under which DF may advance from utility without theory toward a family of artifacts with explicit mechanisms, conditions of applicability, and progressively cumulative evidence.

5. Formalization of the SDF Cycle, Instrumentation, and Methodological Flow

For Design Foresight (DF) to operate as a scientifically evaluable method, it must be articulated at the disciplinary level through a structure that is sufficiently explicit to support replication, traceability, and cumulative assessment. In this sense, the method cannot be identified with a single implementation, platform, or domain-specific architecture. Rather, it must be specified through a set of minimum formal conditions: a defined unit of analysis, a system of operationalized constructs, a generic and replicable procedure, explicit instrumentation that renders the transition from evidence to closure observable, and an evaluation logic oriented toward accumulation, comparison, and attribution [31–33]. This framing is consistent with Design Science Research (DSR), which understands scientific progress as proceeding through the construction and evaluation of artifacts—methods, models, instruments, and technological rules—capable of addressing relevant problems while generating transferable principles [31–33].

Under this logic, the scientific object of Design Foresight is not "the future" as a directly observable empirical entity, but the methodological artifact that links evidence, interpretation, filtering, and closure under conditions of uncertainty. This distinction is decisive. It prevents DF from being reduced to a repertoire of inspirational techniques or situated design practices without cumulative theory, and it responds to the requirement that scientific theorization specify mechanisms, conditions, and evaluable relations rather than merely assembling variables, cases, or procedural repertoires [10,11]. In the present formulation, Singular Design Foresight (SDF) is therefore formalized not as a finalized forecasting system, but as a disciplinary core that can be instantiated through heterogeneous artifacts while preserving a common methodological logic.

The unit of analysis adopted in SDF is the prospective-design decision episode. This is defined as a bounded cycle in which an organization explores possible future configurations, interprets signals and trends, applies viability and criticality filters, and produces a decisional closure materialized in outputs such as prioritized scenarios, roadmaps, prototypes, strategic options, or decision rules under uncertainty conditions typical of wicked

problems [34]. The method is explicitly non-deterministic. Its purpose is not to predict the future correctly in a point-forecasting sense, but to improve, in a justified and traceable manner, the capacity to anticipate, decide, and act. Accordingly, the quality of the method is assessed not by its elimination of uncertainty, but by the validity, utility, and attributable contribution of the artifact through which closure is achieved [6,11–13].

At the disciplinary level, the method is anchored in four core constructs introduced in the previous section. The Stakeholder Viability Principle (SVP) delimits the relevant future domain through criteria of continuity, agency, and axiological coherence. Social singularity identifies the critical transition condition under which ordinary practical reasoning becomes insufficient and closure becomes decisionally urgent [27–29]. Extended futures cones provide a topological representation of the causal and temporal constraints that shape the relevant future space [26]. Semantic projections, finally, provide a quantified representation of the affinity between signals or trends and an expert-defined conceptual universe, thereby rendering discursive material more observable, comparable, and open to evaluation [15–17]. Taken together, these constructs establish the ontological and epistemological infrastructure that allows SDF to move from exploratory practice toward an auditable methodological process.

On this basis, the SDF cycle can be formalized as a generic procedure composed of four interrelated phases: configuration, futurization, de-futurization, and documentation/extraction. These phases should not be understood as a rigid sequentialism, but as a disciplined methodological flow in which each operation produces intermediate objects that can be inspected, revised, and documented.

The first phase, configuration, establishes the conditions of the decision episode. It involves delimiting the scope, horizon, and constraints of the episode; identifying the critical stakeholders relevant to the problem context; eliciting an initial version of the conceptual universe through expert participation; and defining the initial criteria and thresholds that will govern viability and criticality. This phase is foundational because it determines the initial interpretive architecture of the process. Without an explicit configuration stage, later closure risks depending on implicit assumptions that cannot be reconstructed, challenged, or compared across cases.

The second phase, futurization, operates as the divergent operator of the cycle. Its purpose is to expand the future space by generating candidate configurations through a combination of signal extraction, semantic representation, and design-oriented exploratory practices. At this stage, signals and trends are extracted from the evidence corpus and related to the evolving conceptual universe through semantic projections or other traceable representational devices. In parallel, alternative future configurations are generated through scenarios, artifacts, narratives, or design constructs that maximize relevant variation within the boundaries of the episode. The objective is not open-ended proliferation for its own sake, but disciplined expansion: the method seeks to widen the range of futures under consideration sufficiently to avoid default trajectories, path dependence, and premature convergence [20–22].

The third phase, de-futurization, constitutes the formalizable convergent operator of the method. Here, the expanded space of candidate futures is progressively reduced through the application of viability filters, criticality criteria, and contextual constraints. Methodologically, this phase is where SDF differentiates itself most clearly from approaches that remain primarily exploratory. Closure is not treated as tacit agreement or intuitive narrowing, but as a documented reduction in variance governed by explicit rules. The SVP functions here as a primary ontological filter, excluding futures that fail to satisfy minimum requirements of continuity, agency, or axiological coherence for critical stakeholders. At the same time, the analysis of singularity conditions indicates whether the decision episode has

entered a state of critical transition in which continued deferral becomes strategically costly or untenable. De-futurization therefore does not simply reduce multiplicity; it transforms exploratory plurality into a justified shortlist of viable futures under explicit conditions of residual uncertainty.

The fourth phase, documentation and extraction, gives the cycle its cumulative scientific value. At this stage, the process records the successive versions of the conceptual universe, the evidence corpus used, the semantic representations generated, the thresholds applied, the exclusion and prioritization rules adopted, and the final decision outputs. This documentation renders the chain from evidence to closure reconstructible and creates the conditions for inter-case comparison, sensitivity analysis, and the derivation of transferable technological rules in the sense proposed by the design-science tradition [31–33]. Its function is therefore not merely archival, but epistemic: it converts situated decision episodes into analyzable instances of a common disciplinary method.

This procedural architecture requires explicit instrumentation. A scientific method cannot rely solely on conceptual distinctions; it must provide instruments that render its operations observable. In SDF, instrumentation is designed to make the passage from original evidence to decisional closure traceable, discussable, and auditable.

A first instrument is the versioned conceptual universe (CU). The CU is formalized as a structured set of terms, categories, and relations defined through expert participation and revised as the episode evolves. Version control is essential because it allows changes in results to be distinguished according to whether they stem from new evidence, altered thresholds, or conceptual redefinition. In this way, conceptual evolution is not treated as noise, but as a documented dimension of inquiry.

A second instrument is the representation of signals and trends as semantic objects. Each signal or trend is represented within a semantic space and projected onto the conceptual universe through quantified affinities. Operationally, these affinities may be instantiated as bounded coefficients over a documentary corpus, for example through co-occurrence-based measures that estimate the degree to which a signal and a concept share a semantic environment. The resulting semantic projections generate normalized values that make it possible to construct semantic intensity indicators, compare signals across periods or cases, and document the relation between qualitative evidence and anticipatory constructs [16–18]. Their purpose is not to replace interpretation, but to stabilize and render explicit the relation between discursive material and foresight-relevant concepts.

The signal extraction pipeline operationalized in prior empirical applications of SDF-compatible instrumentation follows a structured four-phase sequence: systematic corpus construction and document selection based on temporal relevance, thematic pertinence, and source authority; information systematization through dual-method analysis combining frequency-based tools (Voyant Tools) and large language model processing (NotebookLM); semantic projection computation via the Defly Compass platform, which generates cross-reference matrices quantifying the co-occurrence strength between identified signal-concept pairs across indexed databases such as DOAJ, Google Scholar, and arXiv; and expert-validated synthesis producing integrated megatrend networks [18]. The quantitative outputs of this pipeline include co-occurrence matrices—for example, in the education domain, the pair [AI in Higher Education × Personalized Learning] yields a co-occurrence value of 111, while [Personalized Learning × Learning Analytics] reaches 387, establishing a measurable intensity gradient across the signal landscape [18]—as well as network visualizations in which node centrality and edge thickness encode the relative semantic strength of each trend relation. Signal validity is established through convergence across three methodologically independent procedures, such that a signal is considered robust when frequentist, qualitative, and semantic projection analyses yield

consistent identification patterns. This validation-through-convergence principle has been demonstrated across three distinct application domains—agriculture, education, and public health—providing transferability evidence for the pipeline's core logic [18]. The present article does not reproduce this pipeline in full but incorporates its logic as the empirical basis for the signal extraction step of the SDF cycle.

A third instrument consists of normalized indices and explicit thresholds. Viability, axiological coherence, and related evaluative dimensions are normalized within common intervals and linked to clearly specified thresholds that support inclusion, exclusion, or prioritization rules. Methodologically, the significance of this operation lies in the fact that closure is no longer treated as tacit consensus emerging informally from discussion, but as a decision documented through discussable criteria that are, at least in principle, open to sensitivity analysis, recalibration, and inter-case comparison [12,13]. In published mixed expert–corpus applications, similar logic has already been instantiated through structured scoring procedures, dispersion criteria, and threshold-based selection rules, providing summative evidence for the feasibility of this kind of filtering architecture.

A fourth instrument concerns the production of evaluable, non-deterministic outputs. SDF does not seek point predictions of future states. Rather, its instrumentation allows anticipatory outputs to be represented as indicators, configurations, or transition-relevant signals. Depending on the decision episode and the available evidence, outputs may take the form of semantic intensity series for key drivers, shortlisted viable futures, ranked intervention windows, prioritized scenarios, or signals associated with critical transitions and singularity conditions. This is what makes SDF compatible with forecasting in a methodological sense: it generates anticipations that are explicit, bounded, and assessable under defined horizons and assumptions, without requiring deterministic prediction as the criterion of value [7,9,18].

At the disciplinary level, the methodological flow of SDF can be summarized through the following generic pseudocode (Algorithm 1).

Algorithm 1 Singular Design Foresight (SDF) Cycle

```
Input:
decision episode D
stakeholders S
evidence corpus E
expert panel P
Output:
shortlist of viable futures F*
decision closure C
trace log T
Step 1. Configure
Define episode D, horizon h, and constraints K
Identify critical stakeholders S
Elicit initial conceptual universe CU_0 from P
Define viability criteria and thresholds theta
Define singularity criteria sigma
Step 2. Sense and Represent
Extract signals and trends from E
Compute semantic projections pi(.) onto CU_t
Produce semantic indicators I_t
```

Algorithm 1 *Cont.*

```
Step 3. Futurize
Generate candidate futures F
Link scenarios, artifacts, or constructs to I_t and CU_t
Step 4. De-futurize
Apply contextual constraints, SVP filters, and evaluative thresholds
Obtain viable set F_v subset of F
Detect or estimate singularity conditions sigma
Prioritize windows of intervention
Step 5. Close
Derive decision closure C
Select shortlist F* subset of F_v
Record rules, parameter values, and justifications in T
Step 6. Evaluate and Learn
Assess validity, utility, and attribution
Update CU_(t + 1) if required
Extract transferable rules
Append full trace to T
```

For operational reproducibility, the required 'input' for Step 2 consists of a co-occurrence matrix generated from the frequency of terms within the evidence corpus E cross-referenced with the dimensions of the conceptual universe CU. While the specific projection formulas are detailed in Manetti et al. (2021) [16] and López Bordao et al. (2025) [18], the underlying logic—based on the ratio of documentary frequencies to semantic affinities—allows researchers to design calculation engines compatible with the SDF flow described in this algorithm.

This pseudocode should not be interpreted as software code or as a complete implementation schema. Its purpose is disciplinary and methodological: it renders explicit the invariant logic of the method across possible domain instantiations. In other words, it formalizes what must remain conceptually stable even when corpora, instruments, interfaces, and application contexts vary. At the same time, it should be read in light of the summative evidence accumulated across prior published studies, which have already instantiated several of its components—semantic quantification, mixed expert–corpus filtering, hybrid signal extraction, and singularity-related classification—without yet integrating them into a single end-to-end framework. Because SDF is proposed as a scientifically evaluable method, its formalization must also include an evaluation logic. In line with DSR and foresight evaluation literature, this logic is organized around three dimensions: validity, utility, and attribution [11,12,31].

Validity concerns whether the constructs and operations of the method are adequately specified and whether the process remains sufficiently stable under controlled conditions. At the construct level, it requires explicit operational definitions of viability, singularity, trend, and axiological coherence. At the internal level, it concerns process stability under identical inputs and parameters, inter-rater consistency among experts, and the capacity to distinguish substantive variation from variation introduced by changes in the conceptual universe. At the external level, it concerns the possibility of replication across cases and the comparison of resulting patterns across application domains.

Utility concerns the capacity of the method to improve anticipation and decision-making under uncertainty. It includes the clarity of closure rules, the coherence of the resulting shortlist with stakeholder conditions, the robustness of outputs under plausible

parameter variations, and the perceived contribution of the method to reducing ambiguity or strengthening strategic alignment. Crucially, utility is treated as distinct from validity: a method may prove pragmatically useful in decision settings without yet satisfying stronger conditions of theoretical or methodological justification [10,13,14].

Attribution concerns the capacity to estimate which part of the resulting closure is due to methodological choices rather than solely to contextual contingencies, institutional routines, or group dynamics. In SDF, attribution is approached through comparative and sensitivity-oriented designs, such as examining the effects of including or excluding SVP filters, semantic projections, threshold configurations, or singularity criteria. The purpose is not to claim conclusive causal isolation within the present article, but to establish the design conditions under which methodological contribution can be progressively estimated in future empirical work [31–33].

In this sense, the relation between SDF and forecasting must be understood instrumentally. Forecasting does not redefine the purpose of Design Foresight, nor is DF reduced here to predictive performance alone. Rather, forecasting provides a compatible evaluative harness for those components of the method that generate explicit anticipatory representations—such as semantic intensity series, transition detection, or regime-like future configurations—that may be contrasted under defined horizons when appropriate [16–18]. What SDF adds is not deterministic prediction, but a formalized architecture through which the closure of futures becomes more explicit, traceable, and open to disciplined evaluation.

Illustrative Application of the SDF Cycle: Higher Education Strategy in Peru

To clarify the operational logic of the SDF cycle without overstating its current validation status, this section traces the six-step procedure through a prior published application in the domain of higher education strategy [30]. This application is not presented as end-to-end validation of the integrated SDF framework; rather, it illustrates how the cycle's invariant logic was instantiated across the key operations of configuration, futurization, de-futurization, and closure. The purpose is to render the pseudocode concrete through data and procedural decisions that are traceable and replicable within the scope of the prior study.

It is imperative to note that the specific numerical values reported in this walkthrough—such as means, standard deviations, and affinity coefficients—are not intended to constitute a system-level validation of the integrated SDF framework within the Peruvian context. Instead, they are presented exclusively as 'test data' derived from previous modular studies to demonstrate the operational capacity of the SDF cycle to process and represent auditable quantitative outputs.

Step 1—Configure. The decision episode was delimited as a prospective-strategic inquiry into the future of AI-driven lifelong learning in Peruvian higher education, with a ten-year horizon (2025–2035) and a specific institutional constraint: the need to align trajectories with the continuity and capability development of universities facing accelerated digital transformation. Critical stakeholders were identified in relation to the decision problem, combining academic experts from multiple institutional roles—planning directors, postgraduate deans, vice-rectors, coordination staff, and area directors—spanning early-career to senior profiles with explicit gender balance (maximum 60% representation by any single gender). The initial conceptual universe (CU_0) was elicited through structured expert participation and comprised six thematic clusters: AI-driven personalization, lifelong learning architectures, digital equity and access, institutional adaptability, workforce alignment, and ethical governance of educational AI. Viability thresholds were set at a mean expert score ≥ 4.0 and standard deviation < 1.0 on a 0–5 scale, with inter-rater cosine similarity consistently ≥ 0.99 required for signal retention [30].

Step 2—Sense and Represent. Signals and trends were extracted from a documentary corpus combining academic publications indexed in DOAJ, Google Scholar, and arXiv, together with grey literature on Peruvian higher education policy. Signal extraction followed the structured four-phase pipeline operationalized in prior published work: corpus construction, information systematization via frequency-based and LLM-assisted tools, semantic projection computation through co-occurrence analysis, and expert-validated synthesis [18]. Semantic projections were computed as bounded affinity coefficients: for a signal A and a concept B belonging to CU, the projection $\pi(A,B)$ was estimated as the ratio of co-occurring documents to total documents containing A, producing normalized values in. For example, the signal pair AI in Higher Education → Personalized Learning yielded a co-occurrence value of 111, while Personalized Learning → Learning Analytics reached 387, establishing a measurable intensity gradient across the signal landscape [18]. Aggregated across all CU dimensions, each signal produced a composite representation vector $P_U(A)$, enabling the construction of semantic intensity indicators comparable across the six conceptual dimensions of the universe.

Step 3—Futurize. Using the semantic intensity indicators derived in Step 2, the future space was expanded into candidate configurations representing plausible trajectories for the 2025–2035 horizon. These included: (F1) an accelerated AI-personalization pathway in which institutional transformation is driven by adaptive learning systems and workforce-aligned curricula; (F2) a stratified access pathway in which digital equity constraints produce differentiated institutional futures across socioeconomic segments; (F3) a hybrid resilience pathway in which institutions combine AI adoption with strong human-centred governance frameworks; and (F4) a disruption-and-consolidation pathway characterized by institutional fragmentation followed by strategic recomposition around lifelong learning hubs. These configurations were not forecasts but disciplined expansions of the viable future space, designed to maximize relevant variation while remaining anchored to the signals and conceptual universe defined in Steps 1–2.

Step 4—De-futurize. The expanded candidate set was progressively filtered through the SVP and singularity conditions (Figure 2). Applying the viability threshold (mean > 4.0, SD < 1.0) to expert evaluations, all 30 candidate trends exceeded the selection threshold, with inter-rater cosine similarity of 0.9960342 confirming high panel convergence [30]. Within the SVP filtering logic, futures were differentiated not by outright exclusion of trends but by the differential semantic intensity of those trends across the three SVP dimensions. Future configurations weighted toward lower-intensity signals (semantic projection values < 0.20, as observed for terms such as Gen Z [0.12], AI-Generated Content [0.15], and Global Education [0.20]) were assessed as having weaker anticipatory grounding and lower viability scores under the axiological coherence dimension. The singularity analysis identified a critical transition condition in the intersection of AI adoption velocity and digital equity gaps, interpreted as a window of decisional urgency: continued institutional inaction in this domain was assessed as strategically costly given the concentration of transformation signals in the 2024–2026 timeframe [30]. This condition foregrounded F2 and F3 as requiring prioritized strategic attention.

Step 5—Close. In this illustrative exercise, the resulting viable shortlist was configured as follows with three prioritized futures: F3 (hybrid resilience) as the primary viable trajectory, supported by the highest expert convergence across all three SVP dimensions (mean = 4.4, SD = 0.7, cosine similarity = 0.99); F1 (AI-personalization) as a conditionally viable trajectory dependent on institutional capability development; and F2 (stratified access) as a critical boundary condition requiring active mitigation strategy. Decision closure was documented as a set of three strategic orientations: (i) invest in hybrid human-AI governance frameworks; (ii) prioritize digital equity infrastructure as a precondition

for sustainable AI integration; (iii) build adaptive lifelong learning architectures aligned with workforce transitions. All filtering rules, parameter values, expert score distributions, and semantic intensity indicators were recorded in the trace log, rendering the chain from evidence to closure reconstructible (Figure 3).

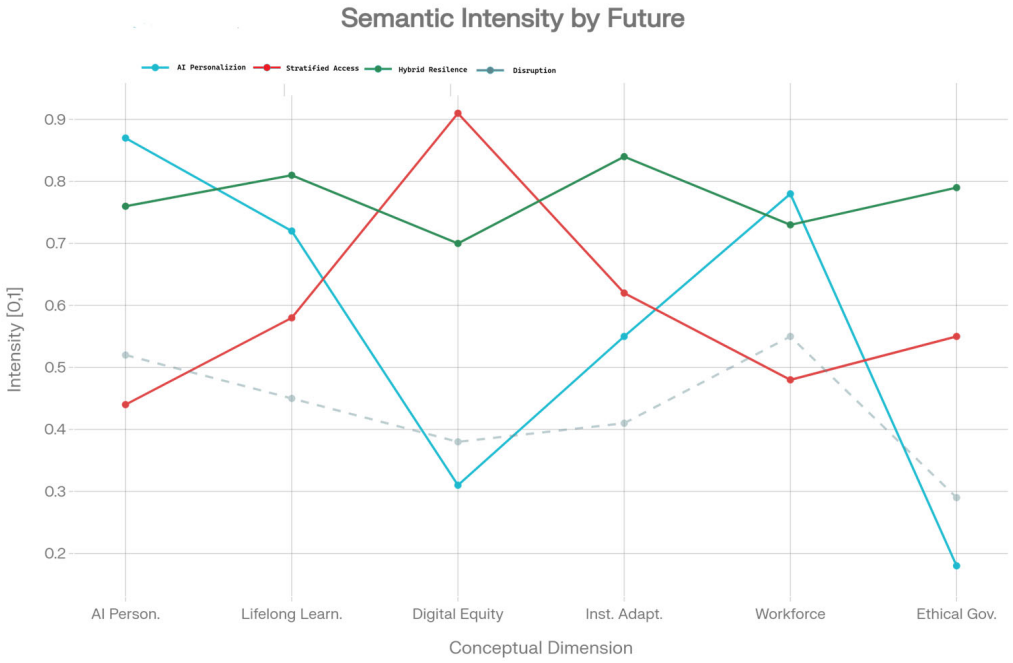


Figure 2. Semantic intensity indicator profiles for the four candidate future configurations (F1–F4) across the six-dimensional conceptual universe (Peru higher education application, 2025–2035). F4 (dashed, low opacity) was excluded by SVP filtering due to insufficient expert convergence on the Agency and Axiological Coherence dimensions. Values derived from co-occurrence-based semantic projections [18,30].

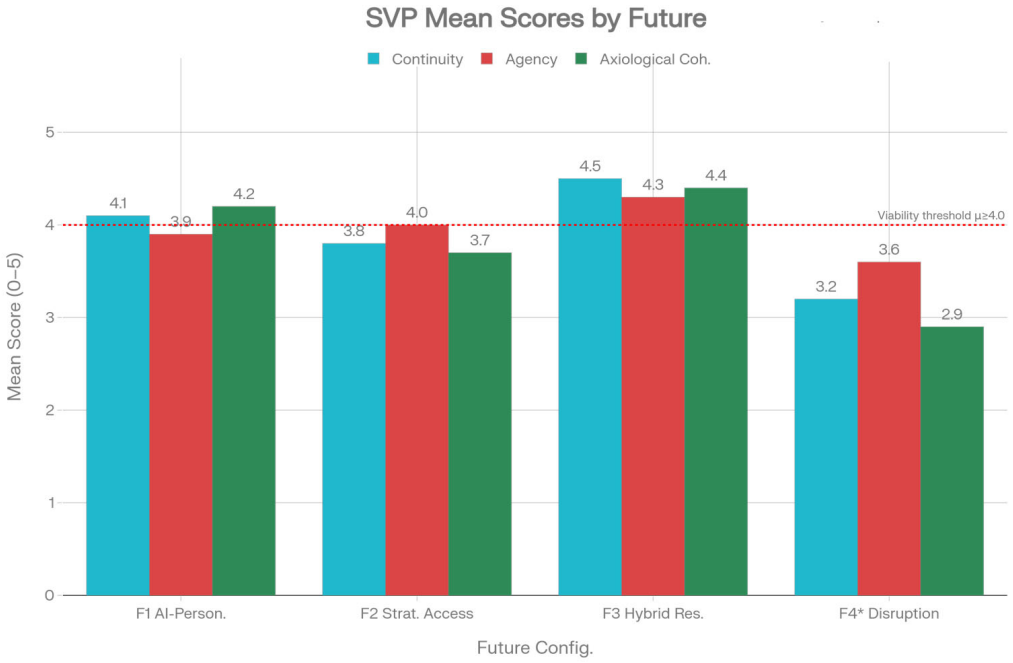


Figure 3. Expert panel mean scores for each candidate future configuration across the three SVP dimensions (Continuity, Agency, Axiological Coherence). Dashed red line marks the viability threshold ($\mu \geq 4.0$, $\sigma < 1.0$). F4 falls below threshold on all three dimensions and is excluded from the viable future set. Scores based on structured expert evaluation on a 0–5 ordinal scale [30].

Step 6—Evaluate and Learn. At the component level, the application confirmed the operational feasibility of semantic projection instrumentation across a multi-dimensional conceptual universe, the discriminatory capacity of SVP thresholds across candidate futures, and the identification of a context-specific singularity condition through observable signal accumulation. The conceptual universe was updated (CU_1) to reflect emergent signal salience around ethical governance and equity, documenting conceptual evolution as a methodological dimension rather than treating it as noise. Transferable rules extracted from this episode include: (i) in high-uncertainty institutional contexts, futures exhibiting low axiological coherence scores are systematically filtered before capability analysis; (ii) singularity conditions are most reliably identified when signal convergence across at least three independent extraction procedures reaches threshold; (iii) composite semantic intensity indices distinguish meaningfully between dominant trends and background noise at co-occurrence differentials above 3:1.

This walkthrough illustrates that the SDF cycle can be traversed with empirically grounded data, explicit parameters, and traceable filtering decisions. It does not constitute full system-level validation of the integrated SDF framework—which remains a future research priority—but demonstrates that the pseudocode is operationally instantiable and that its steps produce inspectable, documentable, and discussable outputs consistent with the forecasting-compatible logic of evaluable anticipatory representation under explicit assumptions.

6. Illustrative Empirical Basis and Programmatic Evaluation of Singular Design Foresight

The formalization of Design Foresight (DF) as a scientifically evaluable method ultimately requires empirical contrast in real decision-making contexts. In the present article, however, such contrast is not introduced as conclusive system-level validation of Singular Design Foresight (SDF), but as a body of illustrative and summative empirical evidence supporting the feasibility of its core mechanisms. In line with the logic of Design Science Research (DSR), empirical applications are relevant not only because they address situated problems, but because they allow methodological artifacts to be examined under varying conditions, thereby contributing to construct refinement, threshold calibration, and the progressive derivation of transferable rules [31–33]. The empirical strategy adopted here therefore relies on distributed evidence already reported in previous studies—including semantic projections for trend forecasting, singularity detection and classification, mixed expert–corpus applications in higher education, and more recent hybrid instrumental implementations. These studies are not treated as proof that the integrated SDF system has already been validated; rather, they are interpreted as modular evidence supporting the plausibility, operational feasibility, and evaluability of the disciplinary framework proposed in this article.

The unit of empirical interpretation remains consistent with the disciplinary definition established above: the prospective-design decision episode. Each episode can be understood as a bounded cycle in which an organization defines an expert-informed conceptual universe, identifies signals and trends, generates candidate future configurations, applies viability and criticality filters, and produces closure in the form of a shortlist and/or other decision artifacts. From the standpoint of the present framework, the central empirical question is not whether a particular scenario later proves to be "true" in a literal predictive sense, but whether the methodological procedure produces closures that are more explicit, more traceable, more robust, and more attributable than reasonable alternatives under explicit conditions [11–13].

From this perspective, the currently available empirical basis can be organized into four complementary layers.

The first is a semantic instrumentation layer, represented by prior work on design trend forecasting through conceptual analysis and semantic projections. This line of research shows that textual and discursive evidence can be transformed into quantifiable and comparable semantic representations, making it possible to derive explicit indicators from qualitative material and to relate signals to an expert-defined conceptual universe. Its importance for the present article lies in grounding the representational and partially forecastable dimension of SDF without displacing its primary orientation toward decision closure [16].

The second is an empirical singularity layer, represented by work on the detection and classification of singularity in higher education. That work operationalizes singularity through observable variables, reduction techniques, and classification logic, thereby supporting the treatment of singularity not merely as a metaphorical construct but as an empirically tractable condition that can be detected, compared, and incorporated into a formal decision framework [29]. Although the present article adopts a broader and more functionally oriented definition of social singularity, this prior evidence is crucial because it demonstrates that rupture-related conditions can sustain explicit empirical treatment.

The third is a mixed expert–corpus application layer, represented by work on AI-driven lifelong learning in Peru and related foresight applications. These studies combine expert judgment and documentary corpora in order to generate trends, challenges, and opportunity structures supported by computational processing and expert validation. Their relevance for the present framework lies in demonstrating that qualitative foresight inputs can be transformed into structured, partially quantifiable, and traceable outputs through hybrid procedures consistent with the architecture proposed for SDF [30].

The fourth is a deployment and transferability layer, represented by more recent instrumental variants—such as Defly Compass—that implement trend-analysis and foresight architectures across broader samples using explicit corpora, semantic analysis, clustering procedures, and expert validation. These studies do not validate SDF as an integrated disciplinary method, but they do provide evidence that several of its core instrumental operations can function in a reproducible and transferable manner across application contexts [18]. Across these four layers, prior empirical applications have collectively covered at least five distinct domains—design education (IED Barcelona), higher education strategy (PUCP, Peru), agriculture, public health, and education technology—providing a preliminary basis for assessing the cross-domain operability of core SDF components [17,18,30].

Taken together, these studies should not be interpreted as having already demonstrated the validity of the full SDF architecture. Rather, they constitute summative component-level evidence. They support the feasibility of semantic instrumentation, the observability of critical-transition constructs, the viability of mixed expert–corpus procedures, and the transferability of computationally supported foresight operations. What they do not yet provide is full end-to-end validation of the integrated cycle proposed in this article. A central contribution of the present work is precisely to articulate how these previously dispersed components can be understood as parts of a common methodological architecture and how future empirical work might evaluate them in a more integrated, cumulative, and comparative manner.

To align empirical contrast with evaluative standards relevant to forecasting—without reducing Design Foresight to predictive performance—SDF is framed as operating at two complementary empirical levels. The first is the level of evaluable anticipatory representation. At this level, semantic instrumentation produces indicators, series, or state-like representations derived from semantic projections that may be examined across time

horizons or compared under defined assumptions. Examples include semantic intensity indicators for key drivers, shifts in conceptual salience, or the identification of emerging regimes. This level is the one most directly compatible with forecasting logic understood as the assessment of anticipatory representations under explicit conditions [7,9,16].

The second is the level of audited decisional closure. At this level, the method produces a shortlist of viable futures under the combined action of the Stakeholder Viability Principle and criticality criteria. Here the central empirical question is not predictive accuracy, but the robustness, traceability, and discussability of closure. A closure is methodologically stronger when the chain linking evidence, semantic representation, filtering, and decision can be reconstructed, and when plausible changes in thresholds, conceptual-universe configurations, or evidence selection do not arbitrarily destabilize the resulting shortlist. This level captures the distinctive contribution of DF within the present framework: the primary output is not a point forecast, but a justified closure under uncertainty.

Within this architecture, evaluation remains organized around the three dimensions already introduced: validity, utility, and attribution. Construct validity concerns the explicit operationalization of viability, singularity, semantic projections, and related indices. Internal validity concerns procedural stability under comparable inputs and parameters, consistency among expert evaluations, and the capacity to distinguish substantive effects from changes introduced by conceptual redefinition. External validity concerns the possibility of replicating decision episodes across different domains and comparing resulting patterns or profiles [11,16,29,30].

Utility refers to the contribution of the method to anticipating and deciding with greater grounding under uncertainty. Relevant indicators include the clarity of closure rules, a reduction in operational ambiguity, the perceived decisional quality among critical stakeholders, robustness under plausible parameter variation, and improved strategic alignment. Importantly, utility is treated as an evaluative dimension distinct from validity. A method may prove pragmatically useful in decision settings without yet satisfying stronger conditions of theoretical specification or methodological stability [10,13,14].

Attribution is the most demanding dimension and, in the present article, remains programmatic rather than conclusively demonstrated. It concerns the possibility of estimating which part of a resulting shortlist or closure is due to methodological choices rather than to contextual contingencies, institutional routines, or group dynamics alone. In future empirical designs, this may be approached through comparative and sensitivity-oriented analyses examining the effects of including or excluding SVP filters, semantic projections, singularity criteria, threshold configurations, or alternative versions of the conceptual universe. Sensitivity analysis will be especially important for distinguishing robust outputs from those that depend strongly on parameter choices [31–33]. To illustrate the type of sensitivity analysis envisioned, consider the threshold configuration used in the Peru application [30]: with selection criteria set at mean > 4 and SD < 1 on a 0–5 scale, 30 macro-trends were retained from an initial set of 276. Had the dispersion threshold been tightened to SD < 0.8, the retained set would have decreased to approximately 21 trends—a reduction of roughly 30%—suggesting moderate but non-trivial sensitivity to threshold calibration. This example does not constitute a formal sensitivity test; rather, it illustrates the kind of parameter-variation reporting that future SDF applications should systematically include, so that methodological choices become discussable, comparable, and progressively calibrated across cases.

For a foundational–methodological article such as this one, the role of empirical material is therefore not to "prove" the method once and for all, but to establish the plausibility of its components and to define a coherent pathway for future contrast. At a minimum, future reports of SDF applications should make reconstructible the full chain

from evidence to closure; report how semantic indicators, filters, and thresholds contributed to the shortlist; examine the sensitivity of closure to controlled variations; and, where appropriate, compare representational outputs or closure patterns against reasonable baselines. Only under such reporting conditions can decision episodes function not merely as local cases, but as cumulative instances of a common methodological artifact.

The scientific purpose of this empirical section is therefore not case display, but disciplinary accumulation. To contribute to cumulative knowledge, each application should document the initial conditions of the episode, the version of the conceptual universe employed, the corpus and filtering criteria used, the parameter settings adopted, the outputs generated, and the traceability and sensitivity conditions of the resulting closure. Only under such conditions can decision episodes become comparable analytical units from which transferable technological rules of the form "if. . . then. . ." may gradually be derived, in line with the design-science tradition [31–33].

7. Discussion

The central argument of this article is that Design Foresight (DF) can consolidate as a cumulative scientific field only if it moves beyond the juxtaposition of methods and explicitly articulates its ontological commitments, epistemic instruments, procedural logic, and evaluative criteria. In action-oriented fields, practical utility is indispensable, but it is not sufficient for the production of cumulative theory. Methods may be useful, engaging, or strategically influential without yet becoming scientifically comparable, attributable, or transferable [10,11]. From this perspective, the contribution of Singular Design Foresight (SDF) does not lie in adding another repertoire of tools to the existing foresight landscape, but in restructuring DF so that practice can progressively become knowledge through the specification of a relevant ontological object, an explicit methodological mechanism, and an evaluation architecture oriented toward validity, utility, and attribution [12,31–33].

This discussion is organized around the five research questions guiding the article.

First, regarding RQ1, the paper has argued that DF can be articulated as a disciplinary framework only if it is treated not merely as an intersection between design methods and foresight methods, but as a field with its own explicit architecture. The theoretical implication is that DF should be understood at the level of a theory of foresight-informed design practice rather than as a loose aggregation of interventions, cases, or methodological repertoires. In this sense, the article responds to a longstanding need in foresight theory for clearer epistemology, explicit process logic, and evaluative criteria capable of supporting cumulative explanation rather than isolated practical success [11]. The disciplinary object proposed here is therefore not "the future" in general, but the production of justified closure over a viable subset of futures under conditions of uncertainty.

Second, in relation to RQ2, the article has proposed that the passage from future-space expansion to decision closure can be formalized through the distinction between futurization and de-futurization. The significance of this distinction is both theoretical and methodological. Theoretically, it clarifies that opening and closing futures are not contradictory gestures, but complementary operations within a single anticipatory logic. Methodologically, it establishes that closure should not remain a tacit consequence of discussion, consensus, or feasibility pressure, but should instead be governed by explicit filters, criteria, and traceable operations. The article therefore advances a mechanism-based understanding of DF in which exploratory divergence is followed by disciplined convergence under explicit viability and criticality conditions. In this sense, SDF contributes not merely an ontology of futures, but a formalizable mechanism for moving from signals and discourses to justified closure under structural uncertainty [10,20–22,24].

Third, with regard to RQ3, the article has specified the role of stakeholders, signals, and contextual ruptures in delimiting the relevant future space. The most important ontological implication here is the introduction of the Stakeholder Viability Principle (SVP). SVP changes the status of stakeholders from external beneficiaries, consultees, or validation audiences to co-constitutive participants in determining which futures count as relevant within a given decision episode. This move reduces an ambiguity frequently present in foresight practice, where desirable, plausible, or imaginable futures may coexist without clear criteria of membership or closure. Under the present framework, a future becomes relevant not merely because it is thinkable or normatively attractive, but because it remains viable in relation to the continuity, agency, and axiological coherence of critical actors in a specific context. In this respect, SDF replaces diffuse future plurality with a more disciplined ontology of viable futures [26,40].

The same research question also clarifies the role of social singularity. Rather than treating singularity as a typology of exceptional events, the article has defined it as a context-specific rupture condition in which existing routines of interpretation and action become insufficient and closure becomes decisionally urgent. This formulation is significant because it introduces a functional criterion of criticality into DF. Social singularity does not determine what the future will be; rather, it indicates when an organization can no longer remain in a purely exploratory mode without incurring strategic cost. In this sense, singularity is not an ornamental addition to the framework, but a mechanism linking signal accumulation to intervention windows and to the timing of closure [27–29].

Fourth, in response to RQ4, the article has argued that qualitative foresight inputs can become more traceable and analyzable through hybrid conceptual–computational instrumentation. The central contribution here lies in the combined use of expert-defined conceptual universes, semantic projections, and versioned traceability structures. This architecture does not eliminate interpretation; rather, it seeks to formalize interpretive operations sufficiently to make them more transparent, comparable, and auditable. By transforming signals and trends into semantic objects related to an explicitly documented conceptual universe, the method establishes a bridge between discursive evidence and evaluable representation. This is particularly important in a field where rich narrative outputs often exceed the available means for comparison, replication, or attribution. In this respect, the article supports the claim that DF becomes methodologically stronger when it renders visible the chain linking corpus, interpretation, filtering, and closure [15–18].

Fifth, concerning RQ5, the article has clarified the types of outputs and evaluative logic through which DF may become forecasting-compatible without being reduced to deterministic prediction. The outputs identified in this paper—semantic intensity indicators, transition detection, and shortlists of viable futures—do not function as point forecasts. Their relevance to forecasting lies instead in the fact that they are explicit anticipatory representations that may, under appropriate conditions, be evaluated across time horizons, compared across cases, or subjected to sensitivity analysis. In that sense, forecasting enters the framework instrumentally, as a source of evaluative discipline rather than as the defining purpose of DF. What becomes scientifically valuable is not a claim to predict the one true future, but the possibility of producing anticipatory outputs whose construction, filtering, and closure are explicit enough to be examined, contrasted, and progressively improved [7,9,13,14].

The forecasting compatibility of SDF rests on a concrete claim: that its outputs—semantic intensity indicators, transition-relevant signals, and shortlists of viable futures—are not merely descriptive but evaluable under defined assumptions and temporal horizons. Prior published work demonstrates that this evaluability is achievable in practice. The Deflexor Ranking Number, formalized as R = (gt + tt + tr)/3—combining Google Trends,

social media frequency, and expert report mentions—constitutes a quantifiable trend-strength indicator updated in real time, with a defined threshold (R = 100) for the transition from trend to megatrend status (Manetti, Lara-Navarra & Sánchez-Navarro, 2022 [17]). The Lipschitz-based Trend Index, derived from computed projections onto a conceptual universe and extended via McShane–Whitney regression, yields normalized scalar values that can be compared across time periods and signal configurations [16]. These instruments do not produce point predictions in the sense of narrow deterministic forecasting; they produce anticipatory representations that are explicit, traceable, and open to evaluation—which is precisely what forecasting, understood as the disciplined production of evaluable anticipations under explicit assumptions, requires [7,9]. The present article establishes the ontological and methodological architecture within which such instruments operate—specifically, the SVP-mediated filtering of viable futures and the singularity-sensitive logic of decision closure—but does not replicate their full technical development here, as this has been established in the cited prior studies. Consequently, the compatibility of SDF with forecasting is not merely conceptual but instrumental: it allows dynamic indicators, such as the DeflyCompass Ranking Number or the Lipschitz-based Trend Index, to transition from isolated metrics to auditable criteria that justify why an organization chooses to close the future space and commit to a specific strategic trajectory.

The contribution of SDF is thus not to produce new forecasting algorithms, but to provide the disciplinary framework that determines when, for whom, and under what viability conditions such instruments should be activated and their outputs interpreted.

Taken together, these responses yield three principal implications: ontological, mechanistic, and evaluative. The ontological implication is that SDF delimits a specific object of inquiry: not generic futures, but viable futures bounded by stakeholder conditions, critical transitions, and contextual constraints. The mechanistic implication is that DF can be described not merely in terms of methods employed, but in terms of an explicit process linking futurization, filtering, singularity-sensitive reduction, and closure. The evaluative implication is that DF becomes a stronger scientific candidate when its outputs and operations are traceable and, at least in principle, attributable, rather than only pragmatically useful. This triad—ontology, mechanism, and evaluability—constitutes the principal theoretical contribution of the article.

Methodologically, the paper also implies a shift in how DF research should be designed and reported. If DF is to function as an applied scientific program, narrative case demonstration is no longer sufficient. Research designs must begin to document explicit cycles of artifact configuration, semantic representation, filtering, closure, and learning. In practical terms, this includes versioning the conceptual universe, documenting parameter and threshold choices, recording evidence selection criteria, and reporting how changes in assumptions affect the resulting closure. Traceability is therefore not a supplementary appendix but a methodological requirement, because only a reconstructible chain from evidence to closure allows attribution, sensitivity analysis, and inter-case comparison [12,31,33].

This methodological reframing also affects empirical design. If social singularity is treated as a critical-transition condition, then decision episodes characterized by signals of disruption, accelerated change, or interpretive instability become especially relevant sites for empirical contrast. Such episodes allow the futurization–de-futurization mechanism to be examined under the conditions for which it is theoretically intended. Methodologically, this suggests the importance of designs capable of observing pre- and post-transition windows, the stability of closure under plausible perturbations, and the marginal contribution of specific components such as SVP, semantic projections, or singularity criteria. In this

sense, the framework points toward a more disciplined family of empirical designs, even though the present article does not claim to have fully executed them.

The discussion also requires an explicit statement of limitations and validity conditions.

First, the framework depends on the quality, diversity, and transparency of both the conceptual universe and the expert panel. Although semantic projections formalize representation, they do not automatically eliminate bias. If the conceptual universe is narrow, homogeneous, or insufficiently documented, the method may formalize and stabilize existing biases rather than counteract them. This makes documentation, diversity of expertise, and, where possible, contrast among alternative panel configurations essential conditions of validity [11,17].

Second, the use of thresholds and normalized indices introduces a risk of reified precision. Thresholds may be necessary for explicit filtering and closure, but they should not be interpreted as universal constants. Their legitimacy depends on contextual calibration, documentation, and sensitivity analysis. Without such precautions, formalization may generate a misleading appearance of exactness rather than a disciplined procedure for handling uncertainty [12].

Third, the framework makes selective use of geometric and critical-transition analogies, especially in relation to futures cones and singularity. These analogies are methodologically useful insofar as they clarify structured relations among trajectories, constraints, and closure conditions. However, they should be treated as disciplined conceptual aids rather than as claims of strong ontological equivalence with physical models. The scientific force of the framework depends not on metaphorical richness, but on whether these analogies improve the explicitness, tractability, and evaluability of the method.

Fourth, generalization in DF cannot plausibly take the form of universal laws. Because the method operates on highly situated decision episodes, its strongest form of transfer lies in context-dependent technological rules of the form "if. . . then. . .", consistent with the design-science tradition. This does not weaken the scientific ambition of the framework; rather, it aligns it with the kind of generalization appropriate to artifact-oriented and intervention-centered fields [32,34].

These limitations lead directly to a minimal research agenda for the consolidation of DF as a scientific discipline. A first priority is construct calibration and validation, especially for SVP-related evaluative dimensions and for the operational treatment of social singularity. A second priority is the development of comparative and attribution-oriented studies, including controlled contrasts between methodological configurations with and without specific components such as SVP filters or semantic projections. A third priority is multi-case replication and sectoral transfer, supported by sufficiently transparent reporting standards to allow comparison of viability profiles, closure patterns, and criticality conditions across domains. A fourth priority is the incorporation of critical and pluralizing approaches capable of preventing premature closure under dominant interpretive frames, thereby strengthening the axiological and epistemic legitimacy of closure itself [14,18,40,41].

In summary, Singular Design Foresight proposes a path through which DF may cease to function merely as a contingent assemblage of practices and begin to operate as an applied scientific program. It does so by delimiting an object of inquiry—viable futures—specifying a methodological mechanism—the movement from futurization to de-futurization under viability and criticality conditions—and establishing an evaluation architecture based on validity, utility, and attribution. Forecasting-compatible instrumentation strengthens methodological contrast without reducing DF to prediction. Its role is more specific, and arguably more important: to make anticipatory representation, closure, and learning sufficiently explicit that knowledge can accumulate across cases in a field that has too often relied on utility without theory [10,11,31].

8. Conclusions

This article has argued that Design Foresight (DF) can develop as a cumulative scientific discipline if three conditions are met: its ontological and epistemological commitments are made explicit, the mechanism linking exploration and closure is formalized, and its applications are evaluated through criteria that support accumulation, comparison, and attribution. On this basis, Singular Design Foresight (SDF) has been proposed not as an additional toolkit for future-oriented practice, but as a foundational methodological framework capable of structuring DF as an applied scientific program.

In response to RQ1, the paper has advanced an ontology centered on viable futures, delimited through the Stakeholder Viability Principle and shaped by context-specific conditions of critical transition. In response to RQ2, it has formalized the passage from futurization to de-futurization as the core mechanism through which Design Foresight moves from exploratory expansion to justified closure. In response to RQ3–RQ5, it has proposed a hybrid architecture combining expert-defined conceptual universes, semantic projections, and an evaluation logic based on validity, utility, and attribution, thereby making DF more explicit, traceable, and compatible with forecasting understood as evaluable anticipation rather than deterministic prediction.

The contribution of the article is therefore not predictive certainty, but auditable closure under uncertainty. By specifying its object of inquiry, its methodological mechanism, and its evaluation architecture, the paper seeks to reposition DF as a field capable of generating cumulative knowledge from situated decision episodes through traceable and comparable processes.

This contribution must be understood within clearly defined limits. The article does not claim conclusive system-level validation of the full SDF framework, nor does it demonstrate comparative superiority over alternative approaches. Instead, it draws on summative component-level evidence reported in prior studies—on semantic instrumentation, singularity detection, hybrid signal extraction, and mixed expert–corpus applications—to establish the plausibility, operational feasibility, and evaluability of the proposed framework, while explicitly delimiting integrated validation as a future research objective.

The next stage of this research program requires systematic construct calibration, the development of comparative and attribution-oriented designs, multi-case replication across domains, and the incorporation of critical and pluralizing approaches capable of preventing premature closure under dominant interpretive frames. Particular attention should be given to sensitivity analysis, to the role of alternative conceptual universes and expert configurations, and to the conditions under which closure remains stable under controlled variation.

If such developments are pursued through transparent, traceable, and comparable empirical designs, Singular Design Foresight may contribute to the consolidation of a scientifically defensible and cumulative discipline of Design Foresight—one in which anticipatory representations, filtering mechanisms, and decisional closure can be progressively examined, contrasted, and improved across cases.

Author Contributions: Conceptualization, P.L.-N., A.F.-S. and E.A.S.-P.; methodology, P.L.-N.; software, E.A.S.-P.; validation, P.L.-N. and E.A.S.-P.; formal analysis, P.L.-N. and A.F.-S.; investigation, P.L.-N., A.F.-S. and E.A.S.-P.; resources, A.F.-S.; data curation, P.L.-N. and E.A.S.-P.; writing—original draft preparation, P.L.-N. and E.A.S.-P.; writing—review and editing, A.F.-S. and E.A.S.-P.; supervision, E.A.S.-P.; project administration, A.F.-S.; funding acquisition, A.F.-S. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Generalitat de València, grant number PROMETEO CIPROM/2023/32.

Data Availability Statement: Dataset available on request from the authors. The raw data supporting the conclusions of this article will be made available by the authors on request.

Acknowledgments: During the preparation of this manuscript/study, the authors used ChatGPT (GPT-5 model) and DeepL Pro for grammar support, language refinement, and spell-checking. The authors have reviewed and edited all generated content and take full responsibility for the final version of this publication.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Bennett, N.; Lemoine, G.J. What VUCA really means for you. *Harv. Bus. Rev.* **2014**, *92*, 27–31.
2. Cascio, J. Facing the age of chaos. *Medium*, 30 April 2020. Available online: https://medium.com/@cascio/facing-the-age-of-chaos-b00687b1f51d (accessed on 5 April 2026).
3. Borja de Mozota, B. *Design Management: Using Design to Build Brand Value and Corporate Innovation*; Allworth Press: New York, NY, USA, 2003.
4. Calabretta, G.; Gemser, G.; Karpen, I.O. (Eds.) *Strategic Design: Eight Essential Practices for Every Business*; BIS Publishers: Amsterdam, The Netherlands, 2016.
5. Verganti, R. *Design-Driven Innovation: Changing the Rules of Competition by Radically Innovating What Things Mean*; Harvard Business Press: Boston, MA, USA, 2009.
6. Bell, W. *Foundations of Futures Studies: Human Science for a New Era, Vol. 1: History, Purposes, and Knowledge*; Routledge: London, UK, 2003.
7. Bishop, P.; Hines, A.; Collins, T. The current state of scenario development: An overview of techniques. *Foresight* **2007**, *9*, 5–25. [CrossRef]
8. Miles, I.; Saritas, O.; Sokolov, A. (Eds.) *The Handbook of Technology Foresight*; Edward Elgar Publishing: Cheltenham, UK, 2016.
9. Voros, J. A generic foresight process framework. *Foresight* **2003**, *5*, 10–21. [CrossRef]
10. Sutton, R.I.; Staw, B.M. What theory is not. *Adm. Sci. Q.* **1995**, *40*, 371–384. [CrossRef]
11. Piirainen, K.A.; Gonzalez, R.A. Theory of and within foresight—What does a theory of foresight even mean? *Technol. Forecast. Soc. Change* **2015**, *96*, 191–201. [CrossRef]
12. Georghiou, L.; Keenan, M. Evaluation of national foresight activities: Assessing rationale, process and impact. *Technol. Forecast. Soc. Change* **2006**, *73*, 761–777. [CrossRef]
13. Rohrbeck, R.; Schwarz, J.O. The value contribution of strategic foresight: Insights from an empirical study of large European companies. *Technol. Forecast. Soc. Change* **2013**, *80*, 1593–1606. [CrossRef]
14. Rohrbeck, R.; Kum, M.E. Corporate foresight and firm performance: A systematic literature review. *Technol. Forecast. Soc. Change* **2018**, *129*, 105–116. [CrossRef]
15. O'Halloran, K.L.; Tan, S.; Pham, D.-S.; Bateman, J.; Vande Moere, A. A digital mixed methods research design: Integrating multimodal analysis with data mining and information visualization for big data analytics. *J. Mix. Methods Res.* **2018**, *12*, 11–30. [CrossRef]
16. Manetti, A.; Ferrer-Sapena, A.; Sánchez-Pérez, E.A.; Lara Navarra, P. Design trend forecasting by combining conceptual analysis and semantic projections: New tools for open innovation. *J. Open Innov. Technol. Mark. Complex.* **2021**, *7*, 92. [CrossRef]
17. Manetti, A.; Lara Navarra, P.; Sánchez-Navarro, J. El proceso de diseño para la generación de escenarios futuros educativos. *Comunicar* **2022**, *30*, 33–44. [CrossRef]
18. López Bordao, M.; Ferrer-Sapena, A.; Reyes Pérez, C.A.; Sánchez-Pérez, E.A. Defly Compass trend analysis methodology: Quantifying trend detection to improve foresight in strategic decision making. *Information* **2025**, *16*, 605. [CrossRef]
19. Bleecker, J. *Design Fiction: A Short Essay on Design, Science, Fact and Fiction*; Near Future Laboratory: San Francisco, CA, USA, 2009.
20. Auger, J. Speculative design: Crafting the speculation. *Digit. Creat.* **2013**, *24*, 11–35. [CrossRef]
21. Dunne, A.; Raby, F. *Speculative Everything: Design, Fiction, and Social Dreaming*; MIT Press: Cambridge, MA, USA, 2013.
22. Candy, S.; Dunagan, J. Designing an experiential scenario: The people who vanished. *Futures* **2017**, *86*, 136–153. [CrossRef]
23. Candy, S.; Potter, C. Introduction to the Special Issue: Design and futures (Vol. I). *J. Futures Stud.* **2019**, *23*, 1–2. [CrossRef]
24. Fry, T. *Defuturing: A New Design Philosophy*; Bloomsbury Visual Arts: London, UK, 2020.
25. Fu, Z.; Xia, Q. Design foresight: A unified approach integrating futurization and de-futurization in design. *J. Futures Stud.* **2025**, *30*, 15–32.
26. Christophilopoulos, E. Special relativity theory expands the futures cone's conceptualisation of the futures and the pasts. *J. Futures Stud.* **2021**, *26*, 83–90.
27. Plowman, D.A.; Baker, L.T.; Beck, T.E.; Kulkarni, M.; Solansky, S.T.; Travis, D.V. Radical change accidentally: The emergence and amplification of small change. *Acad. Manag. J.* **2007**, *50*, 515–543. [CrossRef]

28. Scheffer, M. *Critical Transitions in Nature and Society*; Princeton University Press: Princeton, NJ, USA, 2009.
29. Lara-Navarra, P.; Sánchez-Pérez, E.A.; Ferrer-Sapena, A.; Fitó-Bertran, À. Singularity in higher education: Methods for detection and classification. *Expert Syst. Appl.* **2024**, *239*, 122306. [CrossRef]
30. Lara-Navarra, P.; Ferrer-Sapena, A.; Ismodes-Cascón, E.; Fosca-Pastor, C.; Sánchez-Pérez, E.A. The future of higher education: Trends, challenges and opportunities in AI-driven lifelong learning in Peru. *Information* **2025**, *16*, 224. [CrossRef]
31. Hevner, A.R.; March, S.T.; Park, J.; Ram, S. Design science in information systems research. *MIS Q.* **2004**, *28*, 75–105. [CrossRef]
32. Aken, J.E. Management research based on the paradigm of the design sciences: The quest for field tested and grounded technological rules. *J. Manag. Stud.* **2004**, *41*, 219–246. [CrossRef]
33. Peffers, K.; Tuunanen, T.; Rothenberger, M.A.; Chatterjee, S. A design science research methodology for information systems research. *J. Manag. Inf. Syst.* **2007**, *24*, 45–77. [CrossRef]
34. Buchanan, R. Wicked problems in design thinking. *Des. Issues* **1992**, *8*, 5–21. [CrossRef]
35. Carter, B. Large number coincidences and the anthropic principle in cosmology. In *Confrontation of Cosmological Theories with Observational Data*; Longair, M.S., Ed.; Reidel: Dordrecht, The Netherlands, 1974; pp. 291–298.
36. Hawking, S. *A Brief History of Time*; Bantam Books: New York, NY, USA, 1988.
37. Wenham, M. Singular problems in science and science education. *J. Philos. Educ.* **1987**, *21*, 47–58. [CrossRef]
38. Creswell, J.W. *A Concise Introduction to Mixed Methods Research*; SAGE Publications: Los Angeles, CA, USA, 2015.
39. Amara, R. The futures field: Functions, forms, and critical issues. *Futures* **1974**, *6*, 289–301. [CrossRef]
40. Inayatullah, S. Six pillars: Futures thinking for transforming. *Foresight* **2008**, *10*, 4–21. [CrossRef]
41. Inayatullah, S. Causal layered analysis: Poststructuralism as method. *Futures* **1998**, *30*, 815–829. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.