International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(23s) 2026 pp. 1757-1770
https://doi.org/10.70917/ijcisim-2026-5727
Article

Integrating Social Determinants of Health Data into Population Health Management: Acquisition, Standardization, and Operational Use at Scale

Sushrut Naik

Independent Researcher,India. Email: sushrut.s.naik@gmail.com

Abstract: Clinical and claims data capture what happens inside the healthcare system. Most health outcomes are driven by factors outside it, including housing instability, food insecurity, transportation barriers, and social isolation, none of which appear in an encounter record or a claim. Social determinants data is widely collected and rarely acted upon: it is screened at intervals, aggregated into dashboards and equity scorecards, and left in a reporting layer structurally disconnected from the systems that drive care management decisions. This paper argues that the reasons are specific and addressable rather than cultural. Social risk data has historically lacked a common structure, its governance status is ambiguous enough that organisations default to analytics-only use, it is collected point-in-time and stale by the time it could inform a decision, it sits outside the clinician's workflow, and community-level proxies carry weak individual-level signal that justifies distrust. The approach described treats social risk as a first-class standardised data domain inside the health information exchange, structured through the established social care implementation guide so that it is exchanged and updated alongside clinical and claims data rather than stored beside it. Three properties make it operational: provenance and confidence tiering so care teams know whether a flag is individually confirmed or area-inferred, a precedence hierarchy that resolves disagreement without discarding evidence, and a strict guardrail that area-level indices may elevate a person into a screening queue but never stand alone as an individual determination. No coverage rates, population counts, or model performance figures are reported, because the underlying project data is no longer accessible to the author.

Keywords: Social determinants of health, population health management, health information exchange, FHIR, data standardization, risk stratification, closed-loop referral, community-based organizations, health equity, ecological fallacy, provenance, quality measures

1. INTRODUCTION

1.1 What the record does not contain

A population health model built on clinical and claims data is built on the healthcare system's own record of itself. It knows what was diagnosed, what was billed, what was dispensed, and what was utilised. It does not know whether the patient can reach the pharmacy, afford the food the diet requires, or keep a stable address. Those conditions, the circumstances in which people are born, grow, work, and age, are what the term social determinants of health refers to [1].

The gap is not marginal. A patient can appear low risk in every recorded dimension and be at high risk for reasons the record has no field for.

1.2 The case that illustrates the gap

A published and widely cited programme illustrates the mechanism precisely, and it is external work rather than the author's own. A health system found that extensive and clinically sound work on type 2 diabetes management was underperforming expectation, because patients could not consistently access the food their diet-responsive condition required. No clinical or claims record captured that barrier.

The response was to embed two food insecurity questions directly into the encounter record, asked at every visit, and to link the answers with clinical data to identify patients with uncontrolled diabetes living near a participating food resource. Eligibility combined a clinical threshold with a screening result, the latter being a variable no claims field would surface. Published outcomes from that programme, including a subsequent randomised trial tracking glycaemic control as its primary endpoint across two sites, are reported in the literature.

This is cited as published external evidence for the mechanism, not as a result of the work described in this paper. The distinction matters because the mechanism it demonstrates, that a two-question screening instrument can identify a barrier no structured clinical field contains, is the entire argument for treating social risk as a data domain rather than as context.

1.3 Scope of the work

The author's work with social determinants data spans risk stratification and attribution processes in a healthcare consulting capacity from approximately 2015 to 2016, and risk adjustment and quality measure work at a population health platform from approximately 2022, at which point quality measure sets had begun to include social risk screening measures, notably the social need screening and intervention measure in the HEDIS specifications [2]. Data in that work was sourced in part from third-party social risk data vendors.

1.4 Contribution

1. A four-tier taxonomy of social risk data sources, characterised by the level at which each is available, how complete and current it is, and what it costs.

2. A standardisation approach based on the established social care implementation guide, with a precedence hierarchy that resolves disagreement without discarding evidence.

3. A treatment of the ecological fallacy as an architectural guardrail rather than a caveat, with area-level data structurally prevented from acting as an individual determination.

4. A concrete account of what a care team does differently, and the mechanisms that stop the data becoming information nobody acts on.

What this paper does not do. It reports no coverage rates per data element, no population counts, and no model performance comparison with and without social risk data. The project data required to compute or verify any of those is no longer accessible to the author. Section 11 states the evidence position, and where external published benchmarks are referenced they are labelled as external rather than presented as outcomes of this work.

2. Why the Data Stays in the Reporting Layer

Most health systems and payers collect social risk data through periodic screening embedded in the encounter record or in member surveys, then aggregate it into dashboards, population health reports, and equity scorecards. That data typically lives in a separate module or warehouse used for retrospective analytics rather than being structurally joined to the systems driving clinical or care management decisions.

Five structural reasons keep it there, and none of them is a failure of intent.

No standardised, interoperable format. Unlike clinical data, which has decades of terminology and exchange standardisation behind it, social risk data has historically lacked common structure: different screening instruments, different coding systems, and inconsistent granularity between individual-level and community-level proxy.

Governance ambiguity. Whether social risk data constitutes protected health information, who the responsible entity is, and what consent applies to downstream use are frequently unclear. Organisations facing that ambiguity default to internal analytics rather than operational triggers, which is the conservative and defensible choice.

Point-in-time collection. Screening happens once per visit or per enrolment cycle, so the data is stale by the time it could inform a live decision.

Disconnected from workflow. Even when collected in the encounter record, social risk fields sit outside the clinician's or care manager's actual workflow: visible in a chart tab, not wired into alerts, referrals, or risk scores.

Weak individual-level signal in isolation. Community-level proxies are noisy at the individual level, so organisations do not trust them enough to act on automatically. This reinforces the report-only pattern, and it is a rational response to a real property of the data rather than excessive caution.

A sixth reason is worth adding to these five, because it is organisational rather than technical. Social risk data typically arrives without an owner who is accountable for acting on it. Clinical data has an owner by construction: someone is responsible for the result and answerable if it is not acted upon. Social risk data is frequently collected by a quality or equity function whose remit ends at reporting, and delivered to a care function that did not request it and is not measured on it. Section 5's final mechanism, tying the data to measures the organisation is already accountable for, is the response to this and is arguably the most important of the six.

The last point is the one that connects to the architecture in Section 5. An organisation that cannot distinguish a confirmed individual answer from a neighbourhood inference is right not to act on either automatically. Making that distinction visible is what makes action defensible.

3. Where the Data Comes From

Four tiers differ in the level at which they are available, in completeness and currency, and in cost. Conflating them is the source of most downstream error.

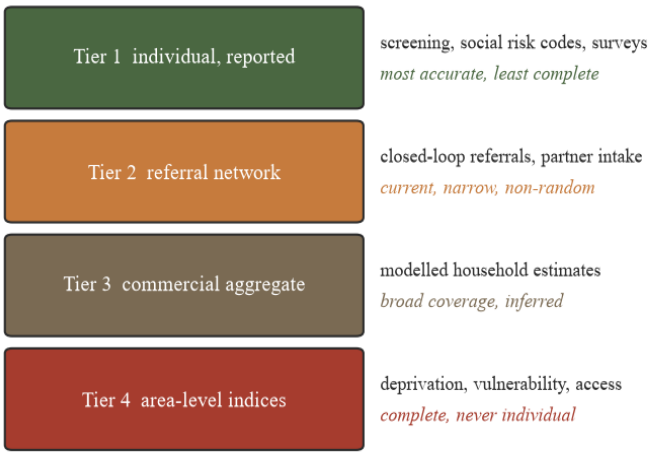

Figure 1. The four source tiers. Coverage and individual accuracy move in opposite directions down the list.

Table 1. Social risk data sources by tier.

Tier	Sources	Level	Completeness and currency	Cost

1. Individual, patient-reported	Encounter-based screening instruments; social risk diagnosis codes in claims; payer member surveys and health risk assessments	Individual	Screening is only as complete as uptake, and goes stale between encounters. Social risk diagnosis codes are under-coded because they do not drive reimbursement. Surveys are annual or per enrolment cycle with modest response rates	Screening and coding carry no licence cost; the real cost is clinical staff time and workflow build. Surveys carry a per-completed-response cost
2. Community and referral network	Closed-loop referral platforms; community organisation intake systems	Individual, but only for those already screened and referred	Very current and specific, with referral status updating in near real time, but coverage is narrow and non-random. Community organisation data is similarly specific and fragmented across systems	Referral platforms licensed to the health system, payer, or exchange rather than per record. Community organisation data often free through partnership; integration is the real cost
3. Commercial aggregators	Consumer and credit-adjacent data modelling services	Individual or household, modelled rather than reported	High population coverage refreshed monthly or quarterly, but modelled from consumer sources rather than self-reported or clinically verified. Complete does not mean accurate for any one person	Licensed per member per month; not publicly priced
4. Public area-level indices	Deprivation, vulnerability, and opportunity indices such as the CDC/ATSDR Social Vulnerability Index [3]; census survey data [4]; food access and housing cost datasets	Neighbourhood only, at tract, block group, or postal level. Never individual or household	Complete national coverage at tract level, but built on survey data typically lagging the present by two to three years	Free and publicly available

Two properties of this table drive everything downstream. Coverage and accuracy move in opposite directions across the tiers. Tier 1 is the most accurate for a given person and the least complete; tier 4 is complete and says nothing about any individual. And tier 2's coverage is non-random by construction, since it exists only for people

already screened and referred, which means any analysis treating it as a sample of the population is biased in a direction that is easy to overlook.

4. Standardisation

4.1 Reconciling sources

The foundation is the established social care implementation guide [5], which defines standardised terminologies for screening questions and conditions and a common set of exchange resources for representing social risk data regardless of origin. It is a constrained profile of the underlying health data exchange standard [6], which matters practically: social risk travels on the same transport, with the same resource types and the same access controls, as the clinical data it has to sit beside. Every incoming source, whether encounter screening, payer assessment, referral platform, community organisation intake, commercial aggregator, or area-level index, is mapped into this common schema before entering the exchange, rather than each source retaining its own proprietary format.

Each data point additionally carries provenance metadata: source type, meaning self-reported, referral-confirmed, modelled, or area-level; collection date; and confidence tier.

Reconciliation is therefore not about picking a winner. All inputs are preserved with their provenance, and a precedence hierarchy applies only when a single value is needed for a downstream decision.

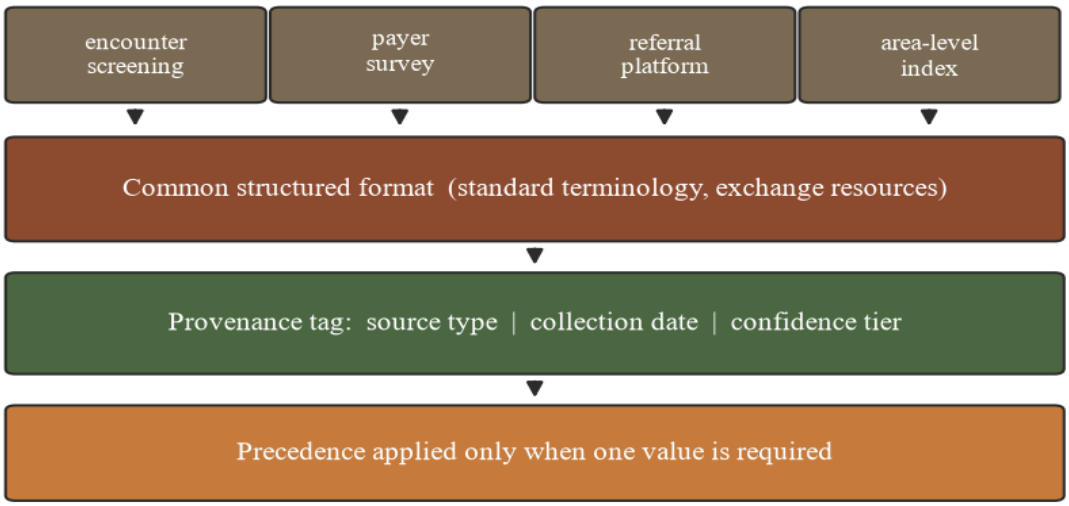

Figure 2. Standardisation and provenance tagging. Every source maps into one format before entering the exchange, carrying its provenance with it.

Table 2. Precedence hierarchy when one value is required.

Rank	Source class	Rationale
1	Self-reported or screened	Individual-specific and directly attested
2	Referral or community-confirmed	Individual-specific and behaviourally confirmed
3	Commercial modelled	Individual or household level, inferred
4	Area-level proxy	Not individual at all; contextual only

The rule is to default to the most individual-specific and most recently verified source available for that patient.

4.2 What the precedence hierarchy costs

The hierarchy resolves disagreement, and it also encodes a value judgement worth surfacing. Ranking self-report above modelled data assumes the person's own account is the most accurate available, which is generally right and is not always right: a patient may under-report a sensitive circumstance, particularly one carrying stigma or perceived consequence, and a modelled estimate may be closer to the truth in exactly those cases.

The design does not resolve this and should not, because the alternative is worse. Weighting a modelled inference above a person's own answer about their circumstances is a position an institution would struggle to defend to that person. The hierarchy therefore accepts a known bias toward under-detection in the most sensitive domains, and the mitigation is repeated screening in a trusted context rather than a change to the ranking.

4.3 Missing and contradictory data

Missing data is never silently imputed as no risk. A null screening result is coded distinctly from a confirmed-negative result, so downstream models and care teams can distinguish "screened, no need identified" from "not yet screened." Collapsing those two is the single most consequential error available in this domain, because it converts absence of information into evidence of absence.

Contradictions trigger a reconciliation flag rather than automatic overwrite. If a payer assessment records stable housing while a recent referral shows a housing assistance request, both records are retained with timestamps and the more recent, more individually verified source is weighted higher, but the conflict itself is surfaced to a care coordinator rather than resolved invisibly.

Staleness is tracked explicitly. Every data point carries a last-verified date, and data older than a defined threshold, varying by domain, is automatically downgraded in confidence rather than treated as current indefinitely.

4.4 The area-level guardrail

This is the most important guardrail in the design and it is handled through explicit tiering, never blending.

Area-level indices attach to a patient's record as a contextual attribute, clearly tagged as neighbourhood-level and not individually confirmed, never merged into the same field as a self-reported answer.

Area-level data is used only for population-level triage and outreach prioritisation, such as flagging a cohort for proactive screening. It is never a standalone trigger for an individual clinical or coverage decision. An individual is never scored as food insecure on the basis of their census tract. The tract data can only elevate them into a screen-this-person queue.

Any model or workflow using area-level data must carry that provenance forward into its output, so a downstream user can see whether a given social risk flag is individually confirmed or area-inferred and weight their action accordingly. Governance rules explicitly prohibit using area-level proxy alone for consequential individual decisions such as benefit determinations or risk scores tied to payment.

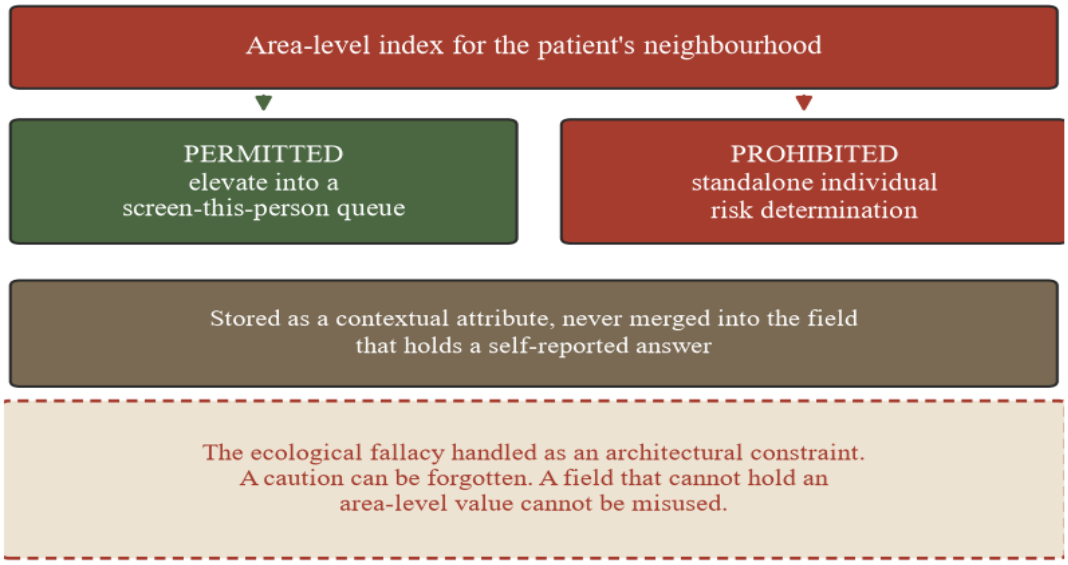

Figure 3. The area-level guardrail. Neighbourhood data can elevate a person into a screening queue and can never become an individual determination.

This is the ecological fallacy treated as an architectural constraint rather than as a caution in a methods appendix. A caution can be forgotten; a field that cannot hold an area-level value cannot be misused.

5. What Changes in the Care Team's Work

Without integrated social risk data, a care manager works from a clinical risk score. They see an elevated glycaemic result and call to reinforce medication adherence and schedule follow-up: the standard clinical playbook.

With individually attached social risk data present, the same care manager's queue instead surfaces that the member is confirmed food insecure, screened and verified recently, and has an open transportation barrier flagged by a recent missed appointment. Four things change concretely.

The outreach becomes a referral trigger. An automatic service request to a food resource programme and a transportation assistance referral are generated in the same workflow as the clinical outreach, not as a separate task.

The conversation changes. The next call includes a specific actionable question about whether the member can reach the pharmacy or needs delivery, rather than a generic adherence reminder, because the system indicated why adherence might be failing.

The risk tier changes. The member moves from moderate clinical risk to high combined risk, which changes outreach frequency and which staff role is assigned, including social work co-assignment rather than nursing alone.

Quality measurement follows automatically. The member is counted in the social risk screening measure and tracked toward closed-loop referral completion, rather than only against clinical outcome measures.

Figure 4. From flag to closed loop. Only loop closure counts as resolution.

Table 3. Mechanisms that prevent this becoming information nobody acts on.

Mechanism	**Why it works**
Embedded in existing workflow, not a separate screen	The referral and risk-tier change occur inside the normal task queue, so acting on it is not extra work layered on the job
Every flagged risk becomes a trackable task with an owner and a due date	A service request rather than a static field
Open referrals escalate	An unresolved referral after a defined window escalates to a supervisor queue, as an overdue clinical task would
Closed-loop tracking makes inaction visible	Bidirectional confirmation with the community partner means "we referred them" is not treated as done; only loop closure is
Tied to something the organisation is already measured on	Screening measures, contract requirements, and risk adjustment accuracy give the data an owner with a real incentive
Provenance and confidence tiering	Care teams act on a flag they can see is confirmed and recent with the same confidence as a laboratory value

The last mechanism is the one that connects the architecture to the behaviour. Section 2 identified distrust of noisy proxies as a reason the data stays in reports. Showing the care team that a specific flag is individually confirmed and recently verified is what makes acting on it reasonable rather than reckless.

6. The Ordered Path from Acquisition to Action

1. Collection. Data enters through one of four channels: encounter-based screening, payer assessment survey, community referral platform, or public and commercial area-level source.

2. Standardisation. Each data point is mapped into the common structured format with standardised terminology for screening questions and conditions [5].

3. Provenance tagging. Source type, collection date, and confidence tier are attached.

4. Reconciliation. Where sources disagree, the most individual-specific and recently verified source is weighted highest; conflicts are flagged for human review rather than silently overwritten.

5. Attachment. Individual-level data attaches directly to the patient record; area-level data attaches as a separate, clearly labelled contextual attribute.

6. Staleness check. Data older than the domain threshold is downgraded in confidence.

7. Risk stratification. Clinical and social risk combine into a combined tier determining outreach frequency and staff assignment.

8. Workflow trigger. A flagged need automatically generates a task in the care team's existing queue.

9. Care team action. The care manager acts on the specific individualised task.

10. Closed-loop tracking. The referral platform confirms whether the service was actually delivered, not merely initiated.

11. Escalation. An open task unresolved after a defined window escalates.

12. Reporting. Outcomes and completion roll up into quality reporting and programme review.

Steps 3, 5, and 6 are what distinguish this from a conventional integration. They exist so that step 9 can be taken with calibrated confidence rather than uniform trust.

7. Risks

Profiling by neighbourhood. Area-level indices describe a tract, not a person, and community-level estimates are imprecise or biased predictors of any individual's circumstances. The danger is that a patient in a high-deprivation tract is silently treated as high risk regardless of their actual situation, while someone in a low-risk tract with real unscreened needs is overlooked. Unchecked, this hardens into geographic redlining inside a clinical or payer system, where postal code rather than the person drives the decision. The mitigation is not better area-level data. It is structural, as set out in Section 4.3.

Consent and sensitivity. Social risk data touches among the most sensitive categories a person can disclose, including housing instability, interpersonal violence, food insecurity, and substance use. Three risks compound. Regulatory ambiguity about whether the data is protected creates gaps in consent protection depending on which system touched it first. Purpose drift means data collected for a food resource referral can end up feeding a risk score or a cost-management model, a use the patient never consented to. And disclosure risk attaches to the person rather than only to the data: unlike a laboratory value, disclosing a social risk status to the wrong person can create immediate harm.

Data that ages badly. Social circumstances change faster than systems refresh. A job loss, an eviction, or a move happens in weeks, while a screening was answered months ago. Two failure modes result: a false negative, where a formerly stable patient in crisis is invisible because their last verified answer was reassuring, and a false positive, where someone who has recovered stability is still treated as high risk. Area indices lag further still. Without explicit staleness tracking, no data and old data both get treated as no risk, which is precisely backwards.

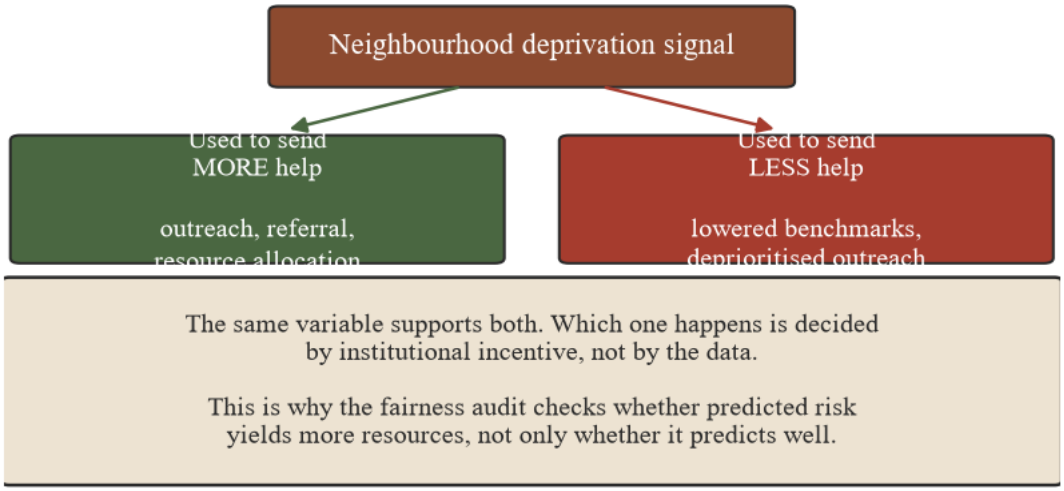

Figure 5. The same signal supports opposite actions, and which one happens is an institutional question rather than a data question.

Models that entrench disadvantage. This is the deepest risk. If a model learns that disadvantaged neighbourhoods correlate with worse outcomes, and that correlation is used to lower expected benchmarks, deprioritise outreach, or justify under-investment, the model does not merely reflect inequity; it automates it. The same variable can justify sending more help or sending less, depending entirely on institutional intent and incentive structure. Any social-risk-informed model therefore needs an explicit fairness audit checking not only accuracy but whether predicted risk translates into more resources for disadvantaged groups or is being used to ration care away from them.

8. Why the Data Ages, and What Follows

Staleness receives a single line in most treatments of this subject and deserves more, because it is the property that most distinguishes social risk data from clinical data and the one most often handled by ignoring it.

The refresh rates and the change rates do not match. A laboratory value describes a physiological state that changes over weeks to months, and it is re-measured on a clinical cadence that roughly matches. A housing situation can change in a week, and it is re-screened annually or per encounter. The mismatch is not a collection failure; it is a structural property of asking a question at a visit about a circumstance that does not wait for visits.

Both error directions are consequential and they are not symmetric. A false negative, where a patient whose circumstances have deteriorated still carries a reassuring answer from months ago, results in someone in crisis being invisible to a system that believes it has screened them. A false positive, where someone who has recovered stability is still flagged, results in misallocated outreach. The first is worse, and it is also the one a system with no staleness tracking produces silently, because a stale reassuring answer is indistinguishable from a current one.

Area-level data ages worst and is trusted longest. Survey-derived indices lag the present by two to three years by construction, because the multi-year rolling estimates they are built from are not published until the collection window closes [4]. They are also the most stable-looking data in the system, since they change slowly and never appear anomalous, which makes them the least likely to prompt anyone to ask when they were last accurate.

The design response has three parts and only two are implemented. Explicit last-verified dating and automatic confidence downgrade past a threshold are described in Section 4 and are straightforward. The third, re-screening triggered by life-event signals already visible in claims or encounter data rather than by the calendar, is proposed in

Section 13 and is where the real improvement lies: an emergency visit, a coverage change, or a run of missed appointments are all observable and all correlate with the circumstances social screening asks about.

The rule that follows. Absence of data and age of data must be represented distinctly from each other and from a negative finding. A system with three states, meaning not screened, screened negative, and screened positive, is missing the fourth that matters most: screened positive or negative long enough ago that the answer no longer carries weight. Collapsing that fourth state into either of the others is the specific error that turns a screening programme into a source of false confidence.

9. Evidence and Its Limits

Table 4. Evidence status by claim class.

Claim class	**Status**
Social risk data used for risk stratification, attribution, risk adjustment, and quality measures	Established from the author's work
Data sourced from third-party social risk vendors	Established
Four-tier source taxonomy, Table 1	Design analysis. Completeness and cost characterisations are practitioner knowledge, not measured from a specific programme
Standardisation and precedence design	Design position
Area-level guardrail	Design position
Workflow and closed-loop design	Design position
Coverage rate per data element	**Not available.** Project data no longer accessible to compute or verify
Population counts	**Not available**
Model performance with and without social risk data	**Not available.** No before-and-after comparison can be produced
Whether coverage analyses or model comparisons exist	**Unknown.** The author no longer has access to determine this
The diabetes and food insecurity programme in Section 1.2	**External published work.** Not a result of the work described here
Screening uptake and survey response rate characterisations	External published benchmarks, labelled as such, not measured here

Three points need making plainly.

The paper has no quantitative results of its own. Every number that could be cited belongs either to external published literature or to project data the author cannot access.

External benchmarks are cited as external. Where published studies report typical screening uptake, survey response rates, or discrimination improvements from adding social risk features to a model, those are properties of the cited studies. They are not outcomes of the work described here, and combining the two registers would misrepresent both.

The reviewer question is unresolved. Independent technical and disclosure review is appropriate for this material, particularly given the protected-information sensitivities in Section 7. Self-review does not satisfy that, and identifying an independent reviewer with relevant domain knowledge remains outstanding.

10. What an Implementing Organisation Faces

Four practical obstacles stand between the architecture and an operating system, and none is technical.

Screening capacity is the binding constraint on tier one. Everything individual-level depends on someone asking the questions, and asking them consumes clinical time in an encounter that is already full. The architecture can make the answers useful and cannot make the asking cheaper, which means the highest-value data tier is gated by the scarcest resource in the setting.

The referral network has to exist before the referrals do. Section 5's workflow generates service requests to community organisations. Establishing those relationships, agreeing data exchange, and confirming capacity is partnership work with a long lead time, and a technically complete system with no partner network produces documented need and no response to it.

Governance sign-off precedes operational use, not the reverse. Section 2 identifies ambiguity about the protected status of social risk data as a structural reason it stays in reporting. An organisation moving it into operational use has to resolve that ambiguity for itself, in advance, in writing, and that is a legal exercise rather than an architectural one.

The measure has to matter to someone. Section 5's final mechanism ties the data to obligations the organisation is already accountable for. Where no such obligation exists, the system depends on goodwill, and goodwill loses to competing priorities on a predictable schedule.

The pattern is that the architecture removes the technical reasons the data stays in reports and leaves the organisational ones untouched. That is the correct division of labour for a technical contribution and it should not be mistaken for the whole problem.

11. LIMITATIONS

Referral capacity is assumed. The workflow in Section 5 generates referrals to community organisations. Whether those organisations have capacity to receive them is outside the system's view, and a screening programme that identifies need faster than the network can meet it produces documented unmet need, which is a different and sometimes worse position than not having screened.

The closed loop depends on partner participation. Bidirectional confirmation requires the community organisation to report back, which requires them to be on a compatible platform and to have reason to use it. Section 3 notes that community organisation systems are fragmented, and the closed-loop mechanism is strongest exactly where that fragmentation is least, which is not where need is greatest.

No measured outcomes. As Section 9 records, the paper reports no coverage, no population, and no model performance. It is a design and practice contribution.

The strongest illustration is not the author's work. The programme in Section 1.2 demonstrates the mechanism convincingly and belongs to someone else. The paper's own contribution is architectural, and a reader should not carry the persuasiveness of the external case over to the untested parts of the design.

Data source characterisations are not measured. Table 1's completeness, currency, and cost entries reflect practitioner knowledge of the market rather than measurements from a specific implementation, and cost figures in particular vary by contract and are not publicly priced.

The precedence hierarchy is unvalidated. Section 4.1 ranks source classes by trustworthiness on reasoning that individual-specific and recently verified sources should dominate. Whether that ranking produces better decisions than alternatives was not tested.

Governance ambiguity is described, not resolved. Section 2 identifies unclear protected-information status as a structural barrier and Section 7 develops the risk. The paper offers design responses in the form of purpose-of-use rules and provenance carry-forward; it does not resolve the underlying legal ambiguity, which is not resolvable by architecture.

The work spans a long period and two roles. Experience runs from around 2015 to 2016 and from 2022, during which the standards landscape changed substantially. The standardisation approach in Section 4 reflects the later position, and the earlier work predates it.

Consent mechanics are not designed. Section 7 identifies purpose drift as a specific risk and Section 12 proposes patient-facing controls as future work. What consent model would actually permit referral use while excluding risk-scoring use, and how that preference propagates to every downstream consumer of the exchange, is not designed here. It is the hardest unsolved problem in the area and the paper does not solve it.

Fairness auditing is named, not specified. Section 7 argues that a social-risk-informed model needs an audit checking whether predicted risk yields more resources or fewer. What that audit measures concretely is not set out, and it is the paper's least developed control.

12. FUTURE WORK

A bias and equity audit layer running alongside the risk model, tracking not only predictive accuracy but whether high-social-risk individuals actually receive more outreach and resource allocation rather than less, and flagging drift toward using social risk data to justify reduced investment. This is the direct answer to the deepest risk in Section 7 and currently the least specified control.

Patient-facing consent and transparency controls, letting individuals see what social data is held, who it is shared with, and for what stated purpose, with granular opt-out for non-clinical uses. This addresses purpose drift at its source rather than through internal policy alone.

Active staleness management, with automatic re-screening triggered by life-event signals already visible in claims or encounter data, such as an emergency visit, a coverage change, or a series of missed appointments, rather than waiting for the next scheduled screening.

Closed-loop outcome tracking beyond referral completion, measuring whether the social intervention actually changed the clinical trajectory, so the system builds its own evidence base rather than relying on external literature. This is what would eventually give the approach results of its own.

A confidence-weighted interface for care teams, surfacing not only a risk flag but its provenance and confidence tier, so frontline staff calibrate trust rather than treating every flag as equally certain.

A note on sequencing. Active staleness management should come first, because it addresses the failure mode in Section 8 that currently operates silently, and because it can be built from signals the estate already holds without new collection. The confidence-weighted interface follows closely, since staleness tracking produces information the care team can only act on if it is shown to them.

Closed-loop outcome tracking is the item that would eventually give this work evidence of its own, and it is placed later because it requires the referral loop to be operating reliably first. The bias and equity audit is listed as answering the deepest risk and is the hardest to specify, which is a reason to start defining it early even though it will complete last.

13. CONCLUSION

Social determinants data is not scarce and is not ignored. It is collected, aggregated, reported, and left in a layer that no operational decision reads.

The reasons are structural. The data lacked a common format, its governance status is ambiguous enough to make analytics-only use the safe default, it is collected point-in-time and ages quickly, it sits outside the workflow, and its community-level form carries weak individual signal that justifies distrust.

The response is to treat social risk as a first-class standardised domain inside the exchange rather than a side system. Three properties make that operational. Provenance and confidence tiering let a care team see whether a flag

is individually confirmed or area-inferred. A precedence hierarchy resolves disagreement when a single value is needed without discarding the evidence that disagreed. And the area-level guardrail is architectural rather than advisory: a neighbourhood index can elevate someone into a screening queue and can never, structurally, become an individual determination.

That last constraint is the one to keep. The same variable that identifies who needs more help can justify giving less, and which of those happens is decided by institutional incentive rather than by the data. Building the system so the inference cannot masquerade as a fact is the part that does not depend on everyone's intentions staying good.

What the paper does not have is evidence of its own. The most persuasive case in it is published work by others, and the coverage, population, and model performance figures that would demonstrate the approach are in project data no longer accessible. The closed-loop outcome tracking in Section 12 is what would change that, because it is the only proposed component that would generate results rather than consume them.

REFERENCES

1. World Health Organization, Social Determinants of Health. [Online]. Available: https://www.who.int/health-topics/social-determinants-of-health
2. National Committee for Quality Assurance, HEDIS Measurement Year Technical Specifications, Social Need Screening and Intervention, NCQA, Washington, DC, USA.
3. Centers for Disease Control and Prevention and Agency for Toxic Substances and Disease Registry, Social Vulnerability Index, CDC/ATSDR.
4. United States Census Bureau, American Community Survey, U.S. Department of Commerce.
5. HL7 International, Gravity Project: Social Determinants of Health Clinical Care FHIR Implementation Guide. [Online]. Available: https://hl7.org/fhir/us/sdoh-clinicalcare/
6. Health Level Seven International, HL7 FHIR: Standard for Health Care Data Exchange, HL7 International, Ann Arbor, MI, USA.