

Expanded Datasets in Support of Targeted Clinical Pathways:

Ethics, Algorithmic Bias, and Data Privacy Imperatives

Healthcare leaders are under pressure to find patient risk earlier, direct scarce resources more precisely, and close gaps in outcomes before they widen. **Targeted clinical pathways** can help—but only when they are built on a more complete picture that accounts for the real-world barriers shaping patient access and emerging health risk.

Evolution Toward Targeted Clinical Pathways

Traditional → EHR + claims; standardized care
Risk-stratified → clinical history + utilization; earlier risk detection

Context-enriched → social, geographic, environmental data; barriers visible

Targeted pathways → integrated data + governance; proactive, equitable support

Goal: move from reactive, episode-based care to proactive, context-aware intervention.

Targeted clinical pathways improve early risk detection and resource allocation by combining clinical data with carefully governed socioeconomic, environmental, geographic, and community-level information. EHRs and claims tend to show what happened *after* a patient entered the healthcare system, not necessarily at the point where early intervention can provide opportunity for the best outcome. (A more in-depth discussion about the evolution toward targeted pathways can be accessed via link to our previous article here: [Next Generation Clinical Pathways](#))

Contextual data can help explain **why** risk is emerging and **where** earlier intervention may be possible, making pathways more proactive and equitable when appropriate safeguards are in place. However, healthcare leaders must prevent historical inequities represented in expanded datasets from being encoded into models, and ensure analytics expand support rather than restrict access. Care must also be taken to minimize data use and otherwise protect privacy of patient-specific information and identifiers.

This paper reviews relevant public-domain datasets, sources of algorithmic bias, and governance safeguards for responsible implementation.

Public Domain Datasets

To construct a robust population health program, data science teams can tap into an extensive ecosystem of federally curated, publicly available datasets. These repositories are typically structured by geographic identifiers—such as Census

tracts, counties, or ZIP Code Tabulation Areas (ZCTAs)—allowing them to be safely mapped to a patient's residential history.

In practice, public-domain datasets can be incorporated into targeted clinical pathways by linking geography-based indicators to patient panels at an appropriate level of aggregation, then using those signals to refine risk stratification, outreach rules, care-team prompts, and resource referrals.

Rather than replacing clinical judgment, these data should be utilized to help identify where barriers such as food access, transportation, housing instability, environmental exposure, or community-level disease burden may require pathway adaptations, earlier intervention, or added social-support navigation.

*Public datasets can add social, environmental, and community context to patient risk. **Caution:** Link geographic indicators only at the aggregation level required for the pathway to limit privacy and re-identification risk.*

Representative Public-Domain Datasets		
Dataset	Signal	Pathway Value
CDC/ATSDR Social Vulnerability Index	Social and transportation vulnerability	Targets proactive support
USDA Food Access Research Atlas	Food access	Supports nutrition-sensitive pathways
EPA EJSCREEN	Environmental exposure	Adds respiratory and cardiovascular context
CDC PLACES	Community disease burden	Benchmarks local risk

Data from Public-Domain Datasets can be incorporated alongside claim and EHR information, as a basis for creating “signals” to refine risk stratification, outreach, care-team prompts, and referrals— but not to replace clinical judgment.

Algorithmic Bias: Enterprise Risk and Guardrails

When machine learning algorithms ingest these expanded datasets, they do not operate in a vacuum. They inherit data reflecting decades of systemic inequities in healthcare access. If left unchecked, algorithms will codify and accelerate these historical disparities into resultant automated clinical workflows.

Why Algorithm Bias Is a Risk

For healthcare executives and provider organizations, algorithm bias should be treated as an enterprise risk—not merely a data science defect. When predictive models influence care management enrollment, clinical prioritization, outreach intensity, or resource allocation; biased outputs can directly affect patient access, quality performance, regulatory exposure, and organizational trust.

The central leadership question is not whether a model is statistically accurate in the aggregate, but whether it performs safely, fairly, and transparently for the populations it affects. A model may appear effective because it predicts utilization, spending, readmissions, or missed appointments, while still disadvantaging patients who have historically faced barriers to care.

In this way, bias can hide inside otherwise strong performance metrics unless leaders require **stratified validation and ongoing monitoring** across race, ethnicity, language, income, geography, disability status, age, payer type, and other relevant cohorts.

For example, a model may appear accurate overall while under-identifying rural patients with limited specialty access, patients whose preferred language is not English, or lower-income patients with fewer claims because care was delayed or unaffordable.

The Healthcare Expenditure Fallacy

Historical utilization or spending is a poor stand-alone proxy for clinical need because it also reflects access to care. In the published example provided below, lower spending led the algorithm to underestimate risk for severely ill patients facing access barriers, excluding some from added care-management support.

Example:

A commercial algorithm was used to identify high-risk patients for an intensive care management program that provided extra nursing support, priority scheduling, and dedicated resources to help vulnerable patients remain stable and avoid hospitalization.

The algorithm's developers used a patient's **historical healthcare spending** as a proxy for how sick they were, reasoning that sicker patients naturally cost the system more money. However, this assumption completely ignored deeply entrenched systemic and economic barriers. Due to disparities in insurance coverage, income levels, and local clinic availability, underserved and minority patients historically spent significantly less on healthcare than wealthier patients.

Because the algorithm equated *spending less money* with *being healthier*, it systematically assigned lower risk scores to severely ill minority patients. Healthy, high-spending individuals were fast-tracked into the supportive care management program, while severely ill patients with lower historical medical spending were excluded from the extra resources they desperately required.

Citation: Obermeyer, Z., Powers, B., Vogeli, C., & Mullainathan, S. (2019). [Dissecting racial bias in an algorithm used to manage the health of populations](#). *Science*, 366(6464), 447–453.


Similar bias can arise when missed visits are treated as disengagement, device data overrepresents patients with reliable technology access, or neighborhood indicators label risk without triggering added support.

Executive Guardrails against Algorithmic Bias

For healthcare executives and providers, preventing the adverse consequences of algorithmic bias requires active governance:

Safeguards may include:

- Mandating pre-deployment and continuous quarterly impact assessments to verify that a model's predictive accuracy and resource routing remain equitable across all racial, geographic, and income cohorts.
- Stripping raw historical financial, insurance, or billing data from clinical recommendation engines. Algorithms can instead be trained to predict risk based on objective clinical markers and verified environmental barriers, not commercial expenditures exclusively.
- Validate performance and resource routing before deployment, and quarterly, across relevant cohorts.



Illustrative Use Case: Type 2 Diabetes

Clinicians and data analysts need to be aware of bias that may emerge when a model treats historical access patterns as clinical need. Lower use of continuous glucose monitors or missed visits may reflect affordability, transportation, work schedules, or food insecurity—not lower need or poor adherence.

For this example, a health plan decided to combine rising HbA1c with food-access and transit data to identify patients who may need support; treating missed visits or limited device use as possible affordability, transportation, work-schedule, or connectivity barriers — not as lower need or poor adherence.

The resulting pathway turns combined data content into appropriate support by:

- **Screening for barriers** when visits are missed or control worsens.
- **Offer targeted outreach**, such as transportation help, nutrition resources, remote monitoring, CGM access review, or community health worker outreach.
- **Auditing outcomes** to confirm lower-income, rural, minority, and historically underserved patients are not excluded from resources.

For this example (and similar efforts utilizing expanded datasets), additional attention needs to be paid to privacy risk that can arise when transit, neighborhood, device and metabolic data are linked to identifiable records. Strong controls are needed to prevent exposure, profiling or non-care related uses.



Data Privacy and Regulatory Safeguards:

Non-traditional data—including geographic indicators, neighborhood vulnerability measures, remote-monitoring telemetry, claims data, and device-generated signals—must be secured, transparent, and treated as high-sensitivity clinical assets. When linked to identifiable patient records, these data can improve outreach and pathway design, but they also increase privacy, consent, and re-identification risk. Expanded data programs therefore need formal safeguards, as detailed here:

Treat geographic, neighborhood, claims, remote-monitoring, and device data as high-sensitivity clinical assets. When linked to identifiable records, they require formal controls for privacy, cybersecurity, transparency, consent, equity, and permissible use.

Safeguard	Executive requirement
Regulatory Compliance	Comply with HIPAA Privacy, Security, and Breach Notification requirements; document risk analyses, access controls, audit logs, encryption, incident response, workforce training, and vendor accountability for electronic protected health information and telemetry.
Minimum Necessary / Purpose Limitation	Limit expanded data to the clinical, operational, or population-health purpose it serves; avoid unnecessary granular location, mobility, retail, employment, or household-level data; exclude, generalize, aggregate, tokenize, or remove unneeded elements.
De-identification / Re-identification Controls	Use formal de-identification methods, such as Safe Harbor or Expert Determination; reassess risk as sources are added; apply small-cell suppression, geographic aggregation, limited data sets, data-use agreements, and role-based access.
Vendor / Third-Party Governance	Require third-party risk review and contracts covering ownership, permitted uses, retention, subcontractors, audit rights, breach duties, model-training limits, and bans on secondary

	commercial use such as advertising, underwriting, premium profiling, or unrelated product development.
Ethical AI Boundaries	Align analytics with trustworthy-AI frameworks; use predictive models only for care-management enrollment, outreach, resource navigation, and clinically appropriate intervention planning; prohibit use for adverse underwriting, discriminatory design, coverage denial, benefit reduction, exclusionary networks, or de-prioritization.
Patient Transparency / Consent	Clearly explain what data is used, why it is used, who can access it, and how it may affect care recommendations; make consent, authorization, and opt-out choices understandable and non-punitive.
Human Oversight / Auditability	Ensure models support—not replace—clinical judgment; keep high-impact recommendations explainable, traceable, reviewable, and monitored; audit inputs, access logs, risk scores, cohort impacts, overrides, and resource allocation to prevent automated inequities.

In summary:

The defining question is no longer whether healthcare organizations can use expanded data and predictive analytics to identify risk earlier, but whether they will do so responsibly. **Targeted clinical pathways** can turn fragmented signals into timely, equitable support—but only when leaders make fairness, privacy, transparency, and human oversight nonnegotiable. Organizations that embed these safeguards from the outset will not merely build better models; they will build a more proactive system of care that earns trust, closes avoidable gaps, and reaches patients before risk becomes crisis.