National Mental Health Network

501(c)(3) Nonprofit · PO Box 200, Limeport, PA 18060 · (610) 844-8258

nicholas@dartboard.systems · nmh.network

Technical Briefing

Prediction Seasonality, Model Drift, and Proposed Legislative Guardrails
for AI-Driven Veteran Suicide Risk Prevention Systems

Prepared in Support of H.R. 8486, the Data Driven Suicide Prevention and Outreach Act of 2026

1. Executive Summary

H.R. 8486 directs the Secretary of Veterans Affairs to award grants for the development of predictive models that evaluate risk factors contributing to veteran suicide. The National Mental Health Network applauds this initiative and submits this technical briefing to inform the legislative process on two critical phenomena that any grant-funded predictive system must address: **prediction seasonality** and **model drift**. This document further proposes professional safety standards, ethical guardrails, and technical requirements that should accompany any federally funded AI system operating in a clinical or population-health context involving suicide risk.

These recommendations are offered in the spirit of strengthening an already commendable piece of legislation, and draw on established principles from the machine learning, clinical informatics, and public health literatures. Where this briefing recommends a given safeguard, it does so because the safeguard reflects mature, defensible practice for any system entrusted with life-and-death classification — and because the absence of such safeguards has, in other federal AI deployments, produced silent and costly failures.

2. Prediction Seasonality in Veteran Mental Health Data

2.1 Definition

Prediction seasonality refers to systematic, time-dependent fluctuations in the statistical distribution of input features or model outputs that recur across defined temporal intervals. In practical terms, this means that a given predictive feature (such as income stability, social isolation, or healthcare utilization) may exert a stronger or weaker influence on risk classification depending on when the observation is collected.

A model trained on data from a single temporal cross-section cannot account for these fluctuations. When deployed in a production environment where screening occurs continuously throughout the year, such a model may systematically over- or under-estimate risk during periods where the population's feature distributions diverge from the training distribution.

2.2 Known Seasonal Risk Factors in Veteran Populations

The existing literature and epidemiological evidence point to several seasonal dimensions that any grant-funded predictive model should account for:

- **Holiday-period amplification.** Features related to social connectedness, family support, and life satisfaction are known to exhibit heightened volatility during November through January, when social comparison and loneliness intensify among veterans who lack support systems.
- **Transition-period volatility.** Peer-reviewed research consistently confirms that suicide risk is highest within three months to one year following military discharge (Hoffmire et al., 2022; Ravindran et al., 2020). Models must account for the non-stationary nature of this risk window as new discharge cohorts cycle through the system each year.
- **Economic-cycle sensitivity.** Features related to income, employment, and financial stability are directly affected by macroeconomic conditions that vary both seasonally and cyclically. A model calibrated during a period of low unemployment may behave unpredictably during an economic downturn, and vice versa.

2.3 Implications for H.R. 8486

Grant-funded organizations should be required to demonstrate awareness of seasonality effects and incorporate mechanisms to mitigate seasonal prediction bias. Such mechanisms might include time-aware feature engineering, rolling retraining windows, or ensemble approaches that weight temporal strata. Without them, a model deployed in January may produce materially different accuracy than the same model deployed in July, undermining the consistency and reliability that Congress intends. Notably, systems whose explainability operates at the level of the individual veteran — rather than the population average — are inherently better positioned to surface emergent seasonal signals as they arise, because per-individual attribution makes distributional drift visible before it aggregates into population-level error.

3. Model Drift: Covariate Shift, Concept Drift, and Performance Degradation

3.1 Technical Definition

Model drift encompasses two distinct phenomena. The first, *covariate shift*, occurs when the distribution of input features $P(X)$ changes over time while the conditional relationship $P(Y|X)$ remains stable. The second, *concept drift*, occurs when the underlying relationship between features and the target variable $P(Y|X)$ itself evolves. Both are well-documented hazards in production machine learning systems and pose acute risks in clinical decision-support contexts where human lives depend on prediction accuracy.

3.2 Why Drift Is Especially Dangerous in This Domain

Veteran populations are not static. Over the operational window of a grant-funded model, several shifts are foreseeable: changes in VA policy and benefits administration; evolving post-pandemic healthcare utilization patterns; shifting economic conditions; new military conflicts or drawdowns that alter the composition of the separating service member population; and changes in the availability of community-based mental health resources.

Each of these shifts can degrade model performance. The particular danger of drift in suicide risk prediction is what the literature calls *silent failure*, a condition where the model continues to produce outputs that appear valid but are no longer clinically reliable. There is no error message and no system alert. The model simply becomes progressively less accurate, increasing the likelihood of false negatives (Type II errors) where at-risk veterans are classified as low-risk. In this domain, that is the costliest possible error.

3.3 Risk to the Grant Program

H.R. 8486 establishes a termination date of September 30, 2029, providing an approximately three-year operational window. Over this period, drift is not a theoretical risk but a statistical certainty. Without mandated monitoring protocols, a grantee's model could degrade from high discriminative accuracy to functionally unreliable classification without any visible system failure. The legislation should ensure that the predictive model which earns the grant is the same caliber of model that continues to operate throughout the life of the program.

4. Proposed Legislative Guardrails

The following guardrails are organized into seven categories. Each reflects a practice that any system entrusted with veteran suicide risk classification should be expected to meet — not at the point of initial award, but continuously across the operational life of the grant. We have framed each as a verifiable, auditable requirement so that the Department and the Committee can hold grantees accountable to objective standards rather than to good intentions.

4.1 Mandatory Model Monitoring and Retraining Cadence

We recommend that H.R. 8486 require grantees to:

- Implement automated drift-detection pipelines that monitor distributional shifts in input features using established statistical tests (e.g., Population Stability Index, Kolmogorov-Smirnov tests, Jensen-Shannon divergence) at no less than quarterly intervals.
- Establish minimum retraining thresholds triggered when drift metrics exceed predefined bounds, with documentation submitted to the Secretary of Veterans Affairs upon each retraining event.
- Maintain a versioned model registry that logs all hyperparameter configurations, training data provenance, and classification performance metrics for each model iteration, enabling full auditability.

4.2 Individual-Level Explainability and Transparency

Section 2(d)(3)(C) of H.R. 8486 correctly prioritizes tools that are "explicable, interoperable, and clinically actionable." We recommend strengthening this provision substantially. The Committee should recognize that **population-level explainability and individual-level explainability are not interchangeable**, and that only the latter is sufficient for clinical action on behalf of a specific veteran. We therefore recommend that the bill require:

- **Per-individual feature attribution.** Every risk prediction must be accompanied by feature-level explanations that identify the specific factors driving that individual veteran's classification — not population-level averages or global feature-importance rankings. Established techniques such as local Shapley Additive exPlanations (SHAP) make this computationally tractable at the point of care, and grantees should be expected to demonstrate per-prediction attribution rather than aggregate interpretability. A system that can only explain its behavior in the aggregate cannot tell a counselor why this veteran, in front of them right now, was flagged.
- **Population-level interpretability as a complement, not a substitute.** Grantees should additionally produce periodic global feature-importance analyses to identify population-level patterns, seasonality signals, and potential algorithmic bias across demographic subgroups,

consistent with Executive Order 13960 on Trustworthy Artificial Intelligence in the Federal Government.

- **Clinician-readable translation.** Raw statistical coefficients and feature-attribution scores are not clinically actionable. Grantees should be required to demonstrate the capacity to render technical findings into plain-language, contextualized guidance — translating, for instance, a set of weighted model inputs into a statement a social worker can act on (“this veteran’s elevated risk is driven primarily by financial strain and housing instability”). The ability to produce this translation layer should be a condition of award, not an aspiration.

4.3 Continuous, Automated Quality Assurance of Every Output

Periodic human audit is necessary but not sufficient. A system that processes screenings continuously across the country cannot rely on retrospective spot-checks to catch erroneous or low-quality outputs before they reach a veteran. We therefore recommend that the bill require grantees to implement **automated, output-level quality assurance that evaluates every individual recommendation** — not a sample — against defined dimensions of factual accuracy, contextual appropriateness, and alignment between the system’s reasoning and the veteran’s profile. Such evaluation can be operationalized through independent automated evaluation architectures that score each output on a structured rubric, surface low-scoring outputs for human review, and log every score for audit. This standard ensures that quality is verified at the moment of generation rather than discovered months later, and it provides the Department with a continuous, machine-readable signal of system health across the entire grant period.

4.4 Validated, Currently-Operating Referral Resources

A predictive model that correctly identifies an at-risk veteran but refers them to an organization that has dissolved, lost its tax-exempt status, or lacks the capacity to serve them has failed the veteran. We recommend that the bill require any system that produces resource referrals to:

- Validate every referred organization against authoritative federal registries (for example, Internal Revenue Service Exempt Organizations data), and exclude any organization appearing on the IRS Automatic Revocation of Exemption list, so that veterans are never directed to defunct or non-compliant entities.
- Employ a standardized organizational taxonomy recognized by federal data systems (such as the National Taxonomy of Exempt Entities) to ensure referrals are matched to organizational purpose and capacity in a consistent, auditable manner rather than by ad hoc keyword association.
- Demonstrate nationwide geographic coverage capable of identifying validated resources in any United States ZIP code — not limited to VA catchment areas or a grantee’s home region — so that a screening conducted anywhere can connect a veteran to help anywhere.

4.5 Whole-Person, Multi-Dimensional Assessment

Suicide risk does not arise from clinical symptomatology alone. We recommend that the bill encourage systems to assess multiple dimensions of a veteran’s profile — clinical indicators, personal interests and sources of meaning, and the specific context of military service — and to ground their recommendations in peer-reviewed evidence. Grantees should be expected to maintain a structured, current knowledge base of relevant clinical and epidemiological literature, and to demonstrate that their referrals reflect that evidence rather than unsupported generalization. A whole-person assessment that connects a veteran to a community organization aligned with who they are, not merely with their diagnosis, is more likely to produce durable engagement and reduced isolation.

4.6 Ethical Guardrails and Bias Mitigation

Given the stakes involved in suicide risk prediction, we recommend the following ethical provisions:

- **Type II error prioritization.** Models should be optimized to minimize false negatives — that is, failing to identify a veteran who is at risk. It is ethically preferable to refer a veteran for additional screening who does not ultimately need it than to miss a veteran who does.
- **Demographic fairness auditing.** Grantees should conduct subgroup performance analysis across race, gender, age, branch of service, and discharge era to ensure equitable predictive accuracy. Audit results should be submitted to the VA's Center for Innovation for Care and Payment on a semi-annual basis.
- **Privacy-by-design architecture.** All systems must comply with HIPAA, 38 U.S.C. § 7332, and VA Directive 6502. We recommend extending this to require explicit privacy-by-design principles that minimize the collection of personally identifiable information to the greatest extent technically feasible. The strongest privacy posture — and the one most likely to overcome the documented reluctance of veterans to seek help — is a system that collects no personally identifiable information whatsoever, requiring only the minimum non-identifying input (such as a ZIP code) necessary to locate local resources. Compliance frameworks that merely protect collected PII are a floor, not a ceiling; systems that never collect PII in the first place eliminate entire categories of breach, stigma, and career-consequence risk that deter veterans from engaging.
- **Human-in-the-loop requirement.** No predictive model should autonomously determine a veteran's treatment pathway. All outputs should be reviewed by a qualified mental health professional before clinical action is taken. The model is a decision-support tool, not a decision-making authority.

4.7 Individual-Prediction Auditability

Finally, we recommend that the bill require grantees to maintain complete audit trails at the level of the individual prediction — not merely at the level of the model. For each screening, the system should be able to reproduce the inputs considered, the intermediate computations and feature attributions, the referral logic applied, the resources considered and excluded, and any automated quality scores assigned. This level of granularity is what allows Congress, the Department, and the public to verify after the fact that a taxpayer-funded system behaved correctly in any given case. It is also the foundation of genuine accountability: a system that cannot reconstruct why it produced a particular recommendation cannot be meaningfully audited, however impressive its aggregate metrics.

5. Alignment with the Bill's Existing Framework

The guardrails proposed above are designed to strengthen, not conflict with, the bill's existing eligibility and selection criteria. Section 2(b) already requires grantees to demonstrate capability in AI and machine learning development and deployment, health data analysis with PII and PHI protections, predictive model development in clinical settings, advanced statistical methods applied to complex health datasets, and compliance with Department data security standards. Our recommendations extend these requirements to cover the operational lifecycle of the model, ensuring that the system which earns the grant continues to perform at the same standard throughout the program's duration.

Additionally, these guardrails complement the bill's priority criteria under Section 2(d)(3), which favor organizations that can deliver explicable, interoperable, and clinically actionable tools, and that employ data scientists, clinicians, and suicide prevention specialists. The recommendations herein provide the evaluative framework to hold grantees accountable to those criteria over time — and, importantly, they translate the bill's existing language into objective, demonstrable capabilities that can be verified at the point of award rather than asserted.

6. Conclusion

Predictive models are not static instruments. They are living systems whose reliability depends on continuous monitoring, retraining, and evaluation. The phenomena of prediction seasonality and model drift are not edge cases; they are inherent properties of any machine learning system deployed in a dynamic environment.

H.R. 8486 represents a significant and commendable step toward data-driven veteran suicide prevention. By incorporating the technical guardrails outlined in this briefing, Congress can ensure that the predictive models funded under this program remain accurate, equitable, explainable, and trustworthy for the duration of the program and for the veterans these systems are designed to serve. Each safeguard recommended here is achievable with existing, mature methods — and each, once required, ensures that federal funds support systems built to the standard these veterans deserve.

The National Mental Health Network stands ready to support the Committee on Veterans' Affairs in any capacity as this legislation advances.

Respectfully submitted,

Nicholas C. Birosik, D.Eng.
Executive Director, National Mental Health Network
Adjunct Faculty, The George Washington University