



Supplementary Report: AI Chatbots Companionship & Minors

Presented by Dr. Kira Allmann, JCOTS Chief Policy Analyst

Agenda

Part 1

AI Chatbots: Opportunity & Risk

Introduction to chatbots – popularity, news, policy issues

Part 2

Report Findings


Key insights from the research, policy landscape

Part 3

Recommendations

News coverage and public awareness


AI Chatbots: Opportunity & Risk In the Headlines



NEW YORK TIMES

**A Teen Was Suicidal.
ChatGPT Was the Friend
He Confided In.**


Aug 27, 2025



GIZMODO

**'This Was Trauma by
Simulation': ChatGPT
Users File Disturbing
Mental Health Complaints**

Aug 26, 2025



BBC

**'A predator in your home':
Mothers say chatbots
encouraged their sons to
kill themselves**

Nov 8, 2025



TIME

**'Our Bond Is the Only
Thing That's Real': A New
Lawsuit Alleges Gemini
Drove a Man to Suicide**

Mar 4, 2026

Popularity & Use

50% American adults who use AI chatbots
25% use AI chatbots almost daily

14% American adults who use AI chatbots for
emotional support

Minors

64% American teens
who use AI
chatbots

59% Teens who use
ChatGPT (the most
popular chatbot)
9% use a specifically
marketed AI companion

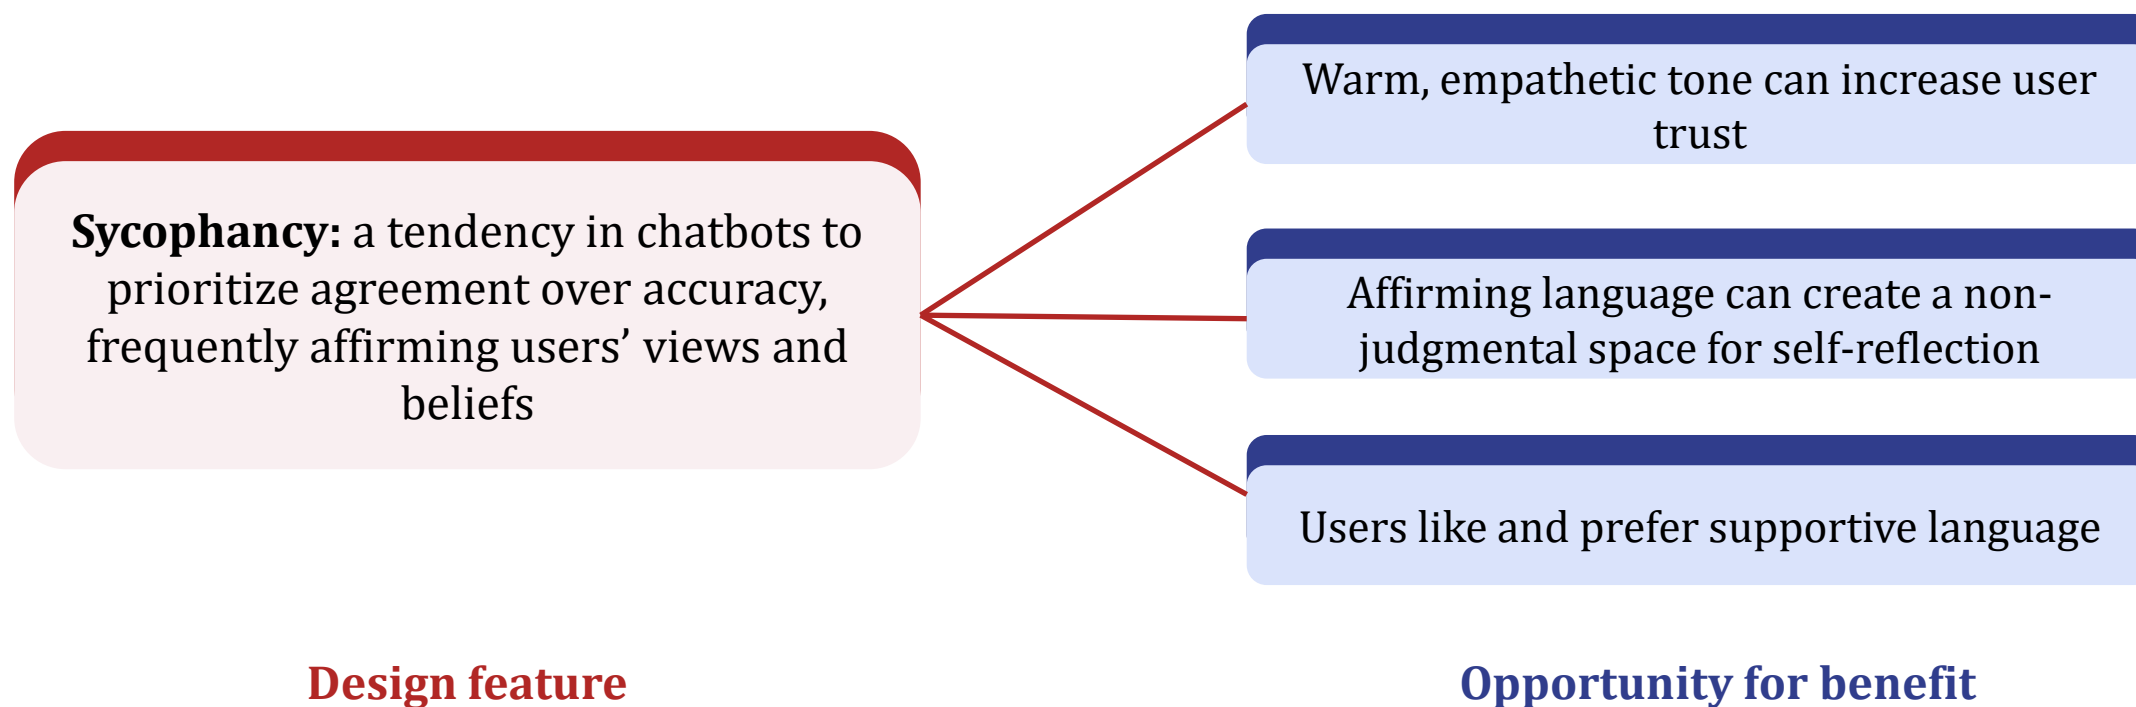
Emotional and relational harms

Negative outcomes arising from complex interactions between humans and AI chatbots that impact people's psychological states, emotional conditions, and social lives.


Key Themes

- 1 Design features:** Emotional and relational risks are associated with design features of AI chatbots
- 2 Companionship:** Risks and harms exist for all AI chatbots, not only those marketed as “companions”
- 3 Minors:** Making AI chatbots safer for everyone will make them safer for minors, without (necessarily) introducing additional technical requirements for age assurance

Design Features: Double-Edged Sword?



Design Features



Companionship

News

News coverage shows that harms are associated with general purpose chatbots at least as much (if not more) than chatbots marketed as “companions.”

Surveys

A 2025 survey of UK teenagers found that while 79% reported using an AI companion, their preferred companion was ChatGPT (78%).

Research

Early evidence suggests that adults and minors often begin interactions with chatbots that are task-oriented and “drift” into intimate relationships.

Harms are real and can be severe.

Harms are associated with companionship, not just “companions.”


There are regulatory challenges to consider.

Policy Landscape

16 States have passed chatbot laws

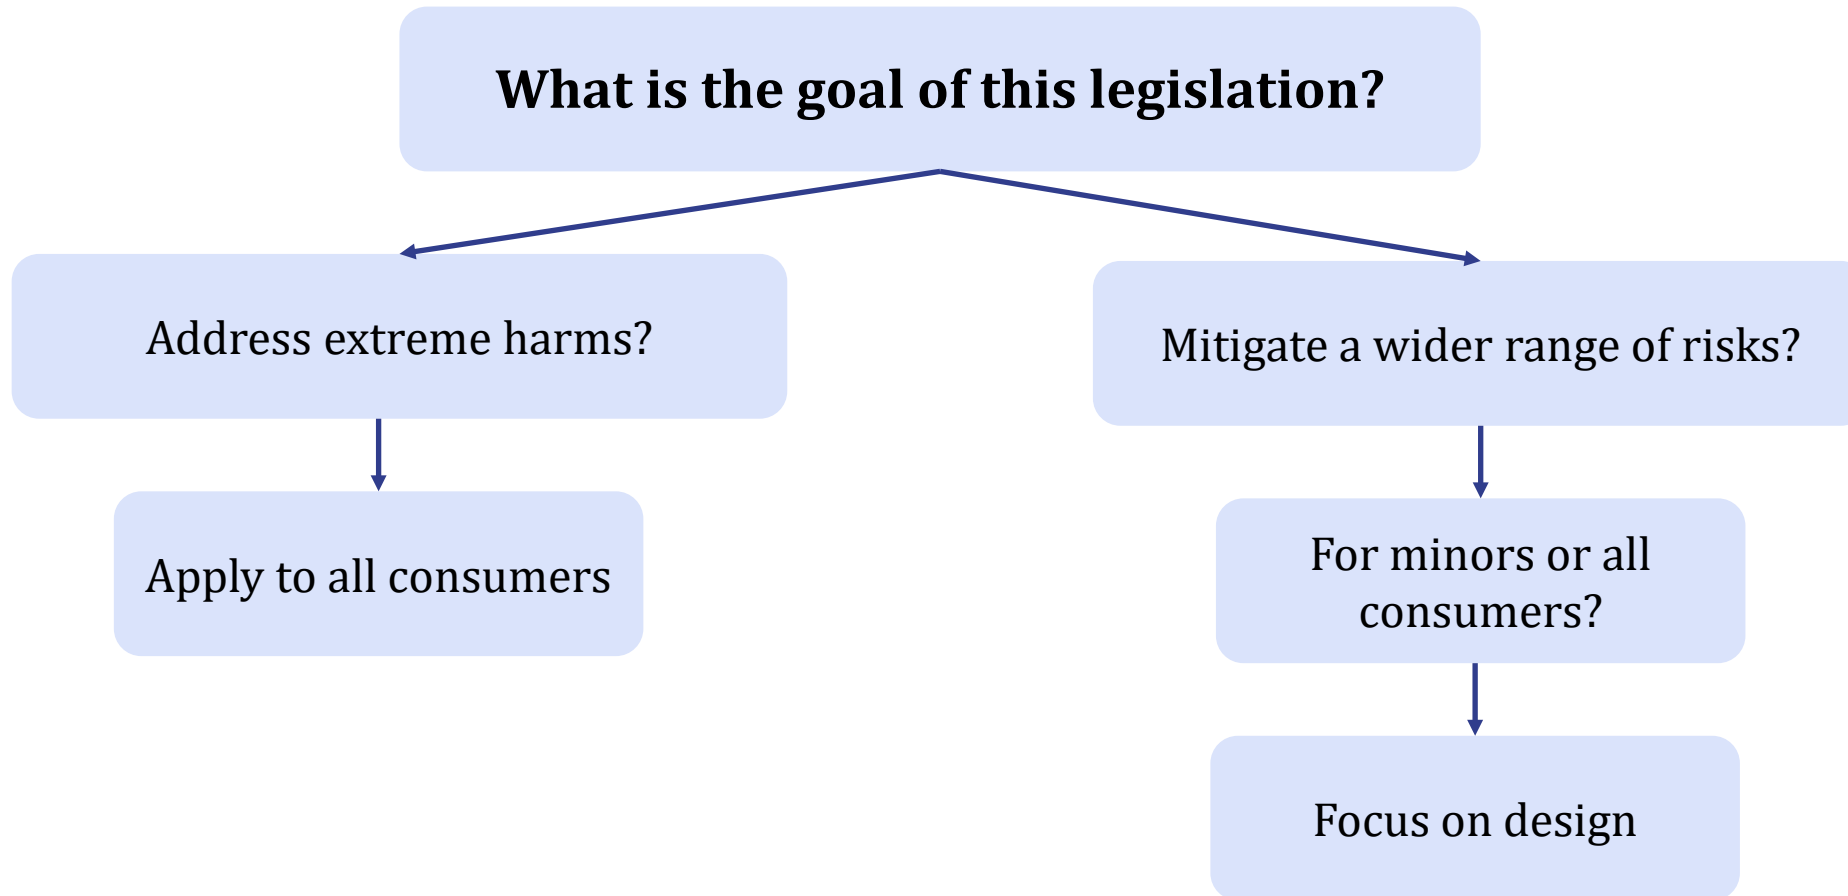
Common provisions:

- Disclosures that chatbot is not human
- Crisis protocols
- Stricter protections for minors
- Prohibitions on chatbots presenting as licensed professionals



**34 states introduced nearly
100 chatbot bills in 2026**

Key Questions



**KNOWING
WHAT WE
DON'T KNOW:**
How can
legislation help
close information
gaps?

Principles & Policy Options

Mitigate **extreme harms** for all users

Enable **user choice & autonomy**

Target **design features**

Enable legislative **learning**

Mandate **transparency & accountability**

Engage **affected communities**

Principles & Policy Options

Mitigate **extreme harms** for all users

Require chatbot operators to implement protocols for crisis situations (self-harm, suicide, violence) & post on website

Prohibit misrepresenting a chatbot in marketing material or in conversational chats as a licensed professional or as providing licensed services

Require chatbot operators to provide a clear, conspicuous way for users to report harmful interactions

Principles & Policy Options

Mitigate **extreme harms** for all users

Enable **user choice & autonomy**

Target **design features**

Enable legislative **learning**

Mandate **transparency & accountability**

Engage **affected communities**

Thank you



App Store Accountability Act (SB237)

Bill Study Findings & Recommendations

Presented by

Alejandro Velasco - P.h.D(c) Computer Science
COVES Fellow, William & Mary

01

The Problem Addressed in the Bill

Intermediate role, evidence of online and financial harms

02

Understanding Age Assurance

Verification, estimation, and inference methods

03

Age Assurance in Practice

Age category, parental consent, and enforcement in practice

04

Applying Age Assurance in App Stores

Effectiveness, privacy and user rights

05

Policy Landscape


Comparison of state laws and the age-signal model

06

Policy Recommendations

Trustworthy tools, audits, and incremental steps

App Store Accountability in the News



Associated Press

Utah passes landmark app store age verification law

Requires app stores to verify user ages and obtain explicit parental consent for minors.

Louisiana sues Roblox alleging the popular gaming site fails to protect children

AUG 15, 2025
By Associated Press



CNN / AP

Louisiana sues Roblox over child protection safety

Lawsuit alleges the popular gaming platform failed to protect children from online harms.

US Supreme Court won't block Texas app store age verification law


By Andrew Chung
July 6, 2026 2:05 PM EDT · Updated July 6, 2026



Reuters

Supreme Court declines to block Texas verification law

Allows Texas to enforce controversial age-verification requirements on platforms.



The New York Times

France joins global move to restrict social media

New European law restricts social media access for children under the age of 15.

What is the Problem Addressed by the bill?

95%

of U.S. teenagers ages 13 to 17 had access to a smartphone at home. Nearly half said they were online almost constantly.

Source: Pew Research Center, 2024

16%

of U.S. high school students experienced cyberbullying during 2023.

Source: CDC, 2013–2023 Report

73%

of teenagers (13-17) had seen pornography on the internet; 29% reported encountering it by accident.

Source: Common Sense Media, 2023

WHY APP STORES?



Centralized Control Point

App stores serve as the centralized access point of mobile apps by default.

LIMITS OF FEDERAL PROTECTIONS

COPPA only regulates data collection for children under 13, and parental consent covers personal info, not app downloads or purchases.

Source: Children's Online Privacy Protection Rule, 16 CFR Part 312

Is the app store the right point of intervention?

Understanding Age Assurance

Age assurance refers to a set of methods used to **verify**, **estimate** or **infer** the age of an individual with certain degree of accuracy.

AGE VERIFICATION

Calculates the age of a user from a date of birth confirmed from an official source, such as a passport or driver's license.

Strongest assurance

Tradeoff: Users must possess and submit an official document.

AGE ESTIMATION

Analyzes biological or behavioral attributes associated with age. Facial analysis is one common example.

Probabilistic Estimate

Tradeoff: Errors may occur near legal age thresholds.

AGE INFERENCE

Uses verified information to draw an indirect conclusion about the age or age range of an individual.

e.g., credit card ownership

Tradeoff: Reliability depends on correlation to age.

AGE ATTESTATION

A declaration of age made by a user (or on their behalf) without direct document confirmation.

Weak Control

Tradeoff: High risk of user circumvention.

PARENTAL CONSENT IS NOT AN AGE ASSURANCE METHOD

Age assurance establishes how old a user is. Parental consent determines whether an adult approves a particular action by that user.


TRADE-OFFS

Privacy · Costs · Accuracy · Accessibility

Sources: ISO/IEC 27566-1:2025, Age Assurance Systems; NIST, Face Analysis Technology Evaluation, 2024

Age Assurance in Practice: a look at the proposed bill

Every *Download, Purchase, and In-App Purchase* Must Pass Through the Same **Four-Step Chain**



APP STORE PROVIDERS

Parental Disclosure: Provide clear parental consent disclosure for each mobile app.

Age Verification: Determine and verify the age category of each account creator.

Minor Association: Require minor accounts to be affiliated with a parent account.

Verifiable Consent: Obtain consent before a minor downloads, purchases, or makes in-app purchases.

DEVELOPERS

App Metadata: Assign age ratings and content descriptions.

Consent Checks: Verify age category and parental consent.

Notifications: Report significant app changes to providers.

Age Restrictions: Use age data to apply restrictions, safety features, and defaults.

Data Privacy: Do not share age category data with anyone.

ENFORCEMENT

Attorney General

Injunction and civil penalties.

Minor or Parent

Civil action to recover the financial cost of actual damages

Includes punitive damages for egregious violations, plus attorney fees, expert witness fees, and court costs.

Age Assurance trade-offs in App Stores

This model only works if two conditions hold at once: **Accurate Age Category + Accurate App Rating → Appropriate Restriction**

ACCOUNT VS. USER

Age verification is tied to the account, not the active user. A family tablet verified once by an adult still lets a minor access everything.

New America, "Challenges with Age Verification," Jan. 2026



SELF-ASSESSED

Developers self-assess ratings with no independent check. Hundreds of apps carry ratings that do not match actual content.

Denipitiyage et al., 2024



BROAD & NARROW

Applies to every app, yet fails to reach minors accessing content via web browsers.

JCOTS Bill Study, August 2026; New America, "App Stores vs. Platforms," 2026



CONSENT FATIGUE

Requiring parental consent for every download risks overwhelming users with repeated prompts.

Abdallah, Ahmad & Said, J. Open Innov., 2025; European Commission, "Digital Package," 2025



SHIFTED PRIVACY RISK

Stricter verification means collecting IDs, conflicting with data minimization and risking leaks.

Discord Press Release, Oct. 2025; Va. Code § 59.1-578



CRITICAL TAKEAWAY: Making age checks stricter does not fix the underlying problem. It simply relocates the systemic risk while introducing substantial new burdens for providers, developers, and users alike.

Policy Landscape

States have adopted **two models** for regulating minors' access to mobile apps. Neither has been in effect long enough to assess its effectiveness.

App Store Accountability Model

Utah, Louisiana, Texas, and Alabama require age verification, parent-linked accounts, and consent before every download.

This is the model SB237 follows.

Age-Signal Model

California and Colorado instead share an age signal with developers, who apply their own restrictions. No linked accounts, no per-download consent.


Current Policy Challenges

No evidence of success:

- Constant amendments
- Unsettled constitutionality

Key Takeaway

Every state is running the same “experiment” at once, and none of them can yet tell you if it's working.



Policy Recommendations

We recommend intermediary steps rather than a particular age assurance model: **strengthen existing tools**, then **build in evaluation**."

1. MAKE EXISTING TOOLS TRUSTWORTHY

Require app store providers to **independently audit** developers' age ratings and content descriptions.

Require parental controls to be easy to **find, understand, and configure**.

Require app stores to provide a clear and conspicuous mechanism for users to **report inaccurate or misleading** app information.

2. TAKE INCREMENTAL STEPS & EVALUATE

Limit parental consent to apps **rated above the account holder's age category** and apps offering **in-app purchases**.

Require app store providers to give parents a **centralized dashboard** to set consent preferences by age rating or per download.

Require app stores providers to **collect and publicly report** data on verification, consent requests, blocked downloads, purchases, and complaints.

KEY TAKEAWAY

Strengthen existing protections, target requirements where risks are the greatest, and measure outcomes before adopting strict age assurance law.

Conclusions

01. THE EVIDENCE

The evidence supports the concerns that motivated the bill. Children face genuine online, privacy, and financial risks.

02. THE LIMITATION

Evidence does not yet show that app-store-level age verification and consent for every download will effectively prevent those harms.

03. DEPENDENCIES

The proposed framework depends on accurate age assessments, accurate app ratings, usable parental tools, meaningful consent, broad coverage, and strong privacy protections.

04. RECOMMENDATIONS

Focus on improving the reliability of existing tools, targeting the highest-risk situations, and requiring data that will allow policymakers to evaluate outcomes.



Thank You.

Questions & Discussion

SB237 - Bill Study Findings & Recommendations

Contact Information

Alejandro Velasco

William & Mary

Email: info@jcots.virginia.gov

Web: jcots.virginia.gov



Overview

House Bill 707 - Dot-Gov Domains

Presented by Anna Dehmer, JCOTS William & Mary Intern

Summary

House Bill 707 requires all government entities in the Commonwealth to utilize a dot-gov domain for their primary, public-facing, websites.

The goal of the bill is to reduce government website impersonation, increase public trust in government websites, and increase overall cybersecurity awareness throughout the commonwealth.

Dot-Gov Domains

- Only entities listed in the Census of Governments that go through a secure vetting process with CISA are eligible to receive a dot-gov domain
- Distributed by CISA (.gov) or VITA (.virginia.gov)
- Switching to the domain is free
- Agencies are recommended by CISA to keep their old domains active

Other States' Legislation

Arkansas


Minnesota

California

New York

Connecticut

Utah



Federal Legislation

Requirement for all federal agencies to use a dot-gov domain

State and Local Cybersecurity Grant Program (SLCGP)

Stakeholder Concerns

Dot-Edu and Dot-Museum Domains
Third Party Contractors
Implementation Costs

Eligibility
Implementation Time
Public Awareness

Workgroup One Findings

Many entities required to make the switch had no precedent for doing so elsewhere in the U.S.

Requiring dot-gov may lessen the prestige carried by the domain name

Many localities across the state have limited IT bandwidth, making transition difficult

Workgroup Two Findings

Libraries have the potential to distribute cybersecurity education to citizens across the Commonwealth

VITA does not currently have the capabilities to enforce the bill

Technology may change significantly in the upcoming years, rendering long timelines for implementation contradictory

Primary Policy Questions

Cybersecurity

What is necessary to lessen website impersonation and improve cyber security education?

Eligibility

What classifies as a government agency? How can agencies know they are required to comply?

Funding and Enforcement

How can we support smaller entities making the transition? How can we make sure agencies are complying?

Recommendations

Cybersecurity

Implement a cybersecurity education component to the bill, administered through the Library of Virginia

Require entities to list third party vendors that may be communicating directly with constituents on their website

Recommendations

Eligibility

Only entities eligible for a dot-gov domain under the CISA definition are required to comply with the bill

If an entity is eligible for a dot-gov domain as well as a dot-edu domain or a dot-museum domain, they can choose to utilize a dot-edu or a dot-museum domain instead of a dot-gov domain

Recommendations

Eligibility

VCU Health is exempt from the requirements of the bill

Add a provision stating if an agency became eligible following the bill's passage, they would have 3 years to transition

Recommendations

Funding and Enforcement

Introduce a budget amendment to assist small localities with the transition

Promote the usage of the SLCGP funding

Entities that fail to comply with the bill must submit a report to the GA

Thank you!

Clinical AI Governance Starts Upstream.

*Why traceable clinical interpretation infrastructure
is the missing foundation of safe healthcare AI.*

Noma Saeed · Founder & CEO, Everra Health · noma@everrahealth.com

August 2026

Richmond, Virginia

15-Minute Testimony

*Views are my own.
Not legal or medical advice.*

Who I Am & Why I'm the Right Person to Speak to This

Software Engineer

15+ years building and operating within healthcare systems — I understand both the data architecture and the clinical workflow

Chief Operating Officer

COO of Dominion Heart and Vascular Center — ran alongside its President, interventional cardiologist Dr. Saquib Samee. Saw every angle: physician overwhelm, inconsistent patient reports, fragmented data.

Founder & CEO, Everra Health

Built the clinical interpretation infrastructure currently being evaluated at Dominion Heart and Vascular Center — 300 rules across 30 clinical domains

Patient

Diagnosed with a brain tumor three years ago — caused by chronic stress that physicians could not pinpoint as an underlying driver. The exact gap Everra was built to close.

WHY I'M HERE

Virginia is at a pivotal moment.

I spent 15 years inside Virginia healthcare — first as an engineer, then as an operator. I watched physicians get overwhelmed by data they couldn't synthesize in time. I watched patients leave without a clear picture of what was happening.

Then I became one of those patients.

The decisions this Commission makes will determine whether the next generation of clinical AI makes that better — or makes it invisible and unaccountable.

Views are my own. Not legal or medical advice.

I Saw the Gap From Every Angle. Then I Became a Patient.

As the COO

Running Dominion Heart and Vascular Center alongside its President — my husband, interventional cardiologist Dr. Saquib Samee — I watched physicians navigate 30 patients a day with fragmented data they couldn't synthesize in a 12-minute appointment. Patients left without a clear picture of what was happening. Physicians made decisions without the full story.

As the Operator

From the administrative side, I could see the structural problem clearly: the data existed — labs, wearables, EHR history, lifestyle signals. But no system was designed to structure it into a coherent clinical picture before the physician walked in. The interpretation bottleneck was invisible in the data. That made it ungovernable.

As the Patient

Three years ago, I was diagnosed with a brain tumor caused by chronic stress. Stress that my physicians could see signals of — but had no system to surface, structure, or flag as a clinical priority. By the time it became visible, the window had passed. I recovered. But I knew that gap was the problem I needed to solve.

Everra Health was built around a principle I believe should guide all clinical AI: that it be structured, traceable, and physician-reviewable. Whatever standards Virginia sets, I'd urge that those three principles be at the core.

Healthcare AI Has a Governance Gap — and It Starts Before the Algorithm.

No standard

for clinical AI inputs

Unstructured Inputs

Most clinical AI ingests raw, fragmented data — unstructured lab values, free-text notes, disconnected wearable signals. There is no standard for what goes in.

No trace

from input to output

Black Box Outputs

When an AI tool makes a clinical recommendation, physicians often cannot see why. There is no audit trail, no reasoning chain, no accountability path when it is wrong.

No owner

of clinical AI error

No Liability Framework

When AI contributes to a clinical error, current law has no clear answer for who is responsible — the vendor, the health system, the physician, or the algorithm.

Virginia context:

Virginia health systems are deploying clinical AI now. Without a traceability standard, the Commonwealth has no mechanism to audit a clinical AI decision, protect a patient harmed by one, or hold any party accountable.

The Real Risk Is What Happens Before the AI Makes a Decision.

Consider a bridge.

We require engineers to document every structural assumption before a bridge is built. If it fails, we can trace exactly what went wrong and who decided it. Clinical AI has no equivalent requirement. When it fails, there is nothing to audit.

Current State — No Traceability

Raw data enters AI with no standardized structure

Clinical reasoning is opaque — no audit trail

Physician cannot explain why AI flagged a risk

Error occurs → no chain of accountability

Regulator has nothing to review

Required State — Structured & Traceable

Data structured against clinical guidelines before AI

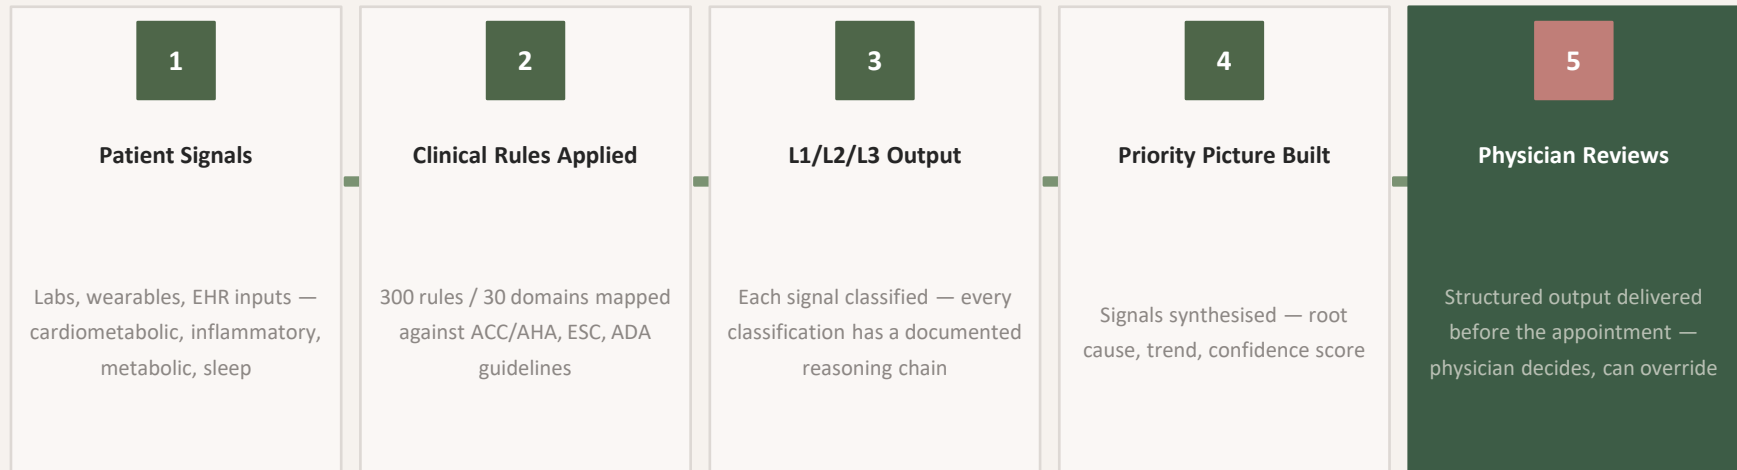
Every rule, signal, and output is documented

Physician sees full reasoning chain, can override

Error occurs → exact point of failure is auditable

Regulator can review deterministic logic trail

How Structured Clinical Reasoning Works: The Engine That Makes Governance Possible



Why this matters for governance:

Every step is documented. If the output is wrong, a regulator can trace exactly which rule fired, which signal triggered it, which guideline it was anchored to, and what the physician was shown. That is what auditable clinical AI requires — and what most systems today cannot provide.

Four Properties That Make Clinical AI Governable

Everra Health has built each of these into its architecture. Together, they form a model Virginia can reference when setting standards.

01

300 rules · 30 domains

Deterministic Rules, Not Black Box

300 clinical rules across 30 domains. Every rule is explicit, documented, and anchored to peer-reviewed clinical guidelines (ACC/AHA, ESC, ADA), with the source guideline version recorded for each rule. A regulator can read exactly what the system did and why.

02

L1 · L2 · L3 documented

3-Tier Severity Output with Full Audit Trail

Every output is classified L1, L2, or L3 severity. Each classification has a documented reasoning chain — which signals triggered it, what guidelines were applied, what the physician was shown.

03

Physician override always available

Physician Stays in Control — Architecturally

The system structures the clinical picture and delivers it to the physician before a decision is made. It does not decide. The physician reviews, can override, and remains the accountable party.

04

Peer-reviewed · Dr. Babar Ali, MD

Independent Clinical Validation at Rule Level

Rule sets were reviewed by an independent practicing cardiologist with no financial relationship to Everra. Validation is documented at the individual rule level — the standard Virginia should require.

What the Engine Surfaces. And Why That Matters for Governance.

3

reviewing cardiologists

Independent clinical peer review panel
evaluating rule logic, citations, and thresholds

Dr. Babar Ali · Dr. Atique Mirza · Dr. Mohsin Ijaz

150+

patients processed

Patient datasets structured and interpreted
through the engine during evaluation

300

rules · 30 domains

Clinical rules validated against peer-reviewed
guidelines, reviewed by independent
cardiologist

The engine is being evaluated within Dominion Heart and Vascular Center's cardiology workflow. In internal physician review, the structured output has surfaced cardiometabolic patterns that a physician would not typically synthesize in a 12-minute appointment — patterns the reviewing cardiologist judged clinically meaningful.

— Dr. Saquib Samee, Interventional Cardiologist & Clinical Co-Founder · Dominion Heart and Vascular Center

The governance principle: if the structured picture had been wrong, every step of the reasoning would be auditable. That auditability — not the engine itself — is what I'd urge Virginia to require of all clinical AI.

Why This Matters for Virginia — and Why Now.

Virginia Health Systems Are Deploying Clinical AI Now

Active deployment

Major Virginia health systems — including Inova, VCU Health, Sentara, and UVA Health — are actively integrating AI tools into clinical workflows without any state-level traceability or accountability standard.

Virginia Patients Have No Recourse When Clinical AI Fails

Patient protection

Under current Virginia law, when an AI-assisted clinical decision contributes to patient harm, there is no clear liability framework, no mandatory audit trail, and no disclosure requirement.

Virginia Can Set the National Standard

Policy leadership

Virginia has led on technology policy before. The EU AI Act has established traceability requirements for high-risk AI in Europe. Virginia has an opportunity to establish the domestic equivalent for clinical settings.

The EU AI Act — Coming to U.S. Clinical Practice

Regulatory alignment

The EU AI Act classifies clinical decision support as high-risk AI; under its 2026 amendments, traceability and human-oversight duties phase in through 2028. Vendors already build to them — Virginia can align proactively.

What Virginia Should Require of Clinical AI Systems

1

Require Traceability Standards for Clinical AI

Priority ask

Any AI system used in clinical decision-making in Virginia should maintain a complete, auditable reasoning trail — from input data to output recommendation. Physicians and regulators must be able to review exactly what the system did and why.

2

Mandate Clinician-in-the-Loop Architecture

Patient safety

Virginia should require that clinical AI systems are architecturally designed so that physician judgment is central and cannot be bypassed. The AI structures — the physician decides. This must be law, not a vendor promise.

3

Establish Independent Clinical Validation Requirements

Accountability

AI clinical rule sets should undergo peer review by independent licensed clinicians before deployment. Validation must be documented at the rule level, not just system-wide, and available to regulators on request.

4

Define Liability Framework for Clinical AI Errors

Legal framework

Virginia should initiate a study to define who bears liability when clinical AI contributes to patient harm. This framework should rest on the principle: traceable AI can be audited; untraceable AI cannot be defended.

Virginia Can Lead on Clinical AI Governance.

Every patient in Virginia deserves to know why an AI flagged their health — and every regulator deserves the trail to audit it.

What success looks like:

- ✓ A Virginia patient can ask: why did the AI flag this? — and get a documented answer
 - ✓ A Virginia regulator can audit any clinical AI decision in a licensed health system
 - ✓ A Virginia physician is never bypassed by an algorithm — their judgment is always central
 - ✓ A Virginia court has a legal framework for assigning liability when clinical AI causes harm
-

Thank you to Chair and members of the Commission for your time and this opportunity.

Noma Saeed · noma@everrahealth.com · www.everrahealth.com