Bipartisan Policy Center

FDA Oversight: Understanding the Regulation of Health AI Tools

Authors: Maya Sandalow, Katie Adams, Gabriel Loud

AI-enabled medical devices are a routine part of clinical practice and many patients' lives. In hospitals, the devices help detect tumors on imaging scans and identify abnormal heart rhythms. At home, wearable electrocardiogram (ECG) monitors and digital therapeutic apps help patients manage atrial fibrillation, diabetes, and other conditions.

The Food and Drug Administration (FDA) is the primary federal agency responsible for regulating medical devices. This issue brief explains how the FDA regulates medical devices enabled by artificial intelligence (AI) across the product lifecycle. It recounts how FDA frameworks were built for static devices with specific indications, describes the steps the agency is taking to modernize oversight for advanced AI technologies, and highlights key challenges ahead.

This is the second policy brief in the Bipartisan Policy Center's series on the regulatory and coverage landscape for health AI. Part one covered oversight of health AI outside of the FDA's jurisdiction, and part three will focus on how the Centers for Medicare & Medicaid Services (CMS) pays for health AI.

To help policymakers keep pace with rapid technological change, BPC launched its AI 101 initiative in 2024. The program provides resources and briefings to build AI literacy on Capitol Hill. Building on this foundation, BPC's Health Program details how AI is transforming health care and the policy landscape. Recent work analyzes AI applications in administrative functions and medical devices; how federal health agencies use AI; the health policy implications of the AI Action Plan; and digital mental health solutions, including generative AI chatbots.

NOVEMBER 2025

TYPES OF HEALTH AI WITHIN THE FDA'S JURISDICTION

Under Section 201(h) of the Federal Food, Drug, and Cosmetic Act [21 U.S.C. 321(h)], AI is considered a medical device—and subject to FDA regulation—if it is intended for use in the diagnosis, cure, mitigation, treatment, or prevention of disease. The FDA does not typically regulate AI used for administrative tasks (e.g., scheduling or billing) or general wellness information unless its "intended use" brings it within scope. For example, an AI tool that merely transcribes clinical notes would likely fall outside FDA oversight, but if that same software generates diagnostic or treatment suggestions, those functions could be regulated as a medical device.

For more information on the types of AI-enabled medical devices and the evolution of AI in clinical care, see BPC's fact sheet "AI 101: AI-Enabled Medical Devices."

AI-ENABLED MEDICAL DEVICES BY SPECIALTY

Radiologists use AI-powered cancer screening to enhance medical images and measure tumors.

Cardiologists use machine learning to detect unseen patterns in heart rhythm data.

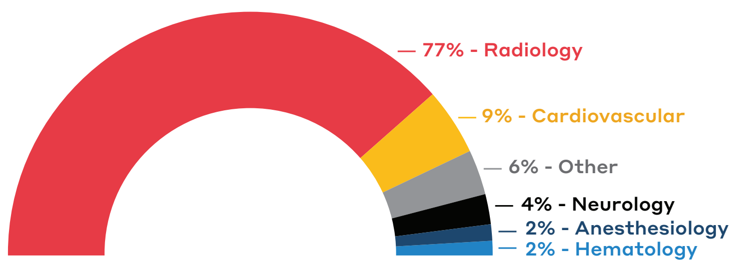

Note: The chart above uses data from the FDA's list of approved AI-Enabled Medical Devices as of July 2025.

The FDA regulates two main categories of medical software, both of which may include AI or machine learning:

- Software **as** a Medical Device (SaMD): This is standalone software intended for medical purposes, but it is not part of a hardware medical device. It operates on general-purpose computing platforms, such as cloud systems, mobile devices, or desktop computers. *Examples: AI-powered software that enhances medical images and measures tumors; machine-learning models that detect unseen patterns in heart rhythm data.*

- Software **in** a Medical Device (SiMD): This software is part of, or drives, a physical medical device and is necessary for the device's intended medical function. *Examples: handheld ultrasound devices with built-in AI to help capture consistent images; digital therapeutic apps that analyze data from glucose monitors and provide personalized guidance.*

As of July 2025, the FDA's public database lists over 1,250 AI-enabled medical devices authorized for marketing in the United States, up from 950 devices recorded as of August 2024. Although generative AI often dominates media headlines, most AI-enabled medical devices rely on predictive models. Researchers are testing generative applications in health care, such as models that generate synthetic chest X-ray and skin-lesion images, or ones that integrate multimodal data—such as imaging and clinical text—to improve diagnostic speed.

KEY AI TERMS

Machine Learning (ML): A subset of AI that enables computers to learn from data and make predictions or decisions.

Predictive AI: Uses machine learning to forecast future outcomes, producing consistent, data-driven results.

Generative AI: A subfield of AI that enables computers to create new content—such as text, images, or data—with outputs that can vary. It is still experimental in clinical settings.

Large Language Models (LLMs): Advanced generative AI systems trained on large volumes of data to summarize, translate, predict, and generate contextual conversations.

SCOPE OF FDA OVERSIGHT AND OPERATIONAL CAPACITY

The FDA applies a risk-based approach to oversight, requiring that devices "demonstrate a reasonable assurance of safety and effectiveness." Higher risk devices undergo more rigorous testing and review than lower risk ones.

In certain cases, the FDA exercises "enforcement discretion" when a tool technically meets the definition of a medical device but poses low risk. For devices under enforcement discretion, the FDA does not expect manufacturers to submit a premarket review application or to register their device with the agency. The FDA often applies this approach to software that supports general wellness or self-management (e.g., weight logging, medication reminders). The FDA's Digital Health Policy Navigator can help developers determine whether their product meets the definition of a device and, if so, whether it falls under enforcement discretion. If the answer is unclear, developers are advised to seek legal or regulatory advice from experts.

Clinical decision support (CDS) software presents a complex regulatory area. The 21st Century Cures Act of 2016 narrowed the FDA's authority over certain health software. It amended Section 520 of the Federal Food, Drug, and Cosmetic Act to exclude some CDS software from being classified as medical devices if they are designed to support—and not replace—clinical decision-making. The law established criteria that CDS software must meet to be outside the definition of "medical device" and therefore outside the FDA's oversight authority. One key criteria is that the software must allow providers to independently review and understand the basis for recommendations.

In 2022, the FDA issued final guidance clarifying which CDS software functions qualify for exclusion and offering examples of how the agency applies the criteria. For example, the guidance notes that CDS tools used in time-sensitive situations or that rely on complex, nontransparent algorithms—particularly those that draw on multiple types of data—may not qualify. Some stakeholders argue that the guidance interprets the statutory criteria too narrowly and discourages innovation. For example, former FDA Commissioner Scott Gottlieb warns the guidance may discourage developers from embedding advanced but low risk AI in electronic health records, because integrating diverse data sources could trigger FDA oversight.

As technologies become more sophisticated and raise complex regulatory questions, the FDA faces workforce and capacity constraints that limit its ability to evaluate AI-enabled medical devices quickly and comprehensively. As of September 2025, staffing levels were down by approximately 2,500, or nearly 15%, from 2023. During recent congressional hearings, lawmakers raised the possibility of using AI to support regulatory functions, such as adverse event monitoring.

In 2025, the FDA rolled out its own generative AI model "Elsa," a chatbot powered by Anthropic's Claude to help staff with reading, writing, and summarizing internal documents. Commissioner Marty Makary has highlighted Elsa's potential to streamline and reduce the time to complete scientific reviews, yet some stakeholders have questioned how the tool might influence decision-making and what would happen if the agency used AI to make decisions that later get challenged in court. For more information, see BPC's brief, "Mapping the Rise of AI in Federal Health Agencies."

For an overview of the FDA's organization structure related to AI/ML, see Appendix A.

AI ACTION PLAN – IMPLICATIONS FOR THE FDA

In July 2025, the Trump administration released a three-pillar AI Action Plan focused on accelerating U.S. innovation, building the infrastructure needed to power it, and leading globally in AI diplomacy and security. Several components intersect with the FDA's oversight role for health AI:

- **Investing in AI-Enabled Science**: Promotes cloud-connected automated labs and encourages data sharing to enable innovators to test ideas and speed up discovery.
- **Building an AI Evaluation Ecosystem**: Seeks to establish a framework for rigorously assessing AI performance. This task involves publishing guidelines and resources for federal agencies to conduct their own evaluations and supporting development of the underlying science for measuring and evaluating models.
- **Enabling AI Adoption**: Encourages AI Centers of Excellence and regulatory sandboxes to rapidly deploy and test AI, while committing to open data sharing. The action plan proposes that these efforts be enabled by agencies such as the FDA. It also calls for the National Institute of Standards and Technology (NIST) to convene public, private, and academic stakeholders to accelerate the development and adoption of national standards for AI.
- **Empowering American Workers in the Age of AI**: Recommends expanding AI literacy and skills development.

For more information, see BPC's blog, "From Vision to Action: Aligning Health Policy and Innovation within the AI Action Plan."

FDA REGULATORY PATHWAYS

Guiding Principles for AI Oversight

The FDA has grounded its oversight of AI-enabled devices in two complementary frameworks, developed and refined across multiple administrations.

The Total Product Life Cycle (TPLC) approach assesses a device across its entire lifespan: design, development, deployment, and postmarket monitoring. This is important for AI, especially adaptive or generative models, that may continue to evolve after they have received authorization.

PREMARKET REVIEW

FDA evaluates safety, effectiveness, and performance, with attention to intended use, training data, and risk management. For AI tools, this includes testing across diverse populations and care settings.

CHANGE MANAGEMENT

Through the PCCP, manufacturers outline expected software updates in advance, allowing certain changes without a new FDA submission.

POSTMARKET MONITORING

Manufacturers must monitor real-world issues such as model drift, bias, or unexpected outputs, and take corrective action as needed.

RISK-BASED OVERSIGHT

Oversight is calibrated to a device's risk, factoring on condition severity, potential harm from errors, and the extend of human oversight.

The Good Machine Learning Practice (GMLP) principles were developed with Canada and the United Kingdom and set 10 guiding principles to ensure safe and effective AI. The recommendations emphasize transparency, data quality, and ongoing model maintenance, as well as the importance of using representative datasets, following solid engineering practices, accounting for human-AI interactions, and monitoring performance over time. GMLP informs regulatory tools, such as Predetermined Change Control Plans (PCCPs) and postmarket monitoring requirements.

The FDA also collaborates with global bodies like the International Medical Device Regulators Forum (IMDRF) to align approaches on change control, validation, and labeling. This coordination helps reduce regulatory fragmentation and supports the safe, consistent adoption of AI across international markets.

The sections that follow explore each stage of the regulatory lifecycle in more detail. For a *full timeline of FDA regulatory actions related to AI/ML, see Appendix B.*

Premarket Authorization

FDA oversight relies primarily on premarket review. The agency regulates medical devices, including ones that integrate AI, through one of several pathways, depending on the device's risk classification and whether a substantially equivalent "predicate" device already exists.

As part of their submission, manufacturers must clearly specify both the intended use (the general purpose or function of the device) and the indications for use (the specific clinical conditions, patient populations, and settings in which the device is intended to be used). The FDA then reviews whether the device is safe and effective for the stated use, is technically accurate, performs consistently across different patients' groups, and is usable in clinical practice.

Risk classifications fall into three categories:

- Class I - Low risk: Devices with minimal potential for harm (e.g., static tools like tongue depressors).
- Class II - Moderate risk: Devices that present some risk and generally require clinical oversight. Many AI-enabled devices fall into this category. These devices typically go through 510(k) or de novo pathways.
- Class III - High risk: Devices that are life-sustaining or pose significant risk. These require the most rigorous review, the Premarket Approval (PMA) process. A growing number of AI-enabled devices are in this category.

PATHWAY	WHEN IT'S USED	KEY FEATURES	AI USE
510(k) Clearance *For an example of an AI-enabled medical device seeking FDA clearance, see Appendix C.*	Devices "substantially equivalent" to a marketed (predicate) device.	Most common, fastest, and least burdensome. Uses device-specific guidelines to demonstrate equivalence. Often does not require clinical trials.	Many AI tools have cleared through 510(k) by showing similarity to existing, marketed software.
De Novo Classification	Novel devices with no predicate but are deemed low to moderate risk.	Establishes a new device type and special controls to ensure safety and effectiveness.	Used for some early AI tools where no clear predicate existed.
Premarket Approval (PMA)	For high risk devices (Class III).	Most stringent pathway. Requires robust clinical evidence and in-depth FDA review of safety and effectiveness.	Less common for AI to date; applies to high risk, life-sustaining technologies.

In 2019, the FDA introduced **Predetermined Change Control Plans** to manage modifications to AI-enabled medical devices postmarket. Companies submit a PCCP alongside the initial device submission. If the changes stay within the scope of an approved PCCP, manufacturers can implement them without undergoing a new FDA review. The FDA has indicated that a PCCP change must operate within a medical device's intended use and indications for use, as well as maintain or improve the safety and effectiveness of the medical device to which it is applied.

PCCPs also require transparency with users, such as version tracking, change documentation, and clear communication about a device's performance and safety. Under Section 3308 of the Food and Drug Omnibus Reform Act of 2022, and refined in the FDA's 2025 final guidance, a PCCP must include three core elements:

1. Description of Modifications: Detailed account of planned changes to the device, including the type of update, frequency, and which parts of a device may change.
2. Modification Protocol: Methods to implement those changes, including data practices, retraining, testing, and user communication. Each change must be tied to a verification and validation plan.

3. Impact Assessment: Evaluation of benefits, risks, and mitigation strategies for the proposed modifications, including how problems will be managed.

Some developers report FDA premarket pathways are unclear, time-intensive, and difficult to navigate, particularly in the context of fast-moving AI innovation. In certain cases, manufacturers have discontinued promising tools, citing the cost and complexity of obtaining FDA market authorization. Others design tools that fall outside the FDA's jurisdiction—for example, by avoiding explicit medical claims or framing products as general wellness tools. However, consumers may not clearly understand that these tools have not undergone FDA review. Further, generative AI presents unique premarket evaluation challenges due to the large number of potential applications and output variability.

In November 2024, the FDA's Digital Health Advisory Committee (DHAC) held its inaugural meeting to discuss how the agency should adapt its regulatory approach for generative AI-enabled devices.

Members recommended the FDA require manufacturers to disclose a model's intended use, training data, and estimates of error rates and "hallucinations," possibly through standardized disclosure forms like model cards. However, members noted that such information may be hard to obtain from commercial models like OpenAI that don't disclose their training databases. Some compared regulating large language models to evaluating doctors—testing with standardized materials and supervised use—but acknowledged key differences, because AI systems don't reason like humans. Members also discussed whether independent bodies should conduct testing and quality assurance, rather than relying solely on the manufacturer.

Finally, the committee recommended that the FDA accelerate its premarket review by expanding its Medical Device Development Tools (MDDT) Program to include AI-specific resources, such as reference datasets and performance benchmarks. Once qualified, any device sponsor could use these tools to generate evidence in their submissions.

These insights informed the FDA's January 2025 draft guidance on marketing submission recommendations for AI-enabled devices.

Marketing Authorization

Once a device is authorized for marketing, the FDA expects device makers to clearly describe how AI functions within a product. Labeling requirements vary by device type. For AI/ML-enabled products, the FDA's 2025 guidance outlines specific labeling expectations, including:

- a clear statement that the device uses AI, with a plain-language description of how it supports the intended use of the device.
- details on the inputs and outputs of the model, how data are collected, and what systems it interacts with.
- performance measures, along with any known risks or potential sources of bias.
- if a PCCP is in place, how performance will be monitored and how updates will occur.
- for patient- or caregiver-facing devices, instructions and explanations need to be written at an accessible level. (Professional societies like the Federation of State Medical Boards and the American Medical Association have also published guidance for clinicians on sharing when and how an AI-enabled device is used in patient decision making.)

Companies must also show that they have designed safeguards against hacking or system failure. In its 2025 premarket guidance, the agency states that manufacturers should show their products are "secure by design," embedding threat modeling, risk assessments, and update mechanisms into development. The FDA requires a Software Bill of Materials (SBOM)—a list of all software components—to enable tracking of vulnerabilities, and it expects clear labeling so that users know how to maintain a device's security.

Postmarket Monitoring

Compared with premarket review, the FDA's postmarket oversight requirements are more limited. Manufacturers of medical devices, including those with AI capabilities, must meet baseline quality requirements under the Quality System Regulation (QSR), which requires companies to design safe products, manage risk, and document updates. In 2026,the FDA will update these rules under the Quality Management System Regulation (QMSR), aligning U.S. oversight with international standards (ISO 13485:2016). The agency recommends that companies gather and analyze real-world performance data and document any corrections they make, especially when a PCCP is in place. Ongoing postmarket monitoring requirements are more common for De Novo and PMA devices than for those cleared through the 510(k) pathway.

The FDA can set additional postmarket requirements for higher-risk products. Section 522 of the Federal Food, Drug, and Cosmetic Act authorizes the FDA to require postmarket surveillance for certain class II or class III devices, including ones that pose a serious risk if they fail, are used frequently in pediatric populations, are implanted in the body for more than one year, or are life-supporting or life-sustaining outside a clinical setting. Manufacturers subject to a 522 order must submit a formal surveillance plan and conduct studies to evaluate long-term safety, effectiveness, and performance. Even when a 522 order is not issued, the FDA can intervene if a product's performance degrades or its marketing claims exceed the evidence. The agency has issued warning letters for products it determines manufacturers marketed as medical devices without appropriate clearance. It has also issued recalls for devices that pose safety risks or cybersecurity concerns.

A cornerstone of FDA's postmarket surveillance system is the Medical Device Reporting (MDR) program, which requires manufacturers, importers, and user facilities to report adverse events, including device malfunctions, serious injuries, or deaths. These reports are published in the Manufacturer and User Facility Device Experience (MAUDE) database. MAUDE has significant limitations, including underreporting, inconsistent data quality, and challenges capturing complex AI-specific risks. Voluntary reporting through the FDA's MedWatch system supplements this information but is not widely used by clinicians or patients. Recent reports indicate the FDA plans to combine vaccine, drug, and device adverse event reporting into one automated platform that uses AI to detect issues with products in real time. However, details on implementation are unclear.

Although these oversight mechanisms provide important safeguards, they do not constitute a comprehensive framework for monitoring AI-enabled devices in real-world settings. The agency has limited mechanisms to track real-world performance, model drift—when a model's accuracy declines over time— and safety over time. While some academic and health centers possess the resources to conduct monitoring, many lower resourced and rural providers do not. Academic and policy experts have proposed postmarket surveillance frameworks, such as a targeted public–private surveillance model focused on higher-risk AI systems and using shared health system registries to periodically revalidate performance and flag emerging risks.

The FDA's Digital Health Advisory Committee discussed ways to establish postmarket monitoring for generative AI, potentially through a central database or registry that would track safety and accuracy over time. The committee also discussed the need for local data reporting to confirm that models are relevant for specific populations. Members further explored how PCCPs could be adapted for generative AI, provided they include clear information on allowable updates, methods for performance monitoring, procedures for updating labeling, and notification protocols for FDA and users if a device no longer functions as intended. Building a centralized monitoring system across diverse health systems is complex and would require significant coordination and data sharing, and members noted that this would be difficult to implement across all health systems.

These insights informed the FDA's September 2025 Request for Public Comment on best practices for measuring and evaluating AI-enabled medical devices' real-world performance. The agency is interested in evidence-based methods to identify and address performance changes over time, especially those currently used at scale in clinical environments.

CONCLUSION

Advancing the responsible adoption of health AI begins with understanding how oversight works today. The FDA regulates AI-enabled medical devices through a risk-based framework and has introduced tools like Predetermined Change Control Plans and updated labeling requirements to account for models that evolve over time. Still, critical questions remain about how to define "intended use" for adaptive or generative algorithms, how to ensure consistent performance after deployment, and what level of postmarket monitoring is feasible across diverse care settings.

Although this brief is focused on the FDA's authority, many AI tools in health care fall outside its jurisdiction. Our companion issue brief explores AI solutions outside FDA oversight and the policy challenges they present. BPC's next issue brief will explore emerging questions about coverage and reimbursement of health AI solutions. This will continue BPC's ongoing work to help policymakers navigate the complexities of health AI and pursue bipartisan solutions that support innovation, safeguard patients, and improve high-value use of health care spending.

Appendices

APPENDIX A. HOW THE FDA ORGANIZES OVERSIGHT OF AI AND DIGITAL HEALTH

The FDA's work on AI spans multiple centers, initiatives, and advisory structures. Key components include:

- Center for Devices and Radiological Health (CDRH): Oversees regulation of AI-enabled devices. Houses AI review teams and regulatory science expertise.
- Digital Health Center of Excellence (DHCoE, housed within CDRH): Coordinates AI and digital health policy across the FDA and supports early-stage engagement with developers.
- Digital Health Advisory Committee (DHAC): An external advisory committee that provides expert input on fast-moving digital health and AI issues; its inaugural meeting was in November 2024 and the next meeting is scheduled for November 2025.

In 2025, the FDA created two cross-agency councils:

- External Policy Council: Establishes principles and policies for AI in regulated products.
- Internal Use Council: Oversees how the FDA uses AI internally to make its work faster and more efficient, including through pilot programs and staff training.

APPENDIX B. FDA AI/ML REGULATORY TIMELINE

December 2017: The FDA published a guidance document, *Software as a Medical Device (SAMD)*: Clinical Evaluation, based on IMDRF principles for global regulators.

September 2019: The FDA published *Changes to Existing Medical Software Policies Resulting from Section 3060 of the 21st Century Cures Act*, a guidance document that explained how amendments to the definition of a device under the 21st Century Cures Act affected FDA policies related to medical device software.

April 2019: The FDA published a discussion paper and request for feedback, *"Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD)*, describing a potential approach to premarket review for AI/ML-driven software modifications.

January 2021: The FDA published *Artificial Intelligence/Machine Learning (AI/ML)-Based Software as Medical Device (SaMD) Action Plan* to summarize stakeholder feedback and describe a five-part action plan for regulating SaMD.

September 2022: The FDA published two guidance documents focused on software. One guidance document, *Policy for Device Software Functions and Mobile Medical Applications*, described the agency's regulatory oversight of device software functions and mobile medical applications (MMAs) that meet the definition of a medical device. The other document, *Clinical Decision Support Software*, clarified the FDA's assessment of the types of CDS software functions excluded from the definition of a device, as amended by the 21st Century Cures Act.

October 2021: The FDA published *Good Machine Learning Practice for Medical Device Development: Guiding Principles*, listing 10 guiding principles to inform the development of GMLP.

April 2023: The FDA published draft guidance, *Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence/Machine Learning (AI/ML)-Enabled Device Software Functions*, describing an approach to regulating ML-enabled medical devices that supports iterative improvement while ensuring device safety and effectiveness.

October 2023: The FDA published *Predetermined Change Control Plans for Machine Learning -Enabled Medical Devices: Guiding Principles* with five guiding principles for predetermined change control plans; the recommendations drew upon overarching GMLP guiding principles.

March 2024: The FDA published *Artificial Intelligence & Medical Products: How CBER, CDER, CDRH, and OCP are Working Together*, describing the agency's coordinated approach to AI.

June 2024: The FDA published *Transparency for Machine Learning-Enabled Medical Devices: Guiding Principle*, establishing guiding principles for transparency for ML-enabled medical devices; the recommendations built upon overarching GMLP guiding principles.

December 2024: The FDA announced the availability of a final guidance, *Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions*, providing recommendations for PCCPs tailored to AI-enabled devices.

January 2025: The FDA published draft guidance, *Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations*, with lifecycle considerations and recommendations to support marketing submissions for AI-enabled medical devices. The agency published additional draft guidance, *Considerations for the Use of Artificial Intelligence To Support Regulatory Decision-Making for Drug and Biological Products*, that included recommendations for sponsors and other stakeholders on the use of AI to produce information or data intended to support regulatory decision-making regarding safety, effectiveness, or quality of drugs.

May 2025: The FDA *began deploying* generative AI internally across all centers to support workflows and reviews.

Mid-2025: The FDA updates its *AI-Enabled Medical Devices List*, adding tagging for devices that incorporate large language models or foundation models.

September 2025: The FDA announced a Request for Public Comment, *Measuring and Evaluating Artificial Intelligence-enabled Medical Device Performance in the Real-World*, to solicit feedback from stakeholders on current, practical approaches to measure and evaluate the performance of AI-enabled medical devices in the real-world.

November 6, 2025: The FDA Digital Health Advisory Committee discussed and made recommendations on the topic of "*Generative Artificial Intelligence-Enabled Digital Mental Health Medical Devices*".

APPENDIX C. LIFECYCLE OF AN AI-ENABLED MEDICAL DEVICE - THE 510(K) PATHWAY

Step 1: Concept and Development
A startup develops an AI tool that spots lung nodules on CT scans. The engineers train it using large imaging datasets.

Step 2: Finding a Predicate
To pursue FDA clearance via the 510(k) pathway, the company identifies a predicate device, a previously cleared device that does something similar or "substantially equivalent."

Step 3: Submission to FDA
The company submits a 510(k) application demonstrating that the new device is as safe and effective as the predicate. The submission includes performance data, validation studies, and a description of any differences.

Step 4: FDA Review
FDA reviewers comb through the submission and may ask questions, request additional information, or require revisions if marketing claims exceed what the evidence supports.

Step 5: Clearance
If the FDA finds the device substantially equivalent to the predicate, it grants clearance for market. 510(k) clearance does not imply endorsement or validation of clinical benefit, only that the device meets the statutory standard of safety and effectiveness relative to the predicate.

Step 6: PCCPs and Labeling
Manufacturers must provide clear labeling, including how the AI functions and any known limitations. If a Predetermined Change Control Plan is included, it must outline how algorithm updates will be managed postclearance.

Step 7: Postmarket

Once on the market, the device enters clinical use. Manufacturers are responsible for monitoring real-world performance, identifying issues such as model drift, and implementing corrections as needed. If a device is found to perform inconsistently or is marketed beyond its cleared scope, the FDA can send warning letters or order a recall. A recent study found that about 6% of AI-enabled devices have faced recalls, often within the first year, usually tied to tools that reach the market without much clinical testing.

ACKNOWLEDGMENTS

The Bipartisan Policy Center would like to thank the Peterson Center on Health Care for its generous support of this work. We also appreciate Jodi Daniel's contributions as a consultant on this project.

HEALTH PROGRAM

BPC's Health Program advances bipartisan policy solutions to build a more cost-effective, evidence-based health care system to improve population health. We work to strengthen and sustain Medicare and Medicaid, accelerate the shift toward value, address inefficiencies and misaligned incentives, and responsibly leverage technology and innovation.