Governing Autonomous AI Agents: Critical Implications for Health Tech Entrepreneurs and Investors

AUG 07, 2025

Disclaimer: The thoughts and opinions expressed in this essay are my own and do not those of my employer.

Abstract

This essay examines Joe Kwon's policy paper "AI Agents: Governing Autonomy i Digital Age" from the Center for AI Policy and analyzes its critical implications health tech entrepreneurs and investors. The paper proposes a comprehensive regulatory framework for autonomous AI agents that includes an "Autonomy Passport" system, continuous monitoring requirements, human oversight manda and workforce impact research. For health tech stakeholders, these proposed regulations present both significant compliance challenges and strategic opportunities. While the regulatory framework may increase development costs time-to-market, it also creates competitive moats for well-capitalized companies standardizes safety practices across the industry. Health tech entrepreneurs mus prepare for substantial regulatory compliance investments, while investors need adjust their due diligence processes and portfolio strategies to account for the emerging regulatory landscape. The paper's emphasis on human oversight in cri domains aligns closely with existing healthcare regulations but may limit the efficiency gains that make AI agents attractive to healthcare organizations.

Table of Contents

2. Key Findings from Kwon's Policy Paper

3. Current State of AI Agents in Healthcare

4. Regulatory Impact Analysis for Health Tech Companies

5. Strategic Implications for Entrepreneurs

6. Investment Considerations and Portfolio Strategy

7. Compliance Framework Development

8. Competitive Landscape Transformation

9. Future Scenarios and Risk Assessment

10. Recommendations and Action Items

Governing Autonomous AI Agents: Critical Implications for Health Tech Entrepreneurs and Investors

In May 2025, Joe Kwon of the Center for AI Policy published a comprehensive p paper titled "AI Agents: Governing Autonomy in the Digital Age" that proposes most detailed regulatory framework yet conceived for autonomous AI systems. F health tech entrepreneurs and investors, this 23-page document represents a pot inflection point that could fundamentally alter the development, deployment, an commercialization of AI-powered healthcare solutions. The paper's recommend if adopted by Congress, would establish mandatory pre-deployment licensing, continuous monitoring requirements, human oversight mandates, and comprehe workforce impact assessments for AI agents operating above basic assistant leve

The timing of this policy proposal coincides with explosive growth in healthcare agent adoption. Industry data from 2025 shows that AI agents are already manag 50-60% of front-office administrative tasks in many healthcare practices, reduci

revenue cycle management costs by up to 70%, and automating complex workflo from appointment scheduling to clinical documentation. Major healthcare techr companies like ServiceNow report USD 325 million in annualized value from autonomous agent deployments, while organizations like Mayo Clinic are pionee "agentic automation architectures" that fundamentally reimagine clinical and operational workflows. This rapid adoption trajectory makes Kwon's regulatory proposals particularly consequential for stakeholders who must navigate the intersection of technological capability and regulatory compliance.

The paper's central thesis argues that AI agents pose three distinct categories of catastrophic misuse through cyberattacks or dual-use applications, gradual hum disempowerment as decision-making migrates to opaque algorithms, and unprecedented workforce displacement affecting up to 300 million jobs globally according to Goldman Sachs projections. Healthcare, with its complex regulator environment, high-stakes decision making, and substantial economic footprint, represents a critical testbed for these regulatory approaches. The proposed fram would require health tech companies to fundamentally rethink their product development cycles, compliance strategies, and go-to-market approaches while potentially creating new competitive advantages for organizations capable of navigating increased regulatory complexity.

Understanding the paper's key findings requires examining the sophisticated fiv level autonomy classification system that Kwon proposes as the foundation for r proportional regulation. Level 1 "shift-length assistants" can work autonomousl roughly eight hours with human oversight, while Level 5 "frontier super-capable systems" operate indefinitely across any domain without human guidance. Most current healthcare AI applications fall into Levels 1-2, but the rapid pace of capa advancement suggests that Level 3-4 systems capable of multi-day autonomous operation may emerge within the next few years. This progression timeline has d implications for health tech companies, as the regulatory requirements become substantially more stringent at Level 2 and above, requiring formal registration, audits, and continuous monitoring.

The paper's "Autonomy Passport" proposal represents the most significant regu innovation, establishing a mandatory federal registration system for AI agents operating at Level 2 or higher. This system would require companies to file deta dossiers documenting their agents' mission envelopes, tool access permissions, autonomy classifications, security validation results, and emergency contact information before deployment. The US AI Safety Institute would set technical standards and maintain a public registry of approved agents, while accredited pr firms would conduct the required safety audits. For health tech entrepreneurs, t represents a fundamental shift from the current largely self-regulated environm a formal pre-market approval process similar to FDA device clearance but poten more complex and time-consuming.

The continuous monitoring and enforcement mechanisms proposed in the paper would create ongoing compliance obligations throughout an AI agent's operatio lifecycle. High-capability agents would need to operate within digital sandboxes enforce pre-approved action lists, attach tamper-evident signatures to every outb action, and remain subject to emergency recall authority by the Cybersecurity ar Infrastructure Security Agency. These requirements mirror some aspects of exis healthcare IT security frameworks but extend far beyond current practice in ter real-time monitoring and federal oversight authority. Healthcare organizations t have invested heavily in AI agent deployments would need to retrofit their syste comply with these new monitoring requirements, potentially creating significan technical and financial challenges.

The human oversight mandate for critical systems represents perhaps the most directly relevant aspect of Kwon's framework for healthcare applications. The pa proposes that when AI agents make recommendations in domains that normally require professional licensing or regulatory approval, qualified humans must rev and approve those recommendations before execution. In healthcare contexts, th would mean that licensed healthcare professionals must approve AI agent recommendations for prescription changes, treatment modifications, diagnostic interpretations, and other clinical decisions. While this aligns with existing med practice standards and liability frameworks, it may limit the efficiency gains tha

make AI agents attractive to healthcare organizations facing staffing shortages a cost pressures.

The workforce impact research mandate represents the final pillar of Kwon's regulatory framework, directing federal agencies to produce annual reports trac job displacement and wage trends related to AI agent adoption. For healthcare, t particularly relevant given projections that up to 40% of jobs in advanced econor show high exposure to AI-driven automation, with college-educated roles facing higher displacement rates. Healthcare organizations are already reporting signif productivity gains from AI agent deployments, with some companies achieving 1 developer productivity improvements and double-digit reductions in call handli times. The proposed research mandate would create systematic tracking of these workforce impacts, potentially informing future regulatory interventions or sup programs.

Current healthcare AI agent deployments demonstrate both the transformative potential and regulatory complexity that Kwon's framework aims to address. Rev cycle management represents one of the most mature application areas, with companies like CodaMetrix using natural language processing to automate medi coding across 200 hospitals and 50,000 providers. These systems continuously le from clinical data while incorporating payer rules and compliance requirements achieving accuracy levels that reduce manual intervention while maintaining regulatory compliance. Under Kwon's proposed framework, these systems would likely require Level 2 classification and full Autonomy Passport registration, giv their ability to make consequential financial decisions autonomously.

Clinical documentation represents another rapidly evolving application area wh agents are fundamentally changing healthcare workflows. Companies like Augm deploy ambient documentation tools that capture natural clinician-patient conversations and convert them into structured medical notes, serving nearly ha million clinicians. These systems operate continuously during patient encounter making real-time decisions about what information to capture and how to struct clinical documentation. The proposed regulatory framework would likely requir human oversight for clinical documentation agents that generate billable diagno

codes or treatment recommendations, potentially limiting their autonomous capabilities while ensuring clinical accountability.

Diagnostic and imaging applications showcase the potential for AI agents to ope at higher autonomy levels while maintaining safety and accuracy. Companies lik Qure.ai have deployed AI systems across 4,500 sites in over 100 countries to aut interpretation of X-rays, CT scans, and ultrasounds, particularly in areas with li radiology expertise. These systems can identify and prioritize conditions like tuberculosis, lung cancer, and stroke, often detecting anomalies that human radiologists might miss. However, under Kwon's framework, such systems woul require human radiologist approval before final diagnostic decisions, maintainir current standard of care while potentially limiting efficiency gains.

The regulatory impact analysis for health tech companies reveals several critical considerations that will shape strategic planning and investment decisions. Compliance costs represent the most immediate concern, as companies would n invest in safety auditing, continuous monitoring systems, human oversight integration, and emergency response capabilities. These costs are not one-time expenses but ongoing operational requirements that scale with system complexi deployment breadth. For early-stage health tech startups, these compliance requirements could represent existential challenges, as the cost of Autonomy Pa registration and continuous monitoring may exceed available funding runway.

Time-to-market implications are equally significant, as the proposed pre-deploy review process could add months or years to product development cycles. Healt AI companies currently benefit from relatively rapid iteration and deployment capabilities, particularly for administrative and operational applications that do require formal FDA approval. The Autonomy Passport system would introduce a formal gate that must be cleared before any Level 2 or higher AI agent can be deployed, potentially slowing innovation cycles and reducing the first-mover advantages that characterize the current competitive landscape.

Market access considerations add another layer of complexity, as the proposed framework would require major cloud providers and app stores to block any AI

that doesn't appear on the federal green list. This creates a binary approval syste where companies either achieve full market access through successful registrati face complete market exclusion. For health tech companies that rely on cloud-ba deployment models, this represents a fundamental shift in go-to-market strategy risk management. Companies would need to build stronger relationships with accredited auditing firms and develop more robust quality assurance processes t ensure successful registration.

The competitive landscape implications of Kwon's regulatory framework are particularly nuanced for health tech companies. Large, well-capitalized organiza with existing regulatory compliance capabilities may benefit from increased bar to entry that limit competition from smaller, more agile startups. Companies lik Epic, Cerner, and other established healthcare IT vendors have extensive experie navigating complex regulatory environments and may be better positioned to ab the costs and complexity of Autonomy Passport compliance. This could accelera market consolidation and reduce the likelihood of disruptive innovation from st companies.

However, the framework also creates opportunities for companies that can devel specialized compliance capabilities or serve as regulatory technology providers. requirement for accredited private auditing firms creates an entirely new service category that doesn't currently exist in the AI industry. Health tech companies w strong regulatory expertise could potentially pivot to serve this emerging marke while specialized consulting firms could emerge to help smaller companies navig the compliance requirements. The standardization aspects of the framework cou also benefit companies by creating clearer requirements and reducing regulatory uncertainty.

Strategic implications for health tech entrepreneurs require fundamental rethin of business model assumptions and development priorities. Product developmen strategies must now account for regulatory compliance from the earliest stages, than treating compliance as a post-development consideration. This means integ safety-by-design principles, building comprehensive audit trails, and designing human oversight capabilities as core system features rather than add-on compon

Entrepreneurs must also consider whether their target applications justify the increased development and compliance costs, potentially shifting focus toward higher-value use cases that can support the additional regulatory burden.

Funding strategy considerations become more complex under the proposed regu framework, as investors will need to factor compliance costs and regulatory risk their investment decisions. Early-stage companies may need larger seed and Seri rounds to fund regulatory compliance activities, while later-stage companies ma longer development cycles that delay revenue generation and extend time-to-exi scenarios. Entrepreneurs must be prepared to articulate clear regulatory strateg potential investors and demonstrate deep understanding of the compliance requirements that affect their specific applications.

Partnership and acquisition strategies may become more attractive under increa regulatory complexity, as smaller companies may find it more efficient to partne or be acquired by larger organizations with established regulatory capabilities. T could lead to earlier exit opportunities for entrepreneurs but potentially at lowe valuations if regulatory uncertainty reduces buyer enthusiasm. Strategic partner with established healthcare IT companies or device manufacturers may become essential for smaller companies that lack the resources to navigate complex regu requirements independently.

Investment considerations for health tech venture capital and private equity firn must evolve to address the new regulatory landscape that Kwon's framework wo create. Due diligence processes must now include detailed assessment of regulat compliance strategies, evaluation of management team regulatory expertise, and analysis of competitive positioning in a more heavily regulated market environm Investors need to develop capabilities to evaluate the technical feasibility and co implications of Autonomy Passport compliance, human oversight integration requirements, and continuous monitoring system development.

Portfolio construction strategies may shift toward companies with stronger regu moats and away from pure-play technology companies that lack healthcare doma expertise. The increased barriers to entry created by regulatory compliance

requirements may make incumbent healthcare IT companies more attractive investment targets, while making early-stage AI startups riskier investments unl they demonstrate exceptional regulatory preparation. Investors may also need to increase reserve requirements for follow-on funding to support compliance activ that were not previously necessary.

Risk assessment frameworks must incorporate regulatory change risk as a prima factor in investment decisions. The Kwon framework represents just one possibl regulatory outcome, and investors must consider scenarios ranging from no addi regulation to even more restrictive requirements. Companies with greater regula flexibility and stronger compliance capabilities will be better positioned to adap various regulatory scenarios, making them more attractive investment targets. Geographic diversification considerations also become relevant, as regulatory requirements may vary significantly across different jurisdictions and create opportunities for regulatory arbitrage.

Compliance framework development represents a critical operational priority fo health tech companies preparing for potential regulatory changes. Legal and regulatory affairs capabilities must be substantially enhanced, with companies needing either in-house expertise or reliable external counsel with deep understanding of AI regulation, healthcare compliance, and emerging technolog governance. The interdisciplinary nature of the proposed requirements means th companies need legal expertise spanning technology law, healthcare regulation, employment law, and federal administrative procedure.

Technical infrastructure development for compliance requires significant engine investment in monitoring systems, audit trail capabilities, human oversight integration, and emergency response mechanisms. These systems must be desigr for regulatory compliance from the ground up, rather than retrofitted onto exist architectures. Companies must also develop capabilities to work with third-part auditing firms and maintain detailed documentation of system capabilities, safet testing results, and operational procedures. The technical complexity of implem tamper-evident signatures, sandbox containment, and real-time monitoring repr substantial development effort that must be factored into product roadmaps.

Quality assurance and risk management processes must be elevated to meet the standards expected by federal regulators and third-party auditors. This includes implementing formal software development lifecycle processes, conducting comprehensive security testing, maintaining detailed change control documenta and developing incident response procedures. Companies must also establish governance structures that can demonstrate appropriate oversight of AI system development and deployment decisions, potentially requiring board-level comm with relevant expertise.

The competitive landscape transformation that would result from Kwon's regula framework creates both risks and opportunities across different market segment Enterprise healthcare IT companies may benefit from increased barriers to entry the opportunity to leverage existing regulatory relationships and compliance capabilities. However, they also face substantial costs to retrofit existing AI syste for compliance and may need to slow innovation cycles to accommodate regulato review processes. The requirement for human oversight in critical applications r limit the automation benefits that justify AI investments, potentially reducing m demand for certain categories of AI solutions.

Startup and emerging company impacts are more complex and depend heavily on specific business models and target applications. Companies focused on non-cri administrative applications may benefit from lower regulatory barriers, while th targeting clinical decision support or autonomous diagnostic applications face n stringent requirements. The requirement for pre-deployment registration could eliminate the rapid iteration and continuous deployment models that many AI startups rely on for competitive advantage. However, successful navigation of the regulatory framework could also create significant competitive moats for compa that achieve compliance.

International competitiveness considerations become relevant as US companies face compliance costs and development delays that don't affect international competitors. However, the proposed framework's alignment with emerging international standards could also position US companies advantageously in glob markets that adopt similar regulatory approaches. The paper specifically mentio

coordination opportunities with the UK AI Safety Institute, the G7 Hiroshima Process, and emerging OECD frameworks, suggesting that regulatory harmoniz may reduce the competitive disadvantage of early compliance.

Future scenario analysis reveals several possible trajectories that health tech stakeholders must consider in strategic planning. The base case scenario assume partial adoption of Kwon's recommendations, with Congress implementing som of AI agent registration and oversight but potentially with modified requiremen longer implementation timelines. This would create a transitional period during which companies can adapt their systems and processes while maintaining competitive dynamics similar to current conditions.

The accelerated regulation scenario envisions rapid adoption of the full framew potentially in response to a significant AI-related incident that creates political pressure for immediate regulatory action. This scenario would create severe challenges for companies that haven't prepared for compliance requirements, potentially triggering market consolidation and significant competitive reshuffli Companies with advanced regulatory preparation would gain substantial compe advantages, while unprepared companies could face existential challenges.

The expanded regulation scenario considers the possibility that initial AI agent regulations could be followed by additional requirements covering broader aspe healthcare AI deployment. This could include integration with existing FDA dev regulation, expansion of human oversight requirements, or additional reporting obligations related to clinical outcomes and patient safety. Companies must con how initial regulatory compliance capabilities could be expanded to address futu requirements and whether their regulatory strategies are sufficiently flexible to a to evolving requirements.

Risk assessment and mitigation strategies must address both direct regulatory compliance risks and indirect competitive and market risks. Regulatory non-compliance risks include the possibility of market exclusion, civil penalties, crin liability for executives, and reputational damage that affects customer relationsh and investor confidence. Companies must develop comprehensive compliance

monitoring systems and legal risk assessment capabilities to identify and addres potential violations before they result in enforcement actions.

Market and competitive risks include the possibility that regulatory requiremen could reduce market demand for AI solutions, increase customer acquisition cos create competitive advantages for non-AI alternatives. Companies must maintai flexibility in their product development strategies and consider pivot options th could reduce regulatory exposure while maintaining market position. Financial include the possibility that compliance costs could exceed revenue generation capabilities, particularly for early-stage companies with limited resources.

Technology and operational risks focus on the possibility that required compliai systems could interfere with product functionality, reduce system performance, create new security vulnerabilities. The integration of human oversight requiren with autonomous AI systems presents particular technical challenges that must carefully managed to maintain both regulatory compliance and operational effectiveness. Companies must also consider the risks associated with third-part dependencies, including auditing firms, cloud providers, and monitoring system vendors.

Recommendations and action items for health tech entrepreneurs center on proa preparation for potential regulatory requirements while maintaining operationa flexibility. Immediate actions should include comprehensive assessment of curre system autonomy levels using the five-level framework proposed by Kwon, evalu of which systems would require registration and oversight under the proposed r and development of preliminary compliance strategies for high-risk applications Entrepreneurs should also begin building relationships with potential auditing partners and regulatory advisors while monitoring Congressional activity relatec regulation.

Medium-term strategic actions should focus on integration of compliance capab into product development processes, including implementation of audit trail sys human oversight interfaces, and monitoring capabilities that could satisfy regul requirements. Companies should also consider geographic expansion strategies

could provide regulatory diversification and develop contingency plans for diffe regulatory scenarios. Investment in regulatory expertise, either through hiring o external partnerships, becomes essential for companies targeting applications th would face significant oversight requirements.

Long-term strategic positioning requires consideration of how regulatory compl capabilities could become competitive advantages and revenue generators. Comp that develop strong compliance capabilities may be able to serve as regulatory technology providers for other organizations or provide consulting services to he smaller companies navigate complex requirements. The standardization aspects proposed framework could also create opportunities for companies that develop practice approaches that can be scaled across the industry.

Investment recommendations for venture capital and private equity firms empha the need for enhanced due diligence capabilities and risk assessment framework account for regulatory change scenarios. Investors should develop expertise in A regulation and healthcare compliance or partner with specialized advisors who c provide detailed assessment of regulatory risks and opportunities. Portfolio companies should be encouraged to begin compliance preparation activities immediately, even in the absence of final regulatory requirements, to maintain strategic flexibility and competitive positioning.

The health tech industry stands at a critical juncture where technological capabi advancement intersects with emerging regulatory frameworks that could fundamentally reshape competitive dynamics and business model viability. Joe K comprehensive policy framework represents the most detailed roadmap yet prop for AI agent governance, with implications that extend far beyond simple compl requirements. For entrepreneurs and investors in the health tech space, success evolving landscape will require not just technological innovation but also regula sophistication, strategic flexibility, and deep understanding of the complex inter between automation capabilities and human oversight requirements.

The companies that thrive in this new regulatory environment will be those that compliance not as a burden but as a source of competitive advantage, building

regulatory capabilities that enable them to deploy AI agents safely and effectivel while creating barriers to entry for less prepared competitors. The investment fi that succeed will be those that can accurately assess regulatory risks and opportunities, support portfolio companies through complex compliance requirements, and identify emerging opportunities created by the new regulator landscape. As the health tech industry continues its rapid evolution toward AI-powered automation, the intersection of innovation and regulation will increasir determine which companies and investors achieve sustainable success in this transformative market.

← Previous

Nex

Discussion about this post

Comments Restacks

Write a comment...