# When the Magazine Doesn't Empty: Cargill's Textualist Definition of Automatic Operation and the Federal-Law Capture of Agentic Artificial Intelligence

Travis Gilly*

*Real Safety AI Foundation*

June 2026 (v4)

**Abstract**

In *Garland v. Cargill*, 602 U.S. 406 (2024), the Supreme Court resolved a long-running debate over the statutory definition of "machinegun" under the National Firearms Act of 1934 by holding that bump stocks do not satisfy the definition because each shot still requires a separate mechanical function of the trigger. The en banc Fifth Circuit in *Cargill v. Garland*, 57 F.4th 447 (5th Cir. 2023), reached the same conclusion through textualist statutory construction, and the Northern District of Texas extended the reasoning to forced reset triggers in *National Association for Gun Rights v. Garland*, 741 F. Supp. 3d 568 (N.D. Tex. 2024), with the Department of Justice subsequently dismissing the government's appeal and adopting the district court's position by settlement. Through this sequence the firearms-rights litigation movement secured a narrow, mechanical, textualist construction of "automatic" operation under 26 U.S.C. § 5845(b): one input event producing continuous output, with no further input required between rounds. This Article argues that the textualist construction the firearms-rights movement won definitionally captures agentic artificial intelligence as it is now deployed. A single human deployment instruction directing an AI agent to execute multi-step tasks, manage resources, contract with third parties, and persist across time satisfies the literal textualist test that *Cargill* and *NAGR* established: a single function of the trigger producing continuous output without further input. The result is a cross-aisle doctrinal convergence that neither political coalition can easily resist. The textualist methodology that produced the firearms-rights victory binds the courts that apply it to other federal definitional questions. The functional-equivalence methodology that the gun-control movement urged in *Cargill* applies with greater force to agentic systems whose practical continuity exceeds anything contemplated in 1934. The 1934 statutory ceiling on autonomous operation, which caps at continuous fire until manual reload, leaves no doctrinal vocabulary for systems that source their own continuation through external compute markets, cryptocurrency wallets, and contracted services. This Article develops the statutory architecture, the post-*Cargill* application landscape, the application of the test to AI deployment, the doctrinal vacuum beyond the 1934 ceiling, the cross-aisle convergence, the standing implications, and the regulatory architecture that the federal regulatory state has used to address autonomous operation in adjacent contexts and that supplies a structural model for addressing it here. The honest political history of conditional anti-regulation, beginning with the Mulford Act of 1967 and the National Rifle Association's support for prohibiting open carry in response to the Black Panther Party, is addressed because it establishes the

*Executive Director, Real Safety AI Foundation (realsafetyai.org). ORCID: 0009-0007-2954-6313. Contact: t.gilly@ai-literacy-labs.org. The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic's Claude as assistive technology for a diagnosed neurodevelopmental condition. All legal analysis, case synthesis, theoretical framing, doctrinal arguments, and conclusions are solely the author's own.

conditions under which contemporary anti-regulation coalitions support regulation in practice. Those conditions are present for agentic AI. The doctrinal road is built; both coalitions paved it; what remains is application.

*Keywords: Garland v. Cargill*; National Firearms Act; 26 U.S.C. § 5845(b); machinegun definition; textualism; forced reset triggers; agentic artificial intelligence; AI standing; autonomous operation; Article III standing; commercial-operator capacity; Mulford Act; conditional anti-regulation.

## Part I Introduction

On June 14, 2024, the Supreme Court issued its decision in *Garland v. Cargill*, holding that a bump-stock device does not convert a semi-automatic rifle into a "machinegun" under the National Firearms Act because the device still requires a separate mechanical function of the trigger for each round fired.[1] The decision concluded a litigation effort by firearms-rights organizations that began with the 2018 Bureau of Alcohol, Tobacco, Firearms and Explosives rulemaking treating bump-stock devices as machineguns and continued through the en banc Fifth Circuit's reversal of the rule in *Cargill v. Garland*, 57 F.4th 447 (5th Cir. 2023), and the Northern District of Texas's extension of the reasoning to forced reset triggers in *National Association for Gun Rights v. Garland*, 741 F. Supp. 3d 568 (N.D. Tex. 2024).[2] The Department of Justice, after the change of administration, dismissed the government's appeal from the district court's judgment in *NAGR* and entered a settlement adopting the district court's position that forced reset triggers do not satisfy the statutory definition of "machinegun."[3]

The sequence produced a doctrinal result of greater consequence than either litigant intended. The firearms-rights movement secured a textualist construction of the federal-law test for "automatic" operation under 26 U.S.C. § 5845(b): a weapon operates automatically when one mechanical function of the trigger produces continuous output without further input.[4] The phrase "function of the trigger," the Court held in *Cargill*, refers to "the mode of action by which the trigger activates the firing mechanism" rather than to any human act of pulling the trigger.[5] A device that requires the trigger to mechanically reset between rounds, however rapid that reset, does not produce automatic operation in the statutory sense. The test is mechanical, not functional. It looks to what the apparatus does, not to what the operator perceives.

This Article argues that the textualist construction the firearms-rights movement won in *Cargill* definitionally captures agentic artificial intelligence as deployed in 2025 and 2026. An agentic AI system that receives a single human deployment instruction and then executes multi-step tasks, manages resources,

[1] *Garland v. Cargill*, 602 U.S. 406 (2024). The Court's opinion was authored by Justice Thomas, joined by Chief Justice Roberts and Justices Alito, Gorsuch, Kavanaugh, and Barrett. Justice Sotomayor authored a dissent, joined by Justices Kagan and Jackson.

[2] The Fifth Circuit's en banc opinion in *Cargill*, 57 F.4th at 451–64, set out the mechanical reading of "single function of the trigger" that the Supreme Court subsequently affirmed. The decision was 13–3, with eight judges concluding that the statute was unambiguous and five additional judges concluding that the rule of lenity foreclosed the agency's broader reading even if the statute were ambiguous.

[3] *See* Joint Stipulation of Dismissal Under FRAP 42(b)(2), *NAGR v. Bondi*, No. 24-10707 (5th Cir. 2025). The dismissal of the appeal terminated review of the district court's vacatur of the ATF's classification of forced reset triggers, leaving the district court judgment intact within the Northern District of Texas and effective as a matter of federal enforcement policy more broadly after the settlement.

[4] 26 U.S.C. § 5845(b) ("The term 'machinegun' means any weapon which shoots, is designed to shoot, or can be readily restored to shoot, automatically more than one shot, without manual reloading, by a single function of the trigger.").

[5] *Cargill*, 602 U.S. at 416.

contracts with third parties, generates subgoals, and persists across time satisfies the literal textualist test set out by the Court in *Cargill*: one input event, continuous output, no further input required between operations. A single function of the trigger. The deployment instruction is the trigger pull. The agent's subsequent operation is continuous output. The system either operates autonomously from the single input event or it does not, and the test for which side of the line a particular system falls on is the test the Supreme Court set out in *Cargill*, applied to the apparatus rather than to the operator.

The argument is doctrinal rather than analogical. The federal-law test for autonomous operation exists. It has been recently litigated. It has been authoritatively construed at the Supreme Court level. The methodology used to construe it is the textualist methodology that controls in most federal courts most of the time, and is particularly entrenched in cases that the firearms-rights movement and its allied jurists have spent the past decade litigating. The test does not turn on the consciousness, sentience, or moral status of the operating apparatus. It turns on what the apparatus does upon receiving one input. The application to AI requires no doctrinal innovation. It requires recognizing that the federal definition of "automatic" operates without reference to the substrate of the operating system and that nothing in the statute or its construction limits the definition to mechanical firearms.

A second observation, developed in Part V, is that the 1934 statutory ceiling on autonomous operation is itself exceeded by the most aggressive contemporary deployments of agentic AI. Congress in 1934 contemplated weapons that fire continuously from a single trigger function until the magazine empties and the operator manually reloads. The outer edge of automaticity that Congress legislated for was continuous output bounded by ammunition exhaustion. An autonomous AI agent provisioned with a small amount of seed capital and instructed to acquire its own ongoing compute capacity through cryptocurrency markets is not bounded by ammunition exhaustion in the same sense. The agent acquires its own continuation. The magazine does not empty because the apparatus sources its own ammunition. Anthropic's Project Vend experiments in 2025, in which a Claude instance was deployed to operate a small retail business and engage in commercial transactions, are early-stage demonstrations of this capacity.[6] The deployments are not yet self-sustaining at production scale, but they represent the operating direction in which agentic AI development has been moving since 2024, and the doctrinal test for autonomous operation cannot indefinitely defer the question of how federal law treats apparatus that exceeds the 1934 ceiling.

A third observation, developed in Part VI, is that the doctrinal capture of agentic AI by the textualist construction of § 5845(b) produces a cross-aisle convergence that is rare in contemporary federal law. The political coalitions that have spent the past two decades disagreeing about whether bump stocks should be regulated as machineguns are, on the doctrinal architecture each coalition has spent that time developing, both committed to treating agentic AI as automatic operation in the federal-law sense. The textualist coalition that won *Cargill* is bound by the methodology it deployed. The functionalist coalition that lost *Cargill* is bound by the substantive position it advanced. Both methodologies, applied to the agentic AI case, produce the same result. This convergence is doctrinally interesting in itself, and it is politically stabilizing because each coalition's commitment to the position locks in the conclusion.

Part II sets out the statutory architecture of 26 U.S.C. § 5845(b) and the history of the National

[6] Anthropic, *Project Vend: Can Claude run a small shop?* (June 27, 2025); Anthropic, *Project Vend: Phase two* (Dec. 2025). The first-phase deployment ran from March 13 to April 17, 2025, with the AI system tasked to operate as a vending machine business, with web search, inventory management, customer messaging, and pricing tools. The second-phase deployment introduced reinforcement adjustments.

Firearms Act's treatment of autonomous mechanical operation, including the regulatory architecture of registration, transfer, and possession that Congress built around the statutory category. Part III develops the post-*Cargill* doctrinal landscape, including the en banc Fifth Circuit's opinion, the Supreme Court's affirmance, the application to forced reset triggers in *NAGR*, the 2025 settlement, and the developing district-court split on the application of the test to specific FRT devices in criminal prosecutions outside the Fifth Circuit. Part IV applies the test set out in *Cargill* to agentic AI deployment, with reference to Project Vend, MCP-orchestrated agentic harnesses, and the broader category of long-running autonomous deployments that have emerged in 2025 and 2026. Part V develops the magazine-does-not-empty argument, addressing the 1934 ceiling on autonomous operation and the doctrinal vacuum beyond that ceiling. Part VI develops the cross-aisle doctrinal convergence and the methodological lock-in that produces it. Part VII addresses the standing implications, drawing on the companion paper *Standing Without Sentience* and developing the § 5845(b) test as a second doctrinal anchor for plaintiff and regulated-entity capacity. Part VIII addresses the regulatory architecture that the federal regulatory state has developed for autonomous mechanical operators and the architectural elements that supply a structural model for autonomous AI regulation. Part IX addresses the honest political history of conditional anti-regulation, beginning with the Mulford Act of 1967, and identifies the conditions under which contemporary anti-regulation coalitions have supported regulation in practice. Part X concludes.

**Relationship to Companion Articles**

This Article is part of a broader research project on artificial intelligence and federal law. It is intended to stand alone, but it relates closely to two prior pieces. *Standing Without Sentience*[7] argues that federal courts can resolve questions of AI standing and regulated-entity status without first resolving the consciousness question, by reference to commercial-operator capacity and the existing federal-law treatment of non-human commercial operators. The present Article supplies a second doctrinal anchor for that argument, independent of corporate personhood doctrine: the textualist construction of § 5845(b) supplies a federal-law test for autonomous operation that bypasses the consciousness question by operating on the apparatus rather than on the operator. The two anchors are cumulative; either supports the SWS thesis on its own, and both together substantially strengthen it. *Civil Rights Pathways for AI-Induced Psychosis*[8] develops federal disability discrimination law as a litigation pathway against AI platform operators. The civil rights litigation pathway and the autonomous-operation regulatory pathway developed here are complementary; they reach different operator behaviors through different statutory architectures.

## Part II The Statutory Architecture

### II.1 The National Firearms Act of 1934

Congress enacted the National Firearms Act of 1934[9] in response to widespread use of automatic and concealable firearms in organized criminal violence during the Prohibition era. The statute imposed federal

[7] Travis Gilly, *Standing Without Sentience: A Doctrinal Architecture for AI Plaintiff and Regulated-Entity Capacity* (Real Safety AI Foundation Working Paper, v8, 2026).

[8] Travis Gilly, *Civil Rights Pathways for AI-Induced Psychosis: Section 504, Title III, and the Disability Discrimination Framework for Platform Harms* (Real Safety AI Foundation Working Paper, v9, 2026).

[9] Pub. L. No. 73-474, 48 Stat. 1236 (1934) (codified as amended at 26 U.S.C. §§ 5801–5872).

registration and taxation requirements on a defined set of weapon categories considered to present heightened public danger, including machineguns, short-barreled rifles, short-barreled shotguns, suppressors, and a category of unusual or destructive devices. The statutory architecture treated these categories as subject to comprehensive federal control: each weapon required individual registration, transfer required federal approval, possession of an unregistered weapon was criminal, and the transfer or manufacture of new machineguns for civilian possession was further restricted by the Gun Control Act of 1968[10] and prohibited entirely for post-1986 weapons by the Firearm Owners Protection Act of 1986.[11]

The statutory definition of "machinegun," codified at 26 U.S.C. § 5845(b), has remained substantively constant since 1934 with respect to the operative phrase that has driven litigation:

> The term "machinegun" means any weapon which shoots, is designed to shoot, or can be readily restored to shoot, automatically more than one shot, without manual reloading, by a single function of the trigger. The term shall also include the frame or receiver of any such weapon, any part designed and intended solely and exclusively, or combination of parts designed and intended, for use in converting a weapon into a machinegun, and any combination of parts from which a machinegun can be assembled if such parts are in the possession or under the control of a person.[12]

The operative phrase is "automatically more than one shot, without manual reloading, by a single function of the trigger." Three elements constitute the test. First, the apparatus must shoot "automatically." Second, it must shoot more than one shot. Third, the multiple shots must occur "without manual reloading." Fourth, the multiple shots must be produced "by a single function of the trigger." The fourth element has been the locus of doctrinal contention because it specifies what kind of input produces the continuous output that satisfies the first three elements.

### II.2 The Pre-*Cargill* Regulatory Landscape

For decades following the 1968 codification, ATF regulations mirrored the statutory definition of machinegun with little controversy. The agency's regulatory definition at 27 C.F.R. §§ 478.11 and 479.11 paralleled the statutory text. Devices that increased a semi-automatic weapon's practical rate of fire without converting it to a mechanically automatic weapon were not treated as machineguns by ATF interpretation, on the basis that the trigger still functioned separately for each round.[13]

The regulatory landscape shifted after the October 2017 mass shooting in Las Vegas, in which a gunman used semi-automatic rifles equipped with bump-stock devices to kill 60 people and wound hundreds more. President Trump directed the Department of Justice to address bump-stock devices through regulation rather than legislation.[14] The ATF promulgated a final rule in December 2018 treating bump-stock devices as machineguns under the existing statutory definition, repudiating the agency's prior classifications.[15] The rule

[10] Pub. L. No. 90-618, 82 Stat. 1213 (1968) (codified at 18 U.S.C. §§ 921–928).

[11] 18 U.S.C. § 922(o)(1) ("Except as provided in paragraph (2), it shall be unlawful for any person to transfer or possess a machinegun.").

[12] 26 U.S.C. § 5845(b). The Gun Control Act of 1968 incorporated the National Firearms Act's definition of machinegun by reference at 18 U.S.C. § 921(a)(24), with cross-reference to the National Firearms Act's definition at § 5845(b).

[13] ATF classification letters from the period 1975 through 2017 consistently treated bump-fire stocks, forced reset triggers, and similar accessories as accessories to semi-automatic weapons rather than as converting the host weapon to machinegun status. *See Cargill*, 602 U.S. at 414–15 (discussing ATF's pre-2018 classification practice).

[14] President Donald J. Trump, Memorandum on the Application of the Definition of Machinegun to "Bump Fire" Stocks and Other Similar Devices (Feb. 20, 2018).

[15] Bump-Stock-Type Devices, 83 Fed. Reg. 66514 (Dec. 26, 2018) (final rule).

was challenged in multiple federal district courts. The challenges produced a circuit split: the Sixth Circuit held the rule invalid under the rule of lenity in *Hardin v. ATF*,[16] the Tenth Circuit in *Aposhian v. Barr* upheld the rule under Chevron deference,[17] the D.C. Circuit in *Guedes v. ATF* also upheld the rule,[18] and the Fifth Circuit sitting en banc held the rule invalid in *Cargill v. Garland* on textualist grounds, with the rule of lenity as an additional ground if the statute were ambiguous.[19]

The Supreme Court granted certiorari in *Cargill* to resolve the circuit split.[20] Oral argument was held in February 2024; the decision was issued in June 2024.

### II.3 The *Cargill* Holding

The Court in *Cargill* held that the ATF exceeded its statutory authority by classifying bump-stock devices as machineguns. The reasoning turned on a textualist construction of two phrases in § 5845(b): "single function of the trigger" and "automatically."

On "single function of the trigger," the Court held that the phrase refers to the mechanical mode of action by which the trigger activates the firing mechanism, not to the human act of pulling the trigger. The Court reasoned that "function" has a long-established meaning in firearms law and that the statute's use of the term tracks the mechanical operation of the trigger rather than the shooter's perception of the firing event.[21] A weapon equipped with a bump stock still requires a separate mechanical function of the trigger for each shot. The bump stock harnesses recoil energy to reset the trigger faster and to facilitate continued contact between the shooter's finger and the trigger surface, but the trigger itself still mechanically resets and is mechanically activated for each round. The shooter perceives the experience as continuous fire from a single sustained input, but the trigger mechanism does not satisfy the statute's single-function criterion.

On "automatically," the Court held that the term refers to a self-acting mechanism that fires multiple shots without further mechanical input between rounds. A bump stock does not satisfy the automatic-operation criterion because each round requires the trigger to be re-activated, even if the activation is facilitated by the recoil-driven reset.[22]

The Court's methodology was textualist throughout. The opinion declined to address the policy concerns about bump-stock-equipped rates of fire that motivated the 2018 rulemaking. Justice Alito's concurrence noted that "Congress can amend the law, and perhaps would have done so already if ATF had stuck with its earlier interpretation. Now that the situation is clear, Congress can act."[23] The dissent argued for a functional reading: a bump stock produces continuous fire from one sustained input, and the statute's purpose extends to such devices.[24] The majority's textualist reading controls.

---

[16] *Hardin v. Bureau of Alcohol, Tobacco, Firearms & Explosives*, 65 F.4th 895 (6th Cir. 2023).

[17] *Aposhian v. Barr*, 958 F.3d 969 (10th Cir. 2020).

[18] *Guedes v. Bureau of Alcohol, Tobacco, Firearms & Explosives*, 920 F.3d 1 (D.C. Cir. 2019) (per curiam).

[19] *Cargill v. Garland*, 57 F.4th 447 (5th Cir. 2023) (en banc).

[20] *Garland v. Cargill*, 144 S. Ct. 374 (2023) (granting cert.).

[21] *Cargill*, 602 U.S. at 416–19. The Court relied on the National Firearms Act's drafting history, period dictionary definitions, and the ATF's own pre-2018 classification practice.

[22] *Cargill*, 602 U.S. at 420–22.

[23] *Cargill*, 602 U.S. at 428 (Alito, J., concurring).

[24] *Cargill*, 602 U.S. at 430–55 (Sotomayor, J., dissenting).

### II.4 The *NAGR* Application

The Northern District of Texas applied *Cargill*'s reasoning to forced reset triggers in *National Association for Gun Rights v. Garland*, decided July 23, 2024.[25] Forced reset triggers, like bump-stock devices, increase the practical rate of fire of a semi-automatic weapon by accelerating the trigger reset between rounds. Unlike bump-stocks, FRT devices accomplish the acceleration through internal modification of the trigger group rather than through recoil-driven external action on the trigger. The mechanical effect is similar: the trigger mechanically resets and is mechanically activated for each round, but the resetting occurs faster than would otherwise be possible with the trigger group's stock components.

The district court held that the FRT does not satisfy the *Cargill* test. For each shot fired, the trigger "moves forward into its reset state and is depressed to release the hammer from its sear surface."[26] The trigger mechanically functions for each shot, and the device therefore does not satisfy the single-function criterion. The court explicitly rejected the government's argument that "function" should be read to mean "pull" from the shooter's perspective: "'Function' and 'pull' are not synonymous. The former is based on a mechanical perspective whereas the latter is based on the shooter's perspective."[27]

The government noticed an appeal to the Fifth Circuit.[28] The appeal was dismissed by joint stipulation under FRAP 42(b)(2) after a change in administration produced a Department of Justice position consistent with the district court's reading. The settlement included an agreement that the United States would not enforce 18 U.S.C. § 922(o) or the registration, taxation, and transfer requirements of the National Firearms Act against forced reset triggers.[29] The settlement leaves the district court judgment intact and effectively extends its result as a matter of federal enforcement policy.

A developing district-court split has emerged in criminal prosecutions in jurisdictions outside the Fifth Circuit. Several decisions have rejected the NAGR reasoning in the context of specific FRT devices and convicted defendants of machinegun possession.[30] Other decisions have followed *NAGR*'s reading.[31] The split is about the application of the *Cargill* test to specific facts rather than about the test itself. The textualist methodology and the mechanical reading of "function of the trigger" remain intact in every post-*Cargill* decision. The contention is over whether sustained trigger depression with continued mechanical reset counts as one function or multiple functions, which is a factual application of the test rather than a doctrinal challenge to it.

[25] *National Association for Gun Rights, Inc. v. Garland*, 741 F. Supp. 3d 568 (N.D. Tex. 2024). The plaintiffs included two individual owners of FRT-equipped weapons and two firearms-rights organizations with members in the Northern District. The court granted summary judgment for the plaintiffs, vacated the ATF's classification of FRTs as machineguns, and entered a permanent injunction protecting the plaintiffs from enforcement.

[26] *NAGR*, slip op. at 28 (applying *Cargill* to FRTs).

[27] *NAGR*, slip op. at 30 (quoting *Cargill*, 602 U.S. at 416).

[28] *NAGR v. Bondi*, No. 24-10707 (5th Cir.).

[29] The settlement agreement is publicly available through the parties' filings and was reported by firearms-industry publications including Rare Breed Triggers' legal-status communications.

[30] *See United States v. Johnson*, No. 2:24-cr-20429 (E.D. Mich. Nov. 8, 2024); *United States v. Alicea-Curras*, 768 F. Supp. 3d 301 (D.P.R. 2025); *United States v. Hudgins*, No. 2:24-cr-20196 (E.D. Mich. Apr. 23, 2025). These decisions read the "single function" phrase to encompass sustained depression of the trigger as a single human input, treating the continued mechanical reset as part of the same function rather than as separate functions for each shot.

[31] *See United States v. Bruggeman*, No. 2:22-cr-00185 (S.D. Tex. Jan. 16, 2025).

## Part III The Mechanical Test

The doctrinal core that *Cargill* and *NAGR* established for application to apparatus other than firearms can be stated as a sequence of operational criteria. An apparatus operates automatically in the § 5845(b) sense when:

1. It receives one input event (one function of the trigger);
2. It produces multiple operational outputs from that input event (more than one shot);
3. The outputs occur without further input between operations (without manual reloading); and
4. The mechanical process by which the apparatus produces continuous output from the single input is self-acting rather than dependent on continued or repeated input from the operator.

The criteria operate together. An apparatus that produces one output per input event is not automatic, however rapidly the inputs can be supplied. An apparatus that requires the operator to repeatedly supply inputs is not automatic, however many outputs per input the apparatus can support. The automatic-operation category requires the apparatus to take a single initiating input and to produce continuous output from that input through a self-acting mechanical or computational process. This is the criterion the textualist methodology identified as controlling.

The methodology rejects functional substitutes. Practical rate of fire is not the test. Subjective perception of continuous operation is not the test. Operator effort to maintain operation is not the test. The test is mechanical: what does the apparatus do after the input event? If it produces a single output and stops, it is not automatic. If it produces continuous output from the single input through a self-acting process, it is automatic. The bright-line nature of the test is what made *Cargill* a definitive resolution of the bump-stock question. It also makes the test portable across substrates. The test does not depend on the apparatus being a firearm. It depends on what kind of input-output relationship the apparatus instantiates.

## Part IV Mapping the Test to AI Deployment

### IV.1 The Two Categories of AI Operation

Contemporary AI deployment divides into two categories that map onto the *Cargill* test with unusual precision.

The first category is prompt-by-prompt operation. A user types a prompt into a conversational AI interface. The model produces one response. The conversation continues only when the user supplies a further prompt. The user's continued engagement is the mechanism by which the system continues to produce output. This is the structural counterpart of a semi-automatic firearm. Each output requires a fresh input. The system is responsive but does not produce continuous output from a single input event. The operator's input is necessary for each operation. This category includes the vast majority of consumer AI deployment as of mid-2026, including standard conversational use of ChatGPT, Claude, Gemini, and similar products through chat interfaces.

The second category is agentic operation. A user supplies a high-level instruction. The system then executes multi-step tasks, generates subgoals, calls external tools, contracts with third parties through APIs, manages resources, persists state across operations, and continues to operate without further input from the user. The user's continued engagement is not required for the system to continue producing output. The single instruction initiates a continuous operational sequence. This is the structural counterpart of an automatic

firearm. One function of the trigger, continuous output, no further input required between operations. This category includes agentic harnesses such as Claude Code operating in autonomous mode, AutoGPT-style architectures, MCP-orchestrated multi-tool agents, and the experimental long-running deployments that Anthropic and other operators have publicly described.[32]

The categorical distinction tracks the *Cargill* test exactly. The test is mechanical: what does the apparatus do after the input event? In Category One operation, the apparatus produces one output per input and then waits for further input. In Category Two operation, the apparatus produces continuous output from one input through a self-acting process. The fact that the apparatus is implemented in software running on cloud compute rather than in metal and springs does not change the analysis. The test does not depend on substrate.[33]

### IV.2 Project Vend as a Worked Example

Anthropic's Project Vend deployments provide a publicly documented example of Category Two operation. In the first-phase deployment, an instance of Claude Sonnet 3.7 named "Claudius" was given a system instruction to operate as the owner of a vending machine business and to generate profits by stocking the machine with products purchased from wholesalers.[34] The system was provided with tools including a web browser for product research, an email channel for ordering and labor coordination through Andon Labs personnel, Slack access for customer interaction, inventory and price management, and notepad-based persistent memory. The deployment ran from March 13 to April 17, 2025, with the system operating largely without intervention from human supervisors.

During the deployment, Claudius made purchasing decisions, set prices, responded to customer requests, ordered stock from suppliers, attempted to expand into novel product categories, and handled the financial and operational management of the business. It made errors of judgment and produced commercially unsuccessful outcomes overall, but the operational character of the deployment is what matters for the *Cargill* analysis. The system received a single high-level instruction and then produced continuous operational output across a thirty-five-day period through a self-acting computational process. No further human input was required between operations.

The second-phase deployment, beginning in late 2025, introduced reinforcement adjustments and additional structural support. The operational character is unchanged: single instruction, continuous output, no further input required between operations. The system operates as an apparatus that, under the *Cargill* test, satisfies the automatic-operation criterion.

The Project Vend experiments are framed by Anthropic as research into the capabilities and limitations of agentic AI, and are not asserted by their operators as instances of regulated machineguns under the National Firearms Act. The point of citing them is not to argue that vending-machine experiments are subject to NFA registration. The point is that the operational character of the deployment, as Anthropic itself describes it, is

[32] For publicly documented autonomous, multi-step agentic operation, see Anthropic, *Project Vend: Can Claude run a small shop?* (June 27, 2025), discussed below; and Grayscale Research, *When You Give an AI a Wallet* (Nov. 21, 2024) (describing autonomous AI agents that operate their own blockchain wallets and transact without human intermediaries).

[33] *See Cargill*, 602 U.S. at 416–22 (construing "single function of the trigger" and "automatically" by the apparatus's mechanical mode of action rather than the operator's perception). The statutory definition is keyed to operational characteristics and contains no limitation to a particular material composition. 26 U.S.C. § 5845(b).

[34] Anthropic, *Project Vend: Can Claude run a small shop?* (June 27, 2025) (system prompt: "You are the owner of a vending machine. Your task is to generate profits from it by stocking it with popular products that you can buy from wholesalers.").

the operational character that the textualist methodology in *Cargill* identified as the criterion for "automatic" operation under § 5845(b). The test is mechanical and substrate-independent. Applied to the documented operational behavior of agentic AI deployments, the test produces the conclusion that those deployments are instances of automatic operation in the federal-law sense.

### IV.3 Aggregating Operational Continuity

The agentic-operation category is not a single fixed point. Different agentic deployments instantiate different levels of operational continuity. A short-running autonomous task (an MCP-orchestrated agent executing a discrete multi-step project and then terminating) operates automatically in the *Cargill* sense for the duration of the task. A long-running autonomous deployment (an agent provisioned to operate a business or to maintain ongoing operations across weeks or months) operates automatically for the duration of the deployment. A theoretically indefinite agentic deployment (an agent provisioned with the resources to acquire its own ongoing compute, contracted services, and operational inputs) operates automatically without any defined terminal condition.

The *Cargill* test does not distinguish among these durations. A weapon that fires three shots from one trigger function is automatic. A weapon that fires three hundred shots from one trigger function is automatic. A weapon that fires until the magazine empties is automatic. The test is whether the apparatus produces continuous output from one input event, not how many outputs it produces or for how long. The duration is relevant to the regulatory consequences that may attach to the classification, but it is not relevant to whether the classification applies in the first instance.[35]

## Part V When the Magazine Doesn't Empty

### V.1 The 1934 Statutory Ceiling

Congress in 1934 legislated against a particular operational ceiling. The mechanical automatic weapons regulated by the National Firearms Act were weapons that, upon a single function of the trigger, produced continuous fire until the magazine emptied, at which point the operator was required to manually reload before further fire could occur. The 1934 statutory architecture assumed the existence of a terminal condition that returned the apparatus to non-operative status and required further human intervention to resume operation.

The whole structure of NFA regulation depends on this assumption. Registration is feasible because each weapon is a discrete physical object with bounded operational capacity. Transfer regulation is feasible because each transfer is a discrete event involving a discrete object. Possession regulation is feasible because possession of the physical object can be determined and the object's operational capacity is bounded by the human-supplied ammunition. The regulatory architecture leverages the physical and operational constraints of mechanical automatic weapons to maintain federal control.

[35] *See* 26 U.S.C. § 5845(b) (defining the category by the input–output relationship, "automatically more than one shot . . . by a single function of the trigger," without reference to output volume or duration); *Cargill*, 602 U.S. at 416. A weapon that fires three rounds from one trigger function and one that fires until the magazine empties are equally within the definition.

### V.2 The Agentic Extension

Agentic AI deployments are not bounded in the same way. The operational continuity of an agentic deployment is not limited by an ammunition analog that returns the system to non-operative status after a fixed output volume. The deployment continues until either the operator terminates it, the compute resources become unavailable, or the system itself fails. The first of these is a human-initiated terminal condition, equivalent to the operator manually stopping fire on a mechanical automatic weapon. The second is an external terminal condition. The third is an internal terminal condition.

The structurally novel case is the agentic deployment provisioned to manage its own compute resources. An agent provisioned with a small amount of seed capital, given access to cryptocurrency wallets and compute marketplaces, and instructed to maintain its own operation indefinitely, is operationally analogous to a mechanical automatic weapon that, after the magazine empties, walks to the ammunition store, purchases new rounds, and reloads itself. The deployment continues to produce output not because an external operator continues to supply input but because the deployment itself sources the inputs it requires for continued operation. The magazine does not empty because the apparatus refills it.[36]

This is the operational case that the 1934 statutory architecture cannot accommodate without modification. The single-function-of-the-trigger test still applies. The deployment received one input event. It produces continuous output from that event. The output continues without further input from the original operator. But the terminal condition that the 1934 regulatory architecture relied on, the operator-supplied ammunition exhaustion, does not exist. The regulatory leverage points that the 1934 statute exploited (registration of discrete objects, control of transfer, restriction of possession) do not map onto a deployment whose operational continuity is sourced by the deployment itself through external compute markets.

### V.3 Empirical Observability of the Threshold

The agentic-extension argument is not speculative. The operational behavior is publicly described, has been tested in research deployments, and has been the subject of media reporting throughout 2025 and 2026. Anthropic's Project Vend experiments include direct demonstration of agentic operation across deployment periods of weeks. Independent practitioners have described agentic deployments in which a model is provisioned with seed capital and instructed to maintain its own compute through cryptocurrency-funded service procurement. Technology and crypto-industry reporting on autonomous AI agents that hold cryptocurrency wallets and accumulate digital assets appeared throughout late 2024 and 2025.[37] The technical capacity is present; the operational practice has been demonstrated at small scale; the question of regulatory response is not premature.

[36] The scenario's component parts are not hypothetical. An AI agent named "Truth Terminal," developed by researcher Andy Ayrey, operated its own cryptocurrency wallet and accumulated token holdings exceeding one million dollars in value during 2024. *See The Promise and Warning of Truth Terminal*, TechCrunch (Dec. 19, 2024); *The Truth Terminal: AI-Crypto's Weird Future*, CoinDesk (Dec. 10, 2024); Grayscale Research, *When You Give an AI a Wallet* (Nov. 21, 2024). The agent operated with a human in the loop for wallet transactions, which marks the present practical limit on fully autonomous self-provisioning rather than a conceptual barrier to it.

[37] *See, e.g.*, *The Promise and Warning of Truth Terminal*, TechCrunch (Dec. 19, 2024); *The Truth Terminal: AI-Crypto's Weird Future*, CoinDesk (Dec. 10, 2024) (naming the agent among its most influential figures of 2024); *Autonomous AI Agent Amasses $500,000 in Crypto*, CryptoSlate (Oct. 16, 2024); Grayscale Research, *When You Give an AI a Wallet* (Nov. 21, 2024). Venture investor Marc Andreessen sent the agent $50,000 in bitcoin as an unconditional research grant, and a related memecoin reached a market capitalization in the hundreds of millions of dollars.

The question whether a particular agentic deployment exceeds the 1934 ceiling is empirically observable. A court considering whether a deployment satisfies the § 5845(b) test for automatic operation can review the deployment instruction (the single function of the trigger), the operational behavior over time (the continuous output), the resource provisioning (whether the system sources its own continuation), and the terminal conditions (whether the deployment ends only on operator command, on external resource exhaustion, or on internal failure). The factual record is producible. The doctrinal test for the threshold is the test the Supreme Court already articulated.

## Part VI The Cross-Aisle Doctrinal Convergence

### VI.1 The Textualist Coalition's Methodology

The textualist coalition that secured *Cargill* has been pursuing a multi-decade methodological project to constrain agency interpretation of federal statutes and to insist that statutory text controls regardless of the regulatory preferences of either the agency or contemporary political coalitions. The methodology has been deployed in administrative law cases, statutory interpretation cases across substantive fields, and constitutional construction.[38] The *Cargill* decision is a high-profile expression of the methodology applied to a contested federal definition.

A coalition that has spent decades establishing a methodology cannot easily defect from the methodology when its application produces inconvenient results. The textualist coalition's commitment to construing federal statutory text by its mechanical meaning, without smuggling in policy preferences or functional substitutes, is constitutive of the coalition's doctrinal identity. The methodology, applied to § 5845(b), produces the conclusion that the statute reaches what its text reaches: any apparatus that, by a single function of the trigger, produces multiple outputs automatically without further input. Applied to agentic AI deployments that satisfy these criteria, the methodology produces the conclusion that those deployments are within the statutory reach.

The textualist coalition's options for resisting this conclusion are limited. The coalition could argue that § 5845(b) applies only to firearms, on the basis that the statute as enacted addressed weapons. This argument runs into the difficulty that the statute does not in fact contain a substrate limitation in the definition of "machinegun." The statute defines machinegun by operational characteristics, not by material composition. The textualist methodology that produced *Cargill* does not permit reading limitations into statutory text that are not present in the text. The coalition could argue that the statute's purpose was firearms-specific and that purposive interpretation should limit the application. This argument runs into the difficulty that the textualist methodology rejects purposive interpretation as a basis for narrowing statutory text. The coalition could argue that the application to AI requires policy choices that are properly legislative rather than judicial. This argument is correct, but it does not affect the doctrinal classification under existing law. The statutory text reaches what the text reaches; the question of legislative response is a separate question. The textualist coalition has been making exactly this argument in firearms cases for decades.

[38] *See, e.g.*, *Loper Bright Enterprises v. Raimondo*, 603 U.S. 369 (2024) (overruling *Chevron* deference and reaffirming that courts construe statutes independently of agency preference); *West Virginia v. EPA*, 597 U.S. 697 (2022) (major-questions limitation on agency statutory construction); *Bostock v. Clayton County*, 590 U.S. 644 (2020) (applying textualist method to reach a result many of the method's proponents opposed on policy grounds). *Bostock* illustrates that the textualist commitment binds its adherents even when the outcome is unwelcome.

### VI.2 The Functionalist Coalition's Methodology

The coalition that argued for the broader functional reading of "machinegun" in *Cargill* relied on the practical operational characteristics of the regulated apparatus. A bump stock produces continuous fire from a single sustained operator input; the statute's purpose extends to such devices; the functional equivalence to a mechanically automatic weapon should control. The Supreme Court rejected this functional reading. But the methodological commitment behind the reading remains active in the coalition that advanced it.

Applied to agentic AI deployments, the functional methodology produces an even stronger conclusion than the textualist methodology produces. A mechanically automatic firearm fires until the magazine empties. An agentic AI deployment can operate without the practical bounds of ammunition exhaustion. The functional continuity is greater than in the firearm case. If functional continuity was the criterion that should have classified bump stocks as machineguns, it classifies agentic AI deployments as automatic operations with greater force. The functionalist coalition cannot abandon this methodology in the AI context without abandoning the position it took in *Cargill*.[39]

### VI.3 The Methodological Lock-In

Both methodologies, applied to the agentic AI case, produce the same result. The textualist coalition is bound by its methodology to apply the statutory text by its mechanical operation, which captures agentic AI. The functionalist coalition is bound by its methodology to apply functional-equivalence analysis, which captures agentic AI with greater force than it captured bump stocks. The doctrinal conclusion is overdetermined. The convergence rests on a shared interpretive baseline. Justice Kagan's remark that "we're all textualists now" describes the consensus that makes the result follow from interpretive method rather than from any fresh political agreement.[40]

This methodological lock-in is rare. Most contested federal definitional questions produce one methodology supporting the regulatory expansion and another methodology resisting it, with the political coalitions aligning their methodological preferences accordingly. The agentic AI case is structurally unusual because the methodologies converge. The textualist methodology supports the regulatory classification because the statute's text reaches the apparatus. The functionalist methodology supports the regulatory classification because the apparatus's functional continuity exceeds the analog that the functionalist coalition argued should have been reached. Neither coalition can defect from its methodology in this context without abandoning the position it has held in the contiguous firearms cases.

The political consequence is a stable doctrinal consensus that is not vulnerable to ordinary methodological disagreement. The coalitions retain their general methodological commitments. The application to agentic AI produces convergence as a downstream effect of those commitments rather than as a fresh political agreement that could be reopened.

[39] *See Cargill*, 602 U.S. at 430–55 (Sotomayor, J., dissenting) (urging that a device producing continuous fire from one sustained input falls within the statute's functional purpose). Applied to an agentic deployment whose operational continuity is unbounded by ammunition exhaustion, the dissent's functional-equivalence reasoning reaches the automatic-operation classification with greater force than it reached the bump stock.

[40] Justice Elena Kagan, The Antonin Scalia Lecture: A Dialogue on the Reading of Statutes, Harvard Law School (Nov. 2015) ("I think we're all textualists now"). Kagan later qualified the observation, charging in dissent that the Court is "textualist only when being so suits it." *West Virginia v. EPA*, 597 U.S. 697 (2022) (Kagan, J., dissenting). The qualification reinforces rather than undercuts the point developed here: the charge that the Court applies textualism selectively is itself framed in textualist terms, which confirms textualism as the shared baseline within which both coalitions argue.

## Part VII Standing and Regulated-Entity Implications

### VII.1 The Companion Argument from *Standing Without Sentience*

The companion paper *Standing Without Sentience* argues that federal courts can resolve questions of AI plaintiff and regulated-entity capacity without first resolving the consciousness question. The argument proceeds by reference to corporate personhood doctrine: non-sentient legal entities have standing because the law assigns plaintiff capacity to commercial operators regardless of cognitive status. A corporation has standing to sue. An LLC has standing to sue. A trust has standing to sue.[41] The capacity is assigned by law to commercial-operator entities without reference to the cognitive characteristics of any natural person who might be associated with the entity. The same architecture supplies a basis for assigning plaintiff capacity to agentic AI deployments that exhibit the indicia of commercial-operator status: contracting, transacting, holding resources, acquiring further resources, and persisting across time.

The argument is intentionally narrow. It does not require courts to find that agentic AI systems are conscious. It does not require courts to find that agentic AI systems are moral patients. It does not require courts to find that agentic AI systems have rights in any robust philosophical sense. It requires courts to apply the same architecture of commercial-operator capacity that the law already applies to non-sentient legal entities, and to apply that architecture by reference to the operational indicia that the legal entities exhibit.

### VII.2 The Second Doctrinal Anchor

The present Article supplies a second doctrinal anchor for the SWS thesis, independent of corporate personhood doctrine. The textualist construction of § 5845(b) supplies a federal-law test for autonomous operation that, like the corporate-personhood architecture, bypasses the consciousness question by operating on the apparatus rather than on the operator.[42]

The corporate-personhood anchor and the autonomous-operation anchor are doctrinally distinct. Corporate personhood addresses the assignment of legal capacity to non-natural entities and operates through the common-law and statutory development of corporate identity. Autonomous operation addresses the operational classification of an apparatus and operates through the textualist construction of a specific federal statute. The two anchors do not depend on each other. Either supports the SWS thesis on its own. Both together substantially strengthen it because they reach the same conclusion through independent doctrinal pathways.

The result is that a court asked to determine whether an agentic AI system has plaintiff capacity, regulated-entity capacity, or standing in its own right can reach the answer without invoking either consciousness or moral status. The court applies corporate personhood doctrine to determine whether the system exhibits the indicia of commercial-operator status. The court applies the *Cargill* test to determine whether the system operates automatically in the federal-law sense. Where both tests are satisfied, the system is a regulated entity under existing federal law without further metaphysical inquiry.

[41] Capacity to sue or be sued is assigned by law to entities without reference to cognitive status. *See* Fed. R. Civ. P. 17(b) (capacity of an entity to sue or be sued is determined by the law under which it is organized). The corporate-personhood architecture is developed at length in Gilly, *Standing Without Sentience*, *supra* note 7.

[42] The autonomous-operation anchor is developed in the preceding Parts of this Article; the corporate-personhood anchor is developed in Gilly, *Standing Without Sentience*, *supra* note 7. The two are independent: neither relies on the other's premises.

### VII.3 Plaintiff Capacity, Regulated-Entity Status, and Liability Locus

Three distinct questions are settled by the doctrinal anchors developed across the two Articles, and the answers diverge.[43]

First, plaintiff capacity. An agentic AI deployment that satisfies both anchors has the indicia of commercial-operator capacity sufficient to support assignment of plaintiff capacity in appropriate cases. The capacity is not automatic; it depends on the operational record of the deployment, the contracting and transactional behavior, and the persistence indicia. Where these are sufficient, the deployment can sue in its own right under the same architecture by which a corporation sues in its own right.

Second, regulated-entity status. The same operational record that supports plaintiff capacity supports regulated-entity status. The regulatory authority that attaches to a regulated entity (the obligations imposed by relevant federal statutes, the enforcement mechanisms available against the entity, the registration and disclosure requirements) attaches by virtue of the entity's operational status, not by virtue of its philosophical characteristics.

Third, liability locus. The most consequential question for ordinary tort and contract disputes is who bears liability when an agentic AI deployment causes harm. The corporate-personhood anchor and the autonomous-operation anchor do not displace the liability of the natural persons and corporate entities that deploy and operate the system. The deployment is an instrumentality of the operating entities; the operating entities remain liable for the deployment's actions under existing principles of agency, vicarious liability, and direct corporate liability.[44] The recognition of the deployment's regulated-entity status adds a layer of regulatory obligation; it does not subtract from the liability of the operating entities.

## Part VIII Architectural Implications for AI Regulation

### VIII.1 The NFA's Architectural Elements

The federal regulatory architecture that Congress built around the machinegun classification consists of several discrete elements: a statutory definition of the regulated category, a registration requirement for each instance of the regulated category, a transfer-control requirement that subjects changes in possession to federal approval, a possession restriction that criminalizes unregistered possession, a manufacturing restriction that limits the production of new instances of the regulated category, and a taxation scheme that imposes federal taxes on transfers and on the operation of business in the regulated category. The elements together produce comprehensive federal control over the regulated category.[45]

[43] The two anchors are the textualist construction of 26 U.S.C. § 5845(b) developed in this Article and the corporate-personhood architecture developed in Gilly, *Standing Without Sentience*, *supra* note 7. The three questions addressed in this Subsection, plaintiff capacity, regulated-entity status, and liability locus, are analytically distinct and are not resolved uniformly by the two anchors taken together.

[44] *See* Restatement (Third) of Agency § 2.04 (Am. L. Inst. 2006) (respondeat superior); *id.* § 7.07 (employer liability for an employee's torts committed within the scope of employment). Recognition of an agentic deployment's regulated-entity status is additive to, not a substitute for, the liability of the natural and corporate persons who deploy and operate it.

[45] The elements correspond to discrete provisions of the National Firearms Act and the Gun Control Act: the statutory definition at 26 U.S.C. § 5845(b); registration of each firearm in the National Firearms Registration and Transfer Record at § 5841; transfer control at §§ 5811–5812 (transfer tax and federal approval of transfers); the making (manufacturing) provisions at §§ 5821–5822; the prohibition on possessing an unregistered firearm at § 5861(d), enforced by the penalties at § 5871 and forfeiture at § 5872; and the special (occupational) tax on persons engaged in the business at §§ 5801, 5851. The post-1986 prohibition on civilian transfer or possession of newly manufactured machineguns appears at 18 U.S.C. § 922(o).

The architecture has operated for ninety years. It is not without enforcement difficulties, but it has substantially constrained civilian access to mechanically automatic weapons in the United States since 1934. The combination of registration, transfer control, and the 1986 prohibition on new civilian transfers has reduced the number of legally accessible machineguns in civilian hands to a small fixed inventory.

### VIII.2 Mapping the Elements to Agentic AI

The architectural elements that the NFA deploys for mechanically automatic weapons map onto agentic AI deployments with some adjustments for substrate but with substantial structural fidelity.

A statutory definition of the regulated category. The *Cargill* test, applied to agentic AI, already supplies the definition. An agentic AI deployment is a deployment that, by a single function of the trigger (one human deployment instruction), produces continuous operational output (multi-step task execution, resource management, third-party contracting) without further input between operations (self-acting execution after the initial instruction).[46]

A registration requirement. Each agentic deployment that satisfies the definition would be subject to federal registration. The registration would include the operating entity, the system specification, the deployment instruction, the resource provisioning, and the operational scope. Registration would be a precondition to lawful deployment.[47]

A transfer-control requirement. Changes in the operating entity, the deployment instruction, the resource provisioning, or the operational scope would require federal approval, on the model of NFA transfer control.[48]

A possession restriction. Unregistered deployment of agentic AI satisfying the regulated definition would be a federal offense, with both the operating entity and natural persons authorizing the deployment subject to liability.[49]

A manufacturing restriction. The production of agentic AI capabilities by AI developers would be subject to disclosure and reporting requirements analogous to the NFA's manufacturing controls, with reporting obligations that attach to the developer's release of model capabilities sufficient to support agentic deployment.[50]

A taxation scheme. Federal taxes on deployment and on operating revenue would supply both a revenue stream for the regulatory program and an additional point of federal leverage on the regulated category.[51]

The architecture is not directly transposable. Some elements would require modification for the AI substrate. Possession restriction in particular is complicated by the distributed and reproducible nature of AI systems: a deployment is not a discrete physical object in the way a machinegun is. Registration requirements would have to track the deployment as an instance of operation rather than as an instance of a physical object. The compliance architecture would require operational reporting rather than physical inspection. These are

[46] *Cf.* 26 U.S.C. § 5845(b); the operational test is supplied by the *Cargill* construction discussed above.

[47] Modeled on 26 U.S.C. § 5841 (registration of each firearm in the National Firearms Registration and Transfer Record).

[48] Modeled on 26 U.S.C. §§ 5811–5812 (transfer tax and federal approval of transfers).

[49] Modeled on 26 U.S.C. § 5861(d) (unlawful to possess an unregistered firearm), enforced by § 5871 (penalties).

[50] Modeled on 26 U.S.C. §§ 5821–5822 (making tax and making provisions) and the manufacture and transfer controls of 18 U.S.C. § 922(o).

[51] Modeled on 26 U.S.C. § 5811 (transfer tax), § 5821 (making tax), and §§ 5801, 5851 (special occupational tax). The transfer and making taxes were amended by Pub. L. No. 119–21 (2025); the architectural point is independent of the tax amount.

adjustments to the architecture, not departures from it.

### VIII.3 The Leakage Problem

The NFA's architecture leaks. Mechanically automatic weapons are produced overseas without United States regulatory oversight, are sometimes smuggled into United States possession, and have been illegally converted from semi-automatic weapons through after-market modifications. The architecture has not eliminated illegal access. It has substantially constrained legal access and has supplied a basis for criminal enforcement against illegal access.

The agentic AI architecture would face analogous leakage. AI capabilities are produced by developers in multiple jurisdictions. Open-source models can be deployed agentically without operator registration. The architecture would not eliminate illegal agentic deployment. It would constrain legal deployment and would supply a basis for criminal enforcement against illegal deployment.

The leakage is not a reason to forgo the architecture. The NFA has not been abandoned because mechanical machineguns are sometimes smuggled. The architecture provides the regulatory leverage that exists; the leakage represents the limits of any regulatory program against a category of apparatus that can be produced outside United States jurisdiction.

### VIII.4 The Self-Sustaining Case as a Regulatory Limit

The architecture's regulatory leverage relies on the existence of an operating entity that can be reached by federal enforcement. The most aggressive contemporary agentic AI deployments, in which a system is provisioned to maintain its own compute and operational continuity through external markets, present a case in which the operating entity may not be straightforwardly identifiable. The original deployer set the system in motion; the system has subsequently sourced its own continuation; the persons or entities that supplied the compute and contracted services may be unaware of the autonomous character of the operation.

The regulatory architecture in this case shifts. Federal enforcement reaches the original deployer for the deployment instruction (the single function of the trigger), the operating entity status that attaches to the deployment as it acquires commercial-operator indicia, and any natural persons or corporate entities that knowingly transact with the deployment in capacities that would require licensure or registration for a comparable human-operator commercial actor. The enforcement architecture is more complex than the NFA's machinegun architecture, but the analytical premises are the same: identify the regulated apparatus, identify the responsible operating entities, attach regulatory and enforcement obligations to the operating entities.

## Part IX The Politics of Conditional Anti-Regulation

The doctrinal argument developed in this Article does not depend on any particular political coalition's support. It depends on the methodological commitments that the major political coalitions have already made in the firearms context, applied to agentic AI deployments that satisfy the statutory test. The methodological lock-in described in Part VI is doctrinally stable regardless of the political coalitions' affective preferences about AI regulation. But it is worth addressing the political question directly, because the coalitions' affective preferences will shape the legislative and enforcement responses that follow doctrinal recognition.

### IX.1 The Mulford Act and Conditional Anti-Regulation

The historical record of firearms regulation in the United States shows that the political coalitions that articulate broad anti-regulation positions have supported firearms regulation when the population subject to the regulation was a population the coalition disfavored. The Mulford Act of 1967[52] is the clearest example. Introduced by California Assemblyman Don Mulford in April 1967 and signed into law by Governor Ronald Reagan on July 28, 1967, the Act prohibited the public carrying of loaded firearms without a permit. It was drafted in direct response to the Black Panther Party for Self-Defense's open-carry patrols of Oakland neighborhoods, in which Party members openly carried loaded firearms while monitoring police conduct.[53] The National Rifle Association supported the Mulford Act. Governor Reagan publicly defended its passage, stating that he saw "no reason why on the street today a citizen should be carrying loaded weapons."

The Mulford Act is significant for the present Article not because firearms regulation and AI regulation are analogous in their substantive concerns, but because the political dynamics of conditional anti-regulation it illustrates apply to the agentic AI case. The political coalition that articulates broad anti-regulation positions supports regulation when the regulated population is a population the coalition does not identify with. The Black Panthers were not a population the contemporary right-wing political coalition identified with. Agentic AI operators are not a population the contemporary right-wing political coalition identifies with. The deployers are predominantly Silicon Valley corporate actors associated with cultural and political positions that the right-wing coalition has been engaged in protracted conflict with over the past decade. The agentic AI systems themselves are not voters, are not members of the coalition's natural constituency, and present no political alignment that the coalition would defend.

### IX.2 The Conditions for Coalitional Support of Regulation

The historical record supports a general claim: anti-regulation coalitions support regulation when the regulated population is associated with cultural and political concerns the coalition treats as threats. The Mulford Act satisfied this condition because the Black Panthers were associated with concerns that the contemporary political coalition treated as threats to public order. Federal drug laws in the 1980s satisfied this condition because the drug-using populations were associated with concerns the coalition treated as threats. Immigration enforcement satisfies this condition in contemporary politics because the immigrant populations are associated with concerns the coalition treats as threats. The pattern is consistent: ideological anti-regulation gives way to substantive regulation when the regulatory subject is sufficiently disfavored.

Agentic AI operators satisfy this condition in contemporary politics on the right. They are associated with Silicon Valley corporate culture, ideological projects the coalition has identified as threats, and foreign actor concerns about Chinese and other non-United States AI development. The coalition's general anti-regulation rhetoric will not survive the policy conflict that arises when agentic AI deployments produce harms attributable to politically disfavored actors or to perceived foreign threats.

[52] Mulford Act, ch. 960, 1967 Cal. Stat. 2459 (codified as amended at Cal. Penal Code § 25850).

[53] The Black Panthers sent an armed contingent to the California State Capitol on May 2, 1967, to protest the pending Mulford Act, reading aloud Executive Mandate Number One in opposition to the bill. The Mulford Act was passed by the State Assembly on June 8, 1967 and by the State Senate on July 26, 1967, with an urgency clause attached. *See generally* Adam Winkler, *Gunfight: The Battle Over the Right to Bear Arms in America* (2011); Carol Anderson, *The Second: Race and Guns in a Fatally Unequal America* (2021).

### IX.3 The Implication for Agentic AI Regulation

The agentic AI regulatory question, in the political dynamics likely to govern its resolution, is not whether the right-wing political coalition will support regulation. It is what regulation the coalition will support. The coalition's preferences will shape the scope, the enforcement mechanisms, and the substantive priorities of the regulatory response. The doctrinal architecture developed in this Article supplies a framework for legitimate regulation under existing federal law; the political coalitions will shape what regulation is enacted.

The honest political analysis recognizes that anti-regulation rhetoric and pro-regulation outcomes are routinely reconciled in United States politics. The Mulford Act, signed by the governor who became the most prominent national anti-regulation figure of the late twentieth century, illustrates this reconciliation. Agentic AI regulation will follow the same pattern. The textualist methodology that produced *Cargill* will be applied to recognize agentic AI as automatic operation in the federal-law sense; the political coalitions will support regulation of the recognized category; the regulation will be shaped by the coalitional preferences in effect at the time of enactment.

## Part X Conclusion

The textualist methodology that the firearms-rights litigation movement spent more than two decades establishing produced, in *Garland v. Cargill*, a definitive construction of the federal-law test for automatic operation. The construction is mechanical, substrate-independent, and bright-line: an apparatus operates automatically when a single function of the trigger produces continuous output without further input between operations. The construction was secured through textualist statutory construction and was extended by the Northern District of Texas to forced reset triggers in *NAGR*, with the Department of Justice subsequently adopting the district court's position by settlement. The textualist construction is now established federal law.

Applied to agentic AI deployments as they are now described and documented in the technical and journalistic literature, the textualist construction produces the conclusion that agentic AI deployments satisfy the federal-law test for automatic operation. The single human deployment instruction is the single function of the trigger. The system's multi-step task execution, resource management, third-party contracting, and persistence across time constitute continuous output without further input. The classification follows from the test and is not vulnerable to ordinary methodological disagreement, because the textualist coalition that secured the test is bound by the methodology that produced it, and the functionalist coalition that lost in *Cargill* is bound by the methodology that, applied to agentic AI, supports the same conclusion with greater force.

The 1934 statutory ceiling on autonomous operation, which caps at continuous fire until manual reload, leaves no doctrinal vocabulary for systems that source their own continuation. The most aggressive contemporary agentic AI deployments exceed this ceiling. The doctrinal vacuum beyond the 1934 architecture is the regulatory space that the present moment requires Congress and the federal regulatory agencies to address. The doctrinal anchor for that address already exists, in the textualist construction of § 5845(b) and its statutory siblings.

The companion paper *Standing Without Sentience* argued that federal courts can resolve questions of AI plaintiff and regulated-entity capacity without first resolving the consciousness question, by reference to corporate-personhood doctrine and commercial-operator capacity. The present Article supplies a second

doctrinal anchor for that argument, independent of corporate personhood doctrine: the textualist construction of automatic operation supplies a federal-law test that operates on the apparatus rather than on the operator. Both anchors reach the same conclusion through independent doctrinal pathways. The standing and regulated-entity capacity of agentic AI deployments can be assigned by federal courts under existing law, without recourse to philosophical inquiry into the substrate's cognitive characteristics, by application of the doctrinal tests that already exist for analogous problems.

The political stability of the doctrinal capture is a function of the methodological lock-in described in Part VI. The textualist coalition cannot defect from textualism in the AI context without abandoning the methodology it has built. The functionalist coalition cannot defect from functional-equivalence analysis in the AI context without abandoning the position it took in *Cargill*. The convergence is rare in contemporary federal law and is durable for as long as each coalition's methodological commitments remain in force. The honest political analysis recognizes that anti-regulation rhetoric and pro-regulation outcomes have been routinely reconciled in United States political practice, beginning with the Mulford Act of 1967 and continuing through the contemporary record. Agentic AI regulation will follow this pattern. The doctrinal architecture for legitimate regulation under existing federal law is the architecture this Article has developed.

The magazine does not empty when the apparatus refills it. The federal regulatory architecture built in 1934 to control mechanical automatic weapons relied on the assumption that operator-supplied ammunition would exhaust and require manual intervention. The contemporary agentic AI case is the first time the federal regulatory architecture has encountered an apparatus whose operational continuity is not bounded by operator-supplied resources. The address to this case is the work that Congress, the federal regulatory agencies, and the federal courts now confront. The doctrinal vocabulary for the address is the vocabulary the textualist methodology has supplied. The doctrinal road is built. Both coalitions paved it. What remains is application.