Public Comment on FTC Advance Notice of Proposed Rulemaking — Negative Option Rule

Submitted by: Empty Set LLC (Nevada Series LLC) **Contact:** Brice Love, Co-Founder **Email:** brice@emptysetllc.com **Date:** March 20, 2026 **Re:** Advance Notice of Proposed Rulemaking — Rule Concerning the Use of Prenotification Negative Option Plans, Project No. P064202, 91 Fed. Reg. Doc. No. 2026-04952, Docket No. FTC-2026-0265 **Submitted via:** Regulations.gov, Docket No. FTC-2026-0265

Executive Summary

Empty Set LLC respectfully submits this comment in response to the Commission's Advance Notice of Proposed Rulemaking on the Negative Option Rule. We write to raise an issue that, to the best of our knowledge, no commenter has addressed in any prior FTC negative option proceeding: the emergence of AI agents that autonomously discover, evaluate, subscribe to, and pay for goods and services on behalf of human consumers.

This is not a hypothetical concern. As of early 2026, multiple commercial platforms enable AI agents to execute purchasing decisions with minimal or no real-time human oversight. OpenAI's ChatGPT Instant Checkout, Google's Universal Commerce Protocol, and the Interactive Advertising Bureau's Agentic Commerce Roadmap represent the first wave of systems designed to let AI agents transact autonomously. These systems will generate subscriptions, free-trial-to-paid conversions, and auto-renewing commitments — the core transactions that negative option rules exist to regulate — at a scale and speed that the Restore Online Shoppers' Confidence Act ("ROSCA") and the Commission's existing Negative Option Rule were not designed to address.

The Commission's prior rulemaking received over 16,000 public comments. None addressed AI agent purchasing. The Eighth Circuit's vacatur of the 2024 Click-to-Cancel Rule in *Chamber of Progress v. FTC* — for failure to conduct the preliminary regulatory analysis required by Section 22 of the FTC Act — presents the Commission with a rare opportunity to rebuild the regulatory framework from the ground up. We urge the Commission to use this opportunity to address the AI agent purchasing problem before the technology outpaces the rules.

Empty Set LLC operates a standards development initiative focused on behavioral safety credentialing for autonomous AI agents, with particular emphasis on the insurance and financial services sectors. We are developing the technical infrastructure for verifiable AI agent identity, consent delegation, and behavioral accountability — the building blocks that any AI-aware negative option framework will require. Our perspective bridges the gap between consumer protection policy and the technical architecture of agentic commerce.

I. The Consent Problem: ROSCA Assumes a Human Consumer

The Restore Online Shoppers' Confidence Act (15 U.S.C. §§ 8401-8405) and the Commission's legacy Negative Option Rule (16 CFR Part 425) rest on a foundational assumption: a human consumer reads disclosures, considers terms, and affirmatively acts to accept a recurring obligation. Every element of the regulatory framework reflects this assumption.

"Express informed consent" under ROSCA requires that the consumer receive, and affirmatively agree to, the material terms of a negative option feature before being charged. The concept presupposes that a human being is reading, understanding, and deliberately accepting those terms.

"Clear and conspicuous" disclosure under ROSCA and the Commission's Dot Com Disclosures guidance is defined in visual and auditory terms — prominence, proximity, placement, font size, contrast. These are design constraints for human perception. They are meaningless to an AI agent parsing a website's DOM structure or interacting via API.

"Simple mechanism for cancellation" — the core consumer protection in any negative option rule — assumes that the consumer who enrolled is the same entity that must cancel. When enrollment is performed by an AI agent but cancellation requires a human to navigate a phone tree, chatbot, or multi-step web form, the symmetry principle is violated by design.

The Commission should use this ANPRM to formally acknowledge that the consent framework underlying negative option regulation was designed for human-to-machine interaction, and to solicit public input on how that framework should evolve to address machine-to-machine and agent-to-machine commerce.

II. The Current Landscape: AI Agent Commerce Is Deploying Now

The Commission need not speculate about whether AI agents will conduct commerce on consumers' behalf. The technology is already in market.

Deployed or Announced Platforms

OpenAI ChatGPT Instant Checkout (2026): Enables ChatGPT to complete purchases on behalf of users within the conversational interface, including subscription sign-ups. The user delegates purchasing authority to the agent through a general instruction (e.g., "find and subscribe to a meal kit under $50/month"), and the agent executes the transaction — selecting the merchant, evaluating terms, and committing the consumer's payment instrument.

Google Universal Commerce Protocol (2026): A technical framework for enabling AI agents to transact with merchants via structured APIs, bypassing traditional web-based checkout flows entirely. Designed for machine-to-machine commerce at scale.

IAB Agentic Commerce Roadmap (early 2026): The Interactive Advertising Bureau — the trade association representing the digital advertising industry — has published a roadmap for agentic commerce that explicitly envisions AI agents navigating advertisements, evaluating

offers, and executing purchases on behalf of consumers. This roadmap contemplates AI agents interacting with negative option marketing materials — promotional trials, subscription offers, bundled add-ons — as a core use case.

IEEE Project 7006: An active standards project addressing personal data AI agents, including agents that manage consumer subscriptions and financial commitments. The project's scope directly encompasses the consent delegation problem that negative option rules must address.

Why This Matters for the ANPRM

Each of these platforms creates scenarios in which the traditional ROSCA safeguards fail:

Scenario 1 — Agent subscribes on consumer's behalf. ROSCA requires express informed consent. *Failure:* Consumer set a general objective ("find a coffee subscription under $50"), not specific consent to the terms the agent accepted.

Scenario 2 — Merchant displays terms visually. ROSCA requires clear and conspicuous disclosure. *Failure:* Agent interacts via API or DOM parsing. Visual prominence is irrelevant to a machine.

Scenario 3 — Free trial converts to paid. ROSCA requires consent to the negative option feature before charge. *Failure:* Agent accepted trial terms; consumer may not know a negative option feature was included.

Scenario 4 — Consumer wants to cancel. ROSCA requires a simple cancellation mechanism. *Failure:* Agent enrolled automatically; cancellation requires a human to call, chat, or navigate a multi-step process.

Scenario 5 — Merchant detects AI buyer. ROSCA requires truthful disclosure of material terms. *Failure:* Consumer never sees the actual terms the agent accepted; dynamic pricing creates a disclosure gap.

Scenario 6 — Agent malfunctions. ROSCA assigns liability for unauthorized charges. *Failure:* Unclear whether liability falls on consumer, agent developer, or merchant.

These are not edge cases. They are the foreseeable and intended operation of agentic commerce platforms.

III. No Regulator Has Addressed This

Empty Set LLC has surveyed the federal, state, and international regulatory landscape for any proceeding that specifically addresses AI agent purchasing in the negative option context. We found none.

Federal Agencies

FTC: The Commission's extensive body of work on AI — including Operation AI Comply (September 2024), the omnibus AI enforcement resolution (2023, adopted 3-0), the AI Companion inquiry (September 2025), and the 2024 Tech Summit — demonstrates a strong institutional commitment to AI consumer protection. However, none of these initiatives has addressed AI agent purchasing in the context of negative option rules. The 16,000+ comments received during the 2023 NPRM on the Click-to-Cancel Rule included zero submissions addressing AI agents as purchasers.

CFPB: The Bureau has explored "agentic payments" and acknowledged that consumer authorization for AI agent transactions is inherently *prospective* rather than contemporaneous — the consumer authorizes a range of future actions, not a specific transaction. Circular 2023-01 addresses negative options but not AI agents. Note: the Bureau's operational capacity has been substantially reduced as of early 2026.

NIST: The AI Agent Standards Initiative (launched February 2026) is developing frameworks for AI agent identity verification, cryptographic authentication, and interoperability. These standards could provide the technical foundation for consent verification in agentic commerce, but NIST has not engaged with FTC negative option policy.

Congress: No legislation currently addresses AI agent commerce in the negative option context specifically.

State Legislatures

Approximately twenty-seven states have enacted automatic renewal or negative option laws. All assume a human consumer. None address AI agent purchasing.

The most advanced state AI laws — Colorado SB 24-205, California SB 243 and AB 489, and the New York RAISE Act — impose behavioral safety, impact assessment, and transparency obligations on AI deployers, but none specifically addresses the scenario in which an AI agent is the *purchaser* rather than the *seller's tool*.

International

The EU AI Act's risk-based framework could apply to AI agents that prioritize affiliate commissions over consumer interests (manipulative practice prohibition), but does not specifically address autonomous purchasing. The EU Payment Services Directive 3 (PSD3), currently in negotiation, is the most relevant international development — it aims to formalize delegated payment initiation and harmonize liability for AI-initiated transactions. The Digital Services Act's prohibition on dark patterns and principles of protocol neutrality are relevant but do not address agent-initiated subscriptions.

The regulatory gap is complete. No jurisdiction has addressed what happens when an AI agent — not a human — is the entity encountering a negative option offer, evaluating its terms, and committing a consumer's payment instrument.

IV. Questions the Commission Should Pose in the Subsequent NPRM

The ANPRM stage is the ideal moment to shape the questions that will define the rulemaking. Empty Set LLC respectfully proposes that the Commission consider the following questions for inclusion in the subsequent NPRM:

A. Definitional Questions

1. **Should AI agents be treated as "consumers" under the Negative Option Rule?** If not, what entity bears the obligations that ROSCA imposes on the "consumer" side of a negative option transaction — the human principal, the AI agent, or the agent's developer?

2. **How should "express informed consent" be defined when the purchasing decision is made by an AI agent operating under general instructions from a human consumer?** Is consent to a general purchasing objective (e.g., "find and subscribe to a meal kit under $50/month") sufficient to constitute "express informed consent" to the specific terms of the subscription the agent selects?

3. **Should the Commission adopt a concept of "delegated consent" that distinguishes between the human's authorization of the agent's purchasing authority and the agent's exercise of that authority in a specific transaction?** Empty Set LLC believes yes. Delegated consent is the correct framework. The human authorizes a scope of purchasing authority (e.g., "subscribe to a meal kit under $50/month"). The agent exercises that authority in a specific transaction. Both layers require documentation, but they serve different functions. The authorization layer — human to agent — should meet a "clear and informed" standard. The exercise layer — agent to merchant — should be verifiable through standardized credential protocols, such as those being developed under the NIST AI Agent Standards Initiative.

B. Disclosure Questions

4. **Should sellers be required to provide machine-readable negative option disclosures — in addition to visual disclosures — so that AI agents can parse and evaluate negative option terms before committing consumers?** Empty Set LLC believes yes. Machine-readable disclosures are the AI-equivalent of "clear and conspicuous." A standardized format — such as JSON-LD, schema.org markup, or a Commission-designated standard — would enable AI agents to parse negative option terms before committing consumers, the same way visual disclosures enable human consumers to read terms before agreeing. This is a low-cost, high-impact requirement that benefits both human and AI-mediated commerce.

5. **Should sellers be required to disclose when they detect that a purchaser is an AI agent rather than a human consumer? Should terms offered to AI agents be identical to terms offered to human consumers?** Empty Set LLC believes yes on both counts. Sellers who detect an AI agent purchaser and offer different

terms — dynamic pricing, modified cancellation procedures, altered trial durations — without disclosing the difference are engaged in a form of deception. The consumer authorized the agent to transact on their behalf and has a reasonable expectation that the terms the agent receives are the same terms a human would receive. Any deviation should be disclosed to the consumer.

6. **Should AI agent developers be required to disclose to consumers what negative option commitments their agents have made on the consumer's behalf, in what format, and within what timeframe?** Empty Set LLC believes yes. Agent developers should be required to provide consumers with a clear, itemized log of every negative option commitment made on their behalf — including the merchant, the material terms, the billing frequency, and instructions for cancellation. This disclosure should be delivered within 24 hours of the commitment in a format the consumer has chosen (email, push notification, in-app dashboard). The obligation should rest on the agent developer, not the merchant, because the developer controls the agent's behavior and has direct access to the consumer's communication channel.

C. Cancellation Questions

7. **If enrollment is automated (agent-initiated), should cancellation also be available through automated means? Should the "simple mechanism" requirement be interpreted to mandate machine-readable, API-accessible cancellation endpoints that AI agents can use on the consumer's behalf?** Empty Set LLC believes yes. The "simple mechanism" requirement under ROSCA should be interpreted to include machine-readable, API-accessible cancellation endpoints. If a seller's system can accept an automated enrollment request from an AI agent, that same system can expose an automated cancellation endpoint. The technical capability already exists — the question is whether the Commission will require it.

8. **Should sellers be prohibited from implementing "human verification" barriers in the cancellation process that are designed to prevent AI agents from canceling subscriptions they enrolled the consumer in?** Empty Set LLC believes yes. The principle is symmetry: if the enrollment path was accessible to an AI agent, the cancellation path must be equally accessible. A seller who accepts automated enrollment but requires human-only cancellation (phone calls, CAPTCHAs, multi-step identity verification designed to block bots) has created exactly the kind of asymmetric friction the Commission targeted in its Amazon Prime enforcement. The rule should prohibit cancellation barriers that exceed the verification requirements imposed at enrollment.

D. Liability Questions

9. **When an AI agent enrolls a consumer in a negative option plan that the consumer did not specifically authorize, who is liable under ROSCA — the consumer, the AI agent developer, or the seller? Should the answer depend on whether the seller knew or should have known it was transacting with an**

AI agent? Empty Set LLC believes liability should be allocated based on fault and knowledge. If the agent exceeded the scope of authority delegated by the consumer, the agent developer bears primary liability — the developer designed the system that overstepped. If the seller detected or should have detected that the agent lacked valid authorization and proceeded anyway, the seller shares liability. The consumer should not bear liability for unauthorized commitments made by an agent operating outside its delegated scope. Knowledge of AI agent status should be a relevant factor: a seller who knows it is transacting with an AI agent has a heightened duty to verify authorization.

10. **Should merchants be entitled to a safe harbor from ROSCA liability when they transact with AI agents that present verified credentials — such as compliance with NIST AI Agent Standards or similar recognized frameworks — confirming that the agent has valid consent delegation from its human principal?** Empty Set LLC believes yes. A credential-based safe harbor creates the right incentive structure. Merchants who verify an AI agent's credentials before accepting a negative option enrollment have exercised reasonable diligence. The safe harbor should require that the credential: (i) was issued by a recognized standards body or credentialing authority, (ii) confirms that the agent has valid consent delegation from its human principal, and (iii) was current at the time of the transaction. This framework incentivizes the development of credentialing infrastructure — which is precisely what Empty Set and the broader AI governance ecosystem are building.

E. Economic Impact Questions

11. **What are the current and projected economic costs of AI agent-initiated negative option transactions?** The Commission's Section 22 obligation requires a preliminary regulatory analysis for rules with annual economic impact exceeding $100 million. Given the scale of platforms already deploying agentic commerce (OpenAI, Google, Amazon), the Commission should affirmatively solicit data on the economic magnitude of this emerging transaction category.

V. Proposed Safe Harbor Framework

Rather than prescribing specific technical standards — which risk obsolescence in a rapidly evolving technology landscape — the Commission should consider establishing a safe harbor framework that incentivizes industry-led solutions while maintaining consumer protection.

We propose that sellers who adopt the following practices receive a presumption of ROSCA compliance when transacting with AI agents:

1. **Machine-readable disclosures.** Sellers publish negative option terms in a standardized, machine-readable format (e.g., JSON-LD, schema.org markup, or a Commission-designated standard) that AI agents can parse before committing consumers. This is the machine equivalent of "clear and conspicuous."

2. **Agent authentication.** Sellers verify that an AI agent presenting a consumer's payment instrument has been authorized by that consumer through a recognized consent delegation protocol — such as the emerging NIST AI Agent Standards framework for cryptographic identity verification.

3. **Symmetric cancellation.** If enrollment was executed by an AI agent via automated means, cancellation is available through equivalent automated means. No human-verification barriers are imposed in the cancellation path that were not imposed in the enrollment path.

4. **Transaction logging and consumer notification.** The AI agent developer maintains a log of all negative option commitments made on the consumer's behalf and provides the consumer with timely, clear notification of each commitment — including the material terms, the merchant, the billing frequency, and instructions for cancellation.

5. **Post-transaction confirmation window.** For agent-initiated negative option enrollments, the consumer has a defined period (e.g., 24-72 hours) to review and rescind the commitment without penalty, providing a human-in-the-loop safeguard against agent errors and unauthorized enrollments.

This framework is technology-neutral, principles-based, and designed to evolve with the agentic commerce ecosystem. It does not mandate specific AI architectures. It establishes *what* must be achieved — verified consent, symmetric access, consumer notification — while allowing industry to determine *how*.

VI. Cross-References to Commission Precedent

The AI agent purchasing problem is not a departure from the Commission's existing enforcement priorities. It is a natural extension of them.

Operation AI Comply (September 2024): The Commission recognized that AI can "supercharge deceptive or unfair conduct." AI agents that enroll consumers in unwanted subscriptions — whether through malfunction, misaligned optimization objectives, or merchant manipulation — represent precisely this kind of supercharged unfairness.

AI Omnibus Enforcement Resolution (2023, adopted 3-0): The bipartisan resolution authorized enforcement against AI practices that cause consumer harm. An AI agent that commits a consumer's payment instrument to a recurring obligation without adequate consent delegation falls squarely within this authority.

AI Companion Inquiry (September 2025): The Commission's inquiry into AI products that exploit consumer vulnerabilities is directly relevant. Consumers who delegate purchasing authority to AI agents are vulnerable to exploitation by merchants who design dark patterns specifically to manipulate AI agent behavior — for example, checkout flows that bury negative option terms in HTML comments visible only to DOM-parsing agents, or API endpoints that return different terms to automated versus human requestors.

Amazon Prime Enforcement ($2.5 billion, late 2025): The largest ROSCA enforcement action in history demonstrates the Commission's capacity and willingness to enforce at scale against subscription dark patterns. The "Iliad Flow" — Amazon's internal name for the multi-step cancellation process designed to frustrate consumers — is a precursor to the AI agent problem. When agents enroll consumers in subscriptions more efficiently than humans, and merchants maintain human-only cancellation barriers, the asymmetry the Commission targeted in Amazon Prime is replicated and amplified.

Uber/Chegg Enforcement: The Commission's enforcement against multi-step cancellation dark patterns establishes that cancellation must be simple and direct. This principle should extend to require that agent-enrolled subscriptions be cancellable through agent-accessible means.

VII. The Section 22 Imperative

The Eighth Circuit's vacatur in *Custom Communications* was procedural, not substantive. The court did not reject the Commission's authority to regulate negative options. It held that the Commission violated Section 22 of the FTC Act (15 U.S.C. § 57b-3) by failing to issue the preliminary regulatory analysis required when a rule's annual economic impact exceeds $100 million.

The Commission must not repeat this error. AI agent-initiated commerce is projected to generate substantial transaction volume within the comment period of this rulemaking. The economic analysis required by Section 22 should affirmatively account for:

- The cost to sellers of implementing machine-readable negative option disclosures
- The cost to AI agent developers of implementing consent delegation verification
- The cost to consumers of agent-initiated negative option commitments that were not specifically authorized
- The economic benefits of reduced consumer harm from unauthorized or inadequately disclosed subscriptions
- The economic benefits of standardized agent-merchant interaction protocols that reduce transaction costs

Soliciting this data now — at the ANPRM stage — is the procedurally correct approach and directly addresses the Eighth Circuit's concern that the Commission's initial cost estimate was artificially low.

VIII. Conclusion

The Commission's decision to restart the negative option rulemaking through an ANPRM — rather than simply reissuing the vacated NPRM — is an act of institutional seriousness. It acknowledges that the regulatory landscape has changed and that a fresh examination is warranted.

We submit that the most significant change in that landscape is not the Eighth Circuit's vacatur itself, but the emergence of AI agent commerce. The vacatur created a procedural reset. AI agent purchasing creates a substantive gap. The existing regulatory framework — ROSCA, the Negative Option Rule, and the Commission's extensive enforcement precedent — was built for a world in which humans encounter, evaluate, and accept negative option offers. That world is ending.

The Commission has an opportunity to build the first regulatory framework in any jurisdiction that addresses AI agent purchasing in the negative option context. No federal agency, no state legislature, and no international regulator has done so. The ANPRM stage is the ideal moment — before proposed rule text is drafted, while the agency is soliciting broad public input on the scope and direction of the rulemaking.

Empty Set LLC urges the Commission to:

1. Formally recognize AI agent purchasing as within the scope of this rulemaking
2. Pose the specific questions outlined in Section IV to solicit public input on definitional, disclosure, cancellation, and liability issues
3. Consider the safe harbor framework proposed in Section V as a principles-based, technology-neutral approach to consumer protection in agentic commerce
4. Include AI agent-initiated transaction costs in the Section 22 preliminary regulatory analysis

We appreciate the opportunity to comment and welcome further engagement with the Commission and its staff on these issues. Empty Set LLC is developing the credentialing and behavioral safety infrastructure that will enable verifiable AI agent identity and consent delegation — the technical foundation for any effective regulation of agentic commerce.

Respectfully submitted,

Brice Love Co-Founder, Empty Set LLC brice@emptysetllc.com