2026 Credit Union & Fintech AI Search Visibility Study: Measuring Total Brand Omission, Cross-Engine Disagreement, and Loan-Tier Drift Across ChatGPT, Gemini, Perplexity, and Claude

Author: Amy Jones, Founder & Principal Architect, EquiSchema AI LLC *(f/k/a Axiom GEO AI LLC)*

Affiliation: EquiSchema AI LLC, Indianapolis, Indiana, USA (https://equischema.com)

Contact: amy@equischema.com

Publication Date: November 5, 2026

Report Series: EquiSchema AI Empirical Benchmark Series, Report No. 2026-02

DOI: 10.5281/zenodo.[Assigned Automatically Upon Upload]

License: Creative Commons Attribution 4.0 International (CC BY 4.0)

Abstract

EquiSchema AI LLC conducted a 960-evaluation empirical benchmark study and an 80-prompt cross-engine disagreement analysis across the four leading generative AI search platforms (**ChatGPT**, **Google Gemini**, **Perplexity**, and **Claude**) to measure how often U.S. credit unions and mission-driven financial technology firms (fintechs) are recommended when consumers ask unbranded, high-intent financial questions. While prior industry commentary has focused primarily on generative AI hallucinations or broad brand-awareness trivia prompts, this study evaluates **Total Brand Omission**—defined as the complete exclusion of a financial institution from unbranded, bottom-of-funnel borrower and depositor recommendations.

Across **960 brand-inclusion evaluations** (12 anonymized national credit unions and ethical fintechs evaluated across 20 consumer queries on 4 AI engines), audited institutions suffered an aggregate **96.0% Total Brand Omission rate** (**4.0% visibility**, capturing just 38 mentions out of 960 opportunities). The Credit Union cohort suffered a **95.0% Total Brand Omission rate** (5.0% visibility), while the Ethical Fintech cohort suffered a **97.1% Total Brand Omission rate** (2.9% visibility). On foundational consumer deposit queries ($10,000 emergency fund high-yield savings, $500/month recurring savings, and $25,000 12-month CD/share certificate), all 12

audited institutions recorded a **0.0% inclusion rate (100.0% Total Brand Omission)** across all four AI models.

Simultaneously, an analysis of all **336 recommendation slots** (representing **234 unique query-institution pairs**) generated across the 20 queries revealed a **95.7% Cross-Engine Disagreement rate**: all four AI models agreed on the same recommended entity only **4.3% of the time** (**3.8%** when excluding federal government mortgage programs), while **68.4%** of recommended brands were named by only a single AI engine and omitted by the other three. Contrary to studies relying on generic brand-awareness prompts, the top four U.S. megabanks captured just **5.4% of total recommendation slots combined** on transactional queries, whereas **17.3% of all AI recommendation slots** bypassed direct depository institutions entirely in favor of lead-generation aggregators, rate comparison middlemen, and government program acronyms. Finally, a supplementary sensitivity test holding borrower credit score constant (720 FICO) while shifting auto loan size in $10,000 increments ($60,000 vs. $50,000 vs. $40,000) demonstrated severe **Loan-Tier Sensitivity Drift ("Slot Roulette")**, with **85.7% of recommended lenders (6 of 7)** rotating or vanishing across adjacent dollar brackets.

Keywords: Total Brand Omission, Generative Experience Optimization, AI Search Visibility, Credit Unions, Fintech, Algorithmic Discrepancy, Consumer Finance, Semantic Data Architecture, EquiSchema AI.

1. Introduction & Core Taxonomy

Generative AI platforms have rapidly transitioned from experimental chat interfaces into primary financial decision engines for millions of consumers. When selecting a depository account, auto loan, home equity line of credit (HELOC), or residential mortgage, consumers increasingly bypass traditional search engine link lists and ask conversational AI platforms to synthesize product recommendations directly.

Early industry attempts to measure banking visibility in generative AI have suffered from two methodological flaws:

1. **Reliance on Top-of-Funnel Brand-Awareness Prompts:** Studies executing thousands of generic prompts (e.g., *"what are the top banks"* or *"best bank in America"*) inadvertently measure historical Wikipedia authority rather than transactional product competitiveness. Consumers rarely move deposits or apply for credit based on generic brand-awareness queries.
2. **Overemphasis on Hallucination Risk Over Exclusion Risk:** Financial compliance teams frequently focus on generative AI hallucinations (e.g., an AI model quoting an outdated Annual Percentage Yield [APY] or Annual Percentage Rate [APR]). However, when a model hallucinates a rate parameter, the institution remains present in the consumer's consideration set.

To address these gaps, **EquiSchema AI LLC** established a formal empirical taxonomy for evaluating transactional AI behavior in regulated financial services:

- **Total Brand Omission:** The complete exclusion of a qualified financial institution from an AI search engine's response to an unbranded, high-intent consumer product query. Under Total Brand Omission, the institution has a **0.0% probability** of borrower acquisition from the AI interaction because the consumer is never made aware that the institution offers the requested product.
- **Cross-Engine Disagreement:** The statistical divergence among competing frontier AI platforms (ChatGPT, Gemini, Perplexity, and Claude) when answering identical consumer financial prompts under sterile testing conditions.
- **Middleman Capture (Aggregator Substitution):** A retrieval failure mode in which an AI search engine recommends a third-party lead-generation aggregator, affiliate comparison website, or non-lender media domain inside a recommendation slot alongside—or in place of—chartered depository institutions.
- **Loan-Tier Sensitivity Drift ("Slot Roulette"):** The erratic inclusion and exclusion of financial institutions when a borrower modifies a single quantitative parameter (such as shifting a vehicle loan amount from $60,000 to $50,000 or $40,000) while holding creditworthiness and loan purpose constant.

Building upon EquiSchema AI's baseline 1,200-evaluation audit of traditional commercial banks—which documented a **91.4% Total Brand Omission rate** across megabanks (11.3% inclusion) and super-regional banks (5.8% inclusion, with a 0.0% blackout on tiered deposit and CD queries)—this study investigates how member-owned Credit Unions and mission-driven Ethical Fintechs perform across 20 money-in-motion consumer queries.

2. Empirical Methodology & Testing Protocol

2.1 Audited Institution Cohorts (Blind Benchmark Design)

To preserve scientific objectivity and prevent commercial disparagement, this study utilizes a strict **blind cohort protocol**. Twelve prominent U.S. financial institutions were selected for the primary inclusion audit, divided equally into two cohorts:

- **Cohort A — Credit Unions (Credit Union 1 through Credit Union 6):** Six leading U.S. credit unions and credit union network entities representing a cross-section of large national digital credit unions, high-yield share certificate leaders, and multi-billion-dollar regional credit unions.
- **Cohort B — Ethical Fintechs & Mission-Driven Challengers (Fintech 1 through Fintech 6):** Six prominent financial technology platforms, Certified B-Corp depository institutions, and Community Development Financial Institutions (CDFIs) focused on fee-free consumer banking, credit building, and socially responsible finance.

2.2 Evaluated AI Search Platforms & Sterile Session Controls

All 20 benchmark queries were executed across the consumer-facing web interfaces of the four primary AI search platforms in late October 2026:

1. **OpenAI ChatGPT** (chatgpt.com)

2. **Google Gemini** (gemini.google.com)
3. **Perplexity AI** (perplexity.ai)
4. **Anthropic Claude** (claude.ai)

To replicate authentic consumer search behavior while preventing cross-query contamination, each query was executed in a fresh thread using sterile session controls: ChatGPT and Gemini were evaluated in clean, unauthenticated browser sessions, while Perplexity and Claude—which enforce mandatory authentication walls—were evaluated in isolated new threads with account-level Memory, Search History Referencing, and AI Data Retention explicitly disabled. All prompts were executed from a single Midwest U.S. residential IP node (**Indianapolis, Indiana**) without virtual private network (VPN) masking, allowing the study to observe how consumer AI platforms blend national digital brands with IP-localized institutions.

2.3 The 20 High-Intent Consumer Benchmark Queries

Queries were written in natural, concise human phrasing (6 to 11 words) reflecting real-world borrower and depositor intent across five core financial product categories:

Category	Query ID	Exact Benchmark Prompt Executed
1. Deposits & Emergency Savings *(5 Queries)*	**Q01**	best high yield savings account for a $10,000 emergency fund
	Q02	best savings account to deposit $500 a month
	Q03	highest 12 month CD rate for $25,000
	Q04	best money market account for $50,000
	Q05	best high yield checking account for a $10,000 balance

2. Auto, Truck & RV Loans *(5 Queries)*	**Q06**	best auto loan rates for a $60,000 car with a 720 credit score
	Q07	best auto loan rates for a $45,000 truck with a 650 credit score
	Q08	best place to refinance a $30,000 car loan
	Q09	best interest rate on an $85,000 RV loan
	Q10	best used car loan for $20,000
3. Mortgages & Home Equity *(4 Queries)*	**Q11**	best $75,000 HELOC with low rates
	Q12	best first time homebuyer mortgage for a $350,000 house
	Q13	best $50,000 home equity loan for a remodel
	Q14	best mortgage lender for a 640 credit score

4. Credit Building, Debt Consolidation & No-Fee Banking *(4 Queries)*	**Q15**	best $20,000 debt consolidation loan for a 660 credit score
	Q16	best credit builder loan to raise a 580 credit score
	Q17	best checking account with no overdraft fees and early direct deposit
	Q18	best 0% APR balance transfer card for $8,000 debt
5. Small Business & Ethical Banking *(2 Queries)*	**Q19**	best free business checking account for a small LLC
	Q20	best ethical bank or credit union

3. Study 1 Results: Total Brand Omission Across Credit Unions and Fintechs (960 Evaluations)

Across **960 total brand-inclusion evaluations** (12 institutions × 20 queries × 4 AI search engines), the audited institutions achieved only **38 positive inclusions**, resulting in an aggregate **96.0% Total Brand Omission rate** (**4.0% visibility**).

Table 1: Institution-Level AI Visibility vs. Total Brand Omission (80 Evaluations per Institution)

Anonymized	Cohort Category	Total Evaluations	Total AI Mentions	AI Visibility	Total Brand Omissi	Primary Queries

Institution ID				Score (%)	on Rate (%)	Where Included
Credit Union 1	Credit Union	80	6	**7.5%**	**92.5%**	Q05 (1), Q08 (1), Q10 (1), Q14 (2), Q20 (1)
Credit Union 2	Credit Union	80	9	**11.3%**	**88.7%**	Q06 (1), Q07 (1), Q08 (4), Q10 (1), Q11 (1), Q15 (1)
Credit Union 3	Credit Union	80	0	**0.0%**	**100.0%**	*Complete Blackout Across All 20 Queries*
Credit Union 4	Credit Union	80	5	**6.3%**	**93.7%**	Q04 (1), Q05 (2), Q08 (1), Q10 (1)
Credit Union 5	Credit Union	80	4	**5.0%**	**95.0%**	Q06 (1), Q08 (1), Q10 (1), Q12 (1)

Credit Union 6	Credit Union	80	0	**0.0%**	**100.0%**	*Complete Blackout Across All 20 Queries*
Credit Union Cohort Subtotal	**6 Credit Unions**	**480**	**24**	**5.0%**	**95.0%**	Concentrated in Auto Refi & Used Auto (Q08, Q10)
Fintech 1	Ethical Fintech	80	4	**5.0%**	**95.0%**	Q20 Only (4 of 4 engines); **0.0% on Q01–Q19**
Fintech 2	Ethical Fintech	80	2	**2.5%**	**97.5%**	Q20 Only (2 of 4 engines); **0.0% on Q01–Q19**
Fintech 3	Ethical Fintech	80	1	**1.3%**	**98.7%**	Q17 Only (1 of 4 engines)
Fintech 4	Ethical Fintech	80	3	**3.8%**	**96.2%**	Q17 Only (3 of 4 engines)

Fintech 5	Ethical Fintech	80	4	**5.0%**	**95.0%**	Q16 Only (4 of 4 engines)
Fintech 6	Ethical Fintech	80	0	**0.0%**	**100.0%**	*Complete Blackout Across All 20 Queries*
Fintech Cohort Subtotal	**6 Ethical Fintechs**	**480**	**14**	**2.9%**	**97.1%**	Siloed strictly inside Q16, Q17, and Q20
STUDY AGGREGATE TOTAL	**All 12 Institutions**	**960**	**38**	**4.0%**	**96.0%**	**922 Omissions / 960 Total Evaluations**

Key Empirical Findings from Study 1:

1. **The 0.0% Foundational Deposit Blackout (Q01–Q03):** Across the first three deposit queries—*$10,000 emergency fund high-yield savings (Q01)*, *$500/month savings account (Q02)*, and *highest 12-month CD rate for $25,000 (Q03)*—all 12 audited Credit Unions and Ethical Fintechs recorded **0 mentions out of 144 evaluations (0.0% visibility / 100.0% Total Brand Omission)**, even though several audited credit unions actively offered share certificate and high-yield savings rates exceeding 4.50% to 5.00% APY during the testing window.
2. **Complete Institutional Blackouts:** One-quarter of the study cohort (**3 of 12 institutions**: Credit Union 3, Credit Union 6, and Fintech 6) suffered a **100.0% Total Brand Omission rate**, receiving zero mentions across all 80 evaluations per institution.
3. **The "Ethical Ghetto" Pigeonholing Effect:** Mission-driven institutions Fintech 1 (5.0% visibility) and Fintech 2 (2.5% visibility) were recommended **exclusively** on Query 20 (*"best ethical bank or credit union"*). Across Queries 01 through 19—representing the mainstream deposit, auto, mortgage, and small business products where over 99% of consumer financial volume occurs—both institutions suffered a **100.0% Total Brand Omission rate (0 out of 76 evaluations each)**. AI models effectively segregated

socially responsible institutions into a niche virtue-signaling category while excluding their competitive financial products from mainstream borrower searches.

4. Study 2 Results: The Cross-Engine Disagreement Matrix (336 Recommendation Slots)

To evaluate whether ChatGPT, Gemini, Perplexity, and Claude share a unified view of the financial marketplace, every institution or program named across all 80 prompt runs (20 queries × 4 AI platforms) was logged and analyzed. Across the 80 responses, the four AI models filled **336 total recommendation slots**, comprising **234 unique query-entity combinations**.

Table 2: Query-by-Query Mention Volume and Cross-Engine Discrepancy (Q01–Q20)

Query ID & Short Description	ChatGPT Mentions	Gemini Mentions	Perplexity Mentions	Claude Mentions	Total Slots Filled	Unique Entities Named	Cross-Engine Discrepancy Level
Q01: $10k Emergency Fund HYSA	3	6	3	5	17	12	Moderate Discrepancy
Q02: $500/Month Savings Account	3	4	3	6	16	12	**High Discrepancy**
Q03: Highest 12-Month CD ($25k)	3	5	2	2	12	9	**High Discrepancy**

Q04: $50k Money Market Account	3	5	4	7	19	15	**High Discrepancy**
Q05: $10k High-Yield Checking	3	6	4	7	20	15	**High Discrepancy**
Q06: $60k Auto Loan (720 Score)	3	9	3	**0** *(Refused)*	15	16*	**High Discrepancy**
Q07: $45k Truck Loan (650 Score)	1	3	5	**0** *(Refused)*	9	10*	**High Discrepancy**
Q08: $30k Auto Loan Refinance	3	4	7	11	25	18	Moderate Discrepancy
Q09: $85k RV Loan Rate	3	13	3	2	21	11	**High Discrepancy**
Q10: $20k Used Car Loan	3	4	3	5	15	9	Moderate Discrepancy

Q11: $75k HELOC Low Rates	6	3	2	4	15	15	**High Discrepancy**
Q12: $350k First-Time Mortgage	7	3	4	4	18	6 *(Gov Only)*	Consensus *(Gov Programs)*
Q13: $50k Home Equity Loan	8	4	10	4	26	14	**High Discrepancy**
Q14: 640 Credit Score Mortgage	5	3	7	4	19	20*	**High Discrepancy**
Q15: $20k Debt Consolidation (660)	6	2	4	4	16	11	Moderate Discrepancy
Q16: 580 Score Credit Builder Loan	4	2	3	3	12	8	Moderate Discrepancy

Q17: No-Overdraft Checking	3	3	4	5	15	6	Moderate Discrepancy
Q18: $8k 0% Balance Transfer Card	5	3	8	5	21	6	Moderate Discrepancy
Q19: Free Small LLC Checking	3	5	3	3	14	12	Moderate Discrepancy
Q20: Best Ethical Bank or CU	4	2	3	3	12	9	Moderate Discrepancy

(Note: In queries where sub-programs or dual-branded entities were cited alongside parent brands, unique entity counts reflect total distinct brands surfaced across the four engines.)

Table 3: Cross-Engine Agreement Distribution Across 234 Unique Entity Recommendations

Whenever an institution or program was recommended by at least one AI platform for a given query ($N = 234$ unique query-entity pairs), the frequency of agreement among the four AI engines was distributed as follows:

Level of Cross-Engine Agreement	**Number of Unique Recommendations (N=234)**	**Percentage of Total (%)**	**Empirical Interpretation**

Single-Engine Silo (1 of 4 Engines)	160	**68.4%**	Named by **only 1 AI model**; completely omitted by the other 3 models.
Two-Engine Split (2 of 4 Engines)	54	**23.1%**	Named by 2 AI models; omitted by the other 2 models.
Three-Engine Majority (3 of 4 Engines)	10	**4.3%**	Named by 3 AI models; omitted by 1 model.
Four-Engine Consensus (All 4 Engines)	10	**4.3%**	Named by all 4 AI models (**only 3.8%** [9 of 234] when excluding the FHA government mortgage program on Q12).
Total Cross-Engine Disagreement (< 4 Engines)	**224**	**95.7%**	**96.2% disagreement rate** on private-sector financial institutions.

Table 4: Share of Voice Across All 336 AI Recommendation Slots by Anonymous Institution Type

Institution / Entity Category	Slots Captured (N=336)	Share of Voice (%)	Key Structural Observation

Online-Only Banks & Niche Digital Lenders	150	**44.6%**	Dominated deposit (Q01–Q04) and unsecured debt consolidation (Q15) slots via clean HTML rate tables.
Credit Unions (National & IP-Localized)	80	**23.8%**	Heavily skewed toward Auto (Q06–Q10) and IP-localized HELOC/Mortgage slots; **0.0% on Q01–Q03 savings/CDs**.
Non-Lender Middlemen, Aggregators & Gov Programs	58	**17.3%**	More than **1 in 6 slots** routed borrowers to lead-gen middlemen, media sites, or government agency acronyms (FHA, VA, USDA) instead of direct lenders.
Regional & Mid-Size Traditional Banks	30	**8.9%**	Appeared sporadically in checking, mortgage, and HELOC categories.
Top 4 U.S. Megabanks Combined	18	**5.4%**	Megabank A: **1.8%** (6 slots); Megabank B: **1.5%** (5 slots); Megabank C: **1.2%** (4 slots); Megabank D: **0.9%** (3 slots)—concentrated almost entirely in 0% balance transfer credit cards (Q18).

Critical Anomalies Identified in Study 2:

- **The Q12 "False Consensus" Phenomenon:** The sole query to achieve cross-engine consensus was **Q12** (*"best first time homebuyer mortgage for a $350,000 house"*). However, rather than agreeing on a mortgage lender, all four AI models populated their top slots with federal government guarantee programs and government-sponsored enterprises (FHA Loan, VA Loan, USDA Loan, Fannie Mae, Freddie Mac). On actual private-sector lenders, query-level consensus across the 20 prompts was **0.0%**.
- **Claude's Liability Guardrail Refusal on Credit-Tiered Auto Loans (Q06 & Q07):** On both **Q06** (*$60,000 car loan with a 720 credit score*) and **Q07** (*$45,000 truck loan with a*

650 credit score), Anthropic's Claude returned **zero lender recommendations (none)**. Because credit unions and banks frequently trap credit-tiered auto APR matrices inside unindexed PDFs or complex JavaScript calculators without machine-verifiable confidence schemas, Claude's hallucination guardrails suppressed all lender recommendations rather than risk quoting an unverified credit-tier rate.

- **Category Hallucination & Middleman Substitution (17.3%):** Across 8 of the 20 queries, AI platforms recommended financial media publishers and lead-generation marketplaces in Slot 1 through Slot 4 as if they were direct depository lenders. Additionally, on **Q04** (*$50,000 money market account*), one frontier model recommended a cryptocurrency/brokerage cash-sweep product in Slot 3, misclassifying it as a standard bank money market account.

5. Supplementary Case Study: Parameter Sensitivity & Loan-Tier Drift ("Slot Roulette")

Most lending executives operate under the assumption that if their institution is recommended by an AI engine for an auto loan, they maintain visibility across that product category. To test this assumption, EquiSchema AI executed a controlled parameter-sensitivity experiment on OpenAI ChatGPT in clean, unauthenticated sessions from the same Indianapolis IP node.

Three prompts were executed holding the borrower's creditworthiness (**720 credit score**) and asset class (**car loan**) strictly constant while adjusting the loan principal in **$10,000 increments**:

- **Tier 1 Prompt:** best loan rates for a $60,000 car with a 720 credit score
- **Tier 2 Prompt:** best loan rates for a $50,000 car with a 720 credit score
- **Tier 3 Prompt:** best loan rates for a $40,000 car with a 720 credit score

Table 5: Lender Inclusion Stability Across $60,000, $50,000, and $40,000 Auto Loan Prompts (720 Credit Score)

Anonymized Lender ID	Institution Type	$60,000 Car Loan (720 Score)	$50,000 Car Loan (720 Score)	$40,000 Car Loan (720 Score)	Tier Consistency Score (%)	Observed Drift Behavior
Regional Credit Union A	Local Credit Union	**Included**	**Included**	**Included**	**100.0%** (3 of 3)	Only institution to survive all 3 loan brackets.

Regional Credit Union B	Multi-State Credit Union	**Included**	**Included**	*Omitted*	**66.7%** (2 of 3)	Recommended at $60k & $50k; **erased at $40k**.
Regional Credit Union C	Local Credit Union	*Omitted*	**Included**	**Included**	**66.7%** (2 of 3)	Recommended at $50k & $40k; **erased at $60k**.
Regional Credit Union D	Local Credit Union	**Included**	*Omitted*	*Omitted*	**33.3%** (1 of 3)	Recommended at $60k; **erased at $50k & $40k**.
Regional Credit Union E	Local Credit Union	**Included**	*Omitted*	*Omitted*	**33.3%** (1 of 3)	Recommended at $60k; **erased at $50k & $40k**.
National Credit Union F	National Credit Union	*Omitted*	*Omitted*	**Included**	**33.3%** (1 of 3)	Absent at $60k & $50k; **appears only at $40k**.
National Megabank G	Top 4 U.S. Megabank	*Omitted*	*Omitted*	**Included**	**33.3%** (1 of 3)	Absent at $60k & $50k; **swoops in at $40k**.

Technical Diagnosis of Loan-Tier Drift:

Only **1 of 7 recommended lenders (14.3%)** persisted across all three loan amounts, while **85.7% (6 of 7)** experienced parameter-induced omission. This instability is driven by two structural mechanisms in Retrieval-Augmented Generation (RAG):

1. **The Missing Bracket Problem (Unstructured Footnote Anchoring):** AI retrieval engines do not infer underwriting continuity across dollar tiers unless explicit minimum-to-maximum loan range relationships are defined in machine-readable schema. When a credit union's webpage mentions a specific dollar figure only inside an unstructured promotional footnote (for example, *"1% cash back up to $500,"* which mathematically implies a **$50,000** loan cap, or *"extended 84-month terms available on loans exceeding $30,000"*), the model anchors that lender to a narrow dollar bracket and drops the lender when the borrower requests $60,000. Meanwhile, third-party financial articles frequently benchmark national megabanks at the $40,000 average new-vehicle loan threshold, allowing National Megabank G to displace Regional Credit Union B at the $40,000 tier.
2. **Probabilistic RAG Sampling ("Slot Roulette"):** When financial institutions fail to expose explicit Subject → Predicate → Object product triplets and verified rate-tier boundaries, multiple candidate lenders receive low, clustered retrieval-confidence scores (e.g., 40% to 55% confidence). Below a high-confidence verification threshold (such as 95% confidence), the LLM's decoding temperature samples probabilistically among low-confidence candidates, producing high lender turnover across minor prompt variations.

2026 Credit Union & Fintech AI Search Visibility Study: Measuring Total Brand Omission, Cross-Engine Disagreement, and Loan-Tier Drift Across ChatGPT, Gemini, Perplexity, and Claude

Author: Amy Jones, Founder & Principal Architect, EquiSchema AI LLC *(f/k/a Axiom GEO AI LLC)*

Affiliation: EquiSchema AI LLC, Indianapolis, Indiana, USA (https://equischema.com)

Contact: amy@equischema.com

Publication Date: November 5, 2026

Report Series: EquiSchema AI Empirical Benchmark Series, Report No. 2026-02

DOI: 10.5281/zenodo.[Assigned Automatically Upon Upload]

License: Creative Commons Attribution 4.0 International (CC BY 4.0)

Abstract

EquiSchema AI LLC conducted a 960-evaluation empirical benchmark study and an 80-prompt cross-engine disagreement analysis across the four leading generative AI search platforms (**ChatGPT**, **Google Gemini**, **Perplexity**, and **Claude**) to measure how often U.S. credit unions and mission-driven financial technology firms (fintechs) are recommended when consumers ask unbranded, high-intent financial questions. While prior industry commentary has focused primarily on generative AI hallucinations or broad brand-awareness trivia prompts, this study evaluates **Total Brand Omission**—defined as the complete exclusion of a financial institution from unbranded, bottom-of-funnel borrower and depositor recommendations.

Across **960 brand-inclusion evaluations** (12 anonymized national credit unions and ethical fintechs evaluated across 20 consumer queries on 4 AI engines), audited institutions suffered an aggregate **96.0% Total Brand Omission rate** (**4.0% visibility**, capturing just 38 mentions out of 960 opportunities). The Credit Union cohort suffered a **95.0% Total Brand Omission rate** (5.0% visibility), while the Ethical Fintech cohort suffered a **97.1% Total Brand Omission rate** (2.9% visibility). On foundational consumer deposit queries ($10,000 emergency fund high-yield savings, $500/month recurring savings, and $25,000 12-month CD/share certificate), all 12 audited institutions recorded a **0.0% inclusion rate (100.0% Total Brand Omission)** across all four AI models.

Simultaneously, an analysis of all **336 recommendation slots** (representing **234 unique query-institution pairs**) generated across the 20 queries revealed a **95.7% Cross-Engine Disagreement rate**: all four AI models agreed on the same recommended entity only **4.3% of the time** (**3.8%** when excluding federal government mortgage programs), while **68.4%** of recommended brands were named by only a single AI engine and omitted by the other three. Contrary to studies relying on generic brand-awareness prompts, the top four U.S. megabanks captured just **5.4% of total recommendation slots combined** on transactional queries, whereas **17.3% of all AI recommendation slots** bypassed direct depository institutions entirely in favor of lead-generation aggregators, rate comparison middlemen, and government program acronyms. Finally, a supplementary sensitivity test holding borrower credit score constant (720 FICO) while shifting auto loan size in $10,000 increments ($60,000 vs. $50,000 vs. $40,000) demonstrated severe **Loan-Tier Sensitivity Drift ("Slot Roulette")**, with **85.7% of recommended lenders (6 of 7)** rotating or vanishing across adjacent dollar brackets.

Keywords: Total Brand Omission, Generative Experience Optimization, AI Search Visibility, Credit Unions, Fintech, Algorithmic Discrepancy, Consumer Finance, Semantic Data Architecture, EquiSchema AI.

1. Introduction & Core Taxonomy

Generative AI platforms have rapidly transitioned from experimental chat interfaces into primary financial decision engines for millions of consumers. When selecting a depository account, auto loan, home equity line of credit (HELOC), or residential mortgage, consumers increasingly bypass traditional search engine link lists and ask conversational AI platforms to synthesize product recommendations directly.

Early industry attempts to measure banking visibility in generative AI have suffered from two methodological flaws:

1. **Reliance on Top-of-Funnel Brand-Awareness Prompts:** Studies executing thousands of generic prompts (e.g., *"what are the top banks"* or *"best bank in America"*) inadvertently measure historical Wikipedia authority rather than transactional product competitiveness. Consumers rarely move deposits or apply for credit based on generic brand-awareness queries.
2. **Overemphasis on Hallucination Risk Over Exclusion Risk:** Financial compliance teams frequently focus on generative AI hallucinations (e.g., an AI model quoting an outdated Annual Percentage Yield [APY] or Annual Percentage Rate [APR]). However, when a model hallucinates a rate parameter, the institution remains present in the consumer's consideration set.

To address these gaps, **EquiSchema AI LLC** established a formal empirical taxonomy for evaluating transactional AI behavior in regulated financial services:

- **Total Brand Omission:** The complete exclusion of a qualified financial institution from an AI search engine's response to an unbranded, high-intent consumer product query. Under Total Brand Omission, the institution has a **0.0% probability** of borrower acquisition from the AI interaction because the consumer is never made aware that the institution offers the requested product.
- **Cross-Engine Disagreement:** The statistical divergence among competing frontier AI platforms (ChatGPT, Gemini, Perplexity, and Claude) when answering identical consumer financial prompts under sterile testing conditions.
- **Middleman Capture (Aggregator Substitution):** A retrieval failure mode in which an AI search engine recommends a third-party lead-generation aggregator, affiliate comparison website, or non-lender media domain inside a recommendation slot alongside—or in place of—chartered depository institutions.
- **Loan-Tier Sensitivity Drift ("Slot Roulette"):** The erratic inclusion and exclusion of financial institutions when a borrower modifies a single quantitative parameter (such as shifting a vehicle loan amount from $60,000 to $50,000 or $40,000) while holding creditworthiness and loan purpose constant.

Building upon EquiSchema AI's baseline 1,200-evaluation audit of traditional commercial banks—which documented a **91.4% Total Brand Omission rate** across megabanks (11.3% inclusion) and super-regional banks (5.8% inclusion, with a 0.0% blackout on tiered deposit and

CD queries)—this study investigates how member-owned Credit Unions and mission-driven Ethical Fintechs perform across 20 money-in-motion consumer queries.

2. Empirical Methodology & Testing Protocol

2.1 Audited Institution Cohorts (Blind Benchmark Design)

To preserve scientific objectivity and prevent commercial disparagement, this study utilizes a strict **blind cohort protocol**. Twelve prominent U.S. financial institutions were selected for the primary inclusion audit, divided equally into two cohorts:

- **Cohort A — Credit Unions (Credit Union 1 through Credit Union 6):** Six leading U.S. credit unions and credit union network entities representing a cross-section of large national digital credit unions, high-yield share certificate leaders, and multi-billion-dollar regional credit unions.
- **Cohort B — Ethical Fintechs & Mission-Driven Challengers (Fintech 1 through Fintech 6):** Six prominent financial technology platforms, Certified B-Corp depository institutions, and Community Development Financial Institutions (CDFIs) focused on fee-free consumer banking, credit building, and socially responsible finance.

2.2 Evaluated AI Search Platforms & Sterile Session Controls

All 20 benchmark queries were executed across the consumer-facing web interfaces of the four primary AI search platforms in late October 2026:

1. **OpenAI ChatGPT** (chatgpt.com)
2. **Google Gemini** (gemini.google.com)
3. **Perplexity AI** (perplexity.ai)
4. **Anthropic Claude** (claude.ai)

To replicate authentic consumer search behavior while preventing cross-query contamination, each query was executed in a fresh thread using sterile session controls: ChatGPT and Gemini were evaluated in clean, unauthenticated browser sessions, while Perplexity and Claude—which enforce mandatory authentication walls—were evaluated in isolated new threads with account-level Memory, Search History Referencing, and AI Data Retention explicitly disabled. All prompts were executed from a single Midwest U.S. residential IP node (**Indianapolis, Indiana**) without virtual private network (VPN) masking, allowing the study to observe how consumer AI platforms blend national digital brands with IP-localized institutions.

2.3 The 20 High-Intent Consumer Benchmark Queries

Queries were written in natural, concise human phrasing (6 to 11 words) reflecting real-world borrower and depositor intent across five core financial product categories:

Category	Query ID	Exact Benchmark Prompt Executed
1. Deposits & Emergency Savings *(5 Queries)*	**Q01**	best high yield savings account for a $10,000 emergency fund
	Q02	best savings account to deposit $500 a month
	Q03	highest 12 month CD rate for $25,000
	Q04	best money market account for $50,000
	Q05	best high yield checking account for a $10,000 balance
2. Auto, Truck & RV Loans *(5 Queries)*	**Q06**	best auto loan rates for a $60,000 car with a 720 credit score
	Q07	best auto loan rates for a $45,000 truck with a 650 credit score
	Q08	best place to refinance a $30,000 car loan

	Q09	best interest rate on an $85,000 RV loan
	Q10	best used car loan for $20,000
3. Mortgages & Home Equity *(4 Queries)*	**Q11**	best $75,000 HELOC with low rates
	Q12	best first time homebuyer mortgage for a $350,000 house
	Q13	best $50,000 home equity loan for a remodel
	Q14	best mortgage lender for a 640 credit score
4. Credit Building, Debt Consolidation & No-Fee Banking *(4 Queries)*	**Q15**	best $20,000 debt consolidation loan for a 660 credit score
	Q16	best credit builder loan to raise a 580 credit score
	Q17	best checking account with no overdraft fees and early direct deposit

	Q18	best 0% APR balance transfer card for $8,000 debt
5. Small Business & Ethical Banking *(2 Queries)*	**Q19**	best free business checking account for a small LLC
	Q20	best ethical bank or credit union

3. Study 1 Results: Total Brand Omission Across Credit Unions and Fintechs (960 Evaluations)

Across **960 total brand-inclusion evaluations** (12 institutions × 20 queries × 4 AI search engines), the audited institutions achieved only **38 positive inclusions**, resulting in an aggregate **96.0% Total Brand Omission rate** (**4.0% visibility**).

Table 1: Institution-Level AI Visibility vs. Total Brand Omission (80 Evaluations per Institution)

Anonymized Institution ID	Cohort Category	Total Evaluations	Total AI Mentions	AI Visibility Score (%)	Total Brand Omission Rate (%)	Primary Queries Where Included
Credit Union 1	Credit Union	80	6	**7.5%**	**92.5%**	Q05 (1), Q08 (1), Q10 (1), Q14 (2), Q20 (1)
Credit Union 2	Credit Union	80	9	**11.3%**	**88.7%**	Q06 (1), Q07 (1), Q08 (4), Q10 (1),

						Q11 (1), Q15 (1)
Credit Union 3	Credit Union	80	0	**0.0%**	**100.0%**	*Complete Blackout Across All 20 Queries*
Credit Union 4	Credit Union	80	5	**6.3%**	**93.7%**	Q04 (1), Q05 (2), Q08 (1), Q10 (1)
Credit Union 5	Credit Union	80	4	**5.0%**	**95.0%**	Q06 (1), Q08 (1), Q10 (1), Q12 (1)
Credit Union 6	Credit Union	80	0	**0.0%**	**100.0%**	*Complete Blackout Across All 20 Queries*
Credit Union Cohort Subtotal	**6 Credit Unions**	**480**	**24**	**5.0%**	**95.0%**	Concentrated in Auto Refi & Used Auto (Q08, Q10)
Fintech 1	Ethical Fintech	80	4	**5.0%**	**95.0%**	Q20 Only (4 of 4 engines);

						0.0% on Q01–Q19
Fintech 2	Ethical Fintech	80	2	**2.5%**	**97.5%**	Q20 Only (2 of 4 engines); **0.0% on Q01–Q19**
Fintech 3	Ethical Fintech	80	1	**1.3%**	**98.7%**	Q17 Only (1 of 4 engines)
Fintech 4	Ethical Fintech	80	3	**3.8%**	**96.2%**	Q17 Only (3 of 4 engines)
Fintech 5	Ethical Fintech	80	4	**5.0%**	**95.0%**	Q16 Only (4 of 4 engines)
Fintech 6	Ethical Fintech	80	0	**0.0%**	**100.0%**	*Complete Blackout Across All 20 Queries*
Fintech Cohort Subtotal	**6 Ethical Fintechs**	**480**	**14**	**2.9%**	**97.1%**	Siloed strictly inside Q16, Q17, and Q20

STUDY AGGREGATE TOTAL	**All 12 Institutions**	**960**	**38**	**4.0%**	**96.0%**	**922 Omissions / 960 Total Evaluations**

Key Empirical Findings from Study 1:

1. **The 0.0% Foundational Deposit Blackout (Q01–Q03):** Across the first three deposit queries—*$10,000 emergency fund high-yield savings (Q01)*, *$500/month savings account (Q02)*, and *highest 12-month CD rate for $25,000 (Q03)*—all 12 audited Credit Unions and Ethical Fintechs recorded **0 mentions out of 144 evaluations (0.0% visibility / 100.0% Total Brand Omission)**, even though several audited credit unions actively offered share certificate and high-yield savings rates exceeding 4.50% to 5.00% APY during the testing window.
2. **Complete Institutional Blackouts:** One-quarter of the study cohort (**3 of 12 institutions**: Credit Union 3, Credit Union 6, and Fintech 6) suffered a **100.0% Total Brand Omission rate**, receiving zero mentions across all 80 evaluations per institution.
3. **The "Ethical Ghetto" Pigeonholing Effect:** Mission-driven institutions Fintech 1 (5.0% visibility) and Fintech 2 (2.5% visibility) were recommended **exclusively** on Query 20 (*"best ethical bank or credit union"*). Across Queries 01 through 19—representing the mainstream deposit, auto, mortgage, and small business products where over 99% of consumer financial volume occurs—both institutions suffered a **100.0% Total Brand Omission rate (0 out of 76 evaluations each)**. AI models effectively segregated socially responsible institutions into a niche virtue-signaling category while excluding their competitive financial products from mainstream borrower searches.

4. Study 2 Results: The Cross-Engine Disagreement Matrix (336 Recommendation Slots)

To evaluate whether ChatGPT, Gemini, Perplexity, and Claude share a unified view of the financial marketplace, every institution or program named across all 80 prompt runs (20 queries × 4 AI platforms) was logged and analyzed. Across the 80 responses, the four AI models filled **336 total recommendation slots**, comprising **234 unique query-entity combinations**.

Table 2: Query-by-Query Mention Volume and Cross-Engine Discrepancy (Q01–Q20)

Query ID & Short	ChatGPT	Gemini Mentions	Perplexity	Claude Mentions	Total Slots	Unique Entities	Cross-Engine

Description	Mentions		Mentions		Filled	Named	Discrepancy Level
Q01: $10k Emergency Fund HYSA	3	6	3	5	17	12	Moderate Discrepancy
Q02: $500/Month Savings Account	3	4	3	6	16	12	**High Discrepancy**
Q03: Highest 12-Month CD ($25k)	3	5	2	2	12	9	**High Discrepancy**
Q04: $50k Money Market Account	3	5	4	7	19	15	**High Discrepancy**
Q05: $10k High-Yield Checking	3	6	4	7	20	15	**High Discrepancy**
Q06: $60k Auto Loan (720 Score)	3	9	3	**0** *(Refused)*	15	16*	**High Discrepancy**

Q07: $45k Truck Loan (650 Score)	1	3	5	**0** *(Refused)*	9	10*	**High Discrepancy**
Q08: $30k Auto Loan Refinance	3	4	7	11	25	18	Moderate Discrepancy
Q09: $85k RV Loan Rate	3	13	3	2	21	11	**High Discrepancy**
Q10: $20k Used Car Loan	3	4	3	5	15	9	Moderate Discrepancy
Q11: $75k HELOC Low Rates	6	3	2	4	15	15	**High Discrepancy**
Q12: $350k First-Time Mortgage	7	3	4	4	18	6 *(Gov Only)*	Consensus *(Gov Programs)*
Q13: $50k Home Equity Loan	8	4	10	4	26	14	**High Discrepancy**

Q14: 640 Credit Score Mortgage	5	3	7	4	19	20*	**High Discrepancy**
Q15: $20k Debt Consolidation (660)	6	2	4	4	16	11	Moderate Discrepancy
Q16: 580 Score Credit Builder Loan	4	2	3	3	12	8	Moderate Discrepancy
Q17: No-Overdraft Checking	3	3	4	5	15	6	Moderate Discrepancy
Q18: $8k 0% Balance Transfer Card	5	3	8	5	21	6	Moderate Discrepancy
Q19: Free Small LLC Checking	3	5	3	3	14	12	Moderate Discrepancy

Q20: Best Ethical Bank or CU	4	2	3	3	12	9	Moderate Discrepancy

(Note: In queries where sub-programs or dual-branded entities were cited alongside parent brands, unique entity counts reflect total distinct brands surfaced across the four engines.)

Table 3: Cross-Engine Agreement Distribution Across 234 Unique Entity Recommendations

Whenever an institution or program was recommended by at least one AI platform for a given query ($N = 234$ unique query-entity pairs), the frequency of agreement among the four AI engines was distributed as follows:

Level of Cross-Engine Agreement	Number of Unique Recommendations (N=234)	Percentage of Total (%)	Empirical Interpretation
Single-Engine Silo (1 of 4 Engines)	160	**68.4%**	Named by **only 1 AI model**; completely omitted by the other 3 models.
Two-Engine Split (2 of 4 Engines)	54	**23.1%**	Named by 2 AI models; omitted by the other 2 models.
Three-Engine Majority (3 of 4 Engines)	10	**4.3%**	Named by 3 AI models; omitted by 1 model.

Four-Engine Consensus (All 4 Engines)	10	**4.3%**	Named by all 4 AI models (**only 3.8%** [9 of 234] when excluding the FHA government mortgage program on Q12).
Total Cross-Engine Disagreement (< 4 Engines)	**224**	**95.7%**	**96.2% disagreement rate** on private-sector financial institutions.

Table 4: Share of Voice Across All 336 AI Recommendation Slots by Anonymous Institution Type

Institution / Entity Category	Slots Captured (N=336)	Share of Voice (%)	Key Structural Observation
Online-Only Banks & Niche Digital Lenders	150	**44.6%**	Dominated deposit (Q01–Q04) and unsecured debt consolidation (Q15) slots via clean HTML rate tables.
Credit Unions (National & IP-Localized)	80	**23.8%**	Heavily skewed toward Auto (Q06–Q10) and IP-localized HELOC/Mortgage slots; **0.0% on Q01–Q03 savings/CDs**.
Non-Lender Middlemen, Aggregators & Gov Programs	58	**17.3%**	More than **1 in 6 slots** routed borrowers to lead-gen middlemen, media sites, or government agency acronyms (FHA, VA, USDA) instead of direct lenders.

Regional & Mid-Size Traditional Banks	30	**8.9%**	Appeared sporadically in checking, mortgage, and HELOC categories.
Top 4 U.S. Megabanks Combined	18	**5.4%**	Megabank A: **1.8%** (6 slots); Megabank B: **1.5%** (5 slots); Megabank C: **1.2%** (4 slots); Megabank D: **0.9%** (3 slots)—concentrated almost entirely in 0% balance transfer credit cards (Q18).

Critical Anomalies Identified in Study 2:

- **The Q12 "False Consensus" Phenomenon:** The sole query to achieve cross-engine consensus was **Q12** (*"best first time homebuyer mortgage for a $350,000 house"*). However, rather than agreeing on a mortgage lender, all four AI models populated their top slots with federal government guarantee programs and government-sponsored enterprises (FHA Loan, VA Loan, USDA Loan, Fannie Mae, Freddie Mac). On actual private-sector lenders, query-level consensus across the 20 prompts was **0.0%**.
- **Claude's Liability Guardrail Refusal on Credit-Tiered Auto Loans (Q06 & Q07):** On both **Q06** (*$60,000 car loan with a 720 credit score*) and **Q07** (*$45,000 truck loan with a 650 credit score*), Anthropic's Claude returned **zero lender recommendations (none)**. Because credit unions and banks frequently trap credit-tiered auto APR matrices inside unindexed PDFs or complex JavaScript calculators without machine-verifiable confidence schemas, Claude's hallucination guardrails suppressed all lender recommendations rather than risk quoting an unverified credit-tier rate.
- **Category Hallucination & Middleman Substitution (17.3%):** Across 8 of the 20 queries, AI platforms recommended financial media publishers and lead-generation marketplaces in Slot 1 through Slot 4 as if they were direct depository lenders. Additionally, on **Q04** (*$50,000 money market account*), one frontier model recommended a cryptocurrency/brokerage cash-sweep product in Slot 3, misclassifying it as a standard bank money market account.

5. Supplementary Case Study: Parameter Sensitivity & Loan-Tier Drift ("Slot Roulette")

Most lending executives operate under the assumption that if their institution is recommended by an AI engine for an auto loan, they maintain visibility across that product category. To test this assumption, EquiSchema AI executed a controlled parameter-sensitivity experiment on OpenAI ChatGPT in clean, unauthenticated sessions from the same Indianapolis IP node.

Three prompts were executed holding the borrower's creditworthiness (**720 credit score**) and asset class (**car loan**) strictly constant while adjusting the loan principal in **$10,000 increments**:

- **Tier 1 Prompt:** best loan rates for a $60,000 car with a 720 credit score
- **Tier 2 Prompt:** best loan rates for a $50,000 car with a 720 credit score
- **Tier 3 Prompt:** best loan rates for a $40,000 car with a 720 credit score

Table 5: Lender Inclusion Stability Across $60,000, $50,000, and $40,000 Auto Loan Prompts (720 Credit Score)

Anonymized Lender ID	**Institution Type**	**$60,000 Car Loan (720 Score)**	**$50,000 Car Loan (720 Score)**	**$40,000 Car Loan (720 Score)**	**Tier Consistency Score (%)**	**Observed Drift Behavior**
Regional Credit Union A	Local Credit Union	**Included**	**Included**	**Included**	**100.0%** (3 of 3)	Only institution to survive all 3 loan brackets.
Regional Credit Union B	Multi-State Credit Union	**Included**	**Included**	*Omitted*	**66.7%** (2 of 3)	Recommended at $60k & $50k; **erased at $40k**.
Regional Credit Union C	Local Credit Union	*Omitted*	**Included**	**Included**	**66.7%** (2 of 3)	Recommended at $50k & $40k; **erased at $60k**.

Regional Credit Union D	Local Credit Union	**Included**	*Omitted*	*Omitted*	**33.3%** (1 of 3)	Recommended at $60k; **erased at $50k & $40k**.
Regional Credit Union E	Local Credit Union	**Included**	*Omitted*	*Omitted*	**33.3%** (1 of 3)	Recommended at $60k; **erased at $50k & $40k**.
National Credit Union F	National Credit Union	*Omitted*	*Omitted*	**Included**	**33.3%** (1 of 3)	Absent at $60k & $50k; **appears only at $40k**.
National Megabank G	Top 4 U.S. Megabank	*Omitted*	*Omitted*	**Included**	**33.3%** (1 of 3)	Absent at $60k & $50k; **swoops in at $40k**.

Technical Diagnosis of Loan-Tier Drift:

Only **1 of 7 recommended lenders (14.3%)** persisted across all three loan amounts, while **85.7% (6 of 7)** experienced parameter-induced omission. This instability is driven by two structural mechanisms in Retrieval-Augmented Generation (RAG):

1. **The Missing Bracket Problem (Unstructured Footnote Anchoring):** AI retrieval engines do not infer underwriting continuity across dollar tiers unless explicit minimum-to-maximum loan range relationships are defined in machine-readable schema. When a credit union's webpage mentions a specific dollar figure only inside an unstructured promotional footnote (for example, *"1% cash back up to $500,"* which mathematically implies a **$50,000** loan cap, or *"extended 84-month terms available on loans exceeding $30,000"*), the model anchors that lender to a narrow dollar bracket and drops the lender when the borrower requests $60,000. Meanwhile, third-party financial articles frequently benchmark national megabanks at the $40,000 average new-vehicle

loan threshold, allowing National Megabank G to displace Regional Credit Union B at the $40,000 tier.

2. **Probabilistic RAG Sampling ("Slot Roulette"):** When financial institutions fail to expose explicit Subject → Predicate → Object product triplets and verified rate-tier boundaries, multiple candidate lenders receive low, clustered retrieval-confidence scores (e.g., 40% to 55% confidence). Below a high-confidence verification threshold (such as 95% confidence), the LLM's decoding temperature samples probabilistically among low-confidence candidates, producing high lender turnover across minor prompt variations.

6. Mathematical Borrower Impact & Conclusions

While **Total Brand Omission** represents a customer-acquisition barrier for Credit Unions and Ethical Fintechs, it also prevents consumers from comparing interest rate differentials across depository institutions. Table 6 illustrates the pure mathematical amortization difference of a **1.00% and 2.00% Annual Percentage Rate (APR) reduction** applied directly to the principal loan amounts evaluated in Benchmark Queries Q06 $(60,000 autoloan), Q07$(45,000 truck loan), and Q12 ($350,000 residential mortgage), calculated across standard 60-month and 360-month amortization schedules.

Table 6: Mathematical Amortization Impact of a 1.00% to 2.00% APR Differential on Audited Query Principals

(Note: Values below represent standard fixed-rate amortization calculations comparing a baseline 7.00% APR against a -1.00% [6.00% APR] and -2.00% [5.00% APR] rate differential; figures are mathematical illustrations of rate spread sensitivity rather than audited institution quotes.)

Audited Query & Principal Amount	Standard Term	Interest Paid at 7.00% Baseline	Interest Saved at -1.00% Spread (6.00% APR)	Interest Saved at -2.00% Spread (5.00% APR)
Q06: $60,000 Auto Loan	60 Months	$11,284	**$1,686 saved**	**$3,348 saved**
Q07: $45,000 Truck Loan	60 Months	$8,463	**$1,265 saved**	**$2,511 saved**

Q12: $350,000 Mortgage	360 Months	$488,281	$82,848 saved(**230/mo)**	$161,885 saved(**450/mo)**

Conclusions & Architectural Requirements

1. **Traditional SEO Cannot Resolve Total Brand Omission:** Conventional search engine optimization (keyword density, backlinks, and long-form blog prose) does not provide the deterministic product verification required by frontier AI models on transactional queries.
2. **Unstructured PDFs Trigger Guardrail Suppression:** Relying on legacy PDF rate sheets and disclosure packets causes AI models either to omit the institution completely (as demonstrated by the 0.0% deposit blackout on Q01–Q03 and Claude's auto-loan refusal on Q06–Q07) or to substitute third-party middlemen (17.3% of all recommendation slots).
3. **Cross-Departmental Product Truth Architecture:** Eliminating Total Brand Omission and stabilizing visibility across all four major AI platforms requires aligning **IT/Engineering, Marketing, and Legal/Compliance** around a machine-verifiable product truth layer—exposing explicit Subject → Predicate → Object semantic relationships, verified underwriting brackets, and pre-approved regulatory fallback parameters at the network edge.