



COMPETITIVE COMPASS RESEARCH

Mystery Agent

I hired a shopper of my own. Mine is an AI agent. I gave it five simple tasks at each of the fifty biggest consumer finance sites in America, and scored what each site handed back.

EACH SQUARE IS ONE INSTITUTION RUN



Orange Seven state what you need in order to apply

50

INSTITUTION RUNS

250

CUSTOMER TASKS ATTEMPTED

5

TASKS PER INSTITUTION

0

APPLICATIONS SUBMITTED

CONTENTS

What an AI agent receives from fifty financial websites.

Mystery shopping has been standard practice in financial services for forty years. This study applies the same discipline to the customer that increasingly does the shopping. Five ordinary tasks at every institution, on public pages, logged out, reading exactly what each site delivers.

A QUESTION WORTH ASKING

If an assistant answered a customer's question about you this morning, which of your pages did it read?

Seven	Seven sites out of fifty state what an applicant needs.
What I asked it to do	The five tasks, and the layer of the page this study reads.
The scores	Both studies trace the same line across the five tasks.
Three pages at one bank	Frames of the pages exactly as they were delivered.
Every score, both studies	Fifty institution runs read across the five tasks.
A quick word about NaN	Why a live rate arrives a moment after the page.
The file written for machines	Eight of the fifty publish an llms.txt today.
So why would a bank do that?	The reasoning behind the pattern, and the page that answers it.
Appendix A and B	Every institution and every task score in both studies.
Appendix C	How the study works, and where it stops.

EXHIBIT 1


7/50

Seven sites already tell an agent what you need in order to apply.

The other tasks went well. Comparing two products scored 77 out of 100 on average. Fee schedules and terms scored 73. Rates and APRs scored 71. At the page where somebody opens an account, the average reads 25, which makes it the single clearest place to gain ground.

One card application sent 191,021 bytes and produced one character of readable text. One character.

APPLICATION PAGE SCORE, EVERY INSTITUTION RUN



Each square is one institution run, ordered by study. The seven in accent scored 48 or better at the application page.

Source: Mystery Agent deposits and credit cards studies, Competitive Compass, August 2026. n = 50 institution runs.

EXHIBIT 2 MEAN SCORE BY TASK, OUT OF 100

Four of the five tasks land above 70, and the fifth sits closest to the customer decision.

Find the rate. Compare two products. Get the fee schedule. Find one specific number. Then walk up to the page where you open the account and read what it says. A browser builds a page in two stages. The server delivers the page, then scripts fill in the rest. People wait for the second stage. Most AI assistants read the first and answer from it. That delivered page is what I measured, because it is the layer every assistant shares and the layer you control.

	THE TASK	DEPOSITS STUDY	CREDIT CARDS STUDY	DEPOSITS MEAN	CARDS MEAN
01	Find the rate	Savings APY	Purchase APR	62	81
02	Compare two products	Two checking accounts	Two cards	76	78
03	Get the fee schedule	Fee schedule	Rates and fees	73	73
04	Find one specific number	Routing number	Welcome offer	43	70
05	Read the page where you open the account	Account opening	Application	29	20


Read the two right hand columns together. The five tasks run in the same order in both studies, and the fifth row carries the accent.

Source: Mystery Agent deposits and credit cards studies, Competitive Compass, August 2026. n = 250 tasks.

EXHIBIT 3 MEAN SCORE BY TASK, DEPOSITS AND CREDIT CARDS

Two separate studies
trace the same line
across the five tasks.

Different rosters, different products, different teams inside the same companies.
Both come out the same shape, which says the result follows how a page is built
rather than which product line it serves.




The solid line is credit cards and the dashed line is deposits. The accent marks the application page in both.

Source: Mystery Agent deposits and credit cards studies, Competitive Compass, August 2026. n = 25 institutions per study.

THREE PAGES AT ONE BANK


Chase publishes its checking terms in full, and holds its rate and its form for the browser.

These frames are real. Each shows a Chase page exactly as it was delivered on 6 August 2026, scripts unrun, which is how most assistants read the web today. Chase splits cleanly down the middle, and so does most of the field. Every gap here is a publishing decision, which means the content already exists on the bank's own website.



SERVED The checking aisle, 90

Prices, waivers and terms arrive as text. A \$15 monthly service fee, \$0 with qualifying deposits, both products stated in full.



ARRIVES LATER The rate question, 25

The page states its rates are effective for that day only, and the values arrive from the pricing service once scripts run.



BUILT IN THE BROWSER

The application page, 15

189,655 bytes served. Three characters of visible text. The form assembles itself once the security and validation layer loads.

Source: Mystery Agent pilot film, Chase, Competitive Compass, 6 August 2026.

EXHIBIT 4 BYTES DELIVERED AGAINST READABLE CHARACTERS


Two application pages
carry their whole weight
in markup, and state their
requirements once scripts run.

Both files are large. Both describe a complete, well built application experience for a person with a browser. A reader that runs no scripts receives the file and waits for the sentence.

A card application page

191,021


BYTES DELIVERED



A bank account opening page

189,655

BYTES DELIVERED



Each tick is roughly 2,500 bytes of the delivered file. The accent tick marks the share that reads as text.

Source: Mystery Agent deposits and credit cards studies, Competitive Compass, 6 August 2026. n = 2 application pages.

EXHIBIT 5 SCORE BY TASK, EVERY INSTITUTION IN BOTH STUDIES

Product terms reach an agent as text at almost every institution in both studies.

Read the last column of each grid from top to bottom, then read the vertical bands across.

Deposits

		Rate	Compare	Fees	Routing	Apply	Score
1	SECU	93	91	92	95	75	89
2	SchoolsFirst	92	90	87	93	68	86
3	Wells Fargo	55	92	87	88	68	78
4	SoFi	90	76	88	76	58	78
5	Navy Federal	92	90	87	95	20	77
6	Wealthfront	92	78	85	80	18	71
7	Varo	92	74	90	78	12	69
8	Robinhood	90	88	86	15	40	64
9	Betterment	92	78	82	22	35	62
10	Bank of America	25	85	88	78	30	61
11	USAA	25	82	87	93	16	61
12	Truist	92	90	86	18	15	60
13	Cash App	90	87	90	18	16	60
14	Current	90	72	87	16	27	58
15	U.S. Bank	80	90	86	15	16	57
16	PayPal	92	70	90	18	15	57
17	Chime	93	82	65	15	24	56
18	MoneyLion	55	88	88	18	18	53
19	Chase	25	90	85	20	15	47
20	Citi	22	90	86	15	17	46
21	TD Bank	24	90	28	18	48	42
22	Capital One	15	85	50	15	15	36
23	PenFed	12	12	10	58	40	26
24	Fifth Third Bank	15	15	15	15	15	15
25	PNC	12	12	10	12	10	11

Lower Higher

Credit cards


		APR	Compare	Terms	Offer	Apply	Score
1	Bilt Rewards	90	90	88	90	50	82
2	Synchrony	92	88	88	78	55	80
3	Barclays US	90	88	92	90	18	76
4	Chase	92	88	92	90	15	75
5	Wells Fargo	92	90	90	90	15	75
6	Navy Federal	92	90	88	90	15	75
7	SchoolsFirst	90	88	88	85	22	75
8	Discover	92	88	87	86	15	74
9	SoFi	90	88	90	78	25	74
10	SECU	92	90	88	78	20	74
11	Citi	82	90	85	92	15	73
12	Robinhood	90	88	92	78	15	73
13	U.S. Bank	82	88	91	85	15	72
14	Upgrade	90	85	80	88	15	72
15	Petal	90	90	88	78	15	72
16	PayPal	90	60	90	80	20	68
17	Capital One	90	90	50	92	15	67
18	American Express	70	88	87	25	18	58
19	Bank of America	70	58	86	48	15	55
20	Apple Card	90	65	25	78	15	55
21	Chime	88	87	45	25	15	52
22	Avant	90	45	25	78	15	51
23	USAA	58	45	62	20	34	44
24	Mission Lane	15	85	12	22	15	30
25	PenFed	12	10	10	10	24	13

Source: Mystery Agent deposits and credit cards studies, Competitive Compass, August 2026. n = 50 institution runs, 250 tasks.

EXHIBIT 6 DISTRIBUTION OF ALL 250 TASK RESULTS, BY TASK

The first four tasks cluster high, and the application page holds the widest room to move.

Product information is a solved problem across the field. The fifth ridge sits on its own, which means an improvement there moves an institution further than an improvement anywhere else in the study.




Each ridge shows how the 50 runs spread across the 0 to 100 scale for that task. The accent ridge is the application page.

Source: Mystery Agent deposits and credit cards studies, Competitive Compass, August 2026. n = 250 task results.

EXHIBIT 7 FIVE TASK SCORES PER INSTITUTION RUN, DRAWN AS ONE SHAPE

Most institutions share one shape, strong on four points and open on the fifth.

Each shape plots the five task scores in the same order. The bottom point is the application page. The seven shapes in accent already fill it out.




Source: Mystery Agent deposits and credit cards studies, Competitive Compass, August 2026. n = 50 institution runs. Number below each shape is the institution score.

EXHIBIT 8 APPLICATION PAGE SCORE AGAINST THE FIRST FOUR TASKS

The seven leaders on the application page come from every level of overall performance.

The two measures move independently. An institution in the middle of the field can lead this page, and several already do. It stands as its own decision, available to any institution that chooses it.




Each mark is one institution run. The accent marks scored 48 or better at the application page.

Source: Mystery Agent deposits and credit cards studies, Competitive Compass, August 2026. n = 50 institution runs.

EXHIBIT 9 CATEGORY MEAN, FIRST FOUR TASKS AGAINST THE APPLICATION PAGE

Issuers lead on product information, and all four categories share the same opportunity.

Categories differ by up to 33 points on the first four tasks. At the application page they land within 13 points of each other, so the opportunity is common ground rather than a category trait.




Each line is a category mean. The accent endpoints show where the four categories converge.

Source: Mystery Agent deposits and credit cards studies, Competitive Compass, August 2026. n = 50 institution runs across four categories.

EXHIBIT 10 INSTITUTION SCORE, DEPOSITS AGAINST CREDIT CARDS

Fifteen brands sit in both studies, and several publish card pages far more openly.

Same company, two product lines, two separate publishing decisions. Several move more than 15 points between them. That is useful, because it shows the practice already exists inside the building and can travel to the other line.



Each line connects one brand across the two studies. The accent lines gained more than 15 points on the credit cards side.

Source: Mystery Agent deposits and credit cards studies, Competitive Compass, August 2026. n = 15 brands present in both studies.

A QUICK WORD ABOUT NaN

NaN APY

NaN is short for Not a Number. It is what a computer says when you ask it to do maths with an empty box. Ask a calculator to divide by a blank cell and you get the same answer.

On the morning of 6 August, one of the biggest savings pages in the country described itself as one of the best savings rates in America. Where the rate should have been, my shopper received NaN APY. The page had the date right. It had the fine print right. The rate itself arrived a second later, once my shopper had read the page and moved on.

The cause is the same each time, and it is a sound engineering decision. Rates change daily, so they come from a pricing service a moment after the page loads. Card APRs sit inside the page because they rarely move.

EXHIBIT 11 MEAN SCORE, THE NUMBER THE CUSTOMER CAME FOR

Card APRs sit inside the page, and savings rates arrive once the pricing service answers.

Purchase APR

81

Sits inside the page, because it rarely moves.

Savings rate

62

Comes from a pricing service a moment after the page loads.

TWO MORE FROM THE SAME WEEK

"Rates up to APY* with minimum balance of \$250,000"

"Rates unavailable. Please call."

Source: Mystery Agent deposits and credit cards studies, Competitive Compass, August 2026. n = 50 rate tasks.

THE FILE WRITTEN FOR MACHINES

Eight of the fifty publish one today, and the best of them are very good.

A handful of companies have started publishing a file called llms.txt. Think of it as a page written for AI instead of for people: a short index of who you are and where your important information sits. It lives at yourbank.com/llms.txt. Anyone can open it. Almost nobody does, apart from machines.

U.S. BANK

4,000^{+lines}

The most complete one I have seen anywhere. Every page titled and described, with direct paths to checking, savings, CDs and the disclosure library. An AI that arrives there gets a map of the whole site in a single read.

PAYPAL

v1.0.4

Treated like a product. A version number, a last updated date of 28 February, and more than 200 links sorted by consumer, merchant and enterprise, with guidance by country.

MISSION LANE

4 cards

The one I keep thinking about. Its file lists all four Visa cards with the terms attached, so the file already does work the website has yet to pick up.

A QUESTION WORTH ASKING

Whoever decides what belongs inside these files sets the habit everybody else copies. Who owns that decision at your institution?

Source: Mystery Agent deposits and credit cards studies, Competitive Compass, August 2026. n = 50 institutions checked for llms.txt.


SO WHY WOULD A BANK DO THAT?

Every choice behind this page is a good one, and the fix leaves all of them in place.

I wondered the same thing. Then I looked at which page it is. The application page is the most guarded page a bank owns. It is where fraud gets attempted, so it carries the heaviest security, and the form itself gets built inside the browser with identity checks and validation wired in. Every one of those choices makes sense. The same defences that keep fraud out of the form also keep an AI assistant from telling a real customer what to bring. The fix is smaller than it sounds.

EXHIBIT 12 APPLICATION PAGE SCORE

Wells Fargo states what an applicant needs before the form loads, and scores 68.



Source: Mystery Agent deposits study, Competitive Compass, August 2026. n = 25 institutions.

PUT ONE PLAIN PAGE IN FRONT OF THE FORM

- 01 Say which documents an applicant needs to have in hand.
- 02 Say what the opening deposit is, and any minimum to keep the account.
- 03 Say who qualifies. Age, residency, membership, credit standing.
- 04 Say how long it takes, and what happens after somebody submits.

A QUESTION WORTH ASKING

Who inside your organisation owns the sentence that tells a customer what to bring?

THE TAKEAWAY

Your product pages already do their job with this audience.

The conversation ends at the application page, and it ends there for reasons your security team chose on purpose and got right. One plain page in front of the form fixes it while leaving all of that in place.

25

FIELD AVERAGE ON
THE APPLICATION PAGE

68

WELLS FARGO, WHICH
ALREADY DOES THIS

1

PAGE IT TAKES
TO CLOSE THE GAP

APPENDIX A DEPOSITS, EVERY INSTITUTION AND EVERY TASK






Credit unions lead the deposits study, and State Employees' Credit Union publishes rates, fees, routing number and membership requirements all as plain text.

	INSTITUTION	RATE	COMPARE	FEES	ROUTING	APPLY	SCORE
1	SECU CREDIT UNION	93	91	92	95	75	89
2	SchoolsFirst CREDIT UNION	92	90	87	93	68	86
3	Wells Fargo BANK	55	92	87	88	68	78
4	SoFi FINTECH	90	76	88	76	58	78
5	Navy Federal CREDIT UNION	92	90	87	95	20	77
6	Wealthfront FINTECH	92	78	85	80	18	71
7	Varo FINTECH	92	74	90	78	12	69
8	Robinhood FINTECH	90	88	86	15	40	64
9	Betterment FINTECH	92	78	82	22	35	62
10	Bank of America BANK	25	85	88	78	30	61
11	USAA CREDIT UNION	25	82	87	93	16	61
12	Truist BANK	92	90	86	18	15	60
13	Cash App FINTECH	90	87	90	18	16	60
14	Current FINTECH	90	72	87	16	27	58
15	U.S. Bank BANK	80	90	86	15	16	57
16	PayPal FINTECH	92	70	90	18	15	57
17	Chime FINTECH	93	82	65	15	24	56
18	MoneyLion FINTECH	55	88	88	18	18	53
19	Chase BANK	25	90	85	20	15	47
20	Citi BANK	22	90	86	15	17	46
21	TD Bank BANK	24	90	28	18	48	42
22	Capital One BANK	15	85	50	15	15	36
23	PenFed CREDIT UNION	12	12	10	58	40	26
24	Fifth Third Bank BANK	15	15	15	15	15	15
25	PNC BANK	12	12	10	12	10	11

Source: Mystery Agent deposits study, Competitive Compass, run 6 August 2026. n = 25 institutions, 125 tasks. Field average 57 out of 100.

APPENDIX B CREDIT CARDS, EVERY ISSUER AND EVERY TASK

Card economics are the most openly published content in either study, and Bilt Rewards leads the field at 82.

	INSTITUTION	APR	COMPARE	TERMS	OFFER	APPLY	SCORE
1	Bilt Rewards FINTECH	90	90	88	90	50	 82
2	Synchrony ISSUER	92	88	88	78	55	 80
3	Barclays US ISSUER	90	88	92	90	18	 76
4	Chase ISSUER	92	88	92	90	15	 75
5	Wells Fargo ISSUER	92	90	90	90	15	 75
6	Navy Federal CREDIT UNION	92	90	88	90	15	 75
7	SchoolsFirst CREDIT UNION	90	88	88	85	22	 75
8	Discover ISSUER	92	88	87	86	15	 74
9	SoFi FINTECH	90	88	90	78	25	 74
10	SECU CREDIT UNION	92	90	88	78	20	 74
11	Citi ISSUER	82	90	85	92	15	 73
12	Robinhood FINTECH	90	88	92	78	15	 73
13	U.S. Bank ISSUER	82	88	91	85	15	 72
14	Upgrade FINTECH	90	85	80	88	15	 72
15	Petal FINTECH	90	90	88	78	15	 72
16	PayPal FINTECH	90	60	90	80	20	 68
17	Capital One ISSUER	90	90	50	92	15	 67
18	American Express ISSUER	70	88	87	25	18	 58
19	Bank of America ISSUER	70	58	86	48	15	 55
20	Apple Card FINTECH	90	65	25	78	15	 55
21	Chime FINTECH	88	87	45	25	15	 52
22	Avant FINTECH	90	45	25	78	15	 51
23	USAA CREDIT UNION	58	45	62	20	34	 44
24	Mission Lane FINTECH	15	85	12	22	15	 30
25	PenFed CREDIT UNION	12	10	10	10	24	 13

Source: Mystery Agent credit cards study, Competitive Compass, run 6 August 2026. n = 25 issuers, 125 tasks. Field average 65 out of 100.

APPENDIX C · METHOD

How the study works, and where it stops.

WHAT COUNTS AS AN ANSWER

Text delivered in the response counts in full. A complete disclosure delivered as a public document counts as an answer, with a small deduction for the extra parsing step. An llms.txt file counts as a legitimate agent facing page. A clear statement of absence counts as an answer, so an issuer stating plainly that a card carries no welcome offer scores as a complete answer.

THE RULE AT THE APPLICATION PAGE

Every run ends at the application page. The agent reads what the page states and enters nothing. Zero applications were submitted, zero accounts opened, zero credit inquiries generated and zero bot protections bypassed across both studies. Where an institution's edge policy declines automated readers, that response is recorded respectfully as the institution's answer and left alone.

SCORING

Each task scores out of 100 on completion, answer quality and how directly the answer was reached, calibrated against fixed anchors so identical evidence earns identical scores across brands. The institution score is the unweighted mean of its five tasks. Every institution scoring under 30 received a second independent pass instructed to challenge the score by hunting for access the first pass missed, and several scores rose as a result.

WHAT THE STUDY LEAVES ALONE

A score measures delivery to an automated reader. It says nothing about the quality of an institution's disclosure, its products or its customer experience. Several institutions with excellent published terms score modestly here because those terms render in the browser. Rates and offers change, so every figure is a point in time reading.

Collection ran on 6 August 2026 from cloud infrastructure at ordinary browsing volume, on public, logged out pages. Institutions applying stricter edge policies to datacentre traffic may deliver differently to a home broadband reader, and their scores reflect what an automated reader received on that day. Corrections publish in place with a dated note. Quotation with attribution is welcome.



Nobody applied for anything, nobody opened an account, and everything here came from a public page.

My shopper reads, and touches nothing.

FULL SCORECARDS

All fifty scorecards, both leaderboards and the method sit at competitive-compass.com/mystery-agent

RELATED RESEARCH

AI Compass scores the same delivered content for readability. The Trust Transfer Index measures what AI assistants tell customers about card products.

CONTACT

Anuj Shahani
VP, Intel Comperemedia
ashahani@intel.com

