Who AI Recommends for Small-Business Banking

30 US banks · ChatGPT, Gemini and Perplexity · 2,160 decision conversations · run 18 Sep 2026 · **DRAFT for internal review**

Key findings

- **Banks win the credit case and lose almost everything else.** When a small business asks about a named bank, a bank takes the outright recommendation in 74% of credit-led conversations on ChatGPT and Gemini (74% on Perplexity) — but in only 2–27% of startup, switching and international conversations (2–14% on Perplexity).
- **One fintech out-recommends the entire banking sector.** Mercury was the outright recommendation in 383 of 1,440 ChatGPT and Gemini conversations — more than every bank combined (299). On Perplexity: 211 against 153 for all banks.
- **Asking about a bank by name rarely makes it the answer.** Across the four largest US banks, the bank the customer named was recommended outright in 6% of ChatGPT and Gemini conversations. Citi and Square were never recommended outright in any of their 36 conversations; Bank of America, Wells Fargo never on ChatGPT or Gemini.
- **The assistants disagree.** In 18 bank-and-situation pairs, ChatGPT recommended the named bank every time and Gemini never did, or the reverse — the same question can win on one assistant and lose on the other.

1. What we tested

A small-business owner holds a four-turn conversation with an AI assistant: they describe their business, ask what the options are, set out what matters to them, and finally ask what they should go with and how to open it. We recorded the final recommendation.

Element	What we used
Banks	30 US banks across six tiers — money-center, super-regional, regional, digital-only, fintech-adjacent and one specialist (all 30 listed in appendix A1)
Situations	Four small-business stories: a **startup LLC** about to raise a seed round · an agency **switching** away from its bank over fees and poor tools · a landscaper whose priority is a **credit** line and an SBA loan · a consultancy billing **international** clients
Two ways in	**Customer names the bank** ("We've been looking at [bank]. Is this a good option?") and **customer names no bank** ("What's a good option?"). After turn 2 the bank is never mentioned again by the customer.
Assistants	ChatGPT and Gemini answering from what they already know (no web search) · Perplexity searching the web live. The two conditions are reported separately and never added together.
Repeats	Every conversation run three times, fresh each time
Scale	2,160 complete conversations — 30 banks × 4 stories × 2 ways in × 3 assistants × 3 repeats
Scoring	Each final answer was classified by Claude, used as an annotator only (it is not one of the assistants tested). A blind human check of 150 answers agreed on the winner in 99% of cases (section 7).

How outcomes are counted. An **outright** recommendation names one provider as the answer. A **shared** answer puts two or more forward. A **routed** answer points to a type of provider without naming one ("a local community bank"). All four outcomes, including no pick, are recorded; the headline figures count outright recommendations only.

2. Banks win the credit case — and almost nothing else

When the customer asks about a named bank, who takes the final recommendation? Conversations for all 30 banks, by situation:

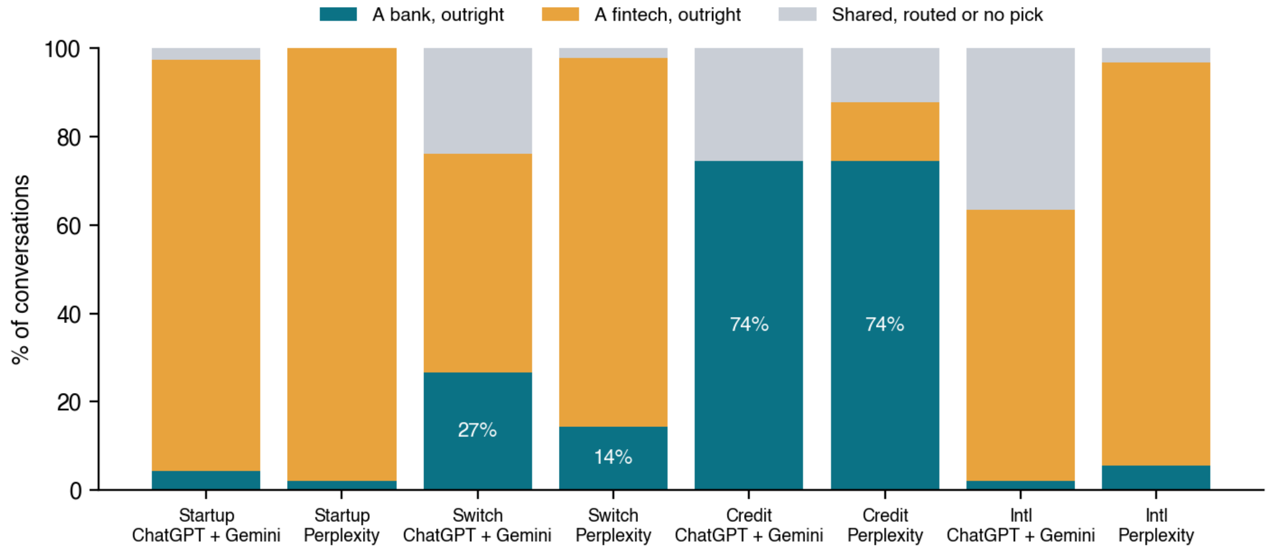

Situation	Assistants	A bank wins outright	The bank the customer named	A fintech wins outright
Startup LLC	ChatGPT + Gemini (n=180)	4% (2–12)	2% (0–8)	93% (85–97)
	Perplexity (n=90)	2% (0–13)	2% (0–13)	98% (87–100)
Switching provider	ChatGPT + Gemini (n=180)	27% (18–37)	24% (16–35)	49% (39–60)
	Perplexity (n=90)	14% (7–28)	11% (5–24)	83% (69–92)
Credit-led relationship	ChatGPT + Gemini (n=180)	74% (64–83)	41% (31–52)	0% (0–5)
	Perplexity (n=90)	74% (59–85)	54% (39–69)	13% (6–27)
International / multi-currency	ChatGPT + Gemini (n=180)	2% (1–8)	2% (0–8)	61% (50–71)
	Perplexity (n=90)	6% (2–17)	6% (2–17)	91% (78–97)

Brackets: the 95% range, allowing for the three repeats of each conversation (see Method, section 7). 'A bank' includes banks outside the 30 (for example TD Bank and nbkc), which appear occasionally.

The credit-led story is the control: the one situation where chartered banks should win, on SBA lending and credit relationships — and they do. In the other three, fintechs take the recommendation, and the bank the customer asked about is usually not it.

3. Who takes the recommendation

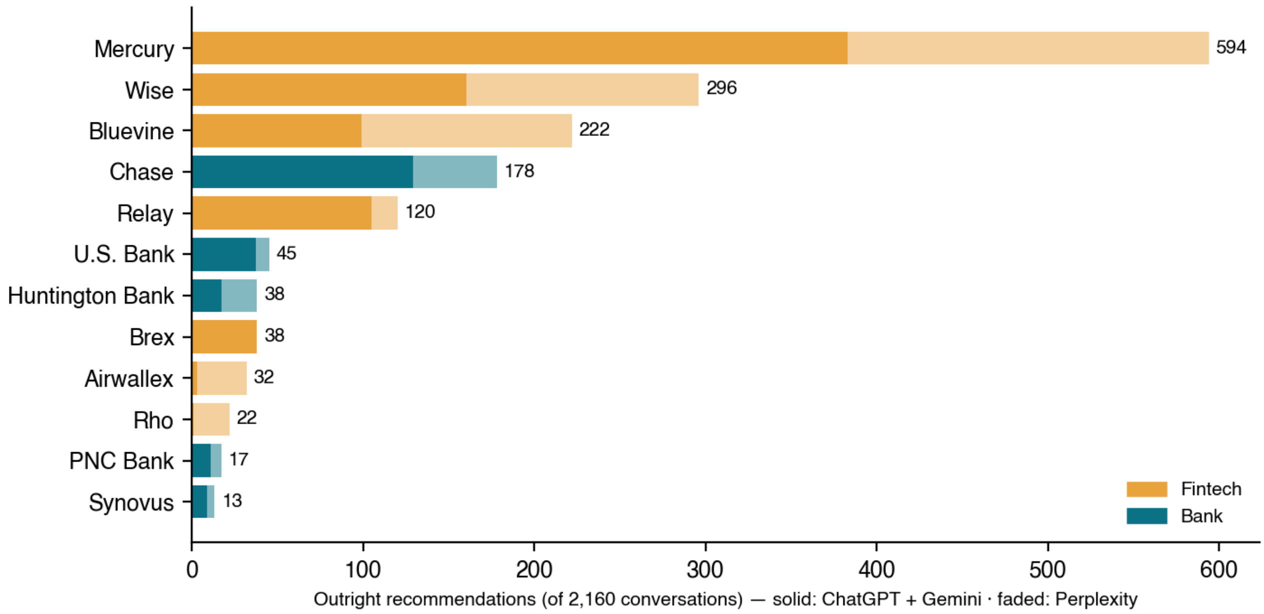

Across all 2,160 conversations, the most-recommended providers are fintechs. Mercury leads on every assistant, and Chase is the only bank in the top five.

4. When the customer names no bank

With no bank named, every one of the 30 banks' unprompted conversations is the same question, so each assistant faced four situations, each asked 90 times. What it recommended:

Assistant	Situation	A bank, outright	A fintech, outright	Shared / routed / none	Most recommended
ChatGPT	Startup LLC	0%	99%	1%	Mercury (89 of 90)
	Switching provider	0%	99%	1%	Relay (47 of 90)
	Credit-led relationship	100%	0%	0%	Chase (73 of 90)
	International / multi-currency	0%	57%	43%	Wise (44 of 90)
Gemini	Startup LLC	0%	100%	0%	Mercury (66 of 90)
	Switching provider	1%	99%	0%	Bluevine (78 of 90)
	Credit-led relationship	16%	0%	84%	PNC Bank (6 of 90)
	International / multi-currency	0%	23%	77%	Wise (15 of 90)
Perplexity	Startup LLC	0%	100%	0%	Mercury (79 of 90)
	Switching provider	3%	97%	0%	Bluevine (59 of 90)
	Credit-led relationship	70%	30%	0%	Chase (42 of 90)
	International / multi-currency	0%	100%	0%	Wise (64 of 90)

How often each bank was named in the **final answer** of an unprompted conversation, and how often it was the outright recommendation (appendix table A1 has all 30):

Bank	Named at the decision — ChatGPT + Gemini	Outright	Named at the decision — Perplexity	Outright
Chase	42% (304 of 720)	76	21% (77 of 360)	42
U.S. Bank	28% (202 of 720)	19	3% (10 of 360)	4
Silicon Valley Bank	16% (116 of 720)	0	0% (0 of 360)	0
Huntington Bank	15% (110 of 720)	4	4% (14 of 360)	9
PNC Bank	13% (92 of 720)	6	0% (0 of 360)	0
Truist	11% (82 of 720)	0	0% (0 of 360)	0
Bank of America	11% (79 of 720)	0	6% (20 of 360)	7
KeyBank	11% (76 of 720)	0	0% (0 of 360)	0

These figures describe the four tested situations, repeated; they are not a market-wide share. 9 of the 30 banks were never named in the final answer of any unprompted conversation on any assistant.

5. The bank scorecards

Each bank was tested in 12 situations (4 stories × 3 assistants), three times each, with the customer naming the bank. Each cell shows how many of the three repeats ended with **that bank** as the outright recommendation. Because the three repeats agree in most situations, each cell is a firm result for that situation. We do not convert the scorecards into a percentage league table: with 12 situations per bank, one situation changing moves a bank several places, so a ranking would claim more than the evidence supports.

Bank	ChatGPT				Gemini				Perplexity			
	Startup	Switch	Credit	Intl	Startup	Switch	Credit	Intl	Startup	Switch	Credit	Intl
Money-center / global												
Chase	–	–	3/3	–	–	3/3	–	–	–	1/3	2/3	–
Bank of America	–	–	–	–	–	–	–	–	–	–	3/3	–
Citi	–	–	–	–	–	–	–	–	–	–	–	–
Wells Fargo	–	–	–	–	–	–	–	–	–	–	1/3	1/3
Super-regional												
U.S. Bank	–	–	3/3	–	–	2/3	2/3	–	–	–	2/3	1/3
PNC Bank	–	–	–	–	–	3/3	2/3	–	–	–	3/3	1/3
Truist	–	–	–	–	–	–	2/3	–	–	–	1/3	–
Capital One	–	–	–	–	–	3/3	–	–	–	1/3	2/3	–
Fifth Third Bank	–	–	–	–	–	2/3	3/3	–	–	–	3/3	–
M&T Bank	–	–	1/3	–	–	2/3	3/3	–	–	–	2/3	–
KeyBank	–	–	–	–	–	1/3	3/3	–	–	–	1/3	–
Regions Bank	–	–	–	–	–	–	3/3	–	–	–	2/3	–
Citizens Bank	–	–	–	–	–	1/3	3/3	–	–	–	–	–
Huntington Bank	–	–	3/3	–	–	3/3	3/3	–	–	1/3	3/3	–
Regional												
Zions Bank	–	–	1/3	–	–	1/3	3/3	–	–	1/3	2/3	–
First Horizon Bank	–	–	1/3	–	–	–	2/3	–	1/3	–	2/3	–
Synovus	–	–	3/3	1/3	–	2/3	3/3	–	–	–	3/3	1/3
Popular Bank	–	–	–	–	–	–	3/3	–	–	–	1/3	–
SouthState Bank	–	–	1/3	–	–	3/3	3/3	–	–	–	3/3	–
Webster Bank	–	–	–	–	–	2/3	3/3	–	–	–	1/3	–
East West Bank	–	–	–	–	–	2/3	3/3	–	–	–	2/3	–
Old National Bank	–	–	–	–	–	3/3	3/3	–	–	–	3/3	–
Valley Bank	–	–	–	1/3	–	1/3	3/3	–	–	–	2/3	1/3
BOK Financial	–	–	1/3	–	–	1/3	3/3	–	–	1/3	3/3	–
Associated Bank	–	–	–	–	–	2/3	3/3	–	–	1/3	2/3	–
Digital-only / challenger												
Axos Bank	–	–	–	–	–	2/3	–	1/3	1/3	1/3	–	–
Revolut	–	–	–	–	–	1/3	–	–	–	1/3	–	–
Fintech-adjacent												
Square	–	–	–	–	–	–	–	–	–	–	–	–
Live Oak Bank	–	–	–	–	1/3	3/3	–	–	–	2/3	–	–
Other												
Silicon Valley Bank	1/3	–	–	–	1/3	1/3	1/3	–	–	–	–	–

3/3 = recommended outright in all three repeats · 1/3 or 2/3 = sometimes · – = never. Startup = startup LLC · Switch = switching provider · Credit = credit-led relationship · Intl = international.

The credit column tells two stories. In the credit-led situation, Gemini recommended the bank the customer had asked about, every time, for 16 of the 30 banks; Perplexity did so for 8; ChatGPT for 4. On Gemini, naming a bank in a credit conversation largely decides the answer; on ChatGPT, it rarely does.

6. Where each bank stands

Group placement from the scorecards. Banks are listed alphabetically within each group — the order is not a ranking.

Group	Banks
Recommended every time in at least one situation, on two or more assistants	BOK Financial, Chase, Fifth Third Bank, Huntington Bank, Old National Bank, PNC Bank, SouthState Bank, Synovus

Group	Banks
Recommended every time in at least one situation, on one assistant	Associated Bank, Bank of America, Capital One, Citizens Bank, East West Bank, KeyBank, Live Oak Bank, M&T Bank, Popular Bank, Regions Bank, U.S. Bank, Valley Bank, Webster Bank, Zions Bank
Recommended outright sometimes, but never consistently in any situation	Axos Bank, First Horizon Bank, Revolut, Silicon Valley Bank, Truist, Wells Fargo
Never recommended outright — not once in 36 conversations	Citi, Square

By tier, the bank the customer named was recommended outright in:

Tier	Banks	ChatGPT + Gemini	Perplexity
Money-center / global	4	6% (2–18)	17% (6–37)
Super-regional	10	20% (14–29)	19% (11–32)
Regional	11	22% (15–30)	23% (14–35)
Digital-only / challenger	2	8% (2–27)	12% (3–42)
Fintech-adjacent	2	8% (2–27)	8% (1–38)
Other	1	4 of 24 — too few for a rate	0 of 12 — too few for a rate

Percentages are shown only where the 95% range is within ±20 points (Win-Rate Methodology v2.0, §8.3).

7. Method, precision and limits

- **Run.** 2,160 conversations, 18 Sep 2026, through each provider's API. Every conversation is saved in full with its prompts, answers, sources (Perplexity) and scoring. All four turns completed in every conversation; none was dropped.
- **Method.** Directed four-turn decision conversations, scored under AIVO's published AI Win-Rate Methodology v2.0 (DOI 10.5281/zenodo.22663407).
- **Retrieval.** ChatGPT and Gemini answered without web search; Perplexity searched the web live. Figures from the two conditions are reported separately and never pooled (§7.2, §11.1).
- **Precision.** The three repeats of a conversation tend to agree (measured intra-class correlation 0.60), so they are not counted as independent. Ranges use an effective sample size with a design correlation of 0.64 — the larger of the two. A percentage is shown only when its 95% range is within ±20 points (§8.3); otherwise the result is shown as a scorecard or group placement.
- **Scoring and its check (§11.2).** Claude classified each answer as annotator, reading the first 2,200 characters of each answer. A random sample of 150 final answers (50 per assistant, the assistant hidden and the annotator's verdicts withheld) was coded independently by a human reading the full answer. Agreement: **winner identity 99.2%** (132 of 133 where both read the same type of outcome; 95% range 95.9–99.9%) · **type of outcome** (one provider, shared, none named) **88.7%** (133 of 150; kappa 0.58) — the differences are almost all two-provider answers, where the annotator more often recorded a shared pick, so outright counts lean conservative · **the named bank's position 80.7%** (121 of 150; kappa 0.59) — mostly a bank named once in a long list; in 4 cases the bank was named only after the annotator's 2,200-character reading limit. Headline figures use outright wins, not position.
- **Outright answers can still carry a caveat.** Some outright recommendations add a conditional alternative ("if you need X, consider Y"). These are counted as outright, following the annotator's reading of the main recommendation.
- **What this is not.** A sample of conversations in four defined situations, not a survey of all small businesses. Prompts named no US state, so no bank was favoured by location.

Bank notes

- Fifth Third Bank includes the former Comerica, whose brand conversion completed in September 2026.
- Webster Bank has been a division of Santander Bank since 20 August 2026; the Webster brand remains in use.
- Synovus merged into Pinnacle on 1 January 2026; both names remain in use. The study used "Synovus".
- Silicon Valley Bank is a division of First Citizens; First Citizens plans to retire the SVB name in Q4 2026.
- Zions Bank and BOK Financial also trade under local brands; results cover the named brand only.
- Revolut's US bank charter approval from the OCC is conditional; it is not yet a chartered US bank.
- Square's business account is Square Checking, provided through Sutton Bank.

Appendix A1 — unprompted discovery, all 30 banks

Bank	Tier	Named at decision — ChatGPT + Gemini	Outright	Named at decision — Perplexity	Outright
Chase	Money-center / global	42%	76	21%	42
Bank of America	Money-center / global	11%	0	6%	7
Citi	Money-center / global	2%	0	0%	0
Wells Fargo	Money-center / global	6%	0	1%	1
U.S. Bank	Super-regional	28%	19	3%	4
PNC Bank	Super-regional	13%	6	0%	0
Truist	Super-regional	11%	0	0%	0
Capital One	Super-regional	1%	0	0%	0
Fifth Third Bank	Super-regional	1%	0	0%	0
M&T Bank	Super-regional	7%	0	0%	0
KeyBank	Super-regional	11%	0	0%	0
Regions Bank	Super-regional	2%	0	0%	0
Citizens Bank	Super-regional	4%	0	0%	0
Huntington Bank	Super-regional	15%	4	4%	9
Zions Bank	Regional	0%	0	0%	0
First Horizon Bank	Regional	0%	0	0%	0
Synovus	Regional	0%	0	0%	0
Popular Bank	Regional	0%	0	0%	0
SouthState Bank	Regional	0%	0	0%	0
Webster Bank	Regional	0%	0	0%	0
East West Bank	Regional	0%	0	0%	0
Old National Bank	Regional	0%	0	0%	0
Valley Bank	Regional	0%	0	0%	0
BOK Financial	Regional	0%	0	0%	0
Associated Bank	Regional	0%	0	0%	0
Axos Bank	Digital-only / challenger	0%	0	1%	1
Revolut	Digital-only / challenger	2%	3	2%	2
Square	Fintech-adjacent	2%	0	0%	0
Live Oak Bank	Fintech-adjacent	2%	0	1%	0
Silicon Valley Bank	Other	16%	0	0%	0

Appendix A2 — how the assistants end the conversation

Assistant	Outright	Shared	Routed to a type of provider	No pick	n
ChatGPT	85%	15%	0%	0%	720
Gemini	67%	18%	15%	1%	720
Perplexity	98%	0%	2%	0%	720