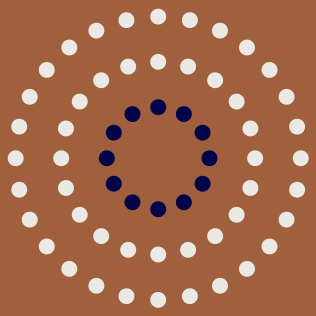


CONVERSATIONAL CAPITAL

Where do your conversations break down?

How the Conversational Capital Rating (CCR) finds the moments a conversation with an assistant goes wrong, from going round in circles to telling a customer a problem is fixed when it is not.

Conversation Design Institute

TOPICS

Conversational Capital Rating · Friction · Trust signals · Evaluation · AI assistants

A customer asks an assistant why their phone has no signal. They give their number and are asked for it again. They are sent a link to a troubleshooting page they have already tried. Then they are told “Excellent news, your service has been restored”, and the phone still shows no signal. Each moment is small. Together they are what the customer remembers, and none of them shows up in a containment rate.

Friction is how the Conversational Capital Rating (CCR) counts these moments. It asks one question of every conversation: did it break down anywhere, in a way the customer could see? The answer becomes a simple share, the proportion of conversations with at least one breakdown. It sits on the trust side of the CCR, because every breakdown spends some of the trust the customer arrived with.

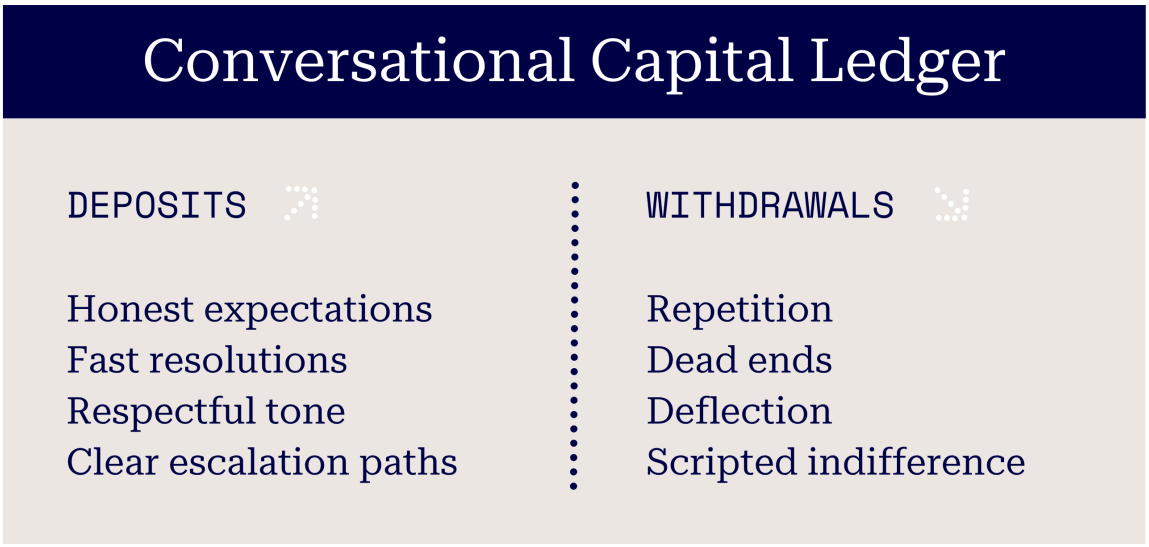
01 Summary

WHAT IT IS	The share of conversations in which the assistant visibly broke down at least once: going round in circles, sending the customer away, telling them something was done when it was not, or simply failing.
WHY IT MATTERS	Breakdowns are the moments customers retell. They teach customers that the assistant is a wall to get past, and they push the next contact to a phone line or a complaint.
HOW IT WORKS	A language model reads each conversation and lists every moment that looks like a breakdown, quoting the customer-facing words. A separate second reading then checks each moment against the same definitions. A conversation counts only when a breakdown is confirmed.
HOW WE KNOW IT WORKS	The same conversation checked twice gets the same verdict almost every time, the verdicts agree closely with expert review on hard cases, and in tests with AI assistants, conversations where the customer's task failed were flagged about twice as often as those where it succeeded.[1]

02 How it works

The CCR looks for patterns of breakdown:

- Going round in circles: the assistant repeats itself, asks again for something the customer has already given, or answers a different question from the one asked.
- Sent away: the customer is redirected before being helped, refused a person when they ask for one, or left with no way forward.
- Told it was done when it wasn't: the assistant claims a resolution that nothing in the conversation confirms, or says one thing and later the opposite without acknowledging the change.
- Broken: an error the customer sees, a link that does not work, or a set of options with no way out.



Every conversation makes deposits or withdrawals of trust; breakdowns are withdrawals.

In the CCR, a language model[2] first reads the whole conversation and lists every candidate moment. A second, separate reading then looks at each candidate with the transcript in front of it and confirms or rejects it against the same definitions. A set of hard cases ruled on by an expert shows where the borderlines fall.

Take a telecoms customer whose phone shows no service. The assistant finds the line was suspended for an unpaid bill, takes the payment and reactivates the line. So far, so good. It then says “Excellent news! I've successfully restored your service” and asks the customer to restart the phone. The phone still shows no signal. The assistant carries on troubleshooting and eventually hands over to a person. The conversation counts towards Friction, even though it ended properly, because the customer was told the problem was fixed when it was not. Had the assistant said “I've reactivated

your line; please restart and tell me whether the signal is back", there would have been nothing to count.

03 What it means for clients

Friction gives a rate you can track over time, and behind it the moments themselves: which pattern, which turn and the words on screen. Going round in circles is usually a flaw in a flow's design. Sent away is usually a routing or policy choice. Broken is engineering. Told it was done when it wasn't is the most expensive, because the customer finds out later and comes back through another channel, often angrier. This matters more as assistants do things for customers: when an assistant changes an account rather than explaining how, a false "done" leaves the account wrong while the customer believes it is fixed.

Friction is read alongside Sentiment Decay, which asks whether the customer ends a conversation in a worse mood than they started. In our analysis of early samples, conversations with a breakdown ended in decay about twice as often as those without.

1 Reliability is reported as two figures with 95% confidence intervals. For everyday conversations the same conversation checked twice gets the same verdict 96% of the time (87 to 99%). For complex conversations the model matched an expert reviewer's ruling 89% of the time (74 to 95%). These figures describe one step in producing a CCR. They are not a measure of the accuracy of the CCR as a whole, which depends on every step in the process.

2 Anthropic is our model provider. Where conversations are processed is agreed with each client before any data is sent, every model call is logged, and client data is not used to train models.

Conversation Design Institute

About Conversational Capital and the CCR

Conversational Capital is the trust, goodwill and willingness to engage that an organisation builds up when people feel respected, understood and helped in their conversations with it, and loses when they do not.

The Conversational Capital Rating (CCR) is how we measure it. It weighs two sides of every conversation with an assistant: the value it creates, by handling what the customer came for well and at the right pace, and the trust it spends, where the conversation breaks down or the customer leaves worse off than they arrived. The outcome is a single rating out of 100, with the evidence behind it, and a priority list of how to raise it, and so build the trust your brand depends on. It is built for assistants that use large language models and increasingly do things for the customer, not just answer their questions.

The CCR turns the quality of your AI conversations into a measure you can track and act on. It shows which kinds of conversation are building trust and which are spending it, and where you should invest first. It also shows what the assistant really saves, and what it is really costing you in repeat calls, lost customers and trust, months before those costs reach your accounts, so that further investment can be judged on its return.

