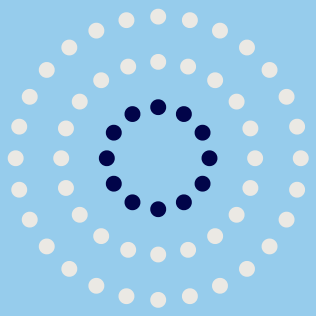


CONVERSATIONAL CAPITAL

Is your assistant handling the conversations that matter?

How the Conversational Capital Rating (CCR) judges whether an assistant takes on the conversations that matter most to customers, rather than sending them elsewhere.

Conversation Design Institute

TOPICS

Conversational Capital Rating · Stake Coverage · Interaction value · Evaluation · AI assistants

A question about opening hours and a call from a customer who cannot pay this month's bill are both conversations. An assistant that answers the first and sends the second to a phone queue has handled half its conversations. It has not handled half of what mattered to its customers.

Stake Coverage is how the CCR measures this. It does not count conversations. It asks how much of what was at stake for customers the assistant took on itself, rather than passing elsewhere.

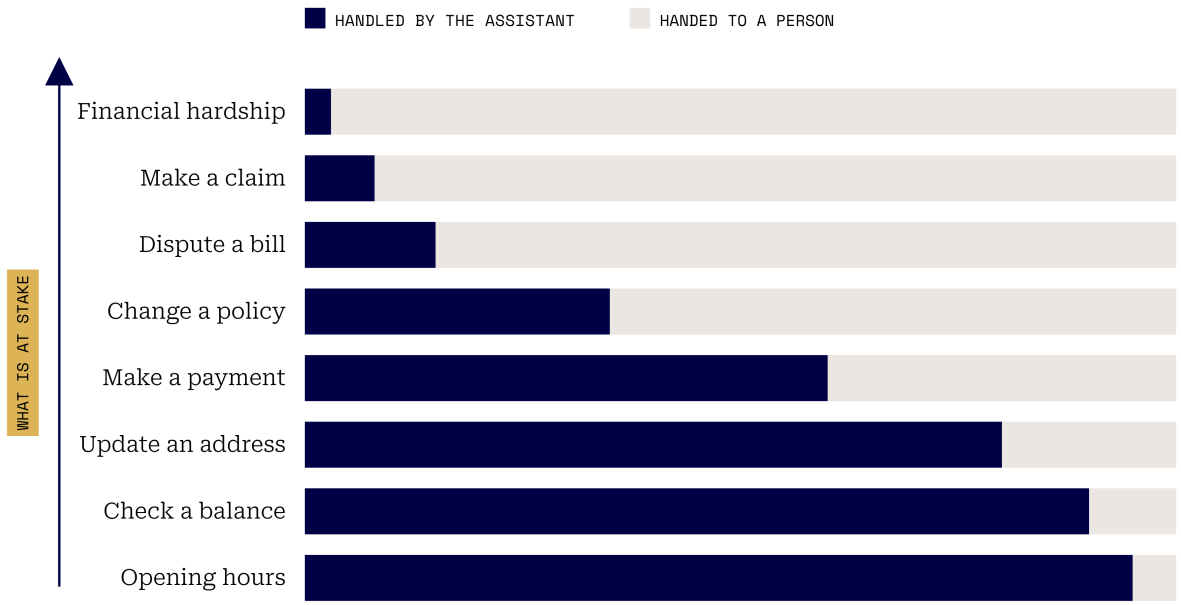
01 Summary

WHAT IT IS	A score out of 100 for how much of the weight of customers' conversations the assistant handles itself, where each conversation counts for what it was worth to the customer.
WHY IT MATTERS	Containment rates treat every conversation alike, so an assistant can look capable by keeping the easy ones and passing on the ones that matter. Customers notice which ones were passed on.
HOW IT WORKS	Every kind of conversation carries a weight set from written criteria and public complaint and regulatory evidence. A language model reads each conversation and says whether the assistant handled it, handled part of it or sent the customer elsewhere, quoting the turn that did so. A separate program checks the quote and does the arithmetic.
HOW WE KNOW IT WORKS	The weights were reviewed one by one by an expert. The model's readings agree with an independent reading by a separate review model on the hardest conversations, and the same conversation judged twice gets the same reading almost every time.[1]

02 How it works

Every reason for contact in our Conversational Capital Rating (CCR) carries a weight that reflects what a typical customer has at stake: their own money, how exposed they are likely to be, whether mishandling could let someone else into their account or lock them out of it, and how much the outcome decides the relationship with your

brand. A conversation about a disputed charge, a suspected fraud or a missed payment weighs more than one about a promotion, for example. Each weight is set from our criteria and checked against existing complaint data that shows what customers talk about. If a customer discloses a bereavement, a serious illness or a vulnerability, for example, that conversation takes the top weight, whatever it was about.



A familiar pattern: the assistant handles the easy conversations and hands off the ones that matter most.

A language model[2] then reads each conversation and looks at how the request was dealt with in the conversation. Take a telecoms provider's assistant that handles eight of its ten conversations: plan questions, an offer, a delivery date, a router reset. The two it passes to the phone line are a customer who cannot pay this month's bill and a customer disputing a charge. By count it handled eight in ten. Its Stake Coverage is well below that, because the two conversations it passed on carried more weight than any it kept. That is the pattern Stake Coverage exists to find: an assistant that looks capable on volume while the conversations customers most needed help with go elsewhere.

03 What it means for clients

Stake Coverage shows where your assistant steps back from the conversations your customers care about most. Because every conversation it passes on comes with the route and the words it used, the evidence shows which kinds of request go to the phone, which reach a dead end, and which are handed to a person. Those are the

places where extending what the assistant can do pays back most. The score also guards against a familiar trap: a containment rate that rises because the assistant has been tuned to keep the easy conversations. As with every part of the CCR, the first run on your conversations is treated as calibration, with a sample of readings checked by an expert before any CCR is reported, and your own complaint and repeat-contact data can adjust a weight where it shows a different picture.

1 Reliability is reported as two figures with 95% confidence intervals. For complex conversations, the hardest cases in development (every one the model judged as partly handled or sent elsewhere, every vulnerability disclosure, and every reading it was unsure of), the model's readings matched an independent reading by a separate review model on 74% (54 to 87%). The same conversation judged twice, independently, got the same reading 96% of the time (87 to 99%). 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.

