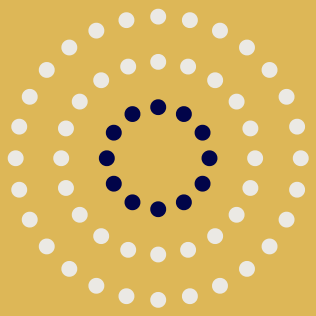


CONVERSATIONAL CAPITAL

What is your assistant worth if customers go round it?

How the Conversational Capital Rating (CCR) stops giving an assistant credit for what it can do once customers no longer rely on it, and why the point at which that happens differs by industry.

Conversation Design Institute

TOPICS

Conversational Capital Rating · Trust and capability · Trust signals · Industry · Evaluation · AI assistants

An insurer's assistant handles claims, policy changes and renewals. It is quick, polite and well built, and its containment rate is high. But many customers who change their policy through it phone the call centre two days later, because the change never saved. On every measure of what the assistant can do, it looks like a success. On the measure that matters, whether customers can rely on it, it is failing, and they have learned to go round it.

The Conversational Capital Rating (CCR) has a rule for this. What an assistant can do only counts while customers trust it enough to use it. Once enough of them have stopped relying on it, the CCR stops giving the assistant credit for its capability and holds its score down until trust is repaired.

01 Summary

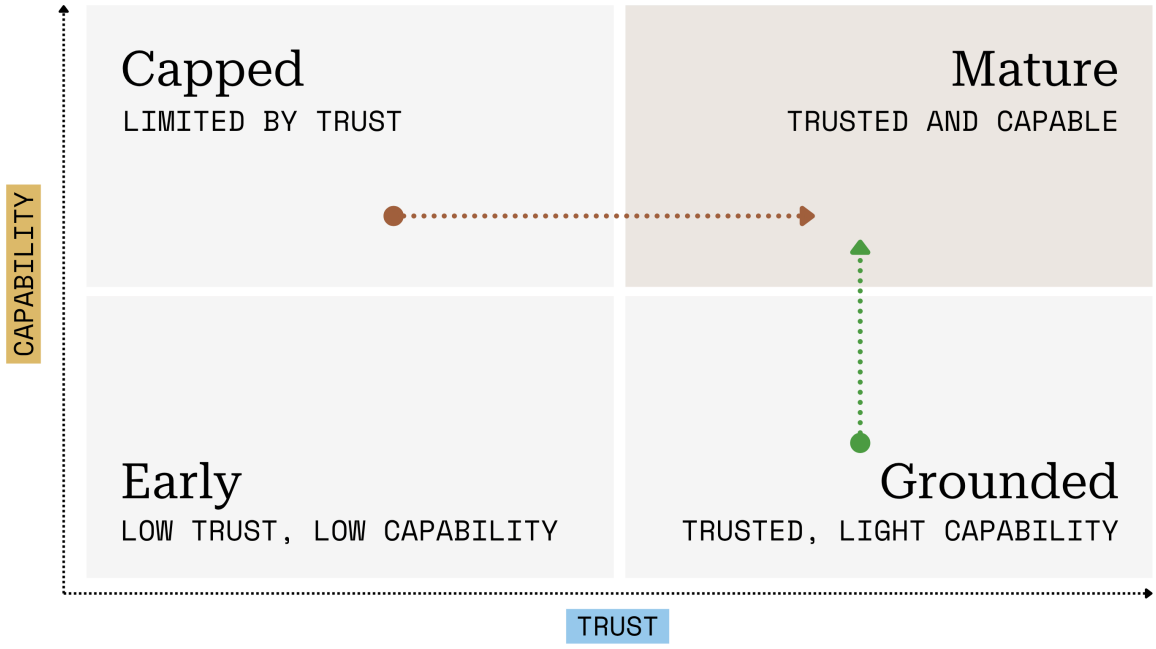
WHAT IT IS	A limit on the score. When customers have stopped relying on an assistant, its score is held down, however capable the assistant is.
WHY IT MATTERS	A clever assistant that customers go round is the most expensive kind: the investment has been made and its value is lost. Without this limit, capability could hide broken trust behind a respectable score.
HOW IT WORKS	Every failure takes a little off the score. Once failures pass a set point, the score is also held below a limit that no amount of new capability can lift until customers come back. The point is set for each industry.
HOW WE KNOW IT WORKS	The point for each industry rests on public evidence about how readily customers switch providers and what a failure costs them, reviewed by CDI.[1]

02 How it works

Trust affects a company's CCR in two ways.

Small failures cost a little. A conversation that goes round in circles, or a customer who ends it in a worse mood than they started, takes something off the score, and the more failures there are, the more each one costs.

But customers do not lose trust gradually for ever. They put up with a failing assistant for a while, then stop trying and go round it: they phone, queue or complain instead. From then on, whatever the assistant can do no longer reaches them, so the CCR stops crediting it. The score is held below the level at which an assistant counts as building trust, and stays there until customers come back. The limit only ever holds a score down; it never lifts one. Research on service quality supports this shape: customers tolerate a range of performance and react sharply once service falls below it,[2] and a good recovery can win a customer back after one failure but not after two.[3]



A capable assistant that customers do not trust is capped; the way to Mature runs through trust.

The clearest sign that customers have reached that point is how often they contact a person about something the assistant counted as handled.

Where that point sits matters, because failure is common. In a Gartner survey of 5,728 customers, only 14% of service issues were fully resolved in self-service, and only 36% even of issues customers described as very simple.[4]

Take an insurer whose assistant handles claims, policy changes and renewals, and on what it can do it scores well. But policy changes made through it often fail to save, so many customers phone afterwards to check. Enough of them go round the assistant

that its score is held down, and the report says why: the assistant is capable, and broken trust is what is holding it back. New features would not move the score, because they would land on customers who have stopped relying on it. Making policy changes save, and confirming them only once they have, would. When customers stop going round the assistant, the limit lifts and the score rises towards what its capability already justifies.

03 Why the point differs by industry

The same failure does not cost the same everywhere. For each industry the CCR sets how soon the limit applies and how far it holds the score down. Both follow from one question: when an assistant fails a customer, what happens next?

Where customers can walk away easily, as in most of retail, they do. A failing assistant is abandoned rather than tolerated, and the damage shows quickly as lost business. So the limit applies early.

Where customers cannot walk away, the damage does not go away; it lands on them. A council tenant reporting a repair, a patient booking care or someone claiming a benefit has no other provider to turn to. Their only way out of a failing assistant is to go round it, and they keep coming back to the same service, so each failure adds to the last. Because nobody leaves, the usual warning sign of customers drifting away never appears. A failing assistant does more damage where customers can't walk away. Where a failure puts health, housing, livelihood or a legal right at risk, the limit applies early here too, and holds the score down furthest.

In between are services customers can leave, but only with effort: moving a bank account, or switching insurer or mobile provider at renewal. Customers there tolerate a little more before they give up, so the limit applies a little later.

04 What it means for clients

When a score is held down this way, the report shows it separately, so a capable assistant is never mistaken for a weak one. The message is specific: the assistant can do the work, and broken trust is holding the score back. The order of work follows. Repair trust first, then add capability, because new features only pay off with customers who still rely on the assistant.

An assistant that customers go round creates repeat contacts: the business pays for the conversation, and again for the phone call or visit that follows. Every one of those repeat contacts is a saving the containment figures reported but the business

never banked. Where customers can leave, some of them also do, and they take the rest of their custom with them, not just the cost of one call. Lost trust becomes repeat calls first and lost customers later, which is why the limit applies before either shows in the accounts.

Take the insurer again, with illustrative figures. If its assistant handles 100,000 policy changes a year and a third of those customers phone afterwards to check, that is over 33,000 calls the business believes it has saved and is still paying for. If even a small share of those customers then leave at renewal, each takes at least a year's premiums with them, and often several. In a service customers cannot leave, the same failure shows up as repeat contact, complaints and the staff time to handle them.

What should a client see after fixing what the CCR flags? First, contact with people about the topics the assistant handles should fall, because customers stop checking that the assistant did its job; the savings the dashboard was already reporting start to become real. After that there should be fewer complaints and, where customers can leave, fewer leaving after a failure. The CCR itself rises as trust returns, towards what the assistant could already do.

Having these figures also changes the conversation about the return on the assistant itself. Most organisations know what their assistant cost; few can say what it returns, because containment counts conversations ended, not problems solved. The CCR identifies what the assistant saved by finishing work customers did not have to repeat, and what it lost on work they did. That gives a baseline against which the next investment, in fixes or in new features, can be judged, and it shows after each fix what that fix returned.

For services customers cannot leave, this shows something their own figures cannot. Those organisations rarely see customers go, so their usual warning never sounds. The CCR looks instead at what customers do when they cannot leave: go round the assistant, come back and complain. It shows the damage while there is still time to repair it.

1 The point for each industry is set from public evidence (switching and churn rates, complaint and ombudsman volumes, duty-of-care rules) reviewed by CDI.

2 Stodnick, M. and Marley, K.A. (2013), “A longitudinal study of the zone of tolerance”, Managing Service Quality, 23(1), 25–42.

3 Maxham, J.G. and Netemeyer, R.G. (2002), “A longitudinal study of complaining customers' evaluations of multiple service failures and recovery efforts”, Journal of Marketing, 66(4), 57–71.

4 Gartner (2024), “Gartner Survey Finds Only 14% of Customer Service Issues Are Fully Resolved in Self-Service”, press release, 19 August.

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.

