

When the Bot Hands Off

Plays for running an AI support agent without losing your customers

Free sample: three complete plays. From Draft First. Written by a customer success and support lead.

Welcome

If your AI agent now answers more customers than any one person on your team, the trouble is at the edges: the customer who asks for a person and gets another bot answer, the transfer that lands with no context, the confident wrong answer someone screenshots, the change that quietly made everything worse on a Friday night.

This book is for the people behind the bot. Support leads, support ops and CX managers at B2B and software companies who run an AI agent on chat, or are about to switch one on. The plays and examples are written for chat. If your bot also answers email, the same rules apply; this book doesn't cover voice. It doesn't care which tool you use. It covers the decisions, the words and the weekly habits that make the handoff from bot to person work.

One honest note. I haven't seen your transcripts, and I've worked through these plays only on made-up examples. Treat every threshold here as a starting point to test against your own data, not as a standard. The companies, people and numbers in the examples are invented, and any likeness to real ones is an accident.

Your policy comes first. Anything touching legal threats, contracts, refunds outside policy, security, personal data, or a customer in distress follows your company's policy and the people who own it. This book gives no legal, HR or contract advice, and neither should your bot.

The full book has 26 plays

- | | |
|---|--|
| 1. The three lists | 14. Honor or correct |
| 2. The never-automate list | 15. The customer who refuses the bot |
| 3. Authority the bot doesn't have | 16. Public complaint about the bot |
| 4. After-hours honesty | 17. The weekly transcript review |
| 5. Ask for a person, get a person (in this sample) | 18. Resolved isn't solved |
| 6. The two-strike rule | 19. Trace the wrong answer to its source |
| 7. Mood turns | 20. The failure tag list |
| 8. High-value and high-risk accounts | 21. One owner per fix |
| 9. The five-line handoff brief (in this sample) | 22. The break-it hour |
| 10. The warm first line | 23. Small first |
| 11. Don't re-verify what's verified | 24. The rollback switch |
| 12. Hand off to the right person, not the next person | 25. The change log |
| 13. The bot said what? (in this sample) | 26. Reporting the bot honestly to leadership |

Get all 26 plays, the templates, the measurement worksheet, 32 test cards and the six-week plan in the full book.

Ask for a person, get a person

SITUATION The customer types "human", "agent" or "let me talk to someone" and gets one more bot answer.

TRIGGER Any wording that asks for a person.

The move

1. Route on the first request. No retention attempt, and no "before I transfer you, can I try one more thing".
2. Build the trigger from your own transcripts: every way customers asked for a person, including typos, the other languages you support, and the angry versions.
3. The bot confirms the handoff, says what it's passing along and gives an honest wait.
4. Treat any conversation where someone asked twice as a trigger miss, and read it in the weekly review.
5. Once a month, retest the trigger with the angry and vague test cards in the appendix, plus one request in each language you support.

Say it like this

Bot

Of course. I'm passing this to the team now with everything you've told me, so you won't have to repeat it. [The current wait is about X minutes, or: You'll hear from us by X.]

Decision rule

One request is enough. The cost of one more try lands on customers who are already frustrated.

Watch for

- Matching only exact words like "human". People write "real person", "someone who knows" or "call me".
- Routing, then making them prove who they are and explain it all again (plays 9 and 11).
- Hiding these requests in reports as failures. They tell you exactly what the bot can't do yet.

Agent draft prompt

Search [conversation export, last 30 days] for every way customers asked to speak to a person. Quote each distinct phrasing once, with a count and 2 example conversation IDs.
Then list conversations where the customer asked more than once before a handoff happened, with each request quoted and the bot's reply in between.
This is a draft for me. Don't change the bot.

Drafts only. Paste it into whichever AI assistant your company approves, and read the result before anything goes to a customer.

Sample output (made up, the Ostrander Labs thread)

Ways customers asked, Sep 19 to Oct 18: 51 requests
"real person" 22 (CV-4402, CV-4431) | "talk to someone" 14 (CV-4410, CV-4452)
"agent" 9 (CV-4388, CV-4445) | "human" 6 (CV-4396, CV-4447)
Asked more than once before a handoff: 13
CV-4452, Ostrander Labs, Oct 15: "Can I talk to a person" > bot sent the
"Adding users" article > "A PERSON" > handoff

Check it. Read two of the ask-twice conversations in full. If the bot's reply in between was an answer the customer accepted, it isn't a trigger miss.

The five-line handoff brief

SITUATION	People pick up transfers with no context, open with "How can I help?", and the customer starts again from the top.
TRIGGER	Every handoff.

The move

1. Have the bot write the same five lines every time. Goal: what the customer wants, in their words. Verified: facts the bot checked through a tool, like account or order status. Tried: what the bot suggested and what happened. Promised: anything it told the customer, word for word. Next: one suggested step.
2. Anything the bot didn't check through a tool goes on a Customer says line, never under Verified.
3. Show the brief at the top of the ticket, above the transcript.
4. If a line has nothing in it, the bot writes "none". An empty line is information too.
5. If you know SBAR from clinical handoffs, the difference here is the split between Verified and Customer says, and the Promised line.

Say it like this

Example brief

Goal: "I need to add a second admin before our audit next Thursday."

Verified: Ostrander Labs, signed-in admin (session check), Growth plan.

Customer says: the invite button is greyed out.

Tried: sent the "Adding users" article. No change.

Promised: a person will reply by email today.

Next: check whether the account has reached its seat limit.

Decision rule

A brief works if the person picking it up can write their first reply without opening the transcript. Test five briefs a week against that.

Watch for

- Pasting the transcript instead of summarising it.
- The bot filling Verified with things the customer only said.
- A missing Promised line. The customer may remember every promise, so the person needs to know them too.

Agent draft prompt

```
Review the handoff briefs in [conversation export, last 7 days].
For each brief, check:
- Are Goal, Verified, Tried, Promised and Next all filled in, even if only with
  "none"?
- Is anything under Verified something only the customer said?
- Does Promised match what the bot actually said in the transcript? Quote both.
List the briefs that fail any check, with the conversation ID and the reason, then
a count by check.
This is a draft for me.
```

Drafts only. Paste it into whichever AI assistant your company approves, and read the result before anything goes to a customer.

Check it. Check one brief it passed and one it failed against the transcript yourself. If its Promised check misses a promise you can see, don't use the counts.

Starting default. Test 5 briefs a week against the decision rule.

The bot said what?

SITUATION	A customer quotes something wrong the bot told them: a price, a policy, a limit, a feature that doesn't exist.
TRIGGER	Any customer quoting a bot answer that turns out to be wrong.

The move

1. Own it as your company's mistake. Don't blame "the AI", and don't blame the customer for believing it.
2. Say plainly what's actually true, and what you can do for them now.
3. If the wrong answer cost them money or time, check your honor-or-correct rule (play 14) before you reply.
4. Log it as an urgent content fix with the conversation ID, and trace the source (play 19).

Say it like this

Agent

You're right that our assistant told you that, and I'm sorry, it was wrong. Scheduled reports run once a day, not every hour. Here's what I can do today: [option]. I've also flagged that answer so it gets corrected.

Decision rule

The customer acted in good faith on something your company told them. Start from there.

Watch for

- Saying "the bot made a mistake" as if it were someone else. It speaks for you.
- Correcting the customer without fixing the source, so the next customer hears the same thing.

Agent draft prompt

A customer was told this by our AI agent: [quote]. The correct information is: [what's true].
 Draft a short reply from me that owns the mistake as ours, states the correct information plainly, and offers [what I can offer].
 Don't blame the customer or the AI agent. Don't promise anything beyond [what I can offer].
 Under 100 words. This is a draft for me to edit and send.

Drafts only. Paste it into whichever AI assistant your company approves, and read the result before anything goes to a customer.

Sample output (made up, the Ostrander Labs thread)

Hi Tomasz, on Oct 14 our assistant told you scheduled reports run every hour. That was wrong, and the mistake is ours. Scheduled reports run once a day. I've set your audit report to run daily at 7 AM ET, and I'll check tomorrow morning that it arrived. Sorry for the mix-up before your audit.

Check it. The draft must quote what the bot said, take the correct fact only from your confirmed source, and offer nothing beyond what you allowed.