

Seven rules for an AI assistant on a live screen.

Run this in any design review of an AI assistant in a trading, wealth, payments or banking product. Every “no” is a finding. 23 questions, about 10 minutes.

01 Never block the data

The data is the product. The assistant is a guest on it.

- ☐ Is every live number still visible with the assistant open?
- ☐ Can it be moved, minimised and dismissed in one action each?
- ☐ Does opening or moving it leave the layout underneath unchanged?

02 Stage, never execute

The consequential button belongs to the person who answers for it.

- ☐ Does every prepared action end in the product’s own form, pre-filled?
- ☐ Is the final, consequential button pressed only by the user?
- ☐ Does the interface say, in plain words, that the assistant never executes?
- ☐ Do staged changes show a before and after?

03 Show your sources

An answer you can’t verify is an answer you can’t act on.

- ☐ Are answers in the product’s own numbers and units, not prose?
- ☐ Does each answer name the data it read, and when it read it?
- ☐ Can the user preview a source and jump from any number to where it lives?
- ☐ Are stale, partial or estimated inputs called out in one line?

04 Suggest what fits the screen

Suggestions come from where the user is, not a fixed menu.

- ☐ Do suggestions change with the current context?
- ☐ Are they limited to actions possible on this screen for this role?
- ☐ Is the current context shown, so users know why they see them?

05 Get noticed only when it matters

Attention is the scarcest thing on a live screen.

- ☐ Does it signal only when there is something new?
- ☐ Does it stay still when the user is focused on it, or idle?
- ☐ Is an unread reply held quietly until the user returns?

06 One pattern, every desk

Only what it can read and stage changes with the role.

- ☐ Is it the same panel, header, opening line and footer across products?
- ☐ Do role and permissions decide what it can read and stage?
- ☐ Does it use the product’s design tokens, in both themes?

07 Stay interruptible

Never leave the user wondering if it’s working.

- ☐ While it works, does it show what it is doing in plain steps?
- ☐ Can the user stop it at once, with the control where Send was?
- ☐ When a source fails, does it say so and offer a retry, with the layout unchanged?

! When to break general chat rules

No side rail: it takes a column of live data. No AI execute button, not even behind a confirm dialog. One quiet signal that means “the assistant” is fine, if it never moves while idle.