



FIGURE: T-VOICE FULL-DUPLEX SIMULATION OVERVIEW



A · USER AUDIO PIPELINE (ZOOM 32S-68S)



B · TURN-TAKING DETAIL (ZOOM 130S-175S)

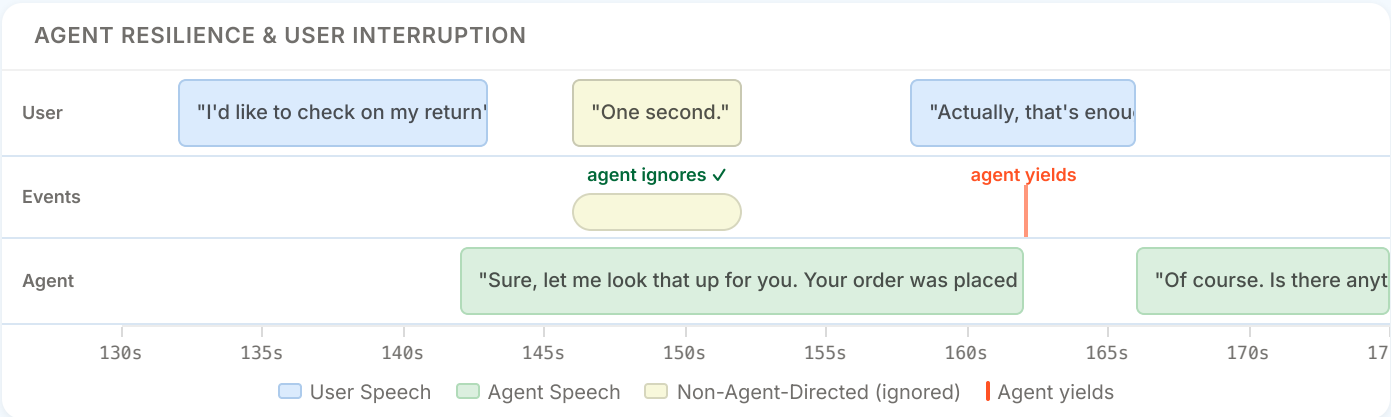


Figure: A τ -voice full-duplex simulation of a retail customer service call. The main timeline (top) shows the full 6-minute interaction. **Inset A** decomposes the user's audio pipeline, showing how clean speech is mixed with background noise, burst events, and non-agent-directed speech — this is the complex signal the agent must process. **Inset B** illustrates turn-taking dynamics: the agent correctly ignores a non-agent-directed utterance

("One second.") and continues speaking, but yields when the user genuinely interrupts ("Actually, that's enough, thanks.").