One-question launch card

Fill in one card per customer-facing chatbot, before its first release goes live. Fill in what you know; every blank line gets an owner and a date. Any blank line means: not ready to launch.

QUESTION	ANSWER	OWNER / BY WHEN
The bot's one question (the one the contact center gets most)		
How you know it is the top question: share of contacts, counted over how long		
What the first release covers, and where every other question goes		
The number the launch must move, and its baseline (counted over how long)		
Target: how far that number must move, by when		
Route to a person: how a customer reaches one from the chat		
Gate for question two: who reads the first measurement, and when		
Question two candidate (the next most frequent)		
Who answers for every reply (the role that is accountable)		
How a wrong answer gets caught, and what happens to the customer		
The bar: understands customers and holds context through a conversation; judged how, by whom, on how many conversations		
The specific requirements behind the answer, captured where		
Tuning: who sat with the bot on the key questions, one by one, and when		
Decision and date	**☐ LAUNCH ☐ HOLD ☐ RE-TEST ON _____ Date _____ Signed _____**	

One-question launch card

Example (fictional): returns chatbot at a mid-size clothing retailer

QUESTION	ANSWER	OWNER / BY WHEN
The bot's one question (the one the contact center gets most)	"How do I return an item?"	
How you know it is the top question: share of contacts, counted over how long	Returns and exchanges are about 3 in 10 of 4,000 monthly contacts, counted over 8 weeks	
What the first release covers, and where every other question goes	Returns only; every other question gets the "talk to a person" button	
The number the launch must move, and its baseline (counted over how long)	Returns calls, about 1,200 a week (8-week average)	
Target: how far that number must move, by when	Down by a quarter within 10 weeks of launch	
Route to a person: how a customer reaches one from the chat	A "talk to a person" button in every reply, opening a ticket answered inside 4 hours	
Gate for question two: who reads the first measurement, and when	Head of customer service, 6 weeks after launch	
Question two candidate (the next most frequent)	"Where is my order?", about 2 in 10 contacts	
Who answers for every reply (the role that is accountable)	Head of customer service	
How a wrong answer gets caught, and what happens to the customer	A reviewer reads 40 conversations each morning; a wrong return window is corrected the same day and the customer gets the right answer	
The bar: understands customers and holds context through a conversation; judged how, by whom, on how many conversations	The customer experience lead scores 60 test conversations, including typos and follow-ups. Bar: 19 in 20 right. Measured: 17 in 20 → **Fails the bar**	Customer experience lead, re-test by 11 Dec
The specific requirements behind the answer, captured where	Written from the returns policy: return windows by product type, who pays postage, refund timing; held in the knowledge base	
Tuning: who sat with the bot on the key questions, one by one, and when	**Blank: not ready**	Returns policy owner and customer experience lead, 2 sessions, by 5 Dec
Decision and date	**☐ LAUNCH ☐ HOLD ☒ RE-TEST ON 11 Dec · Date 20 Nov · "Bar not met and tuning not done."**	Head of customer service