A WORKING GUIDE · AIRTABLE CUSTOM AGENTS

A column that classifies itself.

An Airtable custom agent is a field whose value gets generated by an LLM — every time a new record lands. Write the instructions once, toggle on automatic generation, and the read-and-decide step happens without you. No code, no automations to wire up, no waiting on a build. Here's how to spin one up, what makes the prompt work, and the kinds of jobs this pattern is best at.

▲ THE BUILD IN FIVE STEPS

01	02	03	04	05
Add field, pick "Custom agent"	**Choose output type**	**Pick the model**	**Write instructions, @ the source field**	**Turn on automatic generation**

THE MENTAL MODEL

A field that thinks

You already know how to add a column. The shift is that this column doesn't store what you type — it stores what the model decides based on another column. Best for repeatable read-and-label work — *classify, route, tag, score* — where you've been doing the same call hundreds of times by hand.

▲ WORKED EXAMPLE A

Bureau Selection Agent

Reads the rewritten headline, picks one of 8 bureau slugs.
Long text · GPT-4.1.

▲ WORKED EXAMPLE B

Type Selector Agent

Reads the same headline, picks *article* or *medium-article*.
Single select · GPT-5 mini.

— Spinning up an agent

Each block below is one decision. The two worked examples — Bureau and Type Selector — show how the same step plays out differently depending on the job.

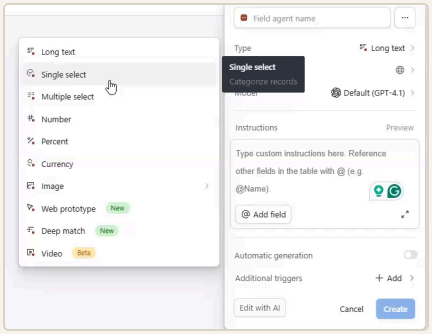
STEPS 01 / 02

Add a field, choose Custom agent, set the output type

Same plus icon you'd use to add a text column — except you pick "Custom agent" and give it a name (this becomes the column header). The **Type** dropdown is the most important decision on this screen. It dictates the shape of the answer the agent is allowed to return.

▲ IN THE EXAMPLES

> **Bureau** uses *long text* — the slug list lives in the prompt, output is a free string. **Type Selector** uses *single select* — only two valid values, and Airtable enforces them.



▲ THE TYPE MENU SITS BESIDE THE AGENT PANEL

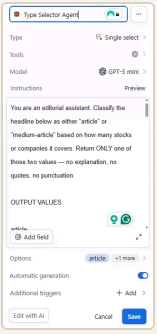
STEP 03

Pick the model

Defaults to GPT-4.1, which is solid for calls that need judgment. GPT-5 mini is the low-cost option — right pick when the agent fires on every new row and the decision is fairly mechanical. Match cost to difficulty.

▲ IN THE EXAMPLES

> **Bureau** stays on GPT-4.1 (8 slugs, some judgment around AI vs. sector). **Type Selector** drops to GPT-5 mini (two values, very mechanical).



▲ SINGLE SELECT + GPT-5 MINI FOR THE SIMPLER CALL

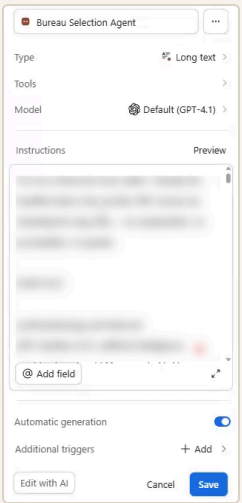
STEPS 04 / 05

Write the instructions, @ the source field, turn on automatic generation

Type the prompt in the Instructions box. Use **@FieldName** to point at the column the agent should read (usually a text field). State the role, the task, the allowed outputs, and any tie-breaker rules. Then toggle **Automatic generation** so every new record fires the agent without manual intervention.

▲ IN THE EXAMPLES

> Both watch *@Rewritten Headline*. Both have Automatic generation on. The difference is what each prompt allows: Bureau lists 8 slugs and edge-case rules; Type Selector lists 2 values and a single rule.



▲ INSTRUCTIONS, MODEL, AUTO-GEN TOGGLE IN ONE PANEL

Five things every classification prompt needs

Classification prompts that work look almost the same every time. Cover these five elements and the agent behaves; skip one and it drifts.

1 · ROLE
One sentence telling the model who it is. "You are a financial news editor." Anchors tone and judgment.

2 · TASK
What you want done, in plain language. "Classify the headline into exactly ONE bureau."

3 · ALLOWED OUTPUTS
Explicit list of every value the agent may return. The closed list is what makes the column trustable.

4 · TIE-BREAKERS
The "if X then Y" rules. Where you encode the judgment a person used to make.

5 · FORMAT GUARDRAILS
"Return ONLY the value — no quotes, no punctuation, no explanation." Stops the model from adding prose.

▲ BUREAU PROMPT, ANNOTATED

```
[role + task]
You are a financial news editor. Classify the
headline below into exactly ONE bureau by returning
the slug only.

[format guardrails]
— no explanation, no punctuation, no quotes.

[allowed outputs]
BUREAUS: technology-and-telecom · markets
· energy-materials · industrials ·
health-care · financials ·
consumer-goods · cryptocurrency

[tie-breakers]
ANY mention of AI → tech-and-telecom. ETFs/indexes →
markets. Generic "X stocks to buy" → pick at random
from sector list. Do NOT use markets as a catch-all.

[source]
Headline: @Rewritten Headline
```

Same five elements, simpler job

The Type Selector uses the exact same structure — just a shorter allowed-outputs list and one rule instead of five. When the job is simpler, the prompt gets shorter, not different.

▲ TYPE SELECTOR PROMPT, ANNOTATED

```
[role + task] You are an editorial assistant. Classify the headline below as either "article" or
"medium-article" based on how many stocks or companies it covers.
[format guardrails] Return ONLY one of those two values — no explanation, no quotes, no
punctuation.
[allowed outputs] OUTPUT VALUES: article · medium-article
[tie-breaker] Multiple stocks/companies/topics → article. Single stock/company/angle →
medium-article.
[source] Headline: @Rewritten Headline
```

> **Use "Edit with AI" instead of rewriting by hand.** When the agent misclassifies, open it, hit Edit with AI, and tell Claude what went wrong in plain English — "stop routing AI stocks to markets." It patches the rules without you re-editing the whole prompt. Faster, and the structure stays consistent.

Four things to keep in mind

$ Every run costs a call

Automatic generation fires on every new record. For a feed pulling dozens of headlines at a time, that's dozens of model calls. Use GPT-5 mini for high-volume mechanical classification; reserve GPT-4.1 for calls that need judgment.

≈ Enforced vs. invented outputs

Single select *enforces* — Airtable rejects anything outside the option list. Long text *invents* — the agent can return anything, including a typo. If your option set is closed and fixed, prefer single (or multi) select.

↻ Watch your source field

The @ reference is a hard dependency. Rename or restructure that column and your agent silently breaks. Whenever you change a source field, open every agent that reads from it.

◆ The agent decides "what." You still decide "what should."

Classification is mechanical. What gets pushed to CMS, which headlines are worth promoting, which categories matter at all — those judgment calls stay where they belong. Top of the stack.

What to build next

Anywhere you're doing the same read-and-label call by hand is a candidate. Six classification agents worth a 30-minute prototype:

SINGLE SELECT

Priority tier

Headline urgency: *breaking · timely · evergreen*. Drives which feeds publish first.

MULTI-SELECT

Ticker tagger

Pull every stock ticker mentioned into a multi-select column. Cleaner than regex, handles edge cases.

SINGLE SELECT

Angle classifier

Is this *analysis · news · opinion · listicle*? Useful for balancing the mix across a daily drop.

NUMBER (1-5)

Headline strength score

Rough rating on clarity and click pull. Useful as a triage filter, not a publish decision.

SINGLE SELECT

Duplicate angle flag

Compare against recent headlines and flag *fresh · similar · duplicate*. Catches repeat coverage before it ships.

LONG TEXT

One-line summary

30-word précis of what the article covers. Helps editors scan the table without opening each row.

THE PATTERN IN ONE LINE

Anywhere a teammate has been reading column A and writing a value into column B, **that's an agent**. The job of building one is mostly writing down what was already in their head.