

An entry, marked

This is a complete entry to the Birdspiracy AI Competition, scored against the published criteria with every point accounted for. It is here so you can see where marks are actually won and lost — not as a standard you are expected to match.

The problem	9 / 10
Solution & AI impact	18 / 20
Presentation	9 / 10
Final score	36 / 40

Read this first

- This entry was written by an adult with time and research tools. A strong entry from a school team will not look like this, and does not need to. The structure and the honesty are what to copy — not the polish.
- Nothing here scores full marks. Four points are lost, and every one of them comes from the same mistake: **stating something more confidently than the evidence supports**. That is the most common way good entries lose marks.
- The deck uses 5 content slides. Four is the minimum and six the maximum. The title slide and the sources slide are both required and neither counts toward that.

Scored against the published judging criteria. birdspiracy.com

ANNADATA

Predict need. Redirect surplus. Deliver with dignity.

A coordination tool for the volunteer groups already redistributing surplus food in Bihar — forecasting where evening meals will fall short, verifying what food is genuinely available, and routing it there before it spoils.

Birdspiracy AI Competition · Exemplar entry



surplus → need

EARNED MARKS	The one-line description says what the thing actually is — a coordination tool for volunteer groups — rather than a slogan. A judge knows what they are about to read.
WHY	Team name, school and class belong here. This example omits them because it is a demonstration.
LOST MARKS	Nothing lost here. A title slide is not scored; it just has to exist and orient the reader.

The problem

Surplus food and unmet need sit within a few kilometres of each other, and nobody knows it in time.

Who this is about. Volunteer redistribution groups operating in Patna district. They collect surplus from wedding halls, hostel canteens and restaurants, and take it to settlements the same evening. Most run on WhatsApp and a phone tree.


What goes wrong. Surplus is offered late and expires within hours. Volunteers guess where to take it, so the same accessible settlements are served repeatedly while others are missed. Nothing records what was delivered where, so nobody learns from last week.

Why it persists. Every group is small, unfunded, and coordinating by hand. The scheduling problem is genuinely hard, and no one group sees enough of the district to solve it.

78 million tonnes

of household food wasted in India in 2022 — second only to China. UNEP, 2024

Bihar is improving, and still highest but one



EARNED MARKS	The population is narrow and specific: volunteer redistribution groups in one district. Not "hunger in India". Almost every entry that scores in the middle band fails exactly here.
EARNED MARKS	Three labelled parts — who this is about, what goes wrong, why it persists. The third is the one most entries skip, and it is what separates a described problem from an examined one.
EARNED MARKS	Both figures are attributed with an organisation and a year, and both are traceable on the sources slide.
EARNED MARKS	The chart shows a decline rather than a single alarming number. Admitting the problem is improving is more credible than implying it is not, and a judge notices.
LOST MARKS	One mark lost. The 78 million tonnes figure is national <i>household</i> food waste, but the problem described is institutional surplus from wedding halls and canteens. Those are different streams. It proves surplus exists at scale; it does not size the surplus this system would actually move.

Proposed solution

A shared dispatch layer the existing groups log into. Four steps, each with a person able to override it.

1

Signal

Donors post surplus in the app: what, how much, cooked when, collect by when. Community workers log served and unserved locations after each run.

2

Forecast

A gradient-boosted model scores each settlement for likely unmet evening demand, from past delivery logs, day of week, weather, festival calendar and ration-shop distribution dates.

3

Verify

Before dispatch, a local worker confirms the need by SMS or call. No delivery is scheduled on a model score alone.

4

Route

A solver matches surplus to confirmed need under a hard time limit for cooked food, and returns a route the driver can follow.

What the AI actually contributes. Ranking settlements by predicted unmet need, and solving the surplus-to-need match under expiry and travel constraints. Both are things a volunteer cannot do quickly by hand at district scale. Collection, confirmation and delivery stay human.

EARNED MARKS	The mechanism is specific: what goes in, what the model does, what comes out, who acts on it. A reader can restate the system after one pass.
EARNED MARKS	The model type and its inputs are named. "A gradient-boosted model scoring settlements from delivery logs, weather and ration-shop dates" is a claim that can be argued with. "An AI system that helps" is not.
EARNED MARKS	Step 3 is a human verification step placed between prediction and dispatch. Building in a limit on the AI's authority is stronger than claiming the AI is accurate.
EARNED MARKS	The band at the bottom states plainly what the AI contributes and what stays human. This is the single most useful sentence in the deck.

Existing approaches

Surplus redistribution in India is not a new idea. It is an uncoordinated one.

Robin Hood Army

Volunteers collect restaurant surplus and distribute it, coordinated over WhatsApp.

Gap. Allocation is by volunteer availability, not by where need is greatest. No record carries between runs.

Feeding India, No Food Waste

Apps that let a donor request a pickup, then dispatch a nearby volunteer or vehicle.

Gap. Demand-side only. The app answers who collects, never which settlement should receive it.

PM POSHAN, PDS, Anganwadi

State programmes delivering scheduled entitlements to enrolled recipients.

Gap. Scheduled and enrolment-based. Cannot absorb a hall of surplus biryani available for the next four hours.

Where Annadata is different

Every existing service optimises the supply side — getting surplus collected. None of them forecast the demand side, so the choice of destination is left to whoever is driving. Annadata adds the missing half and hands it to the groups already doing the work, rather than competing with them for donors.

EARNED MARKS	Four real organisations and schemes named, each with what it does and precisely what it does not do. This is how originality gets <i>evidenced</i> rather than asserted.
EARNED MARKS	The gap is stated in the same terms across all three columns — supply side covered, demand side not — so the argument accumulates instead of repeating.
EARNED MARKS	The closing band says the idea works <i>with</i> existing groups rather than replacing them. Judges are wary of proposals that assume everyone currently doing the work is doing it wrong.
WHY	Leaving this section out is the most common reason a genuinely good idea lands in the middle band. Judges cannot tell an idea is new unless you show what it is being compared against.

Requirements and limitations

Hunger data is thin, late and biased. The system has to be built assuming its own predictions are wrong.

What it needs

- Delivery logs — location, quantity, time, and crucially the runs that found no need. Roughly six months before the model beats a simple rota.
- Donor listings with an honest collect-by time.
- Settlement boundaries at ward level or finer.
- Ration-shop distribution dates, which shift demand sharply.
- One named community worker per settlement who can confirm or refuse a delivery.

Where it fails

Invisibility, not absence. A settlement with no phones and no volunteer history generates no data, and the model reads that as no need. Mitigation: a fixed share of every week's runs goes to settlements with the least data, whatever the model says.

Inherited bias. Training on past deliveries teaches the model who was easiest to reach. Mitigation: audit coverage by settlement monthly and report the gap, not the average.

Dependence. A community that reorganises around the app is worse off if funding stops. Mitigation: the phone tree stays live, and every route is exportable on paper.

EARNED MARKS

Every failure mode is paired with a mitigation that costs something. "A fixed share of every week's runs goes to settlements with the least data, whatever the model says" is a real trade-off, not a reassurance.

EARNED MARKS

The first failure — a settlement with no digital footprint reading as no need — is the sharpest observation in the deck. It identifies a way the system would quietly harm the people it is meant to help.

EARNED MARKS

The left column names what does not exist yet and would have to be collected first. Honest about the gap between the idea and the conditions it needs.

LOST MARKS

One mark lost. "Roughly six months before the model beats a simple rota" is a guess presented with the confidence of an estimate. Naming it as an assumption to be tested would have cost nothing and read as more rigorous.

Implementation plan

One district first. Expansion is earned against published numbers, not assumed.

1

Months 0-6 · short term

Paper pilot, Patna district

Two volunteer groups, 25 settlements, no model. Log every run on a shared sheet to build the training data and prove the logging habit holds. Cost: one part-time coordinator.

2

Months 6-12

Model and app on top of real logs

Train on pilot data, ship the app to the same two groups, keep the phone tree running in parallel. Success is measured against the paper baseline, not against zero.

3

Year 2+ · long term

Second district, only if the gates are met

Cost per delivered meal falls against the paper baseline; median time from surplus posted to delivered drops below three hours; zero food-safety incidents; coverage gap between best- and worst-served settlements narrows.

Funding: corporate CSR budgets for the coordinator, cloud and phones. Government provides ward boundaries and ration-shop schedules through an existing open-data request. No new vehicles are bought — the groups already have them.

EARNED MARKS	Phase 1 contains no AI at all. Six months of paper logging, because the data to train on does not exist yet. This is the strongest single choice in the deck — it shows someone who thought about the plan rather than assuming it.
EARNED MARKS	Expansion is gated on four named, measurable conditions. "Only expand if cost per delivered meal falls" is a commitment; "we will scale nationally" is a wish.
EARNED MARKS	Funding is specific — CSR budgets for the coordinator, government data through an existing request, and no new vehicles because the groups already have them.
LOST MARKS	One mark lost. The four-hour expiry window for cooked food is asserted without a source. Food-safety guidance on this exists; citing it would have settled the question instead of inviting the judge to wonder.

References

- UN Environment Programme and WRAP. Food Waste Index Report 2024. India: 78.2 million tonnes of household food waste in 2022, 55 kg per capita.
- International Institute for Population Sciences. National Family Health Survey, rounds 4, 5 and 6. Stunting in under-fives, Bihar.
- Robin Hood Army — robinhoodarmy.com. Model and coverage.
- Feeding India by Zomato, and No Food Waste — organisation descriptions of pickup dispatch.
- Ministry of Education, PM POSHAN scheme guidelines; Department of Food and Public Distribution, PDS operations.

Figures were checked against the source publication. Photographs of identifiable people were deliberately not used.

EARNED MARKS	Organisation, publication and year for every figure. Not a bare link.
EARNED MARKS	The closing line states that no photographs of identifiable people were used. If you use images of people, you need their permission — say where each image came from.
LOST MARKS	One mark lost, on Presentation. The deck carries speaker notes explaining design decisions to a judge rather than supporting a presenter. In a real entry that reads slightly off.
WHY	Without this slide the entry could not have scored above 5 for Presentation, and both statistics on the problem slide would have counted as unsupported. It costs nothing — it is outside the 4 to 6 limit.

Where the four marks went

All four losses are the same mistake wearing different clothes.

-1	A national figure standing in for a local one. The waste statistic measured a different stream from the one the system moves.
-1	An estimate presented as a finding. Six months of data collection was a reasonable guess, written as though it had been established.
-1	An unsourced constraint. The four-hour food-safety window is checkable, and was not checked.
-1	Speaker notes written for the wrong reader. A small thing, but it shows in presentation.

What to take from this

- **Narrow the problem until it is uncomfortable.** One group of people, one place, one failure. Breadth reads as vagueness.
- **Say what the AI does, in working detail.** If a reader cannot say what goes in and what comes out, the score sits in the middle band no matter how good the idea is.
- **Name what already exists.** You cannot be judged original against nothing.
- **Admit one thing that costs you something.** A limitation you engage with honestly is worth more than a page of advantages.
- **Cite it, or soften it.** Every mark lost here came from a sentence that was more confident than its evidence. Either find the source or write “we estimate”.

Full judging criteria, the presentation template, and the competition packet are all at birdspiracy.com. Questions: birdspiracy.competition@gmail.com