

Lab 5: AI Conversation Log — Discovery Synthesis and Project Proposal

Jordan Blake (Instructor Example) | SE 395-2 | October 8, 2026

Student	Jordan Blake (Instructor Example)
Date	October 6–7, 2026
AI Tool Used	Claude 3.5 Sonnet
Partner Organization	Sunrise Community Kitchen
Project Focus	Weekly meal distribution coordination

This document shows how to use AI productively across all four phases of Lab 5: checking your inventory before you write anything, drafting the proposal with AI as a thinking partner, preparing for peer review, and revising after you get real feedback from another team. The strongest logs show the student pushing back on AI suggestions and using the conversation to sharpen their own thinking, not just generate text.

PHASE 1: INVENTORY CHECK

EXCHANGE 1: ORGANIZING LAB WORK INTO THREE COLUMNS

MY PROMPT

I need to do the inventory check before I start writing the proposal. The lab says to sort everything into three columns: (1) know well enough to build on, (2) think we know but not verified with partner, (3) still guessing. Here is a summary of what we have from Labs 2-4:

Lab 2: SDG 2 (Zero Hunger) and SDG 10 (Reduced Inequalities) are the relevant targets. We identified a bias risk around Maria's historical distribution data embedding her existing knowledge of families. We reviewed AI tools in the food distribution space and found nothing purpose-built for a 3-person operation at this scale.

Lab 3: Maria's biggest pain point is the time it takes to build the weekly distribution plan by hand – about 6-8 hours across the full Monday-Thursday cycle. She said "if I could get that down to two hours I could actually do outreach." We also identified key-person risk: if Maria is out, distribution does not happen because no one else knows the families or the spreadsheet.

Lab 4: Full cold reality: manifest PDF arrives Monday, Maria categorizes items by hand, builds plan in Google Sheet, emails volunteers Wednesday, calls and texts stragglers, distribution Thursday, hand tally at end, report to food bank and United Way by Tuesday. Recipe: AI parses manifest, drafts allocation plan with confidence flagging, generates volunteer checklist, produces report draft. Chef's note: Maria keeps all priority decisions when supply is short, because her crisis knowledge is not in the data. Big open question: manifest format unknown (could be scanned, which breaks Step 1).

Can you help me sort all of this into the three columns?

AI RESPONSE

Here is the inventory sorted into three columns:

Know well enough to build on:

- Maria's weekly workflow from Monday to Thursday (cold reality is solid)
- Time cost: 6-8 hours per coordination cycle, Maria's stated goal is 2 hours
- Key-person risk is real and confirmed by Maria directly
- Chef's note reasoning: Maria holds priority judgment because her crisis knowledge is not in any data source
- SDG 2 and SDG 10 alignment is clear and defensible
- Volunteer coordination spans three channels with no tracking system
- Post-distribution tally is approximate and everyone knows it

Think we know but not verified:

- Family dietary data quality: you have described the Google Sheet, but you have not confirmed the completeness or accuracy of the dietary restriction column
- Report format requirements: you know Maria sends a report to the food bank and United Way, but you do not know whether those funders have specific format or data requirements that would constrain what AI generates
- Volunteer list: you said 15-20 volunteers, but in the Lab 4 cold reality the regular reliable count came up as closer to 12. Is the full list 15-20 and the reliable core 12, or is one of those numbers wrong?

Still guessing:

- Manifest format: digital PDF or scanned image? The entire first step of the recipe depends on this. It is a blocking unknown.
- Whether Maria or anyone on the team could maintain a new AI-assisted workflow without ongoing technical support
- How much time Maria realistically has per week to learn, configure, or troubleshoot a new tool

Two things to address before submitting: the manifest format needs verification, and the volunteer count discrepancy should be resolved. Both show up in the proposal.

EXCHANGE 2: REACTING TO THE FLAGS

MY PROMPT

OK the manifest format being in column 3 is right and I already knew that. But I am a little surprised the family dietary data is in column 2 and not column 1. We saw the spreadsheet during Maria's class visit. Why is that in "think we know but unverified"?

AI RESPONSE

You saw the spreadsheet, but that is different from knowing the data quality. A few questions worth asking:

- How many of the 200 family rows have the dietary restriction column filled in? If 40 out of 200 families have entries and the rest are blank, the blank entries do not mean "no restrictions" — they might mean "we never asked."
- Are the entries consistent? "No pork" and "pork allergy" are different things, but they might both appear in the same column formatted differently.
- When was the data last updated? A family's medical situation changes. If Maria entered dietary information two years ago and never updated it, your AI is building plans on stale inputs.

You have seen the spreadsheet exists and has a dietary column. You have not verified that the data in that column is complete or current enough to use as a reliable AI input. That is the difference between columns 1 and 2.

I assumed that seeing the spreadsheet was the same as knowing the data was good. Those are different things. Column 2 is for things we have observed but not tested. The dietary data needs a data quality check before it becomes a recipe input we can rely on.

EXCHANGE 3: WHAT TO DO WITH COLUMN 3

MY PROMPT

What should I actually do about the things in column 3? The manifest format question is the big one. The maintainability question is also real — I honestly do not know if Maria would be able to keep this running without us.

AI RESPONSE

You have two options for each item in column 3. Either verify it before you submit, or scope the proposal so it does not depend on the guess.

For the manifest format: Verification is fast. Ask Maria to forward you this week's food bank email with the manifest attached. You will know within five minutes whether it is machine-readable. If it is a scanned image, your Step 1 needs to change — AI cannot reliably parse a scanned PDF without OCR preprocessing, and that adds complexity. If you cannot verify before the deadline, scope the proposal honestly: "Step 1 is contingent on the manifest being a digital PDF. We will verify this in Week 8. If the manifest is scanned, the first build milestone shifts to manifest format resolution."

For maintainability: This one is harder to verify quickly. A reasonable scoping move is to constrain the proposal to tools Maria already has access to — Claude, Google Sheets, Gmail — and build processes that do not require her to install software, manage accounts, or learn a new platform. That scoping decision belongs in Section 3 of the proposal, not as a vague risk. State it as a design constraint: "All components use tools Sunrise already has. No new software accounts are required."

Do not leave column 3 items as unresolved background noise. Either close them or name the constraint explicitly in the proposal.

PHASE 2: DRAFTING THE PROPOSAL**EXCHANGE 4: SHARPENING SECTION 3 (PROPOSED SOLUTION)****MY PROMPT**

I am working on Section 3. Here is what I have from the Lab 4 recipe and I want to sharpen it for the proposal:

"We will use AI to help with the distribution workflow. AI will read the manifest, match items to families, help with volunteer coordination, and generate the report. Maria stays in the loop for all final decisions."

I know this is too vague. The Lab 4 recipe has more detail but it is in recipe format and needs to be written for the proposal. Can you help me rewrite this so it is specific? I want to name the actual tools, the inputs, the outputs, and where Maria is in the loop.

AI RESPONSE

Here is a rewrite based on what your recipe established. Fill in any blanks where I am working from inference:

What we will build: A Claude-based prompt workflow that Maria runs each Monday after receiving the food bank manifest. The workflow has four components:

1. **Manifest parser.** Input: the food bank manifest PDF (contingent on it being machine-readable). Output: a structured list of available items by category and quantity, pasted into a Claude prompt. Maria reviews for accuracy before the next step runs.
2. **Draft allocation generator.** Input: the structured manifest list plus an export of the family profile Google Sheet (200 rows). Output: a draft distribution plan in spreadsheet format, with a confidence flag on any family row where the dietary data is incomplete or potentially stale. Maria reviews flagged rows before approving the rest.
3. **Volunteer checklist.** Input: the confirmed volunteer list (12 regulars). Output: a Gmail-ready draft assignment email plus a simple tracking table showing who has confirmed. Maria sends the email and marks confirmations manually; the table shows who is missing at a glance.
4. **Report draft.** Input: tally numbers Maria enters into a simple Google Form after distribution. Output: a formatted report draft with a manifest-vs-tally cross-check flagging any math discrepancies. Maria reviews, corrects, and sends.

What stays human: All priority decisions when supply is short. Any outreach to individual families. Volunteer relationship management. Final sign-off on every output before it leaves Maria's hands.

Minimum viable version by end of semester: Components 1 and 2 working reliably and saving Maria at least 2 hours on the planning step. Components 3 and 4 are the second build milestone.

MY PROMPT

That is much better. I am going to use most of it. I want to change one thing: the manifest parser output should not be "pasted into a Claude prompt" because that sounds weird in a proposal document. Can you rephrase that so it sounds less technical without losing the meaning?

AI RESPONSE

Try: "Output: a structured list of available items by category and quantity that feeds directly into the allocation step." That removes the technical process description and just says what the output is and what it does. The proposal does not need to explain how Claude works internally.

EXCHANGE 5: DRAFTING SECTION 5 (RISKS)

MY PROMPT

Here is my first draft of the risks section:

Technical risks: manifest format is unverified and could break Step 1. Google Sheet

data quality is unverified. If either input is bad, the AI outputs are unreliable.

Adoption risks: Maria may not have time to learn a new tool during the semester. The workflow has to be simple enough to run with no ongoing support.

Limitations: we are not building anything for the tally counting step on distribution day. That problem requires a hardware or process change, not an AI prompt.

Is this OK?

AI RESPONSE

The technical and adoption risks are honest and specific. The limitations section is good. But you have listed only technical and operational risks. You have not listed any ethical risks, and the lab rubric specifically calls out ethical considerations.

Here is one to think through: your allocation generator flags families where dietary data is incomplete. Maria reviews those flagged rows. But she approves the other 150-160 rows without scrutiny. What happens when the AI's suggested allocation is wrong for a family Maria does not know well? She sees the flag and assumes the unflagged rows are fine. The unflagged rows are only as good as the data, which you have already put in column 2 as unverified.

The ethical risk is this: the allocation generator creates a false sense of coverage for the families Maria does not know well. She will focus her review energy on the flagged rows and on the families she already has relationships with. Errors for less-familiar families could pass through unseen. This is not a bug in the code. It is a risk in how Maria will actually use the tool.

Does your risks section address that?

MY PROMPT

No it does not. I thought the confidence flagging design handled this. But I see what you are saying – the flagging only catches data gaps, not judgment gaps. Maria might rubber-stamp allocations for families she does not recognize and trust that if the flag is not there, the AI got it right. That is not the same as the AI actually being right.

How do I write this as a risk without it sounding like I am saying the whole project is too dangerous to do?

AI RESPONSE

Name it as a design constraint you are building around, not a reason to stop. Something like:

"Rubber-stamp risk: even with confidence flagging, Maria may focus her review attention on flagged rows and approve unflagged allocations without close scrutiny. This risk is highest for families Maria does not know well. We will address this in the build phase by limiting the number of allocations presented for approval per session, and by including a periodic prompt asking Maria to spot-check three random unflagged rows before final approval."

You are naming the risk, naming who is affected, and naming your mitigation. That is what the rubric is looking for. You are not saying the project is dangerous. You are saying you know where the risk lives and you have a plan for it.

Listing only technical and adoption risks is not enough. The harder risks are ethical ones, like how a tool changes human behavior in ways that might introduce new problems. The rubber-stamp risk came up in Lab 4 too, but I did not carry it forward into the proposal until Claude pushed back. It needs to be in the proposal explicitly, with a mitigation plan, not just acknowledged in a chef's note.

EXCHANGE 6: FIXING SECTION 7 (TIMELINE)

MY PROMPT

Here is my timeline draft:

Week 8-9: Build the manifest parser and allocation generator

Week 10-11: Test and iterate

Week 12-13: Refine and document

Week 14-15: Onboard Maria and hand off

Week 16: Showcase

Is this OK?

AI RESPONSE

The phase labels are fine. But the milestones are week numbers, not deliverables. I cannot tell from this timeline whether you hit Week 9 or not. What does "build the manifest parser and allocation generator" look like when it is done? What would I check?

Try attaching a testable output to each phase. For example:

- **Week 8-9:** Working Claude prompt that takes a manifest list and family export as inputs and produces a draft allocation plan. Maria has run it against one real week's data and confirmed the output is usable.
- **Week 10-11:** Allocation generator tested against three consecutive weeks of distribution data. Confidence flagging logic reviewed and adjusted based on Maria's feedback on what it catches versus what it misses.

- **Week 12-13:** Volunteer checklist and report draft components complete. Full four-component workflow tested end to end. Documentation written so Maria can run it without our help.
- **Week 14-15:** Maria runs the full workflow independently at least once, with the team observing. Handoff package includes step-by-step guide and a one-page troubleshooting reference.
- **Week 16:** Showcase presentation with live demo or demo recording showing the workflow in action.

The test of a good milestone is: could someone who was not in your meeting look at this and tell you whether you hit it? Week numbers alone cannot pass that test. Testable outputs can.

MY PROMPT

The Week 14-15 milestone about Maria running it independently is good but is it realistic? She does not have much time and we are asking her to run the whole workflow on her own before the showcase. What if she cannot get to it that week?

AI RESPONSE

That is a fair concern. You have two options. One: build in a backup — "Maria runs the workflow independently, or if scheduling does not allow it, the team does a guided walkthrough with Maria present and observing each step." Two: move the independent run to Week 12-13 and use Weeks 14-15 for documentation and any final adjustments Maria requests. Either way, the goal is that someone other than your team can operate this at the end. Write the milestone around the goal, not a specific schedule that may not survive contact with reality.

EXCHANGE 7: TIGHTENING SECTION 8 (SUCCESS CRITERIA)

MY PROMPT

My draft for success criteria:

The project succeeds if Maria's weekly coordination time drops significantly. The tool should be easy enough that Maria can run it without help. The partner should feel the project added real value.

I know this is weak. Help me make it measurable.

AI RESPONSE

You are right, it is weak. Three problems: "drops significantly" is not a number, "easy enough" is not testable, and "feel the project added real value" is a feeling, not evidence. Here is a revision pass:

- **Time reduction:** Maria's weekly coordination time drops from 6-8 hours to 3 hours or less, measured over at least three consecutive distribution weeks.
- **Usability:** Maria can complete the allocation step without requesting help from the team, confirmed by at least one solo run before showcase.
- **Partner validation:** Maria confirms in a brief end-of-semester conversation that the tool is something she would keep using after the course ends, and identifies at least one specific change it made to her workflow.
- **Report accuracy:** The manifest-vs-tally cross-check catches at least one discrepancy during the test period that Maria confirms would not have been caught manually.

But I have a question for you: how will you actually measure the time reduction? Will Maria track her hours? Will you observe a distribution cycle in Week 8 and compare it to Week 14? If you cannot describe how you will collect the evidence, the metric is not real.

MY PROMPT

Good question. I was going to ask Maria to estimate her time after each week. But she will not remember accurately. What would actually work?

AI RESPONSE

A few lightweight options that do not require Maria to log everything manually:

- **Timestamp observation:** have the team note when Maria starts and finishes the allocation step during two distribution cycles — one before the tool is in use, one after. You get a direct before-and-after comparison without relying on memory.
- **Time log in the tool itself:** build a simple one-line entry into the workflow — "How long did the allocation step take today?" as the last prompt before Maria approves. Low friction and captures it in the moment.
- **Use Maria's calendar:** if Maria blocks distribution prep time on her calendar, her calendar entries are more reliable than after-the-fact estimates.

Add your measurement method to the success criteria section. "We will measure time reduction by" followed by one of these. Otherwise it is just a number you said you would hit.

Success criteria need a measurement method, not just a target. Writing "Maria's time drops to 3 hours" is an outcome. Writing "We will measure this by timing one pre-tool and two post-tool allocation sessions" is a plan. The second version is what goes in the proposal.

EXCHANGE 8: FULL DRAFT CONSISTENCY CHECK

MY PROMPT

I have all eight sections drafted now. I am going to paste the full proposal and I want you to check for consistency across sections – contradictions, numbers that don't match, places where one section says something different from another.

[Full proposal pasted here – approximately 1,800 words covering all 8 sections. Key details: Section 1 describes a 3-person organization with Maria as the sole paid coordinator. Section 2 references the workflow and says "she coordinates 15-20 volunteers each week." Section 3 describes the volunteer checklist as built from "the volunteer list of 12 regulars." Section 4 cites SDG 2.1 and SDG 10.2 specifically. Section 5 includes the rubber-stamp risk with mitigation language. Section 6 asks for three things from the partner: manifest sample, 30 minutes to confirm the allocation logic against a past week's data, and a confirmation of funder report format requirements. Section 7 has revised milestones with testable outputs. Section 8 includes the time-tracking method.]

AI RESPONSE

I found one inconsistency worth fixing and two smaller things to review.

Volunteer count mismatch (fix this): Section 2 says Maria "coordinates 15-20 volunteers each week." Section 6 asks for access to "the volunteer list of 12 regulars." These could both be true – 15-20 on the full list, 12 who are reliable regulars – but the reader will see a contradiction unless you connect them. Add a clarifying note in Section 2: "15-20 volunteers are on the contact list, but Maria considers roughly 12 of them reliable regulars for any given week." Then Section 6's reference to "12 regulars" will be consistent.

SDG citation (minor): Section 4 cites SDG 10.2. That target is specifically about the social, economic, and political inclusion of all people regardless of age, sex, disability, race, ethnicity, origin, religion, or economic status. It applies, but your explanation says "equitable food distribution." Make sure your SDG explanation actually connects to the target language, not just the goal title. SDG 2.1 is a stronger fit for this project and your explanation of it is solid.

Section 6 partner ask (worth a second look): You are asking for 30 minutes to confirm the allocation logic against a past week's data. That is reasonable. But you have not mentioned the manifest format verification. Either you closed that question before submitting (in which case add a note saying you confirmed the format) or it is still open and belongs in Section 6 as the first ask, not a background assumption.

Everything else is consistent. The timeline milestones align with the build scope in Section 3. The risks connect to the chef's note logic from Lab 4. The success criteria measurement method is specific. The document reads as one voice.

MY PROMPT

We did get the manifest – Maria forwarded it before class and it is a digital PDF, not a scan. So that blocking unknown is resolved. I should mention that in the proposal somewhere instead of leaving it as an open question.

AI RESPONSE

Yes. Add one sentence to Section 3 after the manifest parser description: "We have confirmed with Maria that the food bank manifest is a digital PDF, making automated parsing feasible without OCR preprocessing." That closes the gap cleanly and shows the reader you did the verification work.

PHASE 3: PREPARING FOR PEER REVIEW**EXCHANGE 9: AI AS PEER REVIEWER****MY PROMPT**

Before the actual peer review in class, I want to stress-test the proposal. Can you role-play as a peer reviewer using the five questions from the lab? Be critical. Do not just tell me what I want to hear.

The five questions are:

1. Can you explain their proposed solution in one sentence? If not, it is not clear enough.
2. Is the "what stays human" section convincing? Would you trust this system if you were the partner?
3. Are there risks they have not considered?
4. Is the timeline realistic? Are the milestones specific enough that you could tell whether they hit them?
5. What is the strongest part of this proposal? What is the weakest?

AI RESPONSE**Question 1: Can I explain the proposed solution in one sentence?**

Yes: "A Claude-based prompt workflow that takes the food bank manifest and family spreadsheet as inputs, generates a draft distribution plan with confidence flags for data gaps, produces a volunteer assignment email and checklist, and drafts the post-distribution funder report — all reviewed and approved by Maria before anything goes out." That is a long sentence but it is specific. The solution is clear.

Question 2: Is the "what stays human" section convincing?

Mostly yes. The reasoning for keeping priority decisions with Maria is the strongest part of the entire proposal — the explanation about crisis knowledge not being in any data source is concrete and persuasive. But I would push back on one thing: the proposal says "final sign-off on every output before it leaves Maria's hands" as a blanket statement. In practice, Maria reviewing 200 allocation rows is not the same as Maria reviewing three priority exceptions. The blanket statement makes the review sound more robust than it is. Be more honest about what Maria will actually scrutinize versus what she will scan.

Question 3: Are there risks not considered?

Two things the proposal does not address. First: what happens during the transition period? In Weeks 8-10 when the tool exists but is not fully trusted yet, Maria will be running two parallel processes — her existing manual workflow and the new AI-assisted one. That is more work, not less, during the build phase. The proposal should acknowledge this. Second: the proposal does not say what happens if the food bank changes the manifest format. If they redesign the PDF structure, the parser breaks. Does the team maintain this after the semester? Does Maria call you? Who owns it?

Question 4: Is the timeline realistic?

The milestones are specific and testable, which is good. But the timeline is front-loaded. Both of the hardest components — manifest parser and allocation generator — are in Weeks 8-9. The easier components (volunteer checklist, report draft) are in Weeks 12-13 when you will also be refining and documenting. If the allocation generator runs into data quality issues in Week 9, you have no buffer before Week 10 testing. Consider whether a smaller first milestone (parser only in Week 8, allocation generator in Week 9) would be more realistic than bundling them together.

Question 5: Strongest and weakest parts?

Strongest: the reasoning in the chef's note section. It is specific, it is honest, and it shows the team actually thought through the ethics. Weakest: the success criteria are good for the team's perspective but they are all team-facing metrics. What does success look like from Maria's perspective? The proposal asks "what does the partner need to see to call this a success?" but the criteria you list are things the team would measure. Add at least one success criterion that is Maria's to define — something like "Maria confirms in her own words that she would keep using this after the semester ends, and states specifically what changed."

MY PROMPT

The partner-facing metric point is fair and I am going to add it. The front-loaded timeline is also something I want to fix. But the question about who owns the tool after the semester — I am not sure how to answer that. It feels out of scope for this proposal. Do I have to address it?

AI RESPONSE

You do not have to solve it, but you should name it. Something like: "Sustainability beyond the semester is outside the scope of this project. The handoff package in Weeks 14-15 is designed so Maria can operate the

workflow independently. Any future modifications to the workflow — for example, if the food bank changes its manifest format — would require the team to be contacted or a new team to take it on." One paragraph. It shows the reader you thought about it and made a deliberate decision to scope it out, rather than missing it entirely. A blank space on this question reads as "we did not think about it." A scoped-out note reads as "we thought about it and decided."

EXCHANGE 10: REVISING BEFORE ACTUAL PEER REVIEW

MY PROMPT

I made the following changes based on the AI peer review:

1. Revised "what stays human" to be honest that Maria will scrutinize the flagged rows and priority exceptions closely, but will scan the bulk of allocations
2. Added a transition period note to Section 5: Weeks 8-10 will require parallel workflows and Maria's time cost may temporarily increase before it decreases
3. Split the Week 8-9 milestone into Week 8 (manifest parser working and tested) and Week 9 (allocation generator working with confidence flagging)
4. Added a partner-facing success criterion: Maria states in her own words one specific change the tool made to her work
5. Added a one-paragraph sustainability note to Section 6

Can you do a final check before I take this into actual peer review?

AI RESPONSE

Those are all good changes. A few quick checks:

- The "what stays human" revision: make sure it reads as honest and not as an apology. You are not undermining the system by acknowledging that Maria will scan most rows. You are being accurate about how humans interact with AI-assisted workflows. Accurate is better than falsely reassuring.
- The transition period note: if you added this to Section 5 (Risks), also make sure Section 7 (Timeline) reflects it. Week 8 should say something like: "Maria runs parallel workflows this week — the existing manual process plus the new parser. This is by design and expected." Otherwise the timeline looks like the old process ends in Week 7 and the new one starts in Week 8 without acknowledging the overlap.
- The partner-facing success criterion: make sure you also note how you will collect it. A brief end-of-semester conversation, written or verbal, and who records it.

The proposal is in good shape for peer review. Go in ready to defend your scoping decisions — specifically why you are not automating the tally step, and why the chef's note boundary is where it is. Expect questions on both.

PHASE 4: POST-PEER-REVIEW REVISION

EXCHANGE 11: ADDRESSING REAL PEER FEEDBACK

MY PROMPT

We just finished the peer review in class. The other team (working with Mills Gallery) gave us this feedback:

"Your solution section is specific and we could follow exactly what you are building, which is good. But your risks section feels like a checklist, not analysis. You named five risks but for most of them you just said what could go wrong without explaining how likely it is, how serious it would be, or what you are actually going to do about it. The rubber-stamp risk is the only one that has a real mitigation. The others feel like you listed them to check a box."

They are right. How do I fix the risks section without just making it longer?

AI RESPONSE

The Mills Gallery team identified a real problem. Length is not the fix. What you need is analysis, not more items. For each risk, you need three things: likelihood (how probable is this?), severity (if it happens, how bad is it?), and mitigation (what specifically will you do?). Not every risk needs a full paragraph, but each one needs all three elements.

Try this structure for each risk:

- **Risk name.** One sentence describing what could happen. Likelihood: high/medium/low, with a brief reason. Severity: high/medium/low, with a brief reason. Mitigation: one specific action you will take or design decision you have already made.

For example, your manifest format risk was a blocking unknown that you have since resolved. You should update it: "Manifest format: resolved. We confirmed the manifest is a digital PDF before submitting this proposal. This risk is closed." That is honest and shows the reader you actually did the verification work.

For the data quality risk on the family spreadsheet, you might write: "Dietary data quality: medium likelihood, high severity. If the dietary restriction column is incomplete or outdated, the AI's confidence flags will miss gaps that matter. Mitigation: in Week 8, we will ask Maria to walk us through five random rows in the spreadsheet and confirm whether the dietary entries are current. We will adjust the confidence flagging logic based on what we find."

The peer reviewers are not saying you need more risks. They are saying the risks you named are not doing analytical work. Rework the section around likelihood, severity, and mitigation for the four risks that currently lack them.

MY PROMPT

That structure is helpful. I am going to apply it. One thing I want to push back on though: the peer reviewer said the risks "feel like a checklist." But some risks really are low likelihood and low severity, and I don't want to pad the analysis just to make every risk look serious. Is it OK to say "this risk is low likelihood and low severity, so no mitigation is needed"?

AI RESPONSE

Yes, that is a legitimate thing to say, and it is actually a stronger move than padding every risk with a mitigation. If you write "this risk is low likelihood and low severity because [reason], and we have decided no mitigation is warranted at this stage," that is an analytical judgment. It shows you thought about it and made a deliberate call.

What the peer reviewers are reacting to is the appearance that you listed items without thinking about them — not that you listed too few mitigation strategies. The fix is to show your reasoning for every item, even if the reasoning concludes "this does not need a mitigation plan." The reasoning is the analysis. The checklist problem is the absence of reasoning, not the absence of action plans.

Peer review revealed that the risks section was item-heavy but analysis-light. The fix is not more items or longer paragraphs. It is showing the reasoning: likelihood, severity, and a deliberate call on whether a mitigation is needed. A well-reasoned "this risk is low enough to accept without mitigation" is stronger than a vague mitigation strategy attached to every item.

EXCHANGE 12: FINAL CONSISTENCY CHECK

MY PROMPT

I have revised the risks section using the likelihood/severity/mitigation structure. I also caught one more thing on my own: Section 2 still said "approximately 30% variance between the manifest and what actually arrives on Thursday" but Section 5 said "manifest accuracy is a known issue but low severity." Those two sentences were inconsistent — 30% variance is not low severity. I changed the severity to medium and added a mitigation: we will track manifest-vs-arrival variance for three weeks during the test period to get a more accurate number than 30%.

Can you run one final check?

AI RESPONSE

Good catch. That was a real inconsistency and fixing it strengthens both sections.

I see two small things in the revised version:

Success criteria: You added "Maria states in her own words one specific change the tool made to her work" as a partner-facing criterion. That is good. But you did not note how you will collect it. Add: "Collected via a brief conversation at the Week 15 handoff session, recorded in writing by the team." One sentence. Otherwise the criterion floats without a collection method.

Section 6 (partner asks): You are asking Maria to walk you through five random spreadsheet rows in Week 8 for the data quality check. That ask is in Section 5 as a mitigation but it is not in Section 6 as a formal partner request. Maria needs to know you are going to ask for this. Add it to Section 6: "One 20-minute session in Week 8 to review five family profile entries for data quality and confirm that the dietary restriction column reflects current information."

Everything else is consistent. The proposal is in good shape. The timeline milestones are testable, the risks show your reasoning, and the solution is specific enough that a reader can picture what you are actually going to build.

Running a final AI consistency check after peer review revisions is worth the 10 minutes it takes. The two small issues Claude caught — no collection method for the partner success criterion and a missing partner ask in Section 6 — would have been easy to miss. Neither is major but both matter because the partner will read this document and the collection method tells Maria what to expect from you.