Apna Hisaab: Building Trustworthy AI by Teaching Women to Trust the Numbers, and Themselves

by Mehza Ayesh

I grew up watching women do invisible labor that never made it into any ledger. My mother managed a household budget she rarely had full control over. Aunts in my extended family quietly absorbed financial decisions made without them. Later, as I moved into classrooms and university hallways, the conversations didn't change as much as I expected them to. They just got more educated. My teachers, my class fellows, now my female university friends: financial anxiety followed almost all of them, regardless of how many degrees were in the room.

What struck me most wasn't that women lacked the ability to manage money. It was that so many of them had never been given the chance to build confidence in their own financial judgment. And the cost of that gap wasn't abstract. I watched it surface, again and again, as one of the most common reasons households fractured: money fights that weren't really about money, but about exclusion, control, and trust. The kind of conflict that doesn't stay contained to two adults arguing in a room, but seeps into how children grow up understanding stability, security, and their own worth. That's not just a personal observation; it tracks with documented research on South Asian households. One study on marital breakdown among Pakistanis found that two-thirds of the divorces examined were attributed primarily to disputes over finances (Qureshi, 2016).

That's the problem I carried into the Bank the Unbanked Ideathon at NIC: not "how do we get more women to use a banking app," but "how do we get women to trust that they're capable of managing money at all."

Why WhatsApp, Why AI?

Apna Hisaab is a WhatsApp-based AI financial assistant built for Pakistani housewives, women who are often financially literate in practice (they run households, stretch budgets, negotiate prices at the bazaar) but have never been formally included in financial systems built around them. We didn't build another banking app because the problem was never really about access to a platform. Most of these women already had smartphones. What they didn't have was a financial tool that spoke their language, asked nothing of them that felt intimidating, and didn't require them to walk into a bank branch and explain themselves to someone who might not take them seriously.

WhatsApp solved the access problem instantly: no new app to download, no unfamiliar interface, just a conversation in a space they already trusted. Trust isn't built by asking users to adopt yet

another platform; in mobile-first countries like Pakistan, trust is often built on the connectivity infrastructure people already use every day, not on convincing them to learn something new. The AI layer solved something subtler: it gave them a private, judgment-free space to ask questions they might never ask a husband, a brother, or even a bank teller. Questions like "is this a fair price," "how do I start saving even a small amount," or "what does this loan term actually mean." We weren't trying to make women into expert investors overnight. We were trying to give them a starting point that didn't come wrapped in shame.

That distinction, between access and trust, is, I think, the real story behind "Women, Trust, and the New AI Reality." AI doesn't automatically make systems more inclusive. It can just as easily replicate the same gatekeeping, the same assumption that financial competence belongs to someone else, unless it's deliberately built to do otherwise.

Building It as the Person It Was For

Pitching Apna Hisaab at the ideathon, I was one of very few women in a room full of male-dominated fintech teams pitching investment trackers and crypto wallets, most of them mixed-gender groups. I went up with a substitute partner after my actual teammate fell ill. I was also, by a fair margin, the youngest person presenting that day. Being young and visibly out of place in that room is its own quiet version of being asked to "check your assumptions with a more technical colleague," the same skepticism women in AI and tech routinely face, just aimed at age instead of gender. I chose to lead with the human story rather than the technical architecture, because I believed the story was about technology. The trust problem wasn't a UX detail to solve after the AI was built; it was the entire design brief. Making it to the top 12 nationally confirmed that being close to the problem wasn't a weakness in my pitch. It was the qualification that mattered most.

That same instinct has shaped the rest of my work too. I'm the first woman in my family to go into tech, and I've taught computer science at the secondary and pre-university level for three-plus years, in a field where only 14% of those enrolling in CS and IT programs in Pakistan are women, against a population that's over 49% female (per an IEEE study on women in Pakistan's computing sector). I see the gap show up in how some of my female students hesitate to commit to an answer even when they're right, the same hesitation I've watched grown women carry into financial decisions at home. My mother, aunts, and university teachers were among the first to ask how Apna Hisaab was going, and kept asking, long before any judge did. That sustained curiosity from women across my life was its own quiet validation.

What "Trust" Actually Means in This Context

If there's one thing I've learned since the ideathon, and I'm still learning it because Apna Hisaab is still actively in development, it's that trust in AI, for the women we're building for, isn't about

explainability dashboards or transparency reports in the abstract. It's much more basic than that. It's: will this thing judge me? Will it tell anyone what I asked for? Will it make me feel stupid for not knowing something I was never taught. Will it actually understand my life, or is it just another system built somewhere else, for someone else, repurposed for me.

But that basic, felt sense of trust still has to be backed by real design decisions, because good intentions don't prevent harm on their own. Building Apna Hisaab meant treating trust as a design constraint, not a feature, through four pillars that any team building AI for an underserved audience could apply.

Inclusive data, not just more data. Financial AI trained on existing transaction or credit data inherits whatever exclusion already exists in that data, and women (especially housewives with no formal credit history) are often invisible in it to begin with. We can't fix that by adding more data from the same broken pipeline. We have to deliberately design for users the underlying data was never collected on, treating their absence from existing systems as a constraint to design around, not a footnote.

Privacy as the default, not a setting. Privacy is what determines whether a woman tells the AI the truth at all. Knowing that no one else (not a husband, not a relative, not an algorithm reporting back to someone) would see what was asked was consistently the single biggest factor in whether women engaged honestly. For our users, the cost of one breach of trust is permanent, so privacy can't be something they have to find and enable.

Conversational transparency over technical transparency. Explaining model architecture is meaningless to someone without a finance or CS background, and pretending otherwise is its own kind of gatekeeping. Real transparency means telling a user plainly why the AI is suggesting what it's suggesting, in language she already uses, so she can disagree with it if her own judgment says otherwise.

Human oversight as a built-in limit, not an afterthought. It's tempting to let AI fully replace the gatekeeper it's meant to bypass. We didn't want that. The AI removes the shame and intimidation of asking a question; it was never meant to replace a woman's own judgment, or the option to escalate to a real person when a decision is too significant for an algorithm to carry alone. Designing AI that earns trust means designing it to know its own limits, not just its capabilities.

Put together, these aren't abstract principles to me. They're the actual difference between a tool that gets used once out of curiosity and one a woman keeps coming back to because it never once made her feel smaller for asking.

In practice, beyond the four principles above, that meant resisting one particular instinct throughout the build: designing for the most sophisticated user first and the most excluded user

last. It's easier to build for someone who already understands finance and already trusts technology. The harder, more necessary work, and the work that actually determines whether an AI tool earns trust rather than just claiming it, is building for the woman with the least reason to trust either.

Where This Goes

We're still building, still learning what these women actually need versus what we assumed they'd need going in, which, honestly, is most of the work. But the long-term vision hasn't changed since that first pitch: a tool that helps women build financial confidence privately, gradually, and on their own terms, so that the financial conversations happening in their households stop being battles over control and start being decisions made together.

I don't think AI alone fixes household financial conflict, and I'd be wary of anyone who claimed it could. But I do think it can do something quieter and more foundational: it can give a woman the words and the confidence to speak up in that conversation in the first place. Financial confidence is, in many ways, a form of digital empowerment: when a woman trusts AI enough to ask questions privately, she becomes more confident not just using financial services, but participating more fully in the digital economy around her. For me, that's what trust in AI should actually mean: not that the system is flawless, but that it finally listens to someone who, for too long, wasn't asked.