

VAMSHI JANDHYALA

How to protect elderly people from financial scams without taking away their independence



January 2026

A product design approach to helping banks protect elderly customers from scams while preserving dignity and autonomy through multi-party authentication and trust networks.

My dad got scammed last year.

He's sharp, educated, has always been careful with money. But he still fell for it. That's when I realized this problem isn't about intelligence, it's about how sophisticated these scams have become, and how isolated elderly people can be when making financial decisions.

Globally, elderly people control trillions in assets but are increasingly vulnerable. The question is: **How do we protect people from financial harm without treating them like children?**

The Scale of the Problem

This is a global crisis that's only getting worse.

The Numbers:

United States:

- $\$28.3$ billion to $\$1.5$ billion lost annually by older adults to financial exploitation (<https://www.cnbc.com/2025/12/13/financial-fraud-seniors-ftc.html>) {2024 FTC estimates}
- Average loss per victim: $\$83,000$
- Over 155,000 suspicious activity reports filed by financial institutions related to elder exploitation in one year alone, totaling $\$27$ billion (<https://www.fincen.gov/news/news-releases/fincen-issues-analysis-elder-financial-exploitation>)
- Only 1 in 44 cases ever gets reported

Japan:

- 36.25 million people aged 65+ (29.3% of population) as of 2024 (<https://www.japantimes.co.jp/news/2025/09/16/japan/society/japans-elderly-population/> Times)

- ¥71.8 billion lost to specialized fraud in 2024, up 58.6% from prior year (<https://www.nippon.com/en/japan-data/ho2424/>) {Nippon.com})
- 65% of fraud victims are seniors
- Average loss per scam: ¥3.5 million

Global:

- 1 in 6 people over 60 experience some form of abuse, including financial (<https://www.consumeraffairs.com/assisted-living/elder-abuse-statistics.html>) {Consumer Affairs})
- World's 60+ population projected to reach 2.1 billion by 2050

The Fraud Landscape:

- **"Grandparent scams"** (Japan's "ore-ore" scams): Criminals pretend to be family in urgent need
- **Investment fraud:** Fake bonds, crypto schemes targeting search for yield
- **Romance scams:** Targeting isolated widows/widowers
- **Digital exploitation:** Phishing, fake bank sites, elaborate social engineering

Who Gets Hurt

The elderly themselves face:

- Fear of losing independence and dignity
- Shame around being scammed (most never report it)
- Difficulty with complex technology
- Social isolation making them vulnerable to manipulation

Their adult children deal with:

- Anxiety about parents' financial safety
- Geographic separation (kids in cities, parents elsewhere)
- Legal limitations on intervention
- Guilt about potentially "taking control" away

Banks are caught between:

- Regulatory pressure to prevent fraud
- Reputational risk when scams happen
- Not wanting to be paternalistic
- Massive operational costs (\$27B in one year in suspicious activity)

Why This Happens

Psychologically: Age-related cognitive changes, optimism bias, social isolation, and scammers who target loneliness with artificial urgency.

Culturally: Stigma around cognitive decline, generational digital divide, and changing family structures (geographic spread).

Economically: Low interest rates drive risky investments, sophisticated criminal networks, branch closures reduce human oversight, and financial products grow increasingly complex.

What Banks Are Doing Today

Financial institutions are already taking measures to protect elderly customers:

Current Protections:

- **Staff training** to <https://www.consumerfinance.gov/consumer-tools/educator-tools/resources/exploitation-red-flags>
- **AI fraud detection** with <https://www.jackhenry.com/fintalk/2025-fraud-trends-protecting-against-emerging-threats/transaction-monitoring>
- **Transaction delays** for suspicious transfers (<https://www.biocatch.com/blog/unspoken-mandate-elder-scams>) (Financial Exploitation Prevention Act)
- **Customer education** on <https://www.wellsfargo.com/privacy-security/fraud/bank-scams/elder-financial-scams/>

Regulatory pressure is mounting. In December 2024, <https://www.occ.gov/news-issuances/news-releases/2024/nr-ia-2024-130.html> (Acting Comptroller Michael J. Hsu stated) that “doing the minimum is not enough” and urged banks to move beyond compliance checkboxes.

Critical gaps:

- Reactive, not proactive
- No family involvement mechanism
- Inconsistent across institutions
- High false positive rates

The trust network approach fills these gaps by combining AI detection with human judgment and enabling family involvement with consent.

Four Possible Solutions

Let me walk through four different approaches, from most promising to least.

Option 1: Tiered Risk System with Trusted Family Network

This is my recommended approach.

The Core Idea: Use machine learning to score transaction risk in real-time. For high-risk transactions, require approval from a trusted family member or friend before the money moves.

Three Risk Tiers:

- **Low Risk:** Immediate execution (utility bills, regular merchants you've paid before)
- **Medium Risk:** 24-hour cooling-off period + SMS verification
- **High Risk:** Someone you trust has to approve it

Setting Up Your Trust Network: You designate 2-4 people as your "watchers" or "peace-of-mind partners." Could be your daughter, a nephew, a close friend, or a professional fiduciary. They're ranked: primary, secondary, tertiary.

What Happens in Practice:

' Elderly customer tries to wire money to an unfamiliar account ↓ System flags it (large amount, new recipient, unusual timing) ↓ Customer's phone: "For your protection, your designated person will review this" ↓ Trusted person gets notification: "Transfer requested. Approve or discuss?" ↓ They call to verify ↓ Either approve if legitimate, or discover it's a scam ↓ Transaction approved or denied accordingly '

The money doesn't leave until verified. The scam is stopped before damage occurs.

The Risk Scoring: The ML model looks at 40+ signals:

- Transaction amount relative to your history
- Is this recipient new or someone you pay regularly?
- Behavioral patterns: Unusual time? Different device? Typing speed changes?
- Velocity: Multiple transfers in short time?
- External data: Is this account on known scam lists?
- Geolocation anomalies

Why This Works:

- Most transactions (90%+) go through instantly, you keep your independence
- Only weird stuff gets flagged
- Human judgment catches sophisticated social engineering that AI misses
- Builds on existing family/community support structures
- You control who's in your trust network
- Can be culturally adapted across different societies

Option 2: Cryptographic Key Splitting

This is the technical purist's approach.

How It Works: Your banking key gets mathematically split into pieces (Shamir's Secret Sharing). To make a transaction, you need 2-of-3 key pieces: yours + one from your trusted parties.

Pros:

- Mathematically guaranteed protection
- No bank can override it
- Perfect audit trail
- Works cross-border

Cons:

- Too technical for most elderly users
- Key loss creates unrecoverable situations
- Zero flexibility for changing circumstances
- Expensive to implement

This might work for crypto enthusiasts, but not for most elderly people.

Option 3: AI Warning System

How It Works: An AI monitors your behavior and flags anomalies. When something seems off, it generates an explanation: "This is unusual because you've never sent money to this country before."

You acknowledge the warning and answer verification questions. No one else involved.

Pros:

- Complete autonomy preserved
- Immediate intervention
- Scales infinitely
- Privacy-preserving

Cons:

- Scammers can coach you through the warnings
- No human judgment
- People develop "warning fatigue" and start ignoring them
- No external accountability

We've seen this fail repeatedly, scammers literally have scripts for bypassing security questions.

Option 4: Age-Based Transaction Limits

This is what some banks are actually doing.

How It Works: Hard caps on transaction amounts if you're 75+. Want to make an exception? Come to the branch, talk to a manager.

Pros:

- Dead simple to implement
- Clear legal protection for banks

Cons:

- Ageist and dignity-destroying
- Punishes competent people indiscriminately
- Likely illegal under discrimination laws
- Drives people to less-protected alternatives

This treats all elderly people like children regardless of capacity.

How to Evaluate These Options

I scored each solution across six dimensions (weighted by importance):

| Dimension | Why It Matters | Weight | ||, , , -, , , , , -, , , -| | **Security** | The whole point | 25% | | **Autonomy** | Can't infantilize people | 20% | | **Cultural Adaptability** | Must work across different societies | 20% | | **Feasibility** | Can we actually build this? | 15% | | **Scalability** | Works for millions of customers? | 10% | | **Legal Compliance** | Regulatory requirements | 10% |

Security (25% weight)

| Solution | Score | Why | ||, , , -, , , -| | Tiered Trust Network | 9/10 | Human judgment catches sophisticated scams; adapts to new threats | | Crypto Key Splitting | 7/10 | Strong technically but can't stop a willing-but-deceived user | | AI Warnings | 6/10 | Vulnerable to adversarial coaching | | Age Limits | 4/10 | Easily circumvented (multiple small transactions) |

Autonomy (20% weight)

| Solution | Score | Why | ||, , , -, , , -| | Tiered Trust Network | 8/10 | You control who's in your network; most transactions unaffected | | Crypto Key Splitting | 5/10 | Loss of unilateral control | | AI Warnings | 9/10 | You make final decision | | Age Limits | 2/10 | Blanket restriction regardless of capacity |

Key insight: Autonomy isn't binary. The UN Convention on Rights of Persons with Disabilities distinguishes between *supported decision-making* (good) vs. *substituted decision-making* (bad). We want the former.

Cultural Adaptability (20% weight)

How the trust network approach adapts to different cultural contexts:

| Culture Type | Adaptation Strategy | Example | | , , , - | , , , , , | , , , | | **Collectivist (East Asia)** | Emphasize family duty, multi-generational obligation | Japan, Korea, China: "watching over" tradition | | **Individualist (US, Australia)** | Emphasize personal control, choice, legal rights | "You choose your protection team" | | **Family-Oriented (Mediterranean, Latin America)** | Extended family networks, community trust | Include cousins, godparents in trust networks | | **High-Trust Institutions (Northern Europe)** | Professional fiduciaries, state partnership | Sweden, Denmark: integrate with social services |

| Solution | Score | Why | | , , - | , , - | - | | Tiered Trust Network | 9/10 | Highly adaptable to cultural contexts | | Crypto Key Splitting | 4/10 | Too mechanistic; culturally alien in most societies | | AI Warnings | 6/10 | Acceptable but impersonal globally | | Age Limits | 3/10 | Discriminatory across all cultures |

Final Scores

| Solution | Total Score | | , , - | , , - | | **Tiered Trust Network** | **8.55/10** | | AI Warnings | 7.45/10 | | Crypto Key Splitting | 5.45/10 | | Age Limits | 4.90/10 |

Competitive Landscape

This isn't a greenfield problem, several companies are already working on elder fraud detection:

BioCatch uses [\{\}href{https://www.biocatch.com/blog/unspoken-mandate-elder-scams}](https://www.biocatch.com/blog/unspoken-mandate-elder-scams) {behavioral biometrics} to detect cognitive decline and coercion in real-time. They analyze how users type, move their mouse, and navigate apps to spot anomalies that indicate fraud or diminished capacity. Strong on detection, but lacks the family network intervention layer.

Feedzai provides enterprise fraud prevention with ML models that flag suspicious transactions. Used by major banks globally. Excellent at pattern recognition, but purely reactive, no proactive family involvement.

Existing bank systems (Wells Fargo, Bank of America, JPMorgan) have in-house fraud teams and rule-based systems. These catch obvious fraud but struggle with sophisticated social engineering where the customer willingly authorizes the transaction.

What's missing across all of these: The human verification layer. Current systems either block transactions (frustrating) or allow them with warnings (ineffective against coached victims). None systematically involve trusted family members in high-risk decisions.

The trust network approach fills this gap by combining behavioral AI (like BioCatch) with human judgment (what banks do manually today) in an automated, scalable way.

Why the Trust Network Approach Wins

1. Behavioral economics principles: Uses default effects, loss framing, cooling-off periods, and commitment devices to make protection feel natural rather than

restrictive.

2. Economic viability as a retention play: This isn't just about preventing losses, it's about preserving the bank's highest-LTV customers. Elderly customers hold the largest deposits and longest relationships. A \$83K scam doesn't just cost the bank reputational damage; it often triggers the customer leaving for a competitor ("They should have protected me"). The real ROI is customer retention: keeping a high-net-worth elderly customer for 10+ more years vs. losing them and their entire family after a scam. Even if trust network setup costs \$50-100 per customer, preventing one account closure from a \$500K deposit customer justifies hundreds of enrollments.

3. Cultural adaptability: Terminology and enrollment strategies adapt to local contexts, "Watcher" in Japan, "Trusted Advisor" in the US, "Financial Supporter" in the UK.

4. Mature ML technology with behavioral UI: Gradient Boosted Trees analyze transaction signals (30%), behavioral patterns (25%), context (20%), history (15%), and external data (10%). Fairness constraints ensure age/gender/race aren't direct features.

Critically, the UI must address the "confidence paradox": Japanese data shows 90% of scam victims believed they were "too smart to be scammed." Traditional warnings ("This looks like a scam") trigger defensive reactions and get ignored.

Better approach, contextual nudges:

- □ "Warning: Potential scam detected"
- □ "Your daughter's voice sounds different than usual. Let's pause for 10 minutes so you can call her back on the number you have saved."
- □ "This recipient is flagged as high risk"
- □ "You've never sent money to Nigeria before. Your trusted contact Sarah will review this, she'll call you within 2 hours."
- □ "Are you sure you want to proceed?"
- □ "This wire can't be reversed. Let's have your son Tom verify the recipient details. We'll send him the last 4 digits of the account only."

The AI doesn't accuse the user of being fooled, it provides a face-saving pause that feels helpful rather than condescending.

5. Failure mitigation:

- False positives → rapid appeal + bank override
- Unavailable trusted party → multi-party hierarchy + 48hr bank review
- Collusion → pattern monitoring + cooling-off for changes
- Low enrollment → voluntary system with AI fallback

Key Risks & Open Questions

Let's be honest about what we don't know and what could go wrong:

Adoption Risk: Will elderly customers actually enroll? We need user research to understand mental models around “giving family access” vs. “getting family protection.” Early data from similar programs suggests 30-40% voluntary adoption in first year with good messaging, but this is hypothetical without real pilots.

Regulatory Uncertainty & Safe Harbor Strategy: Financial regulations vary globally. In the US, GLBA restrictions on information sharing could create liability. In the EU, GDPR’s data minimization principles might conflict with trust network notifications. The <https://www.biocatch.com/blog/unsspoken-mandate-elder-scams> {Financial Exploitation Prevention Act} suggests US regulators are supportive, but banks need legal cover.

The path forward: Partner with the Consumer Financial Protection Bureau (CFPB) or FCA (UK) to create a regulatory sandbox pilot. Frame it as “Supported Decision-Making” (aligned with UN Convention on Rights of Persons with Disabilities Article 12) rather than “monitoring.” Seek explicit Safe Harbor protection: banks that implement trust networks with proper consent can’t be sued for data sharing within the network. This creates regulatory air cover and potentially accelerates adoption if it becomes a best practice standard.

Adversarial Adaptation: Scammers will evolve. Once they know about trust networks, they might:

- Target the trusted parties directly with phishing
- Coach victims to remove family members from trust networks before scamming
- Focus on transactions just below the flagging threshold

How fast can we iterate the ML model? How do we prevent trust network manipulation?

False Positive Costs: What’s the real-world rate? If 10% of legitimate transactions get flagged and require family approval, will customers disable the feature out of frustration? We need precision-recall data from pilots.

Technical Debt: Banks have legacy systems. Integration complexity could be massive. Is this a 6-month or 3-year implementation? What’s the realistic cost per customer?

Family Dynamics: In 5-10% of cases, families might abuse this system, freezing out legitimate transactions for control. How do we detect and prevent this without undermining the system?

Cross-Border Complexity: What happens when the elderly person is in Japan, their trusted party is in the US, and they’re wiring money to the UK? Latency, time zones, and legal jurisdictions all complicate implementation.

I don’t have answers to all these questions. They’d require pilots, user research, and real-world testing. But acknowledging them upfront is critical, overconfident product proposals fail when they hit reality.

Some Counterarguments

“Why not just use AI?” Scammers adapt faster than models retrain and coach victims through security measures. Social engineering exploits human trust, you need human

judgment to counter it. The hybrid approach wins: AI for triage (99%), humans for edge cases (1%).

"What about privacy?" Design with opt-in enrollment, activity logs for transparency, data minimization (trusted parties see risk scores and amounts, not purposes or balances), user control to change networks, and right to disable. Legal basis: consent + legitimate interest (fraud prevention).

"Won't this create family conflict?" Sometimes, but post-scam conflict is worse. This formalizes the process transparently with user-chosen networks. False positives will be ~5-10% of transactions, making verification routine.

"What about weak rule of law countries?" This requires functional banking infrastructure, institutional trust, and consumer protection laws. It's viable for developed and middle-income countries where most elderly wealth is concentrated.

How I'd Validate This

This proposal needs real-world testing. Here's the validation approach:

Qualitative Research: Interview 30-40 elderly customers and adult children to understand mental models around "giving family access" vs. "getting family protection." Test 6-8 different messaging framings.

Prototype Testing: Usability test enrollment flow with 15-20 elderly users across digital literacy levels. Can they complete enrollment? Do they understand what they're agreeing to?

Limited Pilot: Partner with a regional bank for 500-1K voluntary participants. Run in shadow mode first (flag but don't block). Target metrics: 35%+ enrollment, 90%+ review completion, <8% false positives, 4.2/5 satisfaction.

Expansion: Scale to 10K+ customers. A/B test risk thresholds and network sizes. Measure retention and feature abandonment.

Kill criteria: Enrollment <20%, false positives >15%, >10% disable within 3 months, widespread family abuse, or prohibitive technical costs.

Expansion criteria: 40%+ enrollment, <5% false positives, measurable fraud reduction, positive NPS, regulatory support.

Final Thoughts

This isn't just an aging societies problem, it's a preview of the global demographic future. The countries that solve elder financial protection first will protect trillions in assets while building exportable fintech models.

The best product design respects user humanity. For elderly people, autonomy is central to identity. This approach threads the needle: protecting without patronizing, supporting without smothering.

As gerontologist Dr. Bill Thomas says: *"The opposite of independence isn't dependence, it's interdependence."*

We're not taking control away, we're activating existing social networks when risk emerges. Using technology to enable human judgment, not replace it.

The trillions managed by the world's elderly represent decades of life savings, dignified aging, and generational inheritance. We can protect it without infantilizing its owners.

References

Data and concepts from:

- <https://www.cnbc.com/2025/12/13/financial-fraud-seniors-ftc.html> {FTC Elder Financial Fraud Report (2024)} - \$28.3-81.5B annual losses
- <https://www.fincen.gov/news/news-releases/fincen-issues-analysis-elder-financial-exploitation> {Elder Financial Exploitation Analysis} - 155K+ reports, \$27B
- <https://www.japantimes.co.jp/news/2025/09/16/japan/society/japans-elderly-population> {Times - Elderly Population Statistics} - 36.25M aged 65+
- <https://www.nippon.com/en/japan-data/h02424/> {Nippon.com - Japanese Fraud Statistics} - ¥71.8B losses in 2024
- <https://www.consumeraffairs.com/assisted-living/elder-abuse-statistics.html> {Consumer Affairs - Elder Abuse Statistics} - Global prevalence
- <https://www8.cao.go.jp/kourei/english/annualreport/index-wh.html> {Cabinet Office of Japan - Aging Society Report}
- <https://www.ohchr.org/en/instruments-mechanisms/instruments/convention-rights-persons> {Convention on the Rights of Persons with Disabilities - Article 12}
- Thaler & Sunstein's *Nudge: Improving Decisions About Health, Wealth, and Happiness*
- Kahneman's *Thinking, Fast and Slow*