

When AI Starts Spending for Us: From Financial Literacy to Delegation Literacy

Amirah Shazana Magli¹

¹Faculty of Business and Economics, Universiti Malaya, Malaysia

Introduction

In March 2026, an artificial intelligence (AI) agent booked a ride from Kuala Lumpur International Airport to KL Sentral and completed the payment on the passenger's behalf. Conducted in a controlled pilot involving Mastercard, CIMB and RHB, the transaction used tokenised credentials, explicit consumer consent and purchase confirmation through a payment passkey (Mastercard, 2026). The immediate story was technological: Malaysia had completed its first authenticated agentic transaction. The more consequential story, however, was financial. A machine had moved beyond providing information or recommending a choice. It had acted and paid.

The transaction was deliberately bounded and authenticated. Nevertheless, it offers a glimpse of a much wider shift. AI tools are already helping consumers compare products, interpret investment information, prepare budgets and generate financial recommendations. Agentic systems add a further capability: they can plan and execute tasks on the consumer's behalf. The OECD (2026d) describes agentic AI in terms of goal pursuit, autonomy, interaction and the ability to take actions within digital environments. Once connected to payment systems, an agent may not merely say which option appears best; it may select the option, place the order and complete the transaction.

This changes the financial-literacy question. Consumers still need to understand prices, interest, risk, inflation, diversification and affordability. They must now also decide what a machine is authorised to do with their money, which objective it should pursue, where its authority ends and when human

approval remains necessary. An efficient transaction is not necessarily a sound household financial decision. AI can optimise a purchase without optimising financial well-being.

From Financial Assistance to Financial Action

The distinction between generative and agentic AI is important. Generative AI produces content in response to prompts: an explanation, comparison, forecast or suggested budget. The consumer ordinarily decides what to do next. Agentic AI is designed to pursue an objective through a sequence of actions, potentially using external tools and making decisions with limited intervention. The line between recommendation and execution therefore becomes thinner.

This development can create genuine consumer value. An agent could search across providers, reduce information costs, identify unnecessary fees, schedule payments, find an insurance option that meets specified requirements or help a time-poor household manage routine expenses. The OECD (2026a) notes that AI can improve access to financial information, personalisation and decision support. Research on digital finance likewise shows that digital financial literacy is associated with the effective use of digital financial services and improved financial well-being (Choung et al., 2023; Lyons & Kass-Hanna, 2021). Properly designed automation can reduce administrative burden and help consumers act on intentions that might otherwise be neglected.

The benefit, however, depends on what the agent has been told to optimise and what it is able to observe. A household's financial life is not a single-variable problem. Price matters, but so do cash-flow

Corresponding author: ¹ amirahshazana@um.edu.my

timing, existing debt, emergency savings, family commitments, risk tolerance and future goals. An agent asked to find the lowest monthly instalment may produce a technically correct answer while increasing the total financing cost. An agent asked to obtain the cheapest flight may select an inflexible fare that becomes costly when a family's plans change. A subscription agent may renew a service because the price remains within its permitted limit, even though the household no longer values it.

These are not necessarily system failures. They may be examples of a system competently pursuing an incomplete instruction. The gap lies between computational optimisation and household judgement.

When Convenience Is Not Financial Well-Being

The shift towards agentic commerce arrives in an environment where payment is already becoming less visible. Experimental evidence indicates that mobile payment can increase willingness to pay and shorten payment time compared with cash (Liu et al., 2021). This does not mean that every digital payment causes overspending, but it demonstrates why transaction design matters. When searching, choosing and paying are compressed into a nearly frictionless process, the pause during which a consumer might reconsider affordability can also disappear.

AI introduces another layer. Personalisation can make choices more relevant, but it can also make commercial influence harder to recognise. A recommendation may reflect the consumer's stated interest, the platform's ranking logic, a sponsored arrangement or some combination of these. A recent systematic review of AI and consumer financial behaviour identifies consumer interaction, behavioural outcomes and ethical and regulatory challenges as major areas requiring attention (Meng et al., 2025). The OECD (2026a) similarly highlights risks involving inaccurate content, bias, commercial influence, privacy, exclusion and over-reliance, while

noting that the long-term benefits for financial well-being remain uncertain.

Consumers' responses to algorithmic advice are also not uniform. Some people avoid algorithms, especially for tasks they regard as subjective; others give algorithmic recommendations greater weight than human advice. Reviews of this literature show that reliance varies according to the task, the person, the system and the surrounding context (Burton et al., 2020; Mahmud et al., 2022). Trust is therefore not a simple choice between trusting and distrusting AI. The relevant objective is calibrated reliance: using the system where it adds value while preserving scrutiny where the consequences are material.

This is especially important in personal finance because small decisions accumulate. A single unnecessary purchase may be manageable. Repeated automated renewals, frequent optimised offers or a series of low-friction instalment decisions can gradually weaken cash flow. Household financial resilience depends not only on access to technology but also on the capability to use financial resources in ways that protect the household against shocks. Evidence across emerging economies links financial and digital literacy with stronger financial resilience, although the benefits vary across groups and contexts (Kass-Hanna et al., 2022). Digital capability should therefore be judged by what it enables households to achieve, not merely by how many digital tools they use.

Financial Literacy Must Evolve

Financial literacy has never meant knowledge alone. It includes the awareness, skills, attitudes and behaviours needed to make informed decisions and support financial well-being. Digital financial literacy extends this capability to digital products, platforms, risks and transactions (Koskelainen et al., 2023; Lyons & Kass-Hanna, 2021). The OECD's regional framework for adults in ASEAN reinforces the need to understand digital financial products, protect

personal data, identify risks and use digital services safely (OECD, 2026c).

These competencies remain essential, but agentic commerce changes the locus of decision-making. Traditional financial education generally assumes that the person learning is also the person making the final decision. Even when a calculator, comparison platform or robo-adviser is used, the consumer is expected to assess the recommendation and press the final button. Agentic transactions create a different arrangement: the consumer defines some instructions in advance and the system acts later within those instructions.

For this emerging setting, delegation literacy offers a practical extension of digital financial literacy. It refers to a consumer's capacity to define, limit, monitor and revoke the financial authority given to an AI system. The concept is not presented as an empirically validated construct. Instead, it provides a practitioner-oriented framework for translating established concerns about financial literacy, autonomy and human oversight into questions that consumers and advisers can apply.

This distinction is supported by research beyond finance. Experiments on delegating purchasing tasks to AI show that consumers differentiate between autonomy over the available choices and autonomy over the final decision; reducing either can affect their willingness to use AI (Ahmad Husairi & Rossi, 2024). Research also indicates that explainability can influence users' trust and acceptance of AI systems (Shin, 2021). In other words, people require more than an accurate output. They need an intelligible role in the decision process.

Delegation literacy does not mean refusing automation. It means deciding deliberately which functions may be automated, under what conditions and with what safeguards. Just as a financial plan establishes goals and constraints before products are selected, an AI-enabled financial arrangement should

establish authority and boundaries before transactions are executed.

Seven Questions Before an AI Is Allowed to Act

The following questions translate delegation literacy into a practical decision framework.

1. What exactly am I authorising?

Permission to search is not permission to purchase. Permission to pay one bill is not permission to manage an entire category of expenses. Consumers should be able to distinguish among information access, recommendation, selection, commitment and payment. Each represents a different level of delegated authority.

2. What objective is the AI being asked to optimise?

The instruction to find the “best” option is incomplete. Best may mean lowest upfront price, lowest total cost, lowest risk, highest quality, greatest convenience or best alignment with the household budget. Consumers should specify the objective and understand what may be sacrificed when one outcome is prioritised.

3. What financial and behavioural limits have I imposed?

A useful mandate should include transaction ceilings, cumulative limits, permitted product categories and an expiry period. It may also include behavioural safeguards: no purchases after a certain hour, no new instalment commitments, no automatic renewal after a price increase and no transaction that reduces the emergency fund below a defined threshold.

4. Which decisions still require human approval?

Routine, reversible and low-value tasks may be suitable for automation. Complex, long-term, high-value or difficult-to-reverse decisions should trigger explicit approval. Borrowing, investment switching, insurance cancellation, changes affecting joint household finances and transactions involving vulnerable family members deserve a higher threshold.

5. How will I know what the agent has done?

Consent at the beginning is insufficient if subsequent actions are invisible. Consumers need accessible transaction records, timely notifications and plain-language explanations linking each action to the authority previously granted. Monitoring should show not only what was paid but also why the action met the consumer's rules.

6. Can I pause, revoke, challenge or reverse the action?

Control must remain usable after consent is given. A consumer should know how to suspend an agent immediately, withdraw access, dispute an unsuitable transaction and seek redress. Reversibility may not always be possible, which makes approval thresholds more important before an irreversible commitment occurs.

7. Who is accountable when something goes wrong?

Consumers should not be expected to resolve a chain of responsibility involving an AI provider, platform, merchant, bank and payment network. The responsible parties and complaint route should be clear before use. Accountability cannot disappear simply because the transaction

was distributed across several technologies.

A New Assignment for Financial Planners

Agentic commerce will not reduce the need for financial planners. It may change part of their work. Advisers have traditionally helped clients translate life goals into budgets, protection needs, investment strategies and estate plans. They may increasingly need to translate those same goals into rules that AI systems can follow without undermining the wider financial plan.

This begins with an AI authority review. Alongside reviewing income, commitments, assets and risk tolerance, planners can ask clients which applications have access to their financial data, which services can initiate payments and which decisions are automated. Many consumers may know that they approved an application but may not recall the scope or duration of that authority.

Planners can then help clients establish a delegation budget. This is not another spending category. It is a set of limits governing what an AI system may do: a maximum value per transaction, a monthly cumulative amount, permitted merchants or product types, and situations requiring confirmation. These limits should reflect the household's cash flow and resilience, rather than the maximum spending level a platform permits.

A third task is identifying non-delegable decisions. A household may be comfortable automating utility payments and routine reordering but not credit applications, investment liquidation, insurance cancellation or transfers from emergency savings. The boundary should be explicit before a persuasive interface or urgent situation tests it.

Finally, planners should incorporate automation records into periodic reviews. A conventional budget shows where money went. An agentic-finance review

should also ask which decisions were machine-initiated, whether they remained aligned with client goals and whether the rules should be tightened. Human advice is most valuable not when it competes with AI on speed, but when it restores context, challenges an incomplete objective and recognises consequences that a transaction-level system may not see.

The current evidence supports this hybrid role. In an HSBC-commissioned Ipsos survey of about 10,000 affluent and high-net-worth individuals across ten markets, 85% of Malaysian respondents reported using AI for finance and investment. Yet 58% preferred a future approach combining AI with professional human expertise (HSBC Malaysia, 2026). These figures apply to a specific affluent sample and should not be generalised to all Malaysians. Even so, they reveal an important preference: consumers may welcome machine assistance without wanting to surrender human judgement.

Responsibility Cannot Rest on Consumers Alone

Delegation literacy is necessary, but education cannot compensate for poor system design. Financial institutions, payment networks, AI providers and merchants have responsibilities that begin before the consumer grants authority.

Consent should be specific, comprehensible and proportionate to the action. Interfaces should distinguish permission to recommend from permission to transact. Authority should be limited by purpose, amount and time, with conservative defaults rather than open-ended access. High-risk transactions should require renewed human confirmation. Consumers should receive clear records of the agent's actions and have a simple way to pause or revoke access.

Explanations must also be decision-useful. A lengthy technical description of a model may satisfy neither comprehension nor accountability. Consumers need to know which objective was pursued, what information materially influenced the choice, what alternatives were excluded and whether a commercial relationship shaped the recommendation. Research on trust in AI shows that trust is affected by both system characteristics and how the technology is represented and experienced (Glikson & Woolley, 2020). Transparency should therefore help consumers calibrate reliance, not merely reassure them.

Additional safeguards are needed for vulnerable users. Older adults, people with limited digital skills, individuals under financial stress and consumers managing money jointly may face different risks. JWMFP discussions of youth bankruptcy and elder financial abuse already demonstrate that financial vulnerability is shaped by life stage, capability and relationships of trust (Ating et al., 2024; Md. Sum & Nordin, 2025). Agentic systems should not assume that every account holder has equal capacity to understand, monitor or challenge automated actions.

Regulators and financial-education stakeholders should consequently examine agentic commerce as both a consumer-protection and capability issue. The OECD's *Consumer Finance Risk Monitor 2026* warns that opaque algorithms, manipulative digital design, limited human support and unreliable AI assistance can create consumer risks (OECD, 2026b). Policy should encourage innovation while preserving meaningful choice, contestability, redress and access to human support.

Conclusion

The Malaysian pilot was designed around consent, authentication and consumer control. Those safeguards are a sound starting point. The next challenge is ensuring that control remains meaningful as agentic payments move from controlled pilots to

routine purchases and more consequential financial decisions.

For years, financial education has taught consumers to think before they spend. In an agentic economy, that lesson must be extended: consumers must think carefully before authorising a system that may spend later. They need to understand not only the product and price but also the mandate, objective, limits and route for taking authority back.

At the policy level, Malaysia should move early to embed these safeguards in both financial education and consumer-protection practice. The Financial Education Network could incorporate delegation literacy into national learning materials, while relevant regulators could develop risk-based expectations for agentic transactions: purpose-specific and time-limited consent; per-transaction and cumulative limits; renewed human approval for credit, investment, insurance and other high-impact decisions; auditable records; simple pause and revocation functions; and clear liability and redress arrangements. Pilot evaluations should assess not only whether a transaction is secure and successful, but also whether consumers understand the authority granted, can monitor its use and can regain control without undue difficulty. These assessments should include older adults, consumers with limited digital or financial capability and households under financial stress, so that innovation does not widen existing vulnerabilities.

AI can make personal finance faster, more accessible and more responsive. It can also execute an incomplete instruction with remarkable efficiency. Progress should therefore not be measured only by how many decisions can be automated. It should be measured by whether households retain the judgement and control needed to protect their financial well-being.

The next financial-literacy test is not simply whether AI can be trusted with consumers' money. It

is whether consumers understand how much financial judgement they are willing to delegate—and how to take that authority back.

References

- Ahmad Husairi, M., & Rossi, P. (2024). Delegation of purchasing tasks to AI: The role of perceived choice and decision autonomy. *Decision Support Systems*, 179, 114166. <https://doi.org/10.1016/j.dss.2023.114166>
- Ating, R., Palmares-Cruz, M. R. M. J., & Lalisian, A. K. T. (2024). Bankruptcy among Malaysian millennials: Due to lavish lifestyles or low financial literacy as evidenced by PISA examination results? *Journal of Wealth Management & Financial Planning*, 11, 60–66.
- Burton, J. W., Stein, M.-K., & Jensen, T. B. (2020). A systematic review of algorithm aversion in augmented decision making. *Journal of Behavioral Decision Making*, 33(2), 220–239. <https://doi.org/10.1002/bdm.2155>
- Choung, Y., Chatterjee, S., & Pak, T.-Y. (2023). Digital financial literacy and financial well-being. *Finance Research Letters*, 58, 104438. <https://doi.org/10.1016/j.frl.2023.104438>
- Glikson, E., & Woolley, A. W. (2020). Human trust in artificial intelligence: Review of empirical research. *Academy of Management Annals*, 14(2), 627–660. <https://doi.org/10.5465/annals.2018.0057>
- HSBC Malaysia. (2026, August 10). *Malaysia's affluent and high-net-worth investors rank among the highest in adopting AI for finance, but more than half want human-AI synergy in making financial decisions.* <https://www.about.hsbc.com.my/news-and-media/malaysias-affluent-and-high-net-worth-investors-rank-among-the-highest-in-adopting-ai-for-finance>

- Kass-Hanna, J., Lyons, A. C., & Liu, F. (2022). Building financial resilience through financial and digital literacy in South Asia and Sub-Saharan Africa. *Emerging Markets Review*, 51, 100846. <https://doi.org/10.1016/j.ememar.2021.100846>
- Koskelainen, T., Kalmi, P., Scornavacca, E., & Vartiainen, T. (2023). Financial literacy in the digital age—A research agenda. *Journal of Consumer Affairs*, 57(1), 507–528. <https://doi.org/10.1111/joca.12510>
- Liu, Y., Luo, J., & Zhang, L. (2021). The effects of mobile payment on consumer behavior. *Journal of Consumer Behaviour*, 20(3), 512–520. <https://doi.org/10.1002/cb.1880>
- Lyons, A. C., & Kass-Hanna, J. (2021). A methodological overview to defining and measuring “digital” financial literacy. *Financial Planning Review*, 4(2), e1113. <https://doi.org/10.1002/cfp2.1113>
- Mahmud, H., Islam, A. K. M. N., Ahmed, S. I., & Smolander, K. (2022). What influences algorithmic decision-making? A systematic literature review on algorithm aversion. *Technological Forecasting and Social Change*, 175, 121390. <https://doi.org/10.1016/j.techfore.2021.121390>
- Mastercard. (2026, March 4). *Mastercard conducts first live agentic transaction in Malaysia with CIMB and RHB pilot*. <https://www.mastercard.com/news/ap/en/newsroom/press-releases/en/2026/mastercard-conducts-first-live-agentic-transaction-in-malaysia-with-cimb-and-rhb-pilot/>
- Md. Sum, R., & Nordin, N. (2025). The hidden threat: Elder financial abuse in Malaysia. *Journal of Wealth Management & Financial Planning*, 12, 52–57.
- Meng, K., Mahapatra, M. S., & Xiao, J. J. (2025). Artificial intelligence and consumer financial behavior: A systematic literature review and agenda for future research. *Journal of Consumer Behaviour*, 24(4), 1755–1786. <https://doi.org/10.1002/cb.2497>
- Organisation for Economic Co-operation and Development. (2026a). *Artificial intelligence and personal finance* (OECD Artificial Intelligence Papers No. 62). OECD Publishing. <https://doi.org/10.1787/2858fdf4-en>
- Organisation for Economic Co-operation and Development. (2026b). *Consumer finance risk monitor 2026*. OECD Publishing. <https://doi.org/10.1787/61f7dbe0-en>
- Organisation for Economic Co-operation and Development. (2026c). *Digital financial literacy core competency framework for adults in ASEAN*. OECD Publishing. <https://doi.org/10.1787/be938ace-en>
- Organisation for Economic Co-operation and Development. (2026d). *The agentic AI landscape and its conceptual foundations* (OECD Artificial Intelligence Papers No. 56). OECD Publishing. <https://doi.org/10.1787/396cf758-en>
- Shin, D. (2021). The effects of explainability and causability on perception, trust, and acceptance: Implications for explainable AI. *International Journal of Human-Computer Studies*, 146, 102551. <https://doi.org/10.1016/j.ijhcs.2020.102551>