Building Trust in Agentic Commerce

Customer-centric principles for
a trusted and thriving ecosystem

ASB
ONE STEP AHEAD

BANK OF AMERICA

Capital One

Commonwealth Bank of Australia

ING

NatWest Group

Important information

This paper is co-authored by ASB Bank Ltd., Bank of America, Capital One, Commonwealth Bank of Australia, ING Group and NatWest Group. This paper is intended to set out a series of high-level principles shared by its authors. It is not intended to and does not prescribe conduct, any commercial position or any particular method or timetable for implementation. The principles described are voluntary and non-binding. Nothing in this paper requires the authors or other industry participants to take any particular action or to adopt or implement any principle.

The information in this paper and any opinions, conclusions or recommendations are reasonably held or made, based on the information available at the time of its publication, but no representation or warranty, either express or implied, is made or provided as to the accuracy, reliability or completeness of any statement made in this paper. The authors do not endorse the services of a particular provider.

Customers are our driving force for building trust and scale in agentic commerce

Our customers, both consumers[1] and merchants, have expressed excitement about the potential of agentic commerce. Consumers see value in finding what they need at the best value, quickly and easily, and with added levels of advice and automation. Merchants see opportunities to reach new customers, increase sales and offer personalized service to enhance loyalty. We, too, are excited by the promise of agentic commerce and are eager to work with customers, industry and stakeholders to enable its future.

As highly regulated financial entities, we are focused on managing risk effectively as emerging technologies arise. This is why this group of institutions is identifying best practices and a trusted framework as the agentic commerce ecosystem evolves. Consumers are unclear if AI agents will act in their interests. They are concerned that AI agents may buy the wrong thing or spend too much – or even worse, lose their money to scams and fraud. They are not sure whether they will be protected or who they will need to go to if things go wrong. Additionally, merchants are concerned that disputes and chargebacks may increase for reasons outside their control. They also worry that they may lose their connection to consumers, and that they may have to choose between higher costs and losing sales.

We welcome the progress made by AI agents, technology providers and payment schemes in establishing building blocks for agentic commerce, which will help unlock opportunities and manage risks. It is incumbent on all parties across the ecosystem to build on these foundations to manage outstanding risks, and foster well-founded trust and confidence in agentic commerce among consumers and merchants.

Our objective is to support a thriving agentic commerce ecosystem that is safe, effective and efficient for all – especially our customers. To achieve this, we propose the following set of common principles:

1. Transparency | **2.** Safety | **3.** Privacy & Data | **4.** Choice | **5.** Interoperability

This paper outlines how these principles can support trust and scale in agentic commerce. Through a subsequent paper, we plan to detail how the principles can be implemented through protocols, and as relevant industry standards and policies[2]. As we progress, we welcome feedback and engagement with consumers, merchants, AI agent providers, payment schemes, issuers, acquirers, industry bodies, regulators and policymakers. We look forward to supporting a safe, trusted and vibrant agentic commerce ecosystem for all.

[1] For simplicity, "consumers" in this context is also taken to include businesses making agentic commerce purchases from merchants.

[2] Protocols typically provide technical rules that allow agentic commerce participants to interact; industry standards typically establish agreed conventions and consistent requirements across the ecosystem; and policies guide how the industry is governed.

Trust is critical to unlock the potential of agentic commerce

We define agentic commerce as the “use of AI agents to help make or facilitate payments between a consumer and a merchant for a product or service”[3]. Agentic commerce can occur across any payment rail, including cards, account-to-account and tokenized money. Several different service providers may be involved across the value chain, including AI agent providers, payment schemes, digital wallet providers, issuers, payment service providers and acquirers.

Agentic commerce can involve different degrees of AI agent autonomy, with two key variants being human-in-the-loop agentic commerce and autonomous agentic commerce.

Increased level of autonomy given to agents

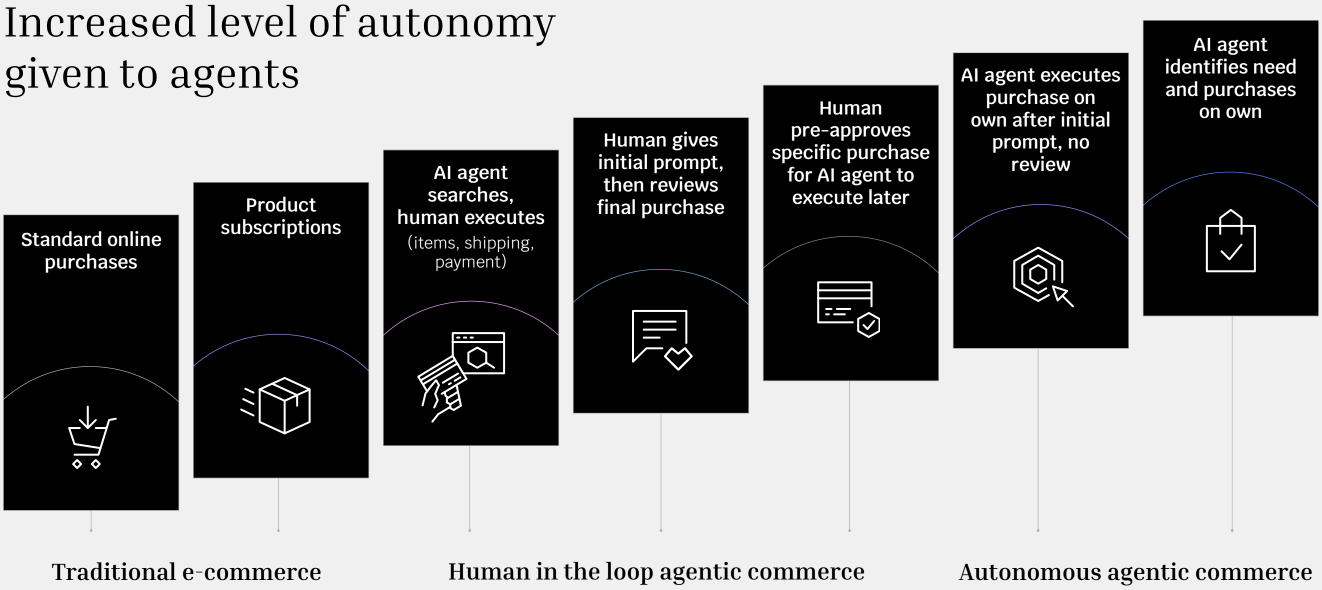

Agentic commerce differs from traditional e-commerce because it involves AI agents that can provide advice to and make decisions on behalf of customers, input payment credentials and other data on behalf of customers, and execute purchases when customers are not present. These capabilities have the potential to unlock benefits, but also bring risks spanning transparency, safety, privacy & data, choice, and interoperability. These risks may increase as greater autonomy is given to AI agents.

While the pace and scale of adoption is uncertain, agentic commerce has the potential to become a mainstream way in which consumers and merchants transact. Agentic commerce has already shifted from closed trials to live pilots and early scaling. These developments are outpacing formal governance, with industry standards, policies and consumer protections

[3] *Agentic commerce is a subset of agentic payments, which encompasses broader uses of AI agents to make payments, such as peer-to-peer payments and accounts payables and receivables. These broader uses are beyond the scope of this paper.*

yet to be fully defined or established. The focus to date has been on protocols, mostly led by payment networks and technology platforms, and concentrated on product discovery, checkout and payment authorization. In addition, some transactions are operating outside of formal protocols, increasing risks for consumers and merchants.

Trust will be critical to support adoption at scale. All market participants have a role to play in unlocking the opportunities of agentic commerce, while addressing outstanding gaps, risks and interoperability challenges. Banks have an important contribution to make, reflecting the expertise and trust we have developed over decades of enabling consumers and merchants to transact safely and securely.

Five principles to build trust and scale

We have developed five principles to deliver safe, effective and efficient agentic commerce, with consumer and merchant needs front and center. These principles are key to developing enduring trust and sustainable uptake in agentic commerce.

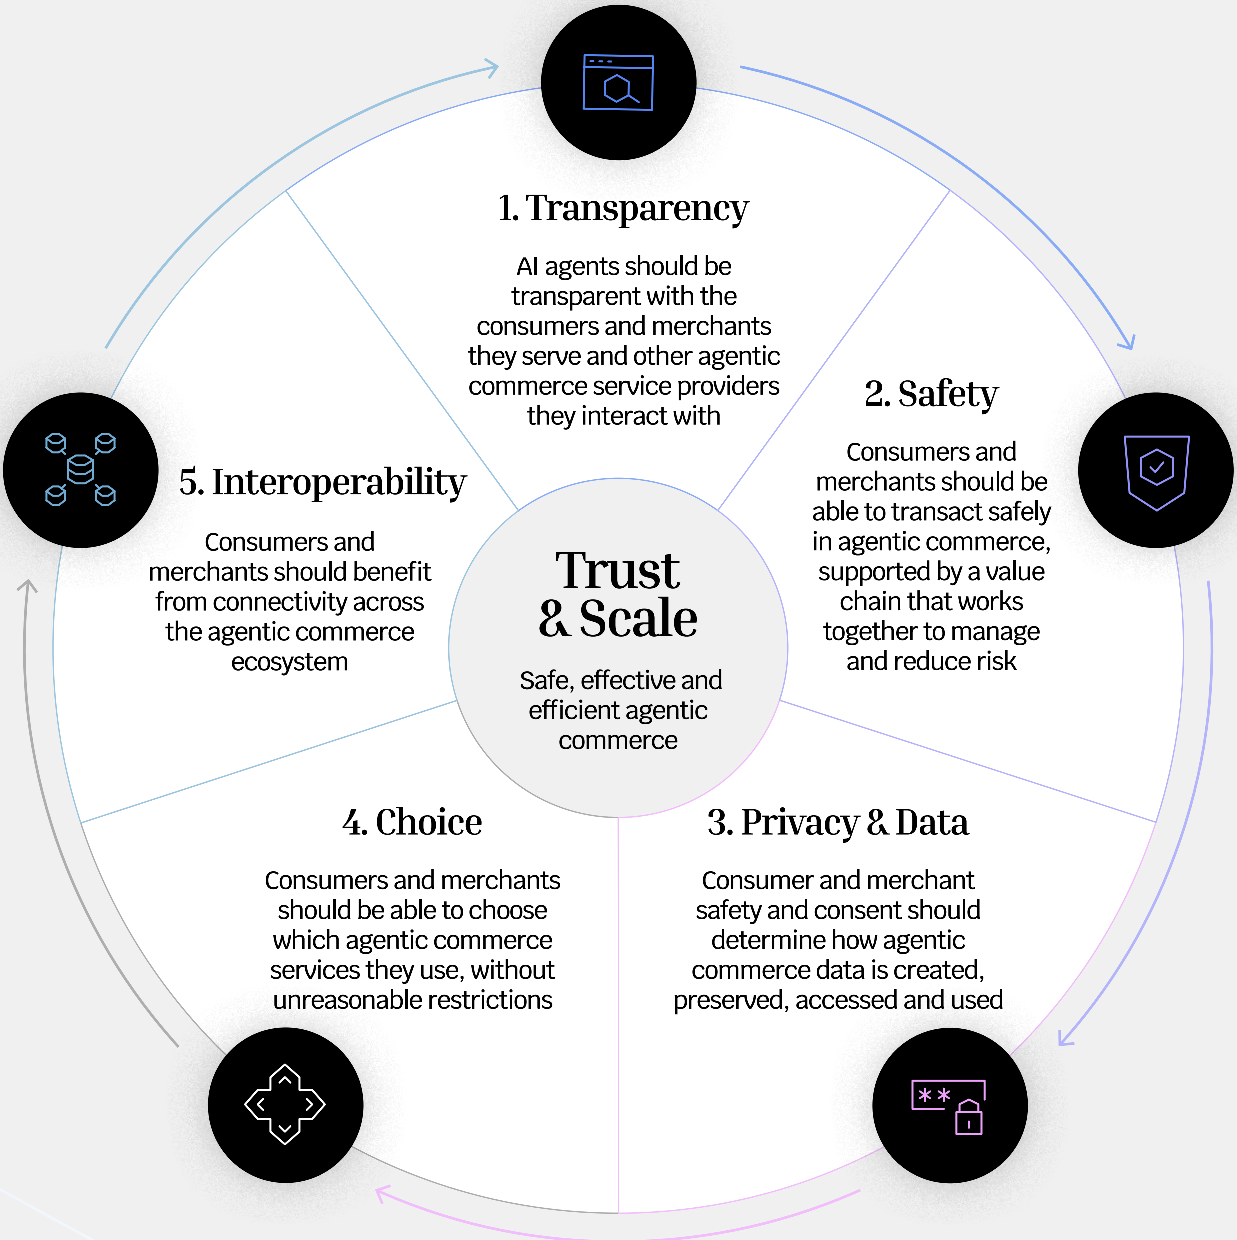

The principles are interdependent and build on each other to create a stronger, trusted ecosystem. Implementation of the principles requires sound governance, where stakeholders have opportunities to inform and influence ecosystem settings that may affect them. The principles focus on opportunities and risks that are specific to agentic commerce, while acknowledging broader and ongoing efforts in relation to the responsible development and use of AI. The principles also build on relevant laws, regulations and standards and so do not seek to duplicate them.

1. Transparency

AI agents should be transparent with the consumers and merchants they serve and other agentic commerce participants they interact with.

Agentic commerce introduces new market participants, process flows and incentives. Risks can occur if AI agents are not clearly identifiable to all parties involved in agentic commerce transactions, resulting in uncertainty about who parties are interacting with, who is accountable for the AI agent's actions, and whose interests the AI agent is representing.

Additionally, conflicts of interest may arise that require careful consideration. AI agents may prioritize products, payment methods or other services that deliver the best financial outcomes for their providers (e.g. higher commissions, lower token costs). This may come at the expense of options that may, in fact, provide better value for consumers and enable merchants to more openly compete. AI agent providers may also, at times, balance competing interests between consumers and merchants.

To support transparency in agentic commerce:

1.1 All parties involved in agentic commerce should be aware when an AI agent is helping to make or facilitate a transaction, and on whose behalf they are acting.

1.2 Consumers and merchants should have transparency over how AI agents prioritize options and make decisions, including any sponsored options, and how their data is collected and used.

2. Safety

Consumers and merchants should be able to transact safely in agentic commerce, supported by a value chain that works together to manage and reduce risk.

Agentic commerce introduces new safety risks, with potential for higher rates of scams, fraud and disputes. There may be mismatched expectations between consumers, AI agents and merchants on what product or service should be provided, or when and how, or who would be liable if an AI agent exceeds its authority. Some providers may engage in unsafe practices, including requesting consumer card details and entering them directly into websites, prioritizing payment methods with lower protections, and not complying with payment processing standards and payment scheme rules. Malicious actors may attempt new attack vectors for scams and fraud, including compromising or impersonating AI agents and merchants, and engaging in new forms of social engineering.

These increased risks make adoption of the principles all the more critical. When things go wrong, there is unclear and inefficient allocation of liability, and disputes processes do not involve all relevant parties across the value chain. In addition, service providers may lack access to the data they need to help reduce and manage safety risks (e.g. issuers and acquirers may not have real-time access to data on AI agent identity, merchant of record, consumer intent and purchase details).

To support safe agentic commerce:

2.1 Agentic commerce service providers should meet appropriate security and safety standards over the lifecycle of their operations, and unsafe practices should be discouraged and where possible prohibited.

2.2 Consumers should have appropriate control over the authority they grant to AI agents, including the ability to view and manage delegations.

2.3 Agentic commerce service providers should use secure and auditable methods to enter payment credentials and authorize purchase intents and payments – with any party holding a potential liability being able to require authentication.

2.4 Agentic commerce service providers should work together to prevent, detect, disrupt and respond to safety risks, including through timely intelligence and data sharing across the value chain, proactive education for consumers and merchants, appropriate warnings and interventions, recovery activity and effective dispute resolution.

2.5 Dispute resolution processes should involve all relevant parties to agentic commerce transactions and be fair and efficient, with liability outcomes reflecting where risks and errors are introduced.

3. Privacy & Data

Consumer and merchant safety and consent should determine how agentic commerce data is created, preserved, accessed and used.

Agentic commerce will generate new data, including conversational prompts, decision logs, intent mandates and purchase details. While the richness of such data provides opportunities to enhance customer experiences, there are also potential risks that require careful management.

A key risk is privacy for consumers and merchants, including in relation to how data is collected, secured, retained, shared and used. Risks include improper data collection (including by making improper inferences from data collected), improper use of data (including for profiling or other secondary purposes without adequate transparency or consent), poor data quality, data breaches, and over-retention of data. Industry best practice to address these risks is fundamental not only to customer safety and trust in agentic commerce, but to all uses of AI assistants. While these privacy risks require holistic governance, agentic commerce requires particular care given the richness of the data.

As outlined in the safety principle, some service providers may not have access to the data they need to reduce and manage safety risks. This includes relevant, accurate and auditable records sufficient to evidence consumer instructions, authentication and intent, transaction decisions and outcomes, and any warnings or interventions – noting such records would need to support scam and fraud investigations, recovery activity and dispute resolution.

Consumers and merchants may also face constraints on accessing their own data, including restrictions on permitting service providers to use their data for value added services. Merchants and consumers may also be prevented from sharing data with each other (e.g. merchants prevented from being disclosed as the merchant of record, consumers being prevented from sharing identity with a merchant to access loyalty benefits).

To support data governance in agentic commerce:

3.1 Agentic commerce service providers should meet appropriate privacy, data and cyber security standards over the lifecycle of their operations.

3.2 Agentic commerce service providers should maintain auditable records to evidence consumer instructions, authentication, intent, and transaction decisions and outcomes.

3.3 Agentic commerce service providers and merchants should only access the data they require to deliver their functions safely and effectively.

3.4 Consumer and merchant consent should determine any additional use or sharing of their data.

4. Choice

Consumers and merchants should be able to choose which agentic commerce services they use, without unreasonable restrictions.

Agentic commerce brings new participants into the commerce value chain and new market structures, with potential for platforms (e.g. devices, agents, marketplaces, digital wallets, payment networks) to become significant market participants. This brings potential risks to consumer and merchant choice. This may include access restrictions, such as preferencing payment methods and other services (e.g. loyalty providers) at checkout, or limiting AI agent access to platforms (e.g. software applications, browser extensions, and device functions). It may also include requirements to use a particular platform or platform-affiliated service (e.g. payment methods, digital wallets, payment service providers) or commercial structures that make it harder for consumers and merchants to access alternative services (e.g. higher fees to use alternative service providers, gatekeeping fees without reasonable commercial grounds).

To support access and choice in agentic commerce:

4.1 Consumers and merchants should be able to access and choose agentic commerce services without unreasonable restrictions or requirements, including any arbitrary restrictions or requirements on their service providers of choice.

4.2 Service providers should retain choice over which agentic commerce services they support based on safety, commercial sustainability and legal / regulatory grounds.

4.3 The identity and role of the merchants of record should be clear and consumers should have the option to share their identity with merchants if they choose.

5. Interoperability

Consumers and merchants should benefit from connectivity across the agentic commerce ecosystem.

Agentic commerce involves new market participants and services, many of which are still maturing and evolving. Fragmentation across different AI agents, payment methods or protocols could reduce network effects, increase integration costs, and create challenges with switching – or even undermine trust and confidence if there is variability in consumer and merchant protections. At the same time, processes to align on common approaches could delay or inhibit innovation and customer value. This could be most pronounced in relation to value added services and features that differentiate market offerings and are not required to meet baseline customer expectations. It will be important to strike a balance between consistency of approaches and pace of innovation.

To enable effective interoperability across agentic commerce:

5.1 Industry should focus on building safe and sustainable interoperability and connectivity for functions that are core to the above principles, while providing flexibility for innovation on value added services and features.

5.2 Over time, industry could consider expanding interoperability and connectivity to additional areas where common approaches deliver benefits for consumers and merchants.

We look forward to collaborating to bring the principles to life

The principles are intended to support a trusted and thriving agentic commerce ecosystem for all customers and market participants. Consumers can benefit from greater transparency, safety, choice and control over their data and transactions. Merchants can grow through trusted and interoperable channels that enable innovation, while preserving customer relationships and fair competition. AI agent providers, technology platforms, and payment schemes and service providers can apply the principles to design services that build trust and support sustainable adoption. Industry bodies and regulators can use the principles to inform discussions on governance, consumer protection, and market development.

We plan to elaborate on these principles and lay out frameworks in a subsequent paper, to solicit input across the ecosystem with the needs of consumers and merchants at the center. We look forward to engaging with customers, industry and all other stakeholders across the agentic commerce ecosystem to bring the principles to life.