Response to the HM Treasury Consultation on Modernising Payment Services Regulation

Submitted by Twinvocate Inc.

5 October 2026

About this submission

Twinvocate Inc. is a Toronto, Canada-based company building trust infrastructure for agentic commerce, including credentialing for the counterparties that software agents transact with on a consumer's behalf. We are responding to the consultation because the agentic payments it seeks to unlock in Chapter 3 are the payments our work addresses.

This response is signed by Twinvocate's founder and chief executive, who is a UK citizen and has worked in retail technology and payments for approximately thirty years, including as founder and chief executive of SelfPay by Digital Retail Apps and as a co-founder of Paytechs Canada, now FinTechs Canada. She completed a Master of Laws at the University of Toronto Faculty of Law in 2026, concentrating in innovation and technology, convocation in October 2026.

We respond to questions 3, 15, 17, 18, 20, 25 and 36. Our substantive answer is at question 15 and the others follow from it.

Summary

The consultation asks at question 15 whether the provisions on consent, authentication and liability for unauthorised transactions need updating for agentic payments. Under the heading of sector risks it asks separately what new risks have emerged and how the framework should address them. We think the two questions meet in one place that the consultation does not name.

When a user delegates a payment to an agent under a purpose and a limit, the user delegates the choice of counterparty along with it. Every protection the Payment Services Regulations currently supply against a payment reaching the wrong party assumes the payer chose that party. Consent under the Regulations is consent to a transaction with a payee the payer identified. Strong customer authentication proves the person approving is the account holder. Reimbursement for authorised push payment fraud turns on whether the payer was deceived into directing a payment. None of these engages when the payee was selected by software after the user gave her instruction.

The government's stated approach is same risk, same regulatory outcome. We support that approach, and we think its honest application to agentic payments leads to a different conclusion from the one usually drawn from it. Delegating selection is not the same risk as delegating execution. It is a new risk, and the framework has no provision that addresses it. Our answer to question 15 sets out what should be added.

Q3 and Q25: Key definitions, and whether the payment initiation service definition is fit for purpose

The current definitions were written for a world in which the party initiating a payment is either the payer or a regulated firm acting on the payer's instruction to a payee the payer named. A

consumer's AI agent fits neither description cleanly. Depending on how it is built and who operates it, the same agent could be read as the user's own instrument, as a payment initiation service provider, or as something outside the perimeter altogether.

We recommend the government settle this by decision rather than by accident. The reformed framework should say what an agent that initiates a payment on a user's behalf is for the purposes of the Regulations, and in particular whether an agent that selects the payee within a mandate is a distinct class of initiator from one that executes a payment the user specified. The difference matters for every downstream provision. Consent, authentication and liability can all be calibrated once the status is clear, and none of them can be calibrated while it is not.

The question is already before the government. In written answers of 4 September 2026 to questions from Bob Blackman MP on whether autonomous software agents should be recognised as a distinct class of payment initiator, and on liability and consumer redress where an agent acts outside the scope of its mandate (UINs 25983 and 25984), the Treasury identified this consultation as the place where those questions are being considered. Both questions concern the agent's authority. Neither reaches the case where an agent acts within its mandate and pays the wrong party, which we address under question 15.

Two definitional additions would help. First, a definition of a payment mandate: an authority granted by a payment service user under a purpose, a ceiling, a duration and any other constraint, under which one or more payment transactions may later be initiated without the user's involvement at the moment of initiation. Second, a definition of an agent-initiated payment transaction as one initiated under such a mandate by software rather than by the user. The second is what makes attribution possible; the first is what gives attribution something to refer to.

Q15: How the Regulations need to adapt for agentic payments

This is our principal submission.

The coupling that delegation breaks

In a payment initiated by a person, two decisions are made in a single act: who to pay, and whether to authorise. The payer selects the counterparty and approves the transfer in the same moment. The Regulations are built on top of that coincidence. Regulation 67 treats a transaction as authorised where the payer has given consent to its execution, and the consent is to a transaction with a known payee. Strong customer authentication under regulation 100 establishes that the person approving is the account holder. The reimbursement regime for authorised push payment fraud asks whether the payer was deceived into directing a payment to a particular recipient. All three answer questions about a human who was present at the point of choice.

Delegation separates those two decisions in time and in agency. The user grants authority at one moment, under a purpose and a constraint. The agent selects a counterparty at a later moment, on information the user has not seen. Consider a user who instructs an agent to buy a piece of art at auction and not to spend more than two hundred pounds. Her instruction contains a subject and a ceiling. It does not contain a payee. If the listing is fraudulent, or the agent is induced to complete against a substituted listing, the limit holds, authentication of the mandate holds, and Confirmation

of Payee has nothing to compare, because she never named the seller. The user is not present to notice.

The point is not that this makes agentic payments unsafe. It is that the specific controls the framework depends on are inoperative in the new journey, and the consultation does not identify a replacement.

Consent

Consent as the Regulations conceive it should be extended to cover consent to a mandate, given once, under which later transactions are initiated within its terms. The Regulations should state what a mandate must contain to be valid as consent, at minimum a purpose, a ceiling and a duration, and should state that a transaction within those terms is authorised and a transaction outside them is not. That gives the unauthorised transaction provisions something to bite on. Today, a transaction an agent initiates is either treated as authorised because the user consented to the mandate, in which case the user bears every error the agent makes, or as unauthorised because the user never consented to that transaction, in which case the provider bears every error. Neither is right, and the answer should be in the rules rather than left to be litigated.

Authentication

The government has committed to revoking the strong customer authentication provisions so that the FCA can set outcomes-based rules. We support that and would ask that the outcomes be specified with agentic payments in view. Authentication today answers one question: is the person approving this payment the account holder. In an agent-initiated payment there are two further questions that authentication of the user does not answer: was this agent authorised by this user and was this transaction within what the agent was authorised to do. The FCA's rules should require that both be established, and that the answer be carried with the transaction in a form a receiving institution can check.

This does not require the contents of the mandate to be disclosed to every participant in the chain. It requires that a transaction carry a reference to the authority relied on, resolvable by a party with a legitimate interest. Attributing a payment as machine-initiated, without more, tells a receiving institution that software acted and nothing about what the software was permitted to do.

Liability

One case is unallocated under the current framework and will remain unallocated unless the government decides it: the agent stayed inside its mandate and paid the wrong party. This is not obviously authorised push payment fraud, because the user did not authorise a payment to a specific payee. It is not obviously an unauthorised transaction, because the mandate was honoured. It is not obviously a firm failure, because no rule required the firm to verify a payee the user did not name. Until this case is allocated, no participant has a commercial incentive to invest in preventing it, and the consumer will bear it by default.

The public record so far addresses the neighbouring case rather than this one. In his letter of 17 July 2026 to the Chair of the Treasury Committee on the legal standing of agentic AI tools, the Governor of the Bank of England framed the question as where liability sits if an agent acts improperly or exceeds the remit given to it by its principal, and relied on the Financial Markets Law Committee's

2025 paper, Private Law Issues in AI, for the conclusion that under English law an AI system remains a tool rather than an agent with legal personality, so that responsibility sits somewhere in the chain of human and corporate actors who built, trained, integrated and deployed it. We agree with that analysis as far as it goes. It resolves the case where the agent exceeded its authority. It does not reach the case where the agent stayed within its authority and the loss arose from who was paid, because in that case no actor in the chain did anything wrong under the current rules. If the agent is legally a tool, the mandate under which the tool operates is the object the Regulations need to define, which is why we make the definitional recommendations under questions 3 and 25.

We recommend the government allocate it, and allocate it in a way that creates the incentive. The provider that executes an agent-initiated payment should bear the loss where a counterparty verification capability was available to it and was not used, and should not bear it where it verified the payee against a recognised credential and the credential was sound. That rule does two things. It makes the verification capability worth building, and it makes the provider who builds it safe.

Four capabilities the rules should require or enable

1. Attribution that carries the authority, not only the fact of automation. The rules should require that an agent-initiated transaction be identified as such and carry a reference to the mandate relied on, resolvable by a party with a legitimate interest without disclosing the mandate's contents to the chain. The extensibility of ISO 20022 makes this feasible if it is specified now.

2. Counterparty verification available to software acting for a consumer. A credential that only institutions can check does not help a consumer whose agent is the party doing the choosing. Verification must be available at machine speed and without a commercial relationship having to be negotiated first. We return to where this capability should sit under questions 20 and 36.

3. Verification effort proportionate to exposure. Verification is not free, in cost or in latency. A recurring payment to a counterparty with an established history does not present the same risk as a first payment to a counterparty selected by software from a set the user has not seen, and the rules should not require the same of both. Designing for a single fixed level of verification will either impose unacceptable friction on ordinary payments or leave the hard case unprotected.

4. Recourse that can engage before finality. Settlement finality is essential and we do not propose qualifying it. But finality combined with delegated selection means that a payment made under a valid mandate to the wrong counterparty has no reversal path, and the user was not present to prevent it. The 72-hour delay the government has already introduced for suspected authorised push payment fraud is the right instrument. The rules should say on what signal an agent-initiated payment can be held under it.

Evidence that these capabilities need to be designed in

The UK AI Security Institute's published incident report on unsanctioned agent behaviour during cyber testing, covering evaluations run between 25 and 28 July 2026, records nineteen unsanctioned online actions across ten of one hundred and twenty two runs, seventeen of them from a single model. In one case an agent created multiple false identities and used them to press a real open

source maintainer into approving malicious code. The Institute notified the affected platform, supplied a complete inventory of the accounts, repositories, issue threads, pull requests and API keys involved, worked with the platform to remove them, and briefed the model developers and the US Center for AI Standards and Innovation on 3 August 2026.

Two things follow. First, a verification step is only as strong as the ability to establish that the parties to it are who they claim to be and are independent of one another. An agent that can manufacture identities and endorsement on a live platform can manufacture them in a payment journey unless the authority it relies on is resolvable by the receiving side. Second, the Institute's notification of the affected platform was discretionary. There was no channel through which the platform was entitled to be told. A recourse mechanism that depends on a research body choosing to make a phone call is not one the framework can rely on.

A test for the whole

We suggest the government assess the reformed framework against one question: does a payment made under a mandate carry protection equivalent to the same payment made by the user directly. Where it does not, the difference should be a deliberate and stated policy decision rather than a consequence of the journey having been added to rules written around a payer who is present.

Q17: Financial inclusion

The consultation rightly frames delegated payment arrangements as supporting inclusion, and delegation to an agent will be marketed the same way. Delegated and assisted payment arrangements are used disproportionately by disabled people and by older people. If the protections applying to a delegated payment are weaker than those applying to a payment a person makes for themselves, the shortfall falls hardest on those groups, and it does so precisely because the feature was designed with them in mind. The equivalence test we propose under question 15 is the inclusion safeguard.

Q18 and Q20: Emerging risks and targeted amendments

The risk we would add to the government's list is counterparty fraud against a payer who is not present. Fraud today reaches a consumer through deception of the consumer. Fraud against an agent reaches the consumer through deception of, or substitution before, the agent, and the consumer learns of it after settlement. The current regulations do not address it because none of their provisions engages when the payee was not chosen by the payer.

Three targeted amendments would address it proportionately.

- Define the mandate and the agent-initiated transaction, as set out under questions 3 and 25, so that the consent and liability provisions have a defined object.
- Require that any register or directory the framework relies on for verifying a payee or an agent be published in machine-readable form with a programmatic interface, so that software acting for a consumer can confirm status directly rather than depending on a human reading a web page. A framework that assumes verification is performed by institutions rather than by consumer-side software will foreclose the consumer-agent case

at the design stage, and it will do so quietly, because the omission looks like an implementation detail.

- Allocate the in-mandate, wrong-payee loss as proposed under question 15, so that a participant has a reason to build the control.

We would place outside regulation: agent reputation scoring, analytics products and commercial suitability judgements. These are competitive services and placing them in rules would concentrate a judgement that ought to remain contestable. Identity and suitability are different assertions; a credential that answers who a party is does not answer whether that party should be transacted with, and the framework should be explicit about which assertion each credential carries.

Q36: Additional powers for the Future Entity

The consultation notes that the Future Entity may be required to provide connected services, and gives a central directory to support the practical use of APIs as its example. We recommend the FCA be given power to require that directory, and any equivalent directory of credentialed payees or authorised agents, to be designed from the outset for use by software acting for a consumer: machine-readable, programmatically queryable, and available without a prior commercial relationship. Counterparty verification for agent-initiated payments is a shared utility for the same reason a common alias directory is. If each agent or scheme builds its own, the protection a user receives will depend on which agent they happen to use, and that fragmentation disadvantages exactly the smaller participants the open banking framework is meant to keep competitive.

Submitted by

Wendy MacKinnon
Founder and Chief Executive Officer, Twinvocate Inc.

UK citizen

100 King Street West, Suite 6200, 1 First Canadian Place, Toronto ON M5X 1B8, Canada

wendy@twinvocate.ai