Legal 500
Country Comparative Guides
Hot Topic | Fintech

Blockchain-based AI-Agents under Swiss Financial Market Law and Data Protection Law

Contributor

Zürcher
Rechtsanwälte AG

Sebastian Hepp

Attorney-at-law, Partner | sebastian.hepp@zurich-law.com

Eric Neuenschwander

attorney-at-law, LL.M. (Columbia), Partner | eric.neuenschwander@zurich-law.com

Blockchain-based AI-Agents under Swiss Financial Market Law and Data Protection Law

1 Introduction

The tech industry is currently undergoing a profound technological transformation driven by the rapid advancement of artificial intelligence ("AI"). The increasing integration of AI into business activities has raised complex legal and supervisory questions.

This evolution has been further accelerated by the growing accessibility of computational resources and the widespread deployment of generative AI applications. As a result, AI is no longer confined to back-office functions or analytical support roles; it is increasingly being deployed in areas that are directly relevant to financial market regulation, including investment decision-making and asset management. At the same time, blockchain technology has introduced a new paradigm for the execution of financial transactions, enabling decentralized, automated, and tamper-resistant processes.

The convergence of these two technologies – AI and blockchain – has led to the emergence of so-called "blockchain-based AI-Agents". These systems combine AI-driven analytical and decision-making capabilities with blockchain-based execution mechanisms, allowing them to autonomously generate and implement transactions. In practice, such agents may analyze market data, identify trading opportunities, and initiate transactions on distributed ledgers without direct human intervention.

Against this backdrop, the key legal question is whether, and under what conditions, the operators of such AI systems fall within the scope of existing Swiss regulations. Rather than introducing entirely new regulatory frameworks for AI, Swiss law adheres to a technology-neutral approach, according to which the decisive factor is the nature of the activity performed rather than the technology used. This principle is often summarized as "same business, same risks, same rules."

Accordingly, the legal assessment of blockchain-based AI-Agents hinges on a functional analysis. The central issue is not whether an AI system is involved, but whether the activities performed through that system correspond to regulated services.

2 Artificial Intelligence and AI-Agents

A precise legal definition of artificial intelligence remains elusive, reflecting the diversity of technologies and methodologies encompassed by the term. AI may include systems based on machine learning, pattern recognition, expert systems, or autonomous planning and decision-making. Despite these differences, most AI systems share a common feature: they process large volumes of data to identify patterns, derive insights, and generate outputs that can be used to support or replace human decision-making.

In economic terms, AI systems are particularly valuable in environments characterized by high complexity, large datasets, and the need for rapid decision-making. Financial markets exemplify such an environment.

Here, the ability to process information and react to market developments within fractions of a second can confer significant competitive advantages. Traditional decision-making processes involving multiple human actors are often too slow and resource-intensive to meet these demands.

AI-Agents represent a specific form of AI application. They are typically understood as software entities that act on behalf of a user, perform tasks autonomously, and interact with their digital environment. Unlike simple rule-based systems, AI-Agents may exhibit a degree of autonomy that allows them to adapt their behaviour, refine their strategies, and even adjust their objectives in response to changing circumstances.

This autonomy, however, also introduces legal challenges. From a regulatory perspective, the key issue is whether the actions of an AI-Agent can be attributed to its operator and whether those actions constitute regulated activities. While AI-Agents may appear to act independently, they are ultimately deployed, configured, and controlled by human actors. Consequently, the legal analysis focuses on the role of the operator rather than the technical autonomy of the system.

From a data protection perspective, the question arises whether existing legal concepts are sufficiently robust to capture the risks created by AI and, in particular, AI-Agents. The following analysis suggests that Swiss data protection law is, in principle, capable of doing so. The challenge lies less in identifying applicable rules than in applying them to systems whose operation is data-driven, dynamic and not always readily transparent. The discussion below therefore focuses on the applicability of the Swiss Data Protection Act to AI-Agents, the main doctrinal and practical friction points, and the implications for compliance in a rapidly evolving technological environment.

3 Structure and functioning of blockchain-based AI-Agents

Blockchain-based AI-Agents are characterized by a dual-layer architecture that separates decision-making from execution. This distinction is essential for understanding both their technical operation and their regulatory classification.

3.1 Off-Chain Decision-Making

The AI component of the agent operates off-chain, typically on centralized infrastructure such as cloud servers. Current blockchain systems are not designed to handle the computational demands of advanced AI models, as their processing capacity is limited and associated costs would be prohibitive. As a result, the “intelligence” of the system – its ability to analyze data and generate decisions – remains outside the blockchain environment.

Within this off-chain layer, the AI system processes a wide range of data inputs. These may include token prices, transaction histories, liquidity metrics, yields from decentralized finance (“DeFi”) protocols, and even qualitative data such as social media sentiment. Based on this analysis, the system generates outputs that take the form of proposed blockchain transactions. Examples include instructions to purchase or sell specific tokens, transfer assets between wallets, or allocate funds to particular DeFi protocols.

Importantly, these outputs do not constitute executed transactions. Rather, they are equivalent to

unsigned transaction proposals that must be validated and authorized before they can be executed on the blockchain.

3.2 On-Chain Execution

The execution of transactions takes place on the blockchain through wallets or smart contracts. In order for a transaction to be executed, it must be signed using a private key associated with a blockchain address. The private key therefore represents the ultimate source of control over the crypto assets held in a wallet.

Wallets function as digital interfaces for managing crypto assets. They store both public and private keys and enable users to initiate transactions. Depending on the configuration, wallets may be self-custodial – where the user retains exclusive control over the private keys – or custodial, where a third-party manages the keys on behalf of the user.

Smart contracts, by contrast, are self-executing programs deployed on a blockchain. They automatically perform predefined actions when specified conditions are met and can, in some cases, function as wallets themselves. By combining AI-generated outputs with smart contract execution, blockchain-based AI-Agents can automate complex operations.

4 Swiss Financial Market Regulation

Based on existing practice of the Swiss Financial Market Supervisory Authority ("FINMA") regarding the attribution of DeFi applications, it can be assumed that the operation of AI systems is attributable to the person who operates the system in a centralized manner (e.g. the licensor or the service provider vis-à-vis the user). The regulatory qualification of AI-Agents under Swiss financial market law is examined in the following analysis, with a focus on the most relevant regulatory frameworks in Switzerland.

4.1 Anti-Money Laundering Act ("AMLA")

4.1.1 General Framework

The Swiss AMLA applies to financial intermediaries who professionally accept, hold, or transfer third-party assets. The concept of financial intermediation encompasses a wide range of activities, including payment services, custody of crypto assets, and certain forms of investment activity. A distinction must be made between financial intermediaries in the narrow sense as defined in art. 2 para. 2 AMLA (such as banks, insurance institutions, licensed payment systems etc.) and financial intermediaries in the broader sense who carry out a financial intermediary activity in a professional manner as defined in art. 2 para. 3 AMLA. The latter encompass providers who:

- carry out credit transactions (in particular in relation to consumer loans or mortgages, factoring, commercial financing or financial leasing, let. a);
- provide services related to payment transactions, in particular by carrying out electronic transfers on behalf of other persons, or who issue or manage means of payment such as credit cards and travellers' cheques (let. b);
- trade for their own account or for the account of others in banknotes and coins, money market

instruments, foreign exchange, precious metals, commodities and securities (stocks and shares and value rights) as well as their derivatives (let. c);
- make investments as investment advisers (let. f);
- hold securities in custody or manage securities (let. g).

Generally, all these activities require some form of power of disposal/control over the assets involved. The analysis of blockchain-based AI-Agents therefore focuses on whether the operator of the AI system has control or influence the transfer or custody of crypto assets in a manner that falls within the scope of AMLA.

4.1.2 Applicability of AMLA on AI-Agents

Where the user retains full control over private keys, the operator does not acquire any disposal power over third-party crypto assets. The AI system merely generates transaction proposals, and the execution of those transactions depends entirely on the user's actions.

As a result, the operator does not engage in payment services, custody, or asset transfer within the meaning of AMLA. Nor does the system provide "assistance" in transferring crypto assets in a regulatory sense, as its role is limited to preparatory activities rather than actual transaction execution. Accordingly, operators in this configuration are not considered financial intermediaries under the Swiss AMLA.

Where the AI system has access to private keys, the operator of the AI system effectively acquires control over third-party assets. This enables the system to independently initiate and execute transactions, thereby fulfilling the criteria for financial intermediation according to the AMLA. In such cases, the operator engages in activities such as asset transfer and custody, both of which fall within the scope of AMLA. Consequently, the operator must comply with anti-money laundering obligations, including affiliation with a self-regulatory organization and adherence to due diligence requirements and KYC-rules.

4.2 Financial Services Act ("**FinSA**")

4.2.1 General Framework

The FinSA primarily establishes the regulatory framework for the provision of financial services. In particular, the act governs the provision of financial services at the point of sale through rules of conduct (art. 6–20 FinSA) and certain organizational requirements (art. 21–27 FinSA). In addition, it promotes transparency regarding financial instruments at the point of production by imposing information obligations on the providers and creators of financial instruments, notably through the prospectus requirement (art. 35 FinSA) and the obligation to prepare a key information document (art. 58 FinSA).

According to art. 2 para. 1 FinSA, the rules apply to financial service providers and to producers and providers of financial instruments, irrespective of their legal form. FinSA itself in art. 3 let. d defines financial service providers as persons who provide financial services on a commercial basis in Switzerland or for clients in Switzerland. Against this background, it follows that the decisive factor for the applicability of the FinSA in this context is, whether a financial service within the meaning of art. 3 let. c no. 1 – 5 FinSA is provided. The provision defines the following activities as financial services, whereby these must always be performed "*for clients*":

- acquisition or disposal of financial instruments (no. 1);
- receipt and transmission of orders in relation to financial instruments (no. 2);
- administration of financial instruments (portfolio management; no. 3);
- provision of personal recommendations on transactions with financial instruments (investment advice; no. 4);
- granting of loans to finance transactions with financial instruments (no. 5).

The classification depends on the functional nature of the activity rather than the technology used.

A critical prerequisite for FinSA applicability is that the service relates to financial instruments. While not all crypto assets qualify as such, tokenized securities and certain derivatives may fall within the definition (so-called asset tokens).

4.2.2 Applicability of FinSA on AI-Agents

AI systems may generate personal recommendations based on user-specific data. Where such recommendations are tailored to the user's financial situation and relate to financial instruments, they constitute investment advice according to FinSA. Although the system does not execute transactions, the provision of personal recommendations qualifies as a financial service according to art. 3 let. c no. 4 FinSA. Operators must therefore comply with FinSA's conduct and organizational requirements.

Where the AI system has control over private keys and can independently execute transactions, the activity could qualify as receipt and transmission of orders in relation to financial instruments (no. 2) and/or portfolio management (no. 3) according to FinSA if financial instruments are involved. In the latter case, the operator of the AI system effectively manages client assets on a discretionary basis, making investment decisions and implementing them without requiring further user approval.

4.3 Financial Institutions Act ("**FinIA**")

4.3.1 General Framework

FinIA governs licensed financial institutions, including the following:

- Portfolio manager (art. 17 para. 1 FinIA)
- Trustees (art. 17 para. 2 FinIA)
- Managers of collective assets (art. 24 FinIA)
- Fund management companies (art. 32 FinIA)
- Securities firms (art. 41 FinIA)

4.3.2 Applicability of FinIA on AI-Agents

The key criterion for determining whether AI-Agents may qualify as financial institutions, particularly as portfolio managers, is the existence of discretionary power over client assets combined with the provision of financial services within the meaning of FinSA.

If the AI-Agent does not have control over the private keys of the user and therefore does not have power

of disposal of assets of the users, such AI systems do not fall under the FinIA and no licensing requirements arise under FinIA.

If the AI-Agent has access to the private keys and thereby exercises power of disposal over client assets, and performs portfolio management activities, this triggers licensing requirements under FinIA, and the operator of the AI system becomes subject to prudential supervision by FINMA.

5 Data Protection Regulation

5.1 Data Protection Act (“DPA”)

5.1.1 General Framework

The Swiss DPA constitutes the principal framework governing the processing of personal data in Switzerland. It is not an AI-specific statute. Rather, it is a technology-neutral regime intended to protect the personality and fundamental rights of natural persons whose personal data is processed (art. 1 DPA). That technology-neutrality is important in the context of AI-Agents: as outlined, Swiss law has traditionally followed a technology-neutral approach. Accordingly, it does not require a bespoke legal category for AI in order to regulate it; instead, existing legal frameworks apply regardless of whether data processing is carried out through conventional software or AI-based systems. The decisive question is whether, and in what manner, the (AI) system processes personal data. Swiss scholarship therefore generally favors applying and, where necessary, selectively refining existing law rather than abandoning the DPA's technology-neutral approach in favor of a wholly separate AI privacy regime.

From a doctrinal perspective, the starting point is conventional: if an AI-Agent processes personal data, the general processing principles set out in art. 6 DPA apply. These include lawfulness (art. 6 para. 1 DPA), good faith and proportionality (art. 6 para. 2 DPA), purpose limitation and recognizability of purpose (art. 6 para. 3 DPA), data minimization and storage limitation (art. 6 para. 4 DPA), and data accuracy (art. 6 para. 5 DPA). In the AI context, those principles are not abstract statements of principle. They operate as concrete design constraints for training, deployment, optimization and re-use of data. In the AI context, these principles are often in structural tension with the logic and operational design of AI applications.

The DPA is complemented by further obligations that are particularly relevant for AI systems. Art. 7 DPA requires privacy by design and privacy by default. Art. 8 DPA requires appropriate data security measures, calibrated to the relevant risk. Art. 19 DPA imposes information duties where personal data is collected. Art. 21 DPA addresses automated individual decisions. Art. 22 DPA provides for data protection impact assessments (“DPIAs”) where processing may result in a high risk to the personality or fundamental rights of data subjects. Taken together, these provisions make clear that the DPA already contains the building blocks needed to assess AI-Agents in practice.

These considerations apply equally to blockchain-based AI-Agents. The use of blockchain technology does not displace the application of the DPA where personal data is processed. Rather, it may intensify certain data protection issues, in particular where (transaction) data is stored in a decentralized and immutable environment, where responsibilities are distributed across multiple actors (e.g. nodes), or where on-chain and off-chain processing interact. The tensions between data protection law and blockchain technology would exceed the scope of this analysis and are therefore not the subject of this

article.

5.1.2 Applicability of DPA on AI-Agents

AI-Agents as data-processing environments

AI-Agents should not be treated as legally exceptional systems falling outside the established logic of data protection law. Rather, they should be understood as intensified data-processing environments. They collect, combine and analyze information across different contexts, generate patterns and correlations, and may materially influence outcomes relating to individuals. Where personal data is involved, the DPA therefore applies across the full lifecycle of the system, including design, training, deployment, feedback and ongoing optimization.

The DPA becomes particularly significant where AI-Agents analyze individuals and shape outcomes affecting them. The statutory concept of profiling in art. 5 let. f DPA is sufficiently broad to capture many AI-driven assessments of behavior, preferences, reliability or likely future conduct. As a result, a considerable number of practically relevant AI applications involving personal data are likely to fall within the scope of profiling.

Where an AI-Agent goes further and materially shapes or determines an outcome affecting an individual, art. 21 DPA on automated individual decisions must also be considered. This does not amount to a general prohibition. It does, however, mean that AI-Agents used in contexts such as creditworthiness, fraud detection, HR screening, client selection or personalized investment handling require closer analysis, particularly where legal effects or comparably significant practical consequences may arise.

5.1.2.1 Challenges in the interaction between AI-Agents and data protection

Purpose limitation and secondary use

One of the clearest pressure points lies in purpose limitation. AI-Agents are attractive because they can generate additional value from existing datasets. That technical capability does not, however, displace the legal requirement that personal data be processed for a specific purpose that is recognizable to the data subject.

The tension becomes particularly acute where customer or operational data are repurposed for model training, feedback loops or product improvement. Under Swiss law, such further use may be permissible, but only where it is appropriately framed, sufficiently transparent and supported by a defensible legal and contractual structure. The issue is especially sensitive where service providers seek to improve their own models using customer data. In practice, the more important legal question is often not whether the AI-Agent functions technically, but whether the surrounding data flows have been properly characterized from a data protection perspective.

Data minimization versus optimization

AI-Agents also sit uneasily with the principles of proportionality and data minimization. From a technical perspective, larger and more diverse datasets will often improve system performance and outcome. From a legal perspective, however, the fact that more data may be useful does not mean that its collection or re-

use is permissible.

This tension is most visible during training and feedback, where there is a natural tendency of providers to collect and retain data on the assumption that it may become useful later. The DPA does not endorse such an open-ended logic of optimization. For operators of AI systems, this means that personal data should be processed only where it is needed for a specific function of the system, and not merely because it may prove useful for future model improvement.

Transparency, information duties and access rights

Transparency is another central point: Data subjects must be informed about the collection of personal data, and more broadly the essential features of the processing must remain recognizable. In the context of AI-Agents, this requirement is complicated by the fact that AI-supported outputs are often difficult to understand from the outside and may rely on probabilistic rather than rule-based logic.

This does not mean that the law requires full technical explainability in every case. A more appropriate standard is one of meaningful intelligibility. Individuals should be able to understand that an AI-supported process is being used, what categories of data are relevant, what the system does in functional terms, and how its output affects the relevant decision chain. This is also relevant in the context of access rights, particularly where automated individual decisions are involved and some explanation of the underlying decision-making process may be required.

Data Subject Rights in Practice

AI-Agents also sharpen the practical significance of data subject rights under the DPA. Access requests may be more difficult to implement where the underlying system operates as a black box and where the output results from probabilistic correlations rather than transparent rule-based logic. That technical complexity does not, however, relieve controllers of their legal obligations.

Instead, it requires that information be provided in a form that is legally meaningful and functionally intelligible, even where full technical explainability is not possible. The same applies to requests for rectification or deletion. Where AI-Agents rely on continuously updated data, feedback loops or derived profiles, organizations must ensure that data subject rights remain effective in practice and are not undermined by the technical architecture of the system.

5.1.2.2 Possible solutions for addressing these challenges

Data protection through technology, data security and DPIAs

AI systems involve new technologies, large-scale data processing, inference-based outputs and decision-making processes that may not be readily transparent. As a result, legal compliance often depends on early design choices rather than on ex post disclosures. Relevant safeguards may include documentation, auditability, data mapping, pseudonymization, anonymization, human oversight, technical segregation of processing environments and access restrictions tailored to the risk profile of the system – also known as “privacy by design” and “privacy by default”.

Data security deserves separate emphasis. In AI environments, privacy risk may not end with the training

dataset. If personal data can later be extracted from a model, or if a model reproduces protected information in a meaningful way, assumptions about anonymization and non-personal outputs may no longer hold. In AI systems, data protection analysis must therefore extend beyond data collection and include the model and inference environment as such.

Anonymization, Pseudonymization and Model Leakage

Finally, anonymization and pseudonymization remain important risk-mitigation tools, but they should not be overstated. Such measures may significantly reduce privacy risk along the AI lifecycle. Whether they do so effectively depend on the design of the specific system and the realistic possibility of re-identification. From a legal perspective, only truly anonymized data falls outside the scope of the DPA. Pseudonymized data, by contrast, remains personal data as long as re-identification remains possible, whether directly or indirectly.

It is therefore not enough simply to label data or models as "anonymous". AI training does not itself amount to anonymization, and a model may remain privacy-relevant if personal data can be extracted from it or if it reproduces elements of the training data in a meaningful way. In the context of AI-Agents, privacy risk may therefore persist not only in the input data, but also in the model and inference environment.

6 Conclusion and Outlook

The analysis demonstrates that for the regulation of AI-Agents under Swiss financial market law the decisive factor is the allocation of control over client assets. Where the AI operator does not have access to private keys, regulatory exposure is generally limited to FinSA obligations, if the investment advice relates to financial instruments. By contrast, where the operator acquires control over private keys of third-parties and therefore has power of disposal over financial instruments or other crypto assets, the full spectrum of financial regulation may apply.

As these technologies continue to evolve, the development of regulatory practice in Switzerland – particularly by FINMA – will play a crucial role in providing clarity. Hybrid models combining elements of automation, decentralization, and traditional financial intermediation are likely to challenge existing classifications and may require further guidance. For market participants, the key takeaway is clear: the technical design of AI systems, particularly the allocation of control and execution authority, will determine their regulatory classification. Careful structuring is therefore essential to ensure compliance while preserving the benefits of technological innovation.

From a Swiss Data Protection Law perspective, AI-Agents do not operate outside the existing data protection framework; rather, they test its limits in a more intensive way. The DPA already provides the key legal tools – in particular the principles of purpose limitation, proportionality, transparency, privacy by design, data security and risk assessment – to assess the use of AI-Agents where personal data is involved. The real challenge therefore lies less in the absence of rules than in their operationalization across the full lifecycle of AI systems. Looking ahead, Swiss law is likely to continue relying on this technology-neutral approach, while expectations regarding governance, documentation, explainability and cross-border data handling will become more demanding in practice. For businesses deploying AI-Agents, compliance will therefore increasingly depend not only on technical performance, but on the ability to

embed data protection into system design from the outset.

Contributors

Sebastian Hepp
Attorney-at-law, Partner
sebastian.hepp@zurich-law.com

Eric Neuenschwander
attorney-at-law, LL.M. (Columbia), Partner
eric.neuenschwander@zurich-law.com