



Agentic Commerce – how can legally binding contracts be created?

Part 1 of 2

Two-part expert opinion on the interplay of legal, technical and business challenges. Supplemented by a glossary of legal-technical terms as well as the assessment of which legal system could pragmatically close the legal loopholes for agentic B2B commerce.

Author: Ass. iur. Oliver M. Merx

As of: October 2026

Part 1: Version 1.3

Language: AI-assisted translation of the original German text

Executive Summary

- **AI agents have so far acted without legal validity:** The action of an AI agent is not bindingly attributed to any recognized legal entity within the framework of agentic commerce. There is no constitutive legal norm that allows the actions of an AI agent to be considered a declaration of intent by the legal entity.
- **The classic figures do not apply:** Messenger, deputy, assistant in performance, assistant agent – all of them require their own will or fault. With autonomous AI agents, neither one nor the other is given so far.
- **The gap is particularly visible in the B2B sector:** Legal requirements such as GoBD, tax law, customs law and money laundering prevention require proof of the legal basis for services. Companies and authorities should be aware of the challenges in this regard, especially in B2B. In addition, there are also questions of risk assessment and reporting.
- **Technology and law must interact:** A new legal norm alone is not enough. For it to work, bidirectional bridges must be built between technology and law. This means that evidence such as identity, mandate and action must be designed in the sense of complementary protocols in such a way that they meet the legal requirements that correlate with them. The legal norm sets the framework. Technology fills it in. Only when both are coordinated will legally binding B2B agentic commerce become possible and sustainably attractive for companies.
- **No "mission impossible":** It is possible to coordinate different protocol types and legal norms in such a way that legally binding agentic commerce is created. It is crucial that a relevant jurisdiction takes up the topic, shapes it with relevant stakeholders and implements it. The first step already has an impact: The country or region that leads this important topic, sets the agenda and brings the relevant players to the table not only shapes the legal norms and protocols, it also distinguishes itself as an innovative architect of the next generation of international trade in the AI age.
- **Takeaways per section:** In order to provide a quick overview of the content that builds on each other, takeaways summarize the content of the sections. In this way, the most important theses and conclusions of this report can be grasped in just a few minutes.



A compass for business, law and technology

Agentic commerce is neither a purely economic nor an exclusively technical or legal issue. Rather, its success is characterized by the coordinated interplay of three worlds, each of which follows its own logics:

- **The business asks:** Where does (added) value arise? What use cases are there? What market acceptance and ROI can be expected?
- **The technology explains:** What can be operationally implemented, executed and documented with which AI tools and protocols?
- **The law regulates:** What rights and obligations arise between the actors? How does the attribution take place? How do declarations of intent and contracts come about?


Each of the three worlds influences the others: Agentic commerce that is not technically feasible does not create any added value and hardly raises any legal questions. But in the meantime, the technology is in place. It works. Now the business asks: "What do I get out of investing in it and using it?" And when looking at ROI in this regard, questions of trust and thus law come into play with regard to market acceptance:

- Agentic commerce is a network. Its value grows with the number of participants. Broad acceptance, in turn, requires that small and medium-sized enterprises participate.
- In the end, the return on participation depends, among other things, on whether legally binding contracts arise from the actions of the AI agents and to whom they are attributed.
- If no civil law contracts are concluded (as before), a lack of legal certainty across the entire network becomes a critical acceptance factor – especially for the large number of small and medium-sized (supplier) companies.

This expert opinion helps to identify and close relevant legal gaps in Agentic Commerce.

It also clarifies the conditions under which stable advantages arise for retailers. Part 1 answers the legal questions in particular: What gaps exist?

Part 2 provides drafts for legal norms and bridging protocols. In addition, there are ROI scenarios. Part 2 is accompanied by an interactive AI prototype that illustrates the interaction of all elements.



In this respect, the report and prototype are intended as a framework for discussion and a compass for cooperation between government, law, B2B stakeholders in trade and logistics, and technical providers of operational agentic commerce infrastructures.

Table of Contents:

I. Legal nature of AI agents	10
1. Growing relevance of agentic AI.....	11
2. Agentic Commerce Despite Growing Risks of Agentic AI?	13
II. What is Agentic Commerce?.....	16
1. The term.....	16
2. Two dimensions: B2C and B2B	16
3. The stages of autonomy	17
4. Differentiation of A2A and H2A	17
III. Three levels of legal binding.....	19
1. The international and dogmatic impasse of stage 2	19
2. The crucial gap	21
3. Why c.i.c. does not replace a warranty for defects.....	21
4. The fault problem	22
5. The consequence	22
IV. Consequences of missing contracts	24
1. The domino effect.....	25
2. The risk levels at a glance	26
3. The risk of disclosure (c.i.c.)	26
4. Risk matrix and risk curve with and without legal norm	27
V. The legal layer model	31
1. The levels and their gaps	31
2. Why the classic constructions don't work	32
3. Why the analogy doesn't help either.....	34
4. The vulnerability is not technical.....	34
5. What this means.....	34
VI. Four precursors, one gap	36
1. Existing International Regulations.....	36
2. First and second generation of automation	37
3. The Crucial Distinction: Proof vs. Validity	39
VII. The current state of the art	42
1. The US standards: Commerce, Payment, Legal Context	42
2. Payment service provider initiatives.....	43
3. The European infrastructure	43
4. The international discussion	45
5. The common gap	45
6. The error in thinking and its causes	46




VIII. The required legal norm	48
1. Requirements for legal norms	48
2. Create an attribution rule	49
3. Enabling the connection to the obligation	50
4. Map the entire chain	50
5. The limits of authority in the case of autonomy	51
6. The requirements for a sustainable legal norm	52
7. The vote of confidence	53
IX. Implementation options	54
1. Option A: Model contractual clause + Choice of Forum	54
2. Option B: Regulatory sandbox	55
3. Option C: Constitutive legal norm.....	55
4. The combination.....	56
X. Quo vadis?	57
1. The wake-up call.....	57
2. The urgency	57
3. Three paths	57
4. The instrument: Choice of Forum	58
5. The business channels and final question	58
6. Conclusion.....	58
XI. Open points	61
1. Attribution and declaration of intent	61
2. Systemic risks	61
3. Interoperability and recognition.....	61
XII. Glossary	62
1. Legal terms with legal definition	62
2. Legal terms without legal definition	62
3. Basic Civil Law Terms (German Law Only)	63
4. Technical Standards and Protocols	64
5. Hybrid legal-technical terms	64
6. Instructions for use.....	65
Appendix	66

Agentic commerce is no longer a vision of the future – the first pilots are already running on the market. Nevertheless, there is a lack of trust and, last but not least, a lack of legal certainty. The question remains open: "How could AI agents make legally binding declarations? To whom are they attributed?" This report outlines the current legal situation and also addresses the necessary interaction between technical processing and legal validity. With regard to the conclusion of an agentic contract, the report discusses various options specifically for agentic B2B trade: from the model contractual clause to the regulatory sandbox. In a separate document, a standardization proposal with matching protocol recommendations is made.

At the WELT AI Summit¹ 2026, VISA announced that it wanted to test the use of AI agents in payment transactions. The agents would suggest products and process transactions – all under the control of the consumer. Albrecht Kiel, Regional Managing Director Central Europe at VISA, says they are confident that it will soon work².

Figure: VISA announcements on Agentic Commerce from the WELT AI Summit 2026



But even the greatest technical confidence faces open questions of trust and legal uncertainty. VISA is also aware of this: Adoption of Agentic Commerce ultimately depends on consumer trust, as well as clear authorization models and a trusted infrastructure that enables secure transactions at scale³.

¹ to the event <https://www.welt.de/ai-summit/>

² see <https://www.welt.de/wirtschaft/ai-summit/video6ab22ffcb1857e86757aa626/welt-ai-summit-2026-VISA-testet-agentic-commerce-ki-agenten-kaufen-kuenftig-fuer-verbraucher-ein.html>

³ See e.g. <https://www.paymentsjournal.com/delegation-with-limits-what-merchants-want-from-agentic-commerce/> and <https://fastcompany.com/fastco-work/ai-could-reshape-commerce-trust-will-drive-it/>

Research shows that user trust remains limited, with only around 1% of respondents wanting fully automated AI shopping, while 56% say they will never allow an AI platform to spend money without their consent.

Trust is thus becoming the decisive currency. VISA India puts it this way: The next phase of digital commerce must be built around a new layer of trust. It should also enable AI agents to reliably complete transactions on behalf of consumers.⁴ VISA is developing a "Trusted Agent Protocol" (TAP) to verify that an agent is authorized by the consumer and to determine the authorizations that have been granted to it.

But one central question remains open:

- On the one hand, global players such as OpenAI, Anthropic, Google, Perplexity or VISA are developing a series of technical protocols that verify the authorization of an agent and document the processing of a transaction.
- Independently of this, a European infrastructure for digital identities is emerging: eIDAS 2.0, the EUDI Wallet for natural persons and the European Business Wallet (EBW) for legal entities. In addition, the EU Digital Power of Attorney regulates cross-border corporate law acts.
- However, neither the protocols nor the infrastructure make it possible for an authorization to be legally valid, for the agent's action to be attributed to a recognized legal entity, and thus for a legally binding contract to be created.

In addition to different protocol types and infrastructure, the legal assignment requires a legal norm that attributes the machine-generated explanatory act of an autonomous AI agent to the legal entity. Such a legal norm with matching bridging protocols is not yet apparent – neither in Europe, nor in America or Asia.

 As shown in Section VI, there are regulations that affect partial aspects: However, they were all designed for the first rule-based and deterministic generation of automation. They do not represent internationally compatible legal norms for the specific phenomena of autonomous, learning agents.

The legal gaps that exist so far, how they would have to be closed and what such a legal norm could look like are outlined in this report.

For this purpose, a glossary was integrated at the end that outlines technical terms and their relationship to legal aspects. Finally, an assessment is also made as to which legal system could implement the required norm most pragmatically in the B2B sector due to various prerequisites – this must be considered first, as infrastructures such as identities and wallets for legal entities (European Business Wallet, EUID, LPID) are

⁴ to this <https://www.exchange4media.com/marketing-news/agentive-commerce-will-make-trust-the-new-layer-of-digital-payments-rishi-chhabra-VISA-158532.html>

already being implemented. The B2B sector is thus a pragmatic starting point for legally binding agentic commerce – less so the consumer sector, where protection needs and consumer law significantly restrict the scope for action.

 The B2B sector is even of particularly high relevance: Mastercard's white paper *"Trust for agentic commerce"*⁵ (September 2026) outlines the trust framework for agentic commerce and specifically addresses both dimensions:

- The consumer sector (B2C) and
- business-to-business transactions (B2B).

The conclusion is even that the strongest growth will come from procurement. Gartner cites a figure of more than 15 trillion US dollars by 2028 – a forecast that depends on the market acceptance of agentic transactions. The decisive factor here is that the B2B sector is not only seen as a central growth area, but that many formal aspects are added in B2B:

- Here, legal evidence vis-à-vis authorities is important. Tax law, customs law, accounting law and money laundering prevention often require binding proof of the legal basis of services.
- So if a company has purchased goods through an agent, the gap often only becomes visible when the VAT return is submitted – when the tax office asks: "Who was the supplier and recipient of the service?"
- According to the German GoBD⁶ for example, business transactions must be recorded individually, including the name of the contractual partner. Exceptions (e.g. for cash sales to unknown customers) also do not apply in the B2B sector.

There is another aspect: B2C agentic commerce is hardly scalable without an already established B2B agentic commerce, as volatility increases overall (flash orders, dynamic bundling): If a customer's B2C agent orders large quantities of goods, the retailer must also be able to deliver in the end. To do this, the retailer sends its own B2B agent to the wholesaler or manufacturer to order the goods.

As long as B2B transactions only generate an "illusion of validity" (model clauses) or only enrichment effects, no supplier will presumably reserve, produce or deliver goods to the retailer in real time. Agentic Commerce's promise is quite right: *"AI agents reduce transaction costs and optimize purchasing."*

However, it should also be noted that transaction costs are not everything:

⁵ Available for download, as of September 2026 <https://www.mastercard.com/global/en/news-and-trends/Insights/2026/mastercard-agentic-commerce-trust-white-paper.html>

⁶ Principles for the Proper Maintenance of Records in Electronic Form https://de.wikipedia.org/wiki/Grundsatz_zur_Ordnungsm%C3%A4%C3%9Figen_F%C3%BChrung_und_Aufbewahrung_von_B%C3%BCchern,_Aufzeichnungen_und_Unterlagen_in_elektronischer_Form_sowie_zum_Datenzugriff

- The number of returns in e-commerce is already enormous⁷: A new study by EHI identifies return transport as the biggest cost driver (70.6%) – and this with an increasing number of returns⁸. Currently, one in four parcels ordered online is returned.
- Without normative guardrails and technical bridging protocols, AI agents optimize transaction costs only on paper for the time being. In physical reality, they could, in turn, generate an even greater flood of returns, which in turn could lead to increased logistics costs, warehouse blockages and losses in value in the legal vacuum of the law of unjust enrichment, and thus to a questionable ROI.

Against this background, the B2B sector is the area in which, according to the assessment of the present report, the legal gap is first visible and noticeable in terms of costs, both economically and in terms of tax and regulatory law. Every company is free to expose itself to the risks in this regard. But a corresponding legal norms, together with the appropriate bridging protocols, could significantly reduce the existing uncertainties and promote agentic commerce first in the B2B sector and then also in B2C.

 Agentic commerce thus potentially optimizes transaction costs, but by no means the costs across the transaction chain as a whole. From this point of view, legally compliant agentic commerce in B2B can certainly be seen as a kind of *conditio sine qua non* for physical agentic commerce in the B2C sector. The target image must be an end-to-end model that actually pays off for all parties involved in the end – including the small and medium-sized players. Corresponding ROI factors or cost/benefit aspects are examined in Part 2 of the report.

Takeaways: Introduction

- **From vision to reality:** Agentic Commerce is no longer a vision of the future. It runs technically reliably in various pilot projects.
- **Legal uncertainty:** AI agents have so far acted without legal attribution. There are no declarations of intent and thus valid contracts.
- **Contracts as proof:** The gap is particularly relevant in the B2B sector, as Tax, customs and GoBD requirements often require a contract as proof.
- **Profitability through legal certainty:** Successful agentic commerce requires the interaction of business (ROI), technology (protocols) and law (legal norms).
- **B2B as a starting point:** B2B is the logical and legally more flexible starting point for legally binding agentic commerce (ahead of the highly regulated B2C).

⁷ more on the EHI study at <https://logistik-heute.de/news/e-commerce-retouren-bleiben-relevantes-thema-286428.html>

⁸ to this <https://logistik-heute.de/news/e-commerce-2026-wird-mit-560-millionen-retouren-hoechststand-rzielt-286383.html>

I. Legal nature of AI agents

As early as the summer of 2025, the author published a section in the book "Grundwissen KI-Recht"⁹ that deals with the interfaces between the AI Regulation and civil law. Agentic shopping was analyzed with regard to the question of whether autonomous AI agents in the business-2-consumer (B2C) sector could conclude legally valid contracts under German civil law or at least de facto "effect" them.

Image: Contribution to Agentic Shopping from summer 2025

126	Grundwissen	KI-Recht
	2. Agentic Shopping	
Agentic Shopping nimmt rasant zu	Ähnlich und doch anders verhält es sich beim sich rasant verbreitenden Phänomen des „Agentic Shopping“ ⁵⁸ . Hier stellt sich die Frage, ob ein „KI-Agent“ auch ohne Abgabe einer KI-generierten Willenserklärung Verträge zumindest faktisch „bewirken“ könnte.	338
	Bsp.: Verbraucher V erhält im Online-Shop A die Möglichkeit, den dauerhaft autonom agierenden KI-Agenten „I-Buy4You“ zu aktivieren, der unter bestimmten Vorgaben selbständig Einkäufe im Shop tätigt. Dazu durchsucht „I-Buy4You“ permanent die Web-Datenbank von A nach vordefinierten Produktprofilen, liest die Produktbeschreibung und den Preis, um darauf basierend Kaufentscheidung zu treffen, z.B. im Fall eines temporären Sonderangebots.	339
	Da ein KI-Agent selbst keine eigenen Willenserklärungen abgeben kann, wäre eine Vertretung oder Botenschaft denkbar. Denkbar wäre z.B., entsprechend der Häufigkeit der Nutzung eines KI-Agenten über die Regeln von Duldungs- und Anscheinsvollmacht einen Rechtsscheinstatbestand zu erwägen.	

At the end of the section was the statement that is still valid today:

"Currently, neither an AI system nor an AI model with a general purpose has legal capacity or can make its own declaration of intent under civil law or conclude a contract."
This leads to the question: How can a legal system create trust in a technology that cannot make a declaration of intent?

 In 2002, the author also published the dmmv-EC handbook as¹⁰ a co-author. The basic work at that time makes it clear that even in the case of e-commerce, which seems to be taken for granted today, a legally binding framework had to be created in order to end an increasing crisis of confidence¹¹.

It was necessary to clarify how *people* can conclude contracts digitally:

- Whether and when declarations of intent can be made automatically.
- When access to a digital declaration of intent is to be assumed.
- How shopping carts are to be designed in a legally compliant manner.

⁹ See the entire section on AI Agents and Agentic Shopping www.grundwissen-ki-recht.de from paragraph 326

¹⁰ for goals and content see e.g. <https://www.etailment.de/magazin/praxishilfe-gestalten-und-betreiben-von-online-shops>

¹¹ worth reading from 2001 <https://www.heise.de/news/Web-Shopping-in-der-Vertrauenskrise-38549.html> and https://www.researchgate.net/publication/242136113_Die_Bedeutung_von_Vertrauen_im_Electronic_Commerce

It was only the Electronic Commerce Act (EGG) and an amendment to the German Civil Code (BGB) that brought clarity. Today, the law in agentic commerce is faced with the far more complex task of regulating the requirements for legally binding contracts that are made for people – on the basis of agent technologies that are developing at breakneck speed.

1. Growing relevance of agentic AI

By mid-2025, the vast majority of people probably didn't even know what autonomous AI agents were and what regulatory challenges they could pose. In this respect, the section in the script was an introduction to a topic that was still new at the time, but was recognizably developing in a highly dynamic way. It consists of two innovations that build on each other:

- Agentic AI in general (technical basis)
- Agentic shopping or commerce in particular (incl. processes and protocols)

a) Rapid development steps from 2023

An intelligent agentic structure is generally characterized by the ability to go beyond the mere reactive response, i.e. to become active in a targeted manner:

- An agent perceives his environment (perception),
- sets goals, independently plans intermediate steps (reasoning/planning)
- and executes actions to achieve these goals (Action Execution)

Intelligent agentic systems are older than you might think: As early as the 70s, the first physical prototypes such as Shakey the Robot¹² or virtual agents such as SHRDLU were created¹³. However, high-performance agentic AI did not emerge until 2023 with frameworks such as BabyAGI or AutoGPT¹⁴. They were based on generative AI models, so they were themselves more like agent track systems. From the FAZ's point of view, the topic even seemed like a kind of science fiction in 2023¹⁵. Agentic AI was quickly declared a market with billions in potential¹⁶. In 2024, the first ideas such as "A bot buys my ticket" emerged. Within two years, this developed into an independent (future) category in online retail, the size of which was already valued at billions in sales in 2030¹⁷. Since 2025,

¹² see [https://de.wikipedia.org/wiki/Shakey_\(Robot\)](https://de.wikipedia.org/wiki/Shakey_(Robot)) with its own sensors, motion planning (STRIPS, A*) and environmental representation.

¹³ see <https://de.wikipedia.org/wiki/SHRDLU> he understood language, could reason logically and perform actions

¹⁴ more on IBM <https://www.ibm.com/de-de/think/topics/babyagi>

¹⁵ see article <https://www.faz.net/aktuell/wissen/babyagi-und-autogpt-was-koennen-autonome-ki-agenten-18855332.html>

¹⁶ Development, among others <https://www.fortunebusinessinsights.com/agentic-ai-market-114233>

¹⁷ e.g. <https://www.marketsandmarkets.com/Market-Reports/agentic-commerce-market-225920956.html>

however, questions about security risks and the trustworthiness of AI agents have also increasingly¹⁸ arisen.

In summary, this shows that Agentic AI and Agentic Commerce have developed at an enormous speed in just a few years, especially technically. At the same time, fundamental legal questions of Agentic AI have remained largely unresolved to this day. This, in turn, is a showstopper for broad market acceptance that should not be underestimated.

b) Legal aspects from the perspective of 2025

The 2025 article first addressed the legal nature of agents as AI systems within the meaning of the AI Regulation. In addition, the characteristics as tools, messengers and representatives, the problems of attribution, the question of the implied conclusion of a contract, the avoidance due to error and deception, and the liability issues within the AI value chain were examined.

Aspect	Topics & Results 2025 (B2C)	Legal Norms
Legal nature	AI systems are neither natural nor legal persons.	§§ 1, 21 et seq. BGB
Attribution	Tool, messenger, deputy, knowledge representation – none relevant.	§§ 164, 166, 120 BGB
Declaration of intent	AI Agent cannot submit its own declaration of intent.	§§ 116 et seq. BGB
Conclusion of contract	Question shifted from "if" to "when" – implied conclusion of the contract conceivable.	§§ 145 et seq. BGB

The text of 2025 also treated the individual aspects as a number of independent individual factors – their systemic interaction with parallel developments such as the eIDAS 2.0 Regulation¹⁹ or technical protocols was not taken into account. In this respect, although many relevant building blocks were already in place (the legal nature, the attribution, the conclusion of the contract, the avoidance, the liability), they stood side by side at the time, without any closer consideration being given to how civil law issues, e.g. of the power of representation, are related to the technical infrastructures of a company ID, a transaction protocol or an agent power of attorney.

The interplay between technology and law has become much more topical and also more concrete within just one year:

- In the meantime, various pilot projects for agent payments have been launched²⁰.
- The industry has developed various new protocols and legal norms²¹.

¹⁸ one of many warning reports from 2025 <https://www.lasso.security/blog/agentic-ai-security-threats-2025>

¹⁹ Learn more <https://www.digitale-verwaltung.de/Webs/DV/DE/digitale-identitaeten/eidas-2-0/eidas-2-0-node.html>

²⁰ Overview e.g. under <https://base.com/de-DE/blog/agentic-commerce-erklaert/> based on

²¹ see e.g. LCP <https://legalcontextprotocol.org/>. More on the standards later

- Government IDs, registers, IDs and wallet solutions have been added²².

Regardless of this, however, the discussion on general trust in AI systems has also reached a new urgency²³. It is becoming increasingly important to synchronize AI security with the speed of further development of AI capabilities²⁴.

2. Agentic Commerce Despite Growing Risks of Agentic AI?

The many examples of agentic AI acting uncontrollably illustrate what can happen when chains of autonomous AI systems pursue goals independently and in a division of labor – regardless of the limits that have been set for them.

a) Agentic Commerce Security

In the meantime, several basic questions arise for Agentic Commerce:

- If even agentic AI as a technical backbone seems insecure in general, are secure transactions still possible at all?
- So would you really entrust your wallet to a "wild" AI agent or entire purchasing and sales processes as a company?
- How expedient is the clarification of the legal basis for agentic commerce when the basic principle of agentic AI seems increasingly questionable?

The question is new in this intensity, because it concerns two independent aspects, namely that of active security (execution of transactions) and that of passive security (protection against cyber attacks by agentic AI).

In this respect, it is important to note that an attacker agent and the purchasing agent can both be agents in B2B. In this respect, they belong to the same genus. But they concern two different categories – similar to how a dog can be a "fighting dog" or a "lap dog". Applied to agentic commerce, this means that active security of a mandated agent in the sense of a "lap dog" is possible – but it must be noted that the best possible passive security is necessary, as a "lap dog" can also be attacked by a "fighting dog" (more on the topic of risk later in section IV)

The degree of autonomy must be seen in relation to the authority to act: So who determines the scope of action – the AI system or the agent itself or a valid, unambiguous external mandate?

²² e.g. from Estonia: <https://edrm.net/2026/06/estonia-aims-to-be-first-to-give-ai-agents-official-digital-ids>

²³ Australian Prime Minister Anthony Albanese, among others, felt compelled to warn against agent AI due to the first known autonomous attack on a government website: <https://www.bbc.com/news/articles/c6vgv0333dppo>

²⁴ among other things. <https://www.linkedin.com/feed/update/urn:li:activity:7510572157697978368/>

b) Differentiation of different agent types

The crucial distinction is not the degree of autonomy as such. A mandated agent can act highly autonomously – he can search through millions of possibilities, compare offers, form alternatives. Agent autonomy does not prevent the principle of control itself.

The decisive difference is the limits of autonomous agency: Who determines the scope of action – the system itself or an external mandate? And this is exactly where it is important to start with the interplay of law and technology as the basis of trustworthy, legally binding agentic commerce!

Aspect	Open Agents	Mandated agents
Objective	System can generate goals and sub-goals itself.	The goal is set externally.
Scope of action	Open – dynamically expandable.	Limited by mandate and policy.
Task Scheduling	Self-determined.	Autonomously within the mandate.
Resources	Potentially self-expandable.	Explicitly authorized resources.
Learning behaviour	Can change target structure and strategy.	Only within defined parameters.
Boundaries	Can be part of your own problem-solving process.	Technically and legally prescribed.
Monitoring	Observation of the system is required.	Actions and policy decisions auditable.
Escalation	If necessary, independent continuation.	In case of policy violation: Stop / Escalation.
Revocation	Conceptually difficult.	Mandate can be withdrawn at any time.
Legal structure	attribution problematic.	Mandate, Action, Attribution and Liability.
Core risk	Goal and action control.	misinterpretation or overstepping the mandate.

Not only does agentic commerce not need fully autonomous AI – on the contrary: it needs a high level of operational autonomy with little normative autonomy. A commerce agent should be able to decide autonomously how to achieve a legitimate goal. However, he should not decide which goal he is allowed to pursue. Ensuring this is both technically possible and legally feasible.

c) Mandated agents

In order to outline the possibilities of legal order structures for agents bound by a mandate, this report first takes up the building blocks of that time again and rearranges them. It shows how a series of individual questions becomes a layered architecture that can be answered neither purely technically nor purely legally, but solely in the interplay of law and technology. Although much is already technically possible today, the US standards

(described later in Section VII and in the glossary) such as UCP, AP2, ACP or LCP do not pay for the legal challenges of making effective declarations of intent.

Rather, law and technology must be closely coordinated. Therefore, this report makes recommendations for bridging standards as well as for legal norms that are aligned with the correlating protocol recommendations. In the end, this can also answer the question of what kind of legal standard is missing so that the overall architecture is viable from a legal point of view.

Takeaways: Legal nature of AI agents

- **No recognized legal entity:** AI and AI agents are not yet recognized legal entities. Classic characters (messenger, representative) do not apply to autonomous AI.
- **Historical parallel:** In the e-commerce of the 2000s, a legal framework was also necessary to end a growing crisis of confidence.
- **Young technology, unclear basics:** Agentic AI (perception, planning, action) did not develop until 2023 – the civil law bases are unclear.
- **Security as a driver:** Agentic commerce is based on agentic AI. Incidents are increasing – control and AI security are becoming more urgent.

II. What is Agentic Commerce?

Before the legal analysis begins, a conceptual clarification is required. This is because "agentic commerce" is not a legally defined term – and it is used differently in the discussion. Depending on which section you look at, the legal situation is different.

1. The term

In the following, agentic commerce is understood to mean the autonomous or delegated conclusion of trade transactions by AI agents – as opposed to AI-supported recommendation, assistance or checkout support. It is crucial that the agent acts himself – not just prepares or suggests.

This understanding is also in line with the official definition, which distinguishes between semi-autonomous and fully autonomous agentic commerce.²⁵ Human involvement is usually limited to initial configuration, periodic monitoring, or exception handling.

2. Two dimensions: B2C and B2B

In the case of agentic commerce, two business-to-consumer (B2C) and business-to-business (B2B) constellations must be differentiated in terms of actors, legal framework and protection needs:

Dimension	Actors	Legal framework	Need for protection
B2C (Agentic Shopping)	Company / Consumers	§§ 13, 14 BGB; §§ 312 et seq., 327 et seq. BGB; EU AI Act	high – consumer protection.
B2B (Agentic Commerce)	Company / Company	§§ 14 BGB; HGB; EU AI Act	– both sides act professionally.

The article from 2025 primarily examined the B2C case of agentic shopping. In order to enable pragmatic legal recommendations, this expert opinion, on the other hand, focuses on the B2B case:

- In the B2B sector, the legal freedom is greater, but the special protective mechanisms such as consumer protection are significantly lower.
- The practical relevance is higher because the first infrastructures such as a European Business Wallet (EBW) or identities for legal entities based on eIDAS 2.0 are in the starting blocks in this area.

²⁵ see the U.S. definition at https://en.wikipedia.org/wiki/Agentic_commerce

3. The stages of autonomy

Within agentic commerce, several decision-making stages can also be distinguished in the area of B2B:

Autonomy level	Who decides?	Identity and mandate context
1. AI Recommendations	People	pure suggestion generation; conventional login.
2. AI-powered checkout	Human confirmed	AI prepares; Humans actively release (SCA/2FA).
3. Delegated Purchasing	Agent within mandate	Agent operates within a fixed framework (budget, supplier); digital mandate.
4. Autonomous Agent	Agent (autonomous)	Agent makes independent decisions; escalation on limits.
5. Agent-to-Agent (A2A)	Agents among themselves	Fully automated transaction chains; Smart contracts / wallets.

The focus is then on levels 3 to 5. Levels 1 and 2 are rather unproblematic from a legal point of view: the person decides – and makes the declaration of intent. Levels 3 to 5 in turn raise the questions that are legally at the center:

- Who makes a declaration of intent if not the human being as a natural person?
- Or the human being as an organ of a legal person or society?

4. Differentiation of A2A and H2A

The combination of the stages results in a matrix that depicts the direction of the agent's actions. The decisive factor is on which side the agent stands, the letters A and H stand for "Agent" and "Human".

Constellation	Who orders?	Who delivers?	Who gives the WE?	Contract?
Person to person (H2H)	People	People	two people.	Contract
Human-to-Agent (H2A)	People	Agent	a person, an agent.	questionable
Agent to Human (A2H)	Agent	People	an agent, a human being.	questionable
Agent to Agent (A2A)	Agent	Agent	two agents.	no contract

The previous matrix illustrates the core of the challenge – it shows:

- As soon as an agent acts on one side and a human on the other, there is a potentially legally valid declaration of intent on at least one side.
- However, since there is no declaration of intent at the same time, it remains questionable under what conditions a contract could be concluded.

- As soon as agents act on both sides, no contract is concluded under current law – because no declaration of intent is made on either side.

In practice, there is another special feature: there is often an intermediary between buyer and seller – a wholesaler, a platform or a hub. He mediates, bundles, coordinates. The intermediary can act in different roles, namely as a messenger (he transmits), as a representative (he acts on behalf of a third party) or as an independent trader (he buys and sells himself). Each role has a different sequence of attribution. The resulting A2A2A constellation with all variants such as H2A2H etc. would not only significantly increase the previous overview, but also the overall complexity. If the foundations for A2A have already been laid, the expansion to three-party relationships is much easier.

 **Wholesale is no exception – rather, it is a very realistic starting point for agentic B2B commerce, because the intermediary could provide coordinated infrastructures and bill them "on demand". However, the following explanations are primarily based on two-party constellations. This not only reduces the complexity of the presentation. It is also sufficient to illustrate that the lack of legally valid contracts in national and even more so in international two-party constellations already has quite relevant consequences. In the end, however, a legal norm recommendation must be formulated in such a way that it also covers multi-level chains. Otherwise, a "new gap in the closed gap" will arise.**

Takeaways: What is Agentic Commerce?

- **No legal term:** "Agentic commerce" is not defined – the focus in the opinion is on the autonomous or delegated conclusion of effective legal transactions.
- **Two dimensions:** B2C (consumer protection, high need for protection) vs. B2B (professional players, greater room for manoeuvre).
- **Five levels of autonomy:** From AI recommendations (human decides) to agent-to-agent (fully automated). Focus on levels 3–5.
- **A2A is the core:** As soon as agents act on both sides, no contract is concluded under current law – because neither side makes a declaration of intent.
- **Intermediaries note:** The A2A2A constellation (wholesaler) is important in practice and expands the A2A attribution issues.

III. Three levels of legal binding

If the contract is missing, in legal terms, in addition to tortious liability law, there is above all the Unjust Enrichment (§§ 812 et seq. BGB) – goods and money were exchanged, but without a legal basis:

- Both sides can demand surrender step by step.
- This is the rescission, but not a contractual settlement.

In addition, the possibility of culpa in contrahendo (§ 311 para. 2 BGB) remains:

- Anyone who violates a duty of protection or information during contract negotiations is liable for the resulting damage.
- In addition, the c.i.c. only compensates for the loss of trust – not the interest in performance. And it presupposes (attributable) fault, which is absent in the case of the AI agent.

This results in a gradation that is important for the entire further contribution:

Level	Legal basis	What is created?	What is missing?
1	§ 812 et seq. BGB	Claim for enrichment.	Every bond.
2	Section 311 (2) of the German Civil Code (BGB)	Obligation with protective duties.	Performance obligations, warranty for defects.
3	§§ 433 et seq. BGB	Contract with full rights.	–

5. The international and dogmatic impasse of stage 2

Stage 1 is the de facto exchange. Stage 2 arises with the business contact. Only at level 3 is there a legally valid contract that confers full rights in the internal relationship and is also recognised as proof in the external relationship.

With regard to level 2, it should be noted that, according to German law, according to Section 311 (2) of the German Civil Code, the initiation of contract negotiations, the initiation of a contract and even business contacts are sufficient to establish an obligation in this regard. If two AI agents from different companies interact with each other, an obligation can certainly be established.

Nevertheless, the difference between an obligation and a contract is important:

- Those who remain at level 2 have protective duties – but no performance obligations.
- It has c.i.c. – but no warranty for defects.
- A contractual relationship arises – but not a contract.

What is presented here using the example of German law is not a special national problem. The culpa in contrahendo is a German figure – but its subject matter is international. Every legal system has a form of pre-contractual liability, even if it has a different name and is constructed differently:

Legal system	Figure	Basis
Germany	culpa in contrahendo	§§ 280, 241, 311 para. 2 BGB
France	responsabilité précontractuelle	Art. 1112 Code civil
Austria	culpa in contrahendo	§ 1295 ABGB
Switzerland	Liability of trust / Art. 41 CO	Judges' Law, Art. 41 CO
Common Law	negligent misrepresentation / promissory estoppel	Case law
UNIDROIT Principles	Termination of negotiations in bad faith	Art. 2.1.15

The following applies everywhere: Anyone who violates an obligation during contract negotiations is liable. But it is also true everywhere: This liability does not replace a contract. And nowhere is there a constitutive legal norm that recognizes the action of an agent as a declaration of intent.

a) First stage: The de facto exchange – §§ 812 et seq. of the German Civil Code (BGB)

Goods and money are exchanged – without there being a legal basis. Both sides can demand the return of what has been obtained. This is the reversal – step by step.

Feature	Significance
Prerequisites	De facto exchange without legal grounds.
Legal consequences	enrichment claim (§ 812 para. 1 BGB).
What is created?	Claim for rescission.
What is missing?	Every bond.

Stage 1 is the rescission. But it is not a settlement. There is no claim for performance, no warranty for defects, no transfer of ownership by contract.

b) Second stage: Business contact – § 311 (2) of the German BGB

An obligation arises – even without a contract. Through contract negotiations, contract initiation or similar business contacts. It establishes duties of protection and consideration (Section 241 (2) of the German Civil Code) – but no obligations to perform.

Feature	Significance
Prerequisites	Business contact.
Legal consequences	Obligation with protective duties.
What is created?	c.i.c. liability, loss of confidence.
What is missing?	Performance obligations, warranty for defects.

Level 2 is basic security. But it is weak – because it only compensates for the damage to trust, not the interest in performance. And because it presupposes fault, which is absent in the autonomous agent.

c) Third stage: The declaration of intent – §§ 433 et seq. BGB

Two concurring declarations of intent – and the contract is concluded. With all rights and obligations: claim for performance, warranty for defects, withdrawal, reduction, damages.

Feature	Significance
Prerequisites	Two declarations of intent.
Legal consequences	Purchase contract (§ 433 BGB).
What is created?	Claim for performance, warranty for defects.
What is missing?	–

Level 3 is full security. However, it requires that two declarations of intent are submitted. The agent cannot do this.

6. The crucial gap

There is no transition between level 2 and level 3. Anyone who stops at level 2 has duties of protection – but no performance obligations. He has c.i.c. – but no warranty for defects. He has an obligation – but no contract.

And this is exactly where the problem lies: The agent cannot reach level 3. So it remains at level 2 – or below.

7. Why c.i.c. does not replace a warranty for defects

The c.i.c. is not a substitute for the warranty for defects. This is shown by a comparison:

	c.i.c. (§ 311 para. 2 BGB)	Warranty for defects (§ 437 BGB)
What will be replaced?	Damage to trust.	Interest in performance.
What does it require?	Breach of duty, culpability.	Defect in the event of a transfer of risk.

What is missing from the agent?	Debt.	Declaration of intent.
Legal consequences	Damages.	Subsequent performance, withdrawal, reduction.

The c.i.c. only compensates for the loss of trust – not the interest in performance. And it presupposes fault, which is absent in the autonomous agent.

8. The fault problem

C.i.c. liability presupposes fault. The agent cannot be at fault – he is not culpable. So only the operator remains. But how is his fault justified?

Construction	Problem
Own breach of duty of care	The operator can exculpate himself.
§ 278 BGB (assistant in performance)	Assuming fault on the part of the assistant – missing.
§ 831 BGB (assistant in performance)	Exculpation possible – and only in tort.

There is no construction that achieves both: attribution without exculpation and without fault on the part of the agent.

9. The consequence

Without a declaration of intent, it remains at level 2 – or below. The episode:

Without a constitutive legal norm	With a constitutive legal norm
Level 1 or 2.	Stage 3.
Reversal or c.i.c.	Full contract.
No warranty for defects.	§§ 434, 437 BGB.
No claim for performance.	§ 433 BGB.
Legal uncertainty.	Legal certainty.



Takeaways: Three levels of legal binding

- **Level 1 – Unjust Enrichment:** Goods and money are exchanged, but without legal grounds. Only reversal – no commitment.
- **Level 2 – c.i.c.:** Obligation with duties of protection – but no performance obligations, no warranty for defects.
- **Stage 3 – Contract:** Two concurring declarations of intent – full rights: claim for performance, warranty for defects, withdrawal, damages.
- **The crucial gap:** There is no transition between level 2 and level 3. The agent cannot reach level 3.
- **Fault problem:** C.i.c. liability presupposes fault – the agent is not culpable. There is no construction that makes attribution without exculpation and without fault.

IV. Consequences of missing contracts

It is important to be aware of the consequences of missing contracts: If a contract is not concluded, then not only is a document missing, but there is no legal basis for many things that have been taken for granted in B2B trade to date. In some cases, for example, there is no recognized proof when companies interact with authorities. On closer inspection, a contract is much more than a binding agreement on goods and price.

Contracts are the basis for, among other things:

Function	Without a contract
Claim for performance	No claim to delivery (§ 433 para. 1 BGB).
Warranty for defects	No subsequent performance, no withdrawal, no reduction (§ 437 BGB).
Obligation to Complain / Commercial Purchase	§§ 377, 378 HGB require a purchase contract.
Transfer of ownership	Section 929 of the German Civil Code requires agreement – which is missing.
Transfer of Risk	Section 446 of the German Civil Code is linked to handover on the basis of a purchase contract.
Tax registration	The invoice and input tax deduction require an exchange of services.
Customs classification	Who is the declarant, who bears the taxes?
Accounting	No activation, no passivation without legal grounds.
Burden of proof	No contractual documentation obligations.

 The contract thus does not only have an effect inter partes between the parties. It also has an effect vis-à-vis the states in which the parties are domiciled: among other things, as recognised proof of tax, customs, money laundering prevention, accounting. The principles of proper accounting (GoB) show how concrete this obligation can be. According to § 146 (1) AO and § 239 (2) HGB, business transactions must be recorded in full and without gaps. According to the case law of the Federal Fiscal Court, this obligation includes the name of the contractual partner and – as far as reasonable – a sufficient description of the business transaction²⁶. If the contract is missing, this obligation can hardly be fulfilled. And anyone who does not fulfil it risks having the accounting system discarded – with considerable tax consequences. The contract is therefore not only an instrument under civil law. It is sometimes a legally prescribed obligation to keep records or a recognised type of document vis-à-vis authorities.

²⁶ See e.g. <https://ao.bundesfinanzministerium.de/ao/2024/Anhaenge/BMF-Schreiben-und-gleichlautende-Laendererlasse/Anhang-33/inhalt.html>

For B2B trade, therefore, obligations are hardly sufficient. Both sides need more than the c.i.c. or the Unjust Enrichment: explicit contract rights!

What the buyer needs, including:	What the seller needs, including:
Defect-free goods.	Claim for performance.
Supplementary performance in the event of defects.	Acceptance of the goods.
Withdrawal in case of incorrect delivery.	Payment of the purchase price.
Damages instead of performance.	Default interest.

All of this is only available at level 3. And level 3 requires that two declarations of intent are made. An agent has not yet been able to do this, which is why no contract has yet been concluded with either the A2A variant. In addition, it is highly questionable with the variants of H2A, A2H if an AI agent only acts on one side.

1. The domino effect

The absence of the contract has a number of consequences beyond civil law:

Field of Law	Results
Tax law	Important formal proof of input tax deduction and business expenses.
Customs law	It often remains unclear who is the declarant and who bears the duties.
Commercial Law	No obligation to give notice, no commercial custom.
Accounting law	No activation, no passivation.
Money laundering law	No recognized proof of economic background.
Private International Law	No choice of law, no Choice of Forum.

The contract is therefore not merely an instrument of civil law. It is a transmission belt for a large number of downstream obligations. If it is removed, the entire formal infrastructure of commercial transactions, which is based on the existence of contracts, wobbles.

From the point of view of companies, agentic commerce that is not legally protected gives rise to a number of economically relevant risks that must be integrated into an ROI analysis and investment planning with regard to agentic commerce (see Part 2 in detail).

 The domino effect has another facet: If an agent violates a legal prohibition, the legal transaction is void according to § 134 BGB. If it violates morality, according to § 138 BGB. Both norms can apply in parallel or alternatively – and in many cases the distinction is blurred: In the sale of radar detectors for road traffic, in the trade in narcotics or in illegal work, there has been a controversial discussion for decades as

to whether a contract is void under § 134 BGB, § 138 BGB or both. There is no uniform line – the decision depends on the individual case²⁷.

This vagueness is transferred to agent trade. In the case of prohibited AI practices under Art. 5 of the AI Regulation, the legal situation is still comparatively clear: Art. 5 of the AI Regulation is a prohibition law within the meaning of Section 134 of the German Civil Code (BGB) – the contract is void. It becomes more difficult in the case of practices that are not covered by Section 134 of the German Civil Code (BGB) but could violate common decency – for example, in the case of particularly aggressive negotiation tactics, the exploitation of emergencies or the systematic circumvention of protective regulations.

2. The risk levels at a glance

As a result, the attribution gap affects three levels – the settlement of the transaction, regulatory law and executive duties:

- **Settlement:** Without a legal norm, it is unclear whether and with what content a contract exists. The claimant bears the burden of proof for the conclusion of the contract – but can hardly fulfil it. With the legal norm, the agent's action is considered a declaration of intent by the legal entity. The contract is concluded. The burden of proof follows the control.
- **Regulatory law:** Without a legal norm, the contractual partner and the recipient of the service can be unclear. The input VAT deduction is at risk. The GoB require the recording of the contractual partner – and this can be unclear. The customs value presupposes a "sale for export" (Art. 70 UCC) – and without an attribution norm, it is unclear who is a party to the sale. With a norm, the legal entity is unambiguous. The input VAT deduction is secured. The recording obligation can be fulfilled. The transaction value is verifiable.
- **Duties of directors:** The management must be aware of and manage the risks associated with the use of agents (Section 93 of the German Stock Corporation Act (AktG), Section 43 of the German Limited Liability Companies Act (GmbHG); for AG: Section 91 (2) of the German Stock Corporation Act). Without a legal norm, a high level of legal uncertainty remains. With a norm, the residual risk is significantly reduced. The duty of care would be much easier to fulfil.

3. The risk of disclosure (c.i.c.)

A risk that is often overlooked: the pre-contractual obligation to provide information if only one side uses an agent (see variants H2A or A2H above).

²⁷ See in detail the paragraphs 343 et seq. of the script "Basic Knowledge of AI Law" prepared by the author

- Anyone who unilaterally uses an AI agent knows – or must know – that the legally binding effect of the agent's action is unclear. If he does not inform the business partner about this, he could violate his pre-contractual duty to provide information.
- The dogmatic basis: Section 311 (2) of the German Civil Code establishes an obligation by the commencement of contract negotiations. The obligations arising from this obligation include the disclosure of circumstances that are obviously important for the formation of the will of the other party. Anyone who appoints an agent without pointing out the lack of attribution could violate this duty.

The situation is different if a legal norm is passed:

- Then the attribution is clarified. The question of the duty to provide information does not arise.
- The legal norm eliminates the c.i.c. risk – not by excluding liability, but by clarity.

 Certainly, a distinction must still be made here according to the different variants, because A2H and H2A constellation require different clarification than A2A: If both sides use agents, both would theoretically have to inform each other with regard to the lack of contractual effect, which does not seem to make much sense.

4. Risk matrix and risk curve with and without legal norm

Summarizing the previous aspects results in a risk matrix that illustrates the risks that exist without a legal norm and what influence such a norm would have.

Dimension	Current risk	What the legal norm changes	What technology and governance Complement
Processing	Unclear whether and with what content a contract exists. Warranty and claim for performance uncertain. Burden of proof on the claimant.	Clear attribution. Burden of proof on the operator for the relief. Full contract.	Evidence (intent, mandate, log).
Tax and Supervisory Law	Contracting parties and recipients of services may be unclear. Input VAT deduction at risk (sec. 15 UStG). GoB: Record of the contractual partner. Customs value: party to the sale unclear (Art. 70 UCC).	Clear legal entity as contractual partner. Input tax deduction secured. Obligation to record can be fulfilled. Transaction value verifiable.	Identity and documentation (e.g. EBW).
Duties of directors	Management must be aware of and manage risks (Section 93 of the German Stock Corporation Act, Section 43 of the German Limited Liability	Reduces residual risk. Due diligence can be fulfilled.	Inventory, Roles, Controls (NIS2/DORA, as applicable).

Dimension	Current risk	What the legal norm changes	What technology and governance Complement
	Companies Act (GmbHG)). Residual uncertainty remains.		
Duty to inform (c.i.c.)	Risk of breach of duty to inform business partners.	Attribution clarified. There is no obligation to provide information.	–

The risks outlined above are due, among other things, to the fact that the tech providers of agentic solutions in general, but also in agentic commerce in particular, operate according to the principle of trial and error.

The principle of "silent versioning",²⁸ which just about works in the development of chatbots in order to achieve a better result through iterative and silent experimentation, is also critical. Agentic AI is reaching its hard limit:

- In B2B commerce, predictable continuity is expected over longer periods of time.
- And in the event of changes, early and clear communication is also necessary.

As industry experts criticize, Agentic AI often degenerates into a mere "stochastic guess" without a formal and reliable architectural foundation.²⁹ For Agentic Commerce in B2B, this is a clear showstopper.

Of course, especially at the beginning, it is not possible without "trial and error": Experience must be built up again and again in order to optimize products and services. A kind of permanent state of "trial and error", in which tomorrow everything can be different from today, but in the end only provides a chain of presumably working, not yet mature demos. For Agentic Commerce in B2B, this does not create a reliable, production-ready system for long-term enterprise use.

 In the tech industry, the paradigm of "trial and error" obviously still dominates the development of autonomous systems³⁰. The implicit premise is that the macro-economic efficiency gains through AI agents are so enormous that systemic "errors" – from hallucinations to unauthorized transactions to algorithmic collusion – should

²⁸ on the challenges in this regard, among other things. <https://daily.dev/posts/the-silent-versioning-problem-in-ai-inference-haeqzbrgl>


²⁹ There are a number of interesting articles on Trial and Error, including: https://www.linkedin.com/posts/davidgouvias_ai-agenti-cai-machinelearning-activity-7431937833436594176-DNcB or learning after mistakes <https://fortune.com/2026/09/29/how-adobe-salesforce-kyndryl-freeport-use-ai-agents/>. On the change in strategy with regard to AP2 in Agentic Commerce: <https://internetwarriors.de/blog/agentic-commerce-und-agentic-shopping>

³⁰ There are several statements on the principle of trial and error in the context of the further development of agentic AI, which are as topical as they are revealing: By Sam Altman <https://www.reuters.com/business/openais-altman-says-ai-benefits-warrant-accepting-some-risks-2026-10-04/> and OpenAI security expert David Robinson: <https://www.theatlantic.com/technology/2026/10/openai-safety-team-resignation/688881/>. Regarding NVIDIA: <https://en.infomaxai.com/news/articleView.html?idxno=74179>. Also interesting: <https://medium.com/illumination/fortune-500-companies-let-ai-make-fatal-mistakes-1e4d8ce7165e>

be accepted as "acceptable frictional losses". The costs of these errors are largely externalized: to the users, to the courts, to society.

The question that arises for B2B commerce: Who bears the "errors" – and who is allowed to bear them at all? A CEO, CFO or compliance officer should not regard the "error" of an autonomous commerce agent as statistical noise in the system. If an agent exceeds the mandate in the procurement process, this error does not affect "society". It directly affects the company and its executive bodies. The opportunities of agentic commerce are real – and so are the risks. However, the latter are of a different nature in B2B trade than in the tech world.

The motto "Move fast and break things" may be a successful innovation strategy in the tech world. However, agentic commerce in B2B should not be a space for tech companies to experiment. Legally compliant guardrails are not a brake on innovation. They are the prerequisite for the real economy to be able to take advantage of the opportunities – without getting involved in incalculable adventures that a responsible retail company would probably not take.



The overview illustrates that the "trial and error" approach in the field of agentic commerce is an asymmetric risk model and that "rationality" is a matter of perspective:

- For the tech industry (providers): Is it rational to bring agents to market quickly and "error-prone", as the benefits (scale, data, market share) outweigh the risk.
- For the B2B company (user): Is it irrational to use these agents without a legal norm, because the risk (compliance, directors' and officers' liability, tax benefits) immediately "eats up" the potential efficiency gain.

A legal norm forces the tech provider to internalize the risk again (by enabling warranty for defects and liability at level 3). Only when the matrix is "symmetrical" by a legal norm (i.e. the provider is liable for its "errors" like any other seller) does agentic commerce become attractive for the B2B market.

Takeaways: Consequences of missing contracts

- **Domino effect:** Without a contract, there is no claim to performance, warranty for defects, transfer of ownership, transfer of risk – and the basis for tax, customs, balance sheet, money laundering, IPR.
- **Risk levels:** Processing (burden of proof), tax and supervisory law (input tax deduction, GoBD, customs value), duties of directors (§ 93 AktG, § 43 GmbHG), disclosure (c.i.c.).
- **c.i.c. risk differentiated:** Only for H2A and A2H – not for A2A, because both sides would have to be aware of the gap.
- **Trial & error:** No neutral waiting, but active risk. The tipping point in B2B has been passed.
- **Benefit-risk curve:** Without a legal norm, the risk exceeds the benefit. With such a norm, the tipping point shifts to the right – the B2B sector becomes manageable.

V. The legal layer model

The question of the conditions under which an agent can conclude a B2B contract depends on the interaction of various factors. The required architecture can be represented in nine levels. Each level requires that the previous layer is fulfilled.

In this way, it can be checked on a successive basis which factors for a contract conclusion by AI agents have already been fulfilled and which still have to be fulfilled. The following layers refer to the German Civil Code and required evidence:

No.	Level	Guiding question	Legal basis
1	Legal entity	Who is behind the agent?	§§ 1, 21 et seq. BGB
2	Agent Identity / Authorization	Who acts – and with what authority?	EBW, IDs (technical)
3	Agent action	What did the agent actually have done?	§ 120 BGB (messenger); Section 164 of the German Civil Code (Representative)
4	Attribution	Whose plot is this legally?	§§ 164, 166 BGB
5	Declaration of intent	If there is a legally relevant explanation?	§§ 116 et seq. BGB
6	Contract / Legal transaction	Is a contract concluded?	§§ 145 et seq. BGB
7	Obligation	What is the legal relationship?	Section 311 of the German Civil Code (BGB)
8	Duties	What are the duties of performance and consideration?	§§ 241, 433 BGB
9	Breach of duty / liability	Who has to stand up for what?	§§ 276, 278, 280, 437 BGB

1. The levels and their gaps

If you want to be liable at level 9, you need level 8 (duties). If you want to establish obligations, you need level 7 (obligation). If you want to establish an obligation, you need level 6 (contract). If you want to establish a contract, you need level 5 (declaration of intent). If you want to establish a declaration of intent, you need level 4 (attribution). And attribution presupposes that there is anything attributable at all at levels 2 and 3.

And this is exactly where the gap becomes apparent. Today, the technical discussion usually starts at level 2 (identity, authorization, EBW) and then jumps to level 6 (contract – ACP, WP2, UCP). Levels 3, 4 and 5 are skipped – or taken for granted. But they are not self-evident. They are the core of the problem.

Conversely, the legal discussion often begins at level 4 (attribution) – and finds that the constructions of the BGB (messenger, representative, assistant in performance) do not

fit the autonomous agent. It then ends with the statement that a legal norm is missing – without saying what it is and what it should look like.

The architecture shows where the norm has to start: at level 4, the transformation point between action and declaration of intent. This level seems quite underexposed in the discussion so far. It is either taken for granted (by technology) or recognized as problematic, but has not yet been solved by law.

And it shows why the current discussion is so unsatisfactory:

- The US standards such as ACP, AP2, UCP and LCP address levels 2 (authorization), 6 (commerce) and – in the case of LCP – the time after the contract is concluded (Dispute Resolution, Terms).
- They do not address levels 3, 4 and 5. They assume that the declaration of intent will come about "somehow" – and leave the clarification to the law. The law, however, has not yet provided the clarification.

The layered architecture thus makes visible what remains hidden in the discussion so far:

- It's not about a single problem. It's about a chain.
- And it is currently cracking at the point where a technical act is supposed to become a legal declaration.
- This passage is called attribution – and it cannot be closed by a protocol, but only by a legal norm.

In the following, it is shown what solutions already exist at the respective levels, where there may be an error in thinking, what the gradation of the legal binding looks like and what a legal norm would have to do to close the gap.

2. Why the classic constructions don't work

The layered architecture illustrates where the problem lies:

- Between the agent action (level 3) and the declaration of intent (level 5).
- This is the point of legally binding attribution (level 4).
- And it is the real breaking point.

Technically, an agent can act. It can formulate an offer, send an acceptance, initiate a payment, place an order. But legally nothing arises from this – as long as it is not clear whose action it is. The agent's action is first and foremost only a technical process. It only becomes legally relevant when it is attributed to a recognized legal entity. And it is precisely this attribution that is still unclear. As already explained, civil law knows several

figures for attributing the actions of others.³¹ But all of them fail with AI agents along the layer model from level 3:

Figure	Prerequisites	Why she fails with the agent
Messenger (§ 120 BGB)	Transmission of a third-party declaration of intent.	The agent does not transmit an external explanation – he creates his own.
Representative (§ 164 BGB)	Submission of one's own declaration of intent in the name of a third party.	The agent has no will of his own – he cannot make a declaration.
Assistant in performance (§ 278 BGB)	Fault of the assistant.	The agent is not culpable.
Assistant Agent (§ 831 BGB)	Order for a performance.	Exculpation possible – and only tortious liability.

All four characters presuppose something that the autonomous agent does not have: a will of his own or a fault. The messenger transmits a foreign will. The deputy forms his own. The assistant in performance and the assistant agent presuppose fault. The AI agent has neither one nor the other.

 The classic constructions fail not only because of a lack of will. They also fail because of a second problem: the visibility of the internal directive. Anyone who grants a power of attorney to an agent defines not only the external framework (what the agent is allowed to do), but also internal boundaries (what the agent is supposed to do). If the business partner knows these internal boundaries, he is no longer in good faith – and the Apparent Authority³² does not apply. This is not a problem among humans: internal instructions are not recognizable to the third party. It is problematic with agentic commerce: In many cases, a verifiable credential is completely legible for anyone to whom it is presented. Anyone who sees the internal directive knows the limit – and can no longer invoke good faith. If the dealer-agent knows limits ("Don't buy a BMW, maximum €50,000"), but the external power of attorney is wider ("May buy cars up to €100,000"), he could even manipulate the offer.

The solution lies in selective disclosure, e.g. via SD-JWT VC.³³ The agent discloses to the business partner only the parts of the power of attorney that are relevant (for the Apparent Authority) – not the internal instructions.

³¹ for details, see the author's script on the Basic Knowledge AI Law, para. 326 et seq.

³² see on the theory of Apparent Authority as well as on external and internal relations, among others. <https://www.mtrle-gal.com/wiki/rechtsscheintheorie/>

³³ on this, among other things: <https://docs.walt.id/concepts/digital-credentials/sd-jwt-vc>

3. Why the analogy doesn't help either

One could try to apply the figure of the deputy by analogy. But an analogy presupposes a gap in the provision that is contrary to plan – and it must not lead to the legislature being replaced by the judge. A provision that regulates the attribution of agent acts cannot be derived from the wording of § 164 of the Civil Code. It is not a case of interpretation, but a case of setting (see also Section VI.2).

Incidentally, even if the analogy were to be affirmed, the problem of the declaration of intent would remain. This is because Section 164 of the German Civil Code requires that the representative makes his own declaration of intent. The agent cannot do this – he has no will in the legal sense.

4. The vulnerability is not technical

At this point, a distinction is important that is often blurred in the current discussion: The gap is not technical. It is legal.

The technical infrastructure is well advanced on several levels:

- Level 1: eIDAS 2.0, EBW, EUDI Wallet create the identity basis.
- Level 2: MCP and A2A create the basis for communication.
- Level 3: The US standards ACP, AP2, UCP create the operational commerce basis.

What is missing is the normative basis including the bridge protocols as the fourth level:

- A legal norm that determines when the agent's act is considered a declaration of intent by the legal entity.
- What is needed are bridge protocols for power of attorney and proof of transaction – they create the documentation basis (e.g. LDM, TAM & TAA, see Part 2 in detail).
- This combination of both elements is not a technical specification. It is a constitutive act – an act of law-making.

5. What this means

Without the combination of legal norm and bridge protocols, the layered architecture remains incomplete. It carries up to level 3.

But it does not carry beyond that:

Without a constitutive legal norm	With a constitutive legal norm
No declaration of intent (level 5).	Declaration of intent is deemed to have been made.
No contract (level 6).	Contract is concluded.
No obligation (level 7).	Obligation arises.

No duties (level 8).	performance obligations arise.
No warranty for defects (level 9).	§§ 434, 437 BGB apply.
Only enrichment law (§ 812 et seq. BGB) or c.i.c. (§ 311 para. 2 BGB).	Full contract law.

The gap has been identified: there is a lack of a legal norm that allows the agent's action to be regarded as a declaration of intent and is flanked by bridging protocols, thus opening up the entire chain of legal binding. Everything else – identity, authorization, documentation, trust – is necessary, but not sufficient on its own.

Takeaways: The legal layer model

- **Nine levels:** From legal entity (1) to liability (9). Each level presupposes the previous one.
- **The gap lies at level 4:** the attribution – the transformation point between agent action (3) and declaration of intent (5).
- **Technology skips levels 3–5:** it begins at level 2 (identity) and jumps to level 6 (contract). Attribution is assumed – but it is not self-evident.
- **Classic characters do not apply:** messenger, deputy, Assistant in performance, Assistant in performance – all of them presuppose will or fault.
- **Selective disclosure as a bridge:** The external scope is disclosed, the internal scope remains hidden – via selective disclosure.

VI. Four precursors, one gap

Before outlining a draft legal norm, it is worth taking a look at what already exists on the normative side. The thesis of this expert opinion is not that there are no regulations at all. It is that there is no sufficient regulation.

1. Existing International Regulations

There are at least four legal norms that regulate the attribution of "automated" actions. Each of the legal norms regulates this more precisely than the previous one in chronological terms:

Legal Norm	Year	Country	Type	Regulatory content
UETA § 14	1999	USA ³⁴	National law (states)	<i>"A contract may be formed by the interaction of electronic agents."</i>
ETA Sec. 15	2010	Singapore ³⁵	National law	<i>"A contract formed by the interaction of automated message systems shall not be denied validity [...] on the sole ground that no natural person reviewed or intervened."</i>
Art. 48 ECL	2018	China ³⁶	National law	<i>"The acts of the parties [...] using automated information systems to conclude or perform contracts shall have legal effect on the parties."</i>
Art. 7 MLAC	2024	UNCITRAL ³⁷	Model Law (Soft Law)	<i>"As between the parties to a contract, an action carried out by an automated system is attributed in accordance with a procedure agreed to by the parties. If paragraph 1 does not apply, an action carried out by an automated system is attributed to the person who uses the system for that purpose."</i>

The line is clear: from contract formation (UETA) to the lack of human verification (ETA) and attribution (ECL) to the general model rule (MLAC). Each norm is more precise than the previous one. Each one moves closer to what this report described as Level 4.

However, even a closer look at the legal definition in UETA § 2(6) reveals the technological gap. UETA defines an "electronic agent" as:

"a computer program or an electronic or other automated means used independently to initiate an action or respond to electronic records or performances in whole or in part, without review or action by an individual."

³⁴ Text of the UETA at <https://euro.ecom.cmu.edu/program/law/08-732/Transactions/ueta.pdf>

³⁵ Text of ETA at <https://sso.agc.gov.sg/act/eta2010>

³⁶ Text of the ECL at https://en.wikisource.org/wiki/Electronic_Commerce_Law_of_the_People%27s_Republic_of_China

³⁷ Text of the MLAC at <https://uncitral.un.org/en/mlac>

This definition describes a deterministic *trigger-response model*³⁸ (e.g. EDI or smart contracts): a pre-programmed script triggers an action without human intervention at that moment.

Singapore's law also illustrates the crucial blind spot: Art. 15 ETA speaks of "automated message systems". These are EDI, if-then scripts, deterministic automatons. In 2010 – based on the UN Convention of 2005 – the legislator saw and solved the problem of the lack of human examination. But it related it exclusively to systems that exchange messages, not to systems that make autonomous decisions.

2. First and second generation of automation

An autonomous AI agent is not a messaging system. It does not receive a predefined message and forward it. It receives a target, independently plans intermediate steps, chooses between alternatives and acts. Its behavior is not deterministic, but probabilistic. It can be emergent – i.e. produce results that the operator did not foresee.

In this respect, the aforementioned norms address the attribution of "automated systems". However, these refer exclusively to the first generation of automation – to rule-based, deterministic systems:

- A transfer to modern agentic commerce fails due to a comprehensive generational leap (see following table).
- The probabilistic autonomy and emergence of today's AI agents was not even foreseeable at the time of the legislation.
- Consequently, the legislature could hardly anticipate this case of application, which created an unplanned regulatory gap for the second generation of automation. However, this cannot be closed by analogous application of these norms, because the regulated cases are not comparable in terms of meaning and purpose, but are essentially different (autonomy vs. automation).

Feature	"Automated System" (Generation 1)	"Autonomous AI Agent" (Generation 2)
Logic	Deterministic (rule-based, if-then logic, EDI, scripts).	Probabilistic (Goal-Oriented, Reasoning, LLM-based).
Behavior	Predictable. A defined input always leads to the same output.	Emergent. The system can choose unforeseen ways to achieve goals within the mandate framework.
Decision	No decision of its own. It only executes pre-programmed commands.	Makes independent decisions based on context and consideration (e.g. dynamic price comparisons).
Reason for attribution	Humans wrote the code exactly like this. The attribution is trivial.	Man has only given the <i>framework</i> (LDM) and the <i>goal</i> . The attribution requires a new dogmatic figure.
Risks	Programming error, failure.	Hallucination, Prompt Injection, Algorithmic Collusion, Model Drift.

³⁸ e.g. <https://nhimg.org/glossary/deterministic-trigger/> and <https://www.opentext.com/what-is/electronic-data-interchange>

The attribution logic of a rule-based system ("The operator is liable because he wrote the code that way") is something different from the case when a system acts emergently. If an AI agent exceeds the mandate due to a hallucination or manipulation (prompt injection), it no longer acts as a mere "automated system", but as an autonomous actor. The legal norms therefore offer neither direct nor analogical answers to the specific phenomena of the second generation: to the possibility of overstepping the mandate through emergence, e.g. the protection of the Apparent Authority in selective disclosure or the burden of proof in model drift.

The four predecessors thus regulate the first generation of automation alone. The second generation – autonomous, learning and negotiating agents – is not addressed. Therefore, none of these legal norms is sufficient. Why is summarized below.

a) They are either soft law or national law.

The MLAC is a model law. It is not a positive law. It recommends, it does not bind. It must also be implemented locally – by national legislators. UETA and Art. 48 ECL are positive law – but they only apply in their respective jurisdictions. None of them is an internationally compatible legal norm.

b) They have been designed for a different generation of technology.

Work on MLAC began in 2021 – two years before AutoGPT and BabyAGI produced the first autonomous AI agents. UETA dates back to 1999, Art. 48 ECL from 2018, Art. 15 ETA from 2010. All four were designed for rule-based, deterministic systems: smart contracts, EDI, automated messaging systems. They address exclusively the first generation of automation. The second generation creates significantly different challenges.

As Art. 48 of China's E-Commerce Law (2018) shows, the legislation ties in with the concept of the "automated information system". This addresses the first generation of automation (EDI, If-Then-Scripts), in which humans retain full control and predictability. For the second generation – autonomous, emergent AI agents – this term does not provide a dogmatic basis for attribution.

c) They do not know the specific phenomena of autonomous agents.

None of them govern the following non-exhaustive challenges:

- Mandate exceeded due to emergence
- Hallucination as a lack of will
- Selective Disclosure (External Scope vs. Internal Scope)
- Model Drift after signing the contract

- Algorithmic Collusion
- Apparent Authority for autonomous agents

The gap can thus be clearly narrowed down: There are already legal norms. There are even international model rules. But there is no legal norm that is intended for the agentic AI of the present and future. The three precursors show that the international community has only seen the problem of deterministic predecessors.

And this is exactly where the decisive distinction comes in:

- What these standards and technical protocols do is "proof". What they don't do is "validity".
- "Validity" requires a legal norm and matching bridging protocols that combine law and technology in a targeted manner.

3. The Crucial Distinction: Proof vs. Validity

This also identifies the distinction that is central to the further content of the report:

Proof	Validity
What happened?	What applies legally?
Provided by technical standards.	Generated by a legal norm.
eIDAS, EBW, LCP, UCP, ACP, WP2.	The lack of a constitutive legal norm.
Available.	Missing.

The technical standards can prove that an agent acted, that he was authorized, that the transaction is documented. They cannot justify that the act is legally valid. Even the four normative predecessors (see Section VI) do not change this: they regulate attribution – not validity. This justification is the constitutive act – and it is missing.

And this is exactly where the further analysis comes in. It shows what the technology already does, where the error in thinking lies, what the gradation of the legal binding looks like and what a legal norm would have to do to close the gap.

The following overview shows what already exists or is being implemented with regard to which topic (marked in green) and what is still missing (marked in red). It is important to understand that the technical protocols and infrastructures will not become obsolete because of what is still required – on the contrary: the aspects marked in green are even important prerequisites for the still missing bridging protocols and a legal norm to be effective.

Theme & Key question	Technology & Infrastructure present – provides the proof	Bridging Protocols – translate the proof into legal form	Legal norm missing – creates the validity
Identity & Mandate: Who acted in whose name and with what authority? (Levels 1–2, 4)	eIDAS 2.0, EBW, EUDI Wallet, LPID; AP2 mandates, Visa TAP, KYA	LDM (Legal Delegation Mandate): digital power of attorney as Verifiable Credential; Selective Disclosure separates External Scope from Internal Scope	The provision provides: Attribution to the legal entity; Mandate as a norm of attribution; Apparent Authority in the event of an incognisable exceedance
Transaction & Agreement: What was done and agreed? (Level 3, 6–8)	ACP, UCP (Commerce); AP2, Agent Pay, x402 (Payment); MCP/A2A (Communication); LCP (Terms, ex post)	TAM (Transaction Acceptance Manifesto): documents the action within the mandate framework; Authorized Representative for the written contract within the meaning of the GoB, customs law and money laundering prevention	Norm provides: The machine-generated act of declaration is considered a declaration of intent by the legal entity → contract with full chains of obligations
Trust & Responsibility: Who can be trusted – and who will stand up? (Level 9)	Visa TAP, KYA, M-Trust, Zero Trust	TAA (Trusted Agent Attestation): Certificate of compliance with a recognized agent specification	Norm links attribution and exculpation: TAA as a basis of trust for the legal entity in the event of a technical malfunction
Attribution & Validity: What applies legally? (Levels 4–5)	– none: no protocol can create validity	LDM + TAM + TAA provide the attributable, conclusive facts	Constitutive legal norm: the constitutive act; transforms proof into validity

 In the second part of the report, the above-mentioned international legal norms and technical protocols will once again play an important role. They help to determine exactly how a legal norms would have to be formulated in detail in order to be valid in conjunction with bridge protocols.

Takeaways: Four precursors, one gap

- **International legal norms, one line:** UETA (1999), ETA Singapore (2010), Art. 48 ECL China (2018), UNCITRAL MLAC (2024) – each more precise than the previous one.
- **Generation 1 vs. Generation 2:** All four govern deterministic, rule-based systems – non-autonomous, probabilistic agents.
- **Six unsolved phenomena:** overstepping of mandate through emergence, hallucination, selective disclosure, model drift, algorithmic collusion, Apparent Authority in autonomous agents.
- **Proof vs. validity:** Technical standards provide evidence ("What happened?"), legal norms create validity ("What applies legally?").
- **The constitutive legal norm is missing:** it is the act of law-making that transforms the proof into validity.

VII. The current state of the technology

Technical development is well advanced. In recent months, a number of protocols, standards and initiatives have emerged to enable agentic commerce. They cover different layers – and they are often presented as "the solution" in the public debate. A closer look shows that they are necessary building blocks. But they are not sufficient in terms of "validity".

1. The US standards: Commerce, Payment, Legal Context

In the USA, various protocols have been created that are intended to technically enable agentic commerce, including ACP³⁹, UCP⁴⁰, AP2⁴¹ and LPC⁴².

- What these protocols do: They standardize the exchange of data and actions. They enable agents to find products, compare prices, initiate payments and process orders. They create the technical infrastructure for Agentic Commerce.
- What they do not do: They do not say whether the agent's action is legally binding. They do not say to whom it is attributed. They do not say what legal consequence occurs. They answer the question: "What happened?" – but not the question: "What applies legally?"

Protocol	Author	Function	Status
ACP (Agentic Commerce Protocol)	OpenAI / Stripe	Checkout: AI agents communicate with e-commerce platforms; product catalog, checkout sessions, payment processing.	September 2025; first implementations (ChatGPT Instant Checkout).
AP2 (Agent Payments Protocol)	Google	Payment: Cryptographic mandates for authorization; Checkout and payment mandates in open/closed stages.	September 2025; v0.2 (April 2026) with "Human Not Present" flow; donated to FIDO Alliance.
UCP (Universal Commerce Protocol)	Google	Commerce: Unified language for the entire commerce journey – discovery, cart, checkout, post-purchase.	January 2026; with Shopify, Etsy, Wayfair, Target, Walmart.
LCP (Legal Context Protocol)	AAA / Integra Ledger	Legal Context: Findability of contract terms, governing law, dispute resolution links.	June 2026; with Google, IBM, Circle, Wayfair.

LCP is the most ambitious attempt to address the legal dimension. It makes contract terms findable, verifiable and signable. But it presupposes that the legal relationships have already been determined. It is an instrument of dispute resolution ex post – not of

³⁹ Note: There are several ACP protocols, but only this one applies to Agentic Commerce <https://www.agenticcommerce.dev/>

⁴⁰ more about UCP as Universal Commerce Protocol: <https://ucp.dev/>

⁴¹ on the Agent Payment Protocol and others <https://cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol>

⁴² More about LCP as a dispute resolution protocol at <https://legalcontextprotocol.org/>

the constitution ex ante. It helps if something has already happened. It does not ensure that validity is established.

2. Payment service provider initiatives

In parallel with the protocols, the major payment providers such as VISA,⁴³ Mastercard⁴⁴ and ANT-International have developed ⁴⁵ their own initiatives and also Know-Your-Agent⁴⁶:

Initiative	Author	Function
Trusted Agent Protocol (TAP)	VISA	verification of agent signatures; differentiation between trusted agents and malicious bots.
Agent Pay / Agent Pay for Machines (AP4M)	Mastercard	Real-time authorization, orchestration, and settlement of agent payments; for high-frequency, low-latent micro-transactions.
Agentic Mobile Protocol (AMP)	ANT International	Agent payments in the mobile ecosystem.
Know-Your-Agent (KYA)	VISA, Mastercard, Ant International	interoperability of agent protocols; common principles for trust and governance.

- What these initiatives do: They create trust at the payment level. They verify that an agent is authorized, that a signature is valid, that a payment may be initiated. They establish a "Know Your Agent" – analogous to "Know Your Customer".
- What they do not do: They do not say whether and under what conditions an authorization is legally effective. They do not say whether the payment results in a contract. They do not say who is liable in the event of an error. They create technical trust – not legal validity.

3. The European infrastructure

With eIDAS 2.0 and other elements, the European Union is building its own infrastructure for digital identities and related legal topics. Especially with regard to wallet solutions, two similar-sounding variants can be distinguished:

The EUDI Wallet (Personal Wallet):

- Focus: B2C & B2G.
- Target group: Natural persons (citizens).

⁴³ <https://developer.VISA.com/capabilities/trusted-agent-protocol>

⁴⁴ <https://www.mastercard.com/global/en/news-and-trends/press/2026/june/mastercard-launches-agent-pay-for-machines.html>

⁴⁵ <https://www.ant-intl.com/en/news/detail/?id=ant-international-launches-open-sourced-agentic-mobile-protocol-to-drive-ai-commerce&tag=AI>

⁴⁶ <https://www.entrust.com/resources/learn/know-your-agent>

And the European Business Wallet (EBW / EUBW):

- Focus: B2B & B2G.
- Target group: Legal entities (companies, organizations).

In addition, there are other plans (overview):

Element	Function	Status
EUDI Wallet European Digital Identity Wallet	Digital identity wallet for natural Persons (e.g. for account opening, online ID, proof of age)	Member states must provide at least one wallet by the end of 2026.
EBW European Business Wallet	Digital identity solution for legal entities and economic actors (e.g. automation of KYB and supply chain processes, signing of contracts)	Council adopted negotiating position in June 2026; political agreement targeted by the end of 2026.
LPID Legal Person Identification Data	Standardised identity data for legal entities.	Issuance by registry authorities planned from 2027.
PoA Power of Attorney	Standardised digital power of attorney for cross-border corporate law acts (incorporation, branch, merger, division, change of legal form).	Directive (EU) 2025/25; transposition by 2027.

- What this infrastructure does: It creates the identity basis for Agentic Commerce. Companies can identify themselves digitally, prove representation authorizations and exchange documents. The EBW will enable the delegation of legal power to representatives.
- What it does not do: It does not say whether an agent's action is legally valid. It creates the prerequisite for attribution – but not the attribution itself.

 The EU Digital PoA is a first step towards digital power of attorney in EU law. However, it is limited to register-related corporate law acts – it is not relevant for everyday commercial transactions and thus for agentic commerce. The decisive factor here is a distinction that is often blurred in the current discussion: "Digital" describes a carrier medium. It says nothing about whether a declaration is effective. The EU Digital PoA regulates the form of the power of attorney – its digitality, its machine readability, its cross-border recognition. It does not regulate the effectiveness of the declaration that is made on the basis of it. But it is precisely this distinction that is central to agentic trade. The technical infrastructure – eIDAS 2.0, EBW, EUDI Wallet, EU Digital PoA – creates technical prerequisites. What is missing is normative recognition: the rule as to when a declaration made by AI agents on this technical basis is considered or recognized as a declaration of intent.

4. The international discussion

The gap is also recognized at the international level, both with regard to NIST⁴⁷ and the Asian M-Trust standard⁴⁸:

Source	Statement
NIST (February 2026)	"How can an agent prove its authority to perform a specific action?" – The concept paper asks about identity and authorization, not about legal validity.
M-Trust (GSMA/CAICT, September 2026)	extends Zero Trust to include Multi-party Trust, Intent-aware Trust, and Cross-domain Trust. Addresses authority, not attribution.

- What these discussions achieve: They show that the gap in binding authorization with legal effect is recognized internationally. However, they only create awareness of the problem.
- What they don't do is they don't provide the solution. They describe the problem – but they don't close the gap.

5. The common gap

All of these initiatives – US protocols, payment provider initiatives, EU infrastructure, international discussions – address important layers:

Layer	What exists	What's missing
Identity	eIDAS 2.0, EBW, EUDI Wallet	–
Authorization	AP2, VISA / TAP, Mastercard Agent Pay	–
Commerce	ACP, UCP	–
Payment	AP2, Agent Pay, x402	–
Legal Context	LCP (ex post)	–
Attribution	–	The constitutive legal norm.

The gap is thus clear:

- There is no normative basis for the latest generation of automation – autonomous decision-making.
- It is a legal norm that determines that the agent's act is considered a declaration of intent by the legal entity.
- This norm is not a technical specification. It is a constitutive act that must be flanked by bridge protocols.

⁴⁷ to NIST: <https://www.nccoe.nist.gov/sites/default/files/2026-02/accelerating-the-adoption-of-software-and-ai-agent-identity-and-authorization-concept-paper.pdf>

⁴⁸ see M-Trust with classification and further links: <https://spherity.github.io/spherity-research/beyond-zero-trust-m-trust-authorised-agent-actors.html>

And this act is missing in all initiatives. The US standards presuppose that the declaration of intent somehow comes about – and leave the clarification to the law. The EU infrastructure creates the identity basis – but not the attribution. The international discussion recognises the gap – but it does not close it.

6. The error in thinking and its causes

From this point of view, technical standards and legal norms are regularly lumped together. The episode:

- It is believed that a protocol or a technical standard could replace a legal norm.
- Or you may believe that technical verification is already a legal obligation.

There are several possible reasons why the thinking error is so common:

Reason	Explanation
Technical precision conceals legal loophole.	A well-specified protocol acts like a solution.
Industry has an interest in equation.	Anyone who sells a protocol benefits from the assumption that it is enough.
Law is underrepresented in the discussion.	On LinkedIn and in tech circles, almost only technicians are discussing.
The standard is missing – so the gap is being filled technically.	What is not regulated by law is "solved" by standards and protocols.

 Technical development is not an end in itself. It is driven by global players – large platforms, payment networks and tech companies. Their interest is not the order or optimization of trade, but the market share. The conviction: Whoever sets the (technical) standard first determines the rules of the game – the legislator should react if necessary. The law has a diametrically opposed meaning to this. It does not create market share – it balances. It not only protects the fast – it also protects the weaker. It does not centralize – it decentralizes. And it does not work in weekly stages – it works in years, because it first has to grasp the conflict situations arising from an innovation.

This asymmetry explains why the gap has remained open: Agentic AI has only been around since 2023: Initially, it was frameworks such as AutoGPT and BabyAGI. The major international players immediately integrated the principle into their own products and services. The question of the law was probably hardly "overlooked" because no one could have recognized it, but because the drivers probably have no great interest in closing it. Anyone who sets a technical AI market standard today usually does not want a legal norm to contain it tomorrow. And for this reason, the question of the legal norm also becomes a question of power: Who determines the

rules of the game of the B2B agentic commerce – the technical drivers or trade in cooperation with suitable legislators?

In the end, the "error in thinking" is not really a coincidence at all – and it is not a coincidence either. It is the consequence of a discussion in which the legal perspective is missing because it does not see the circuit breaker or only wants to accept it late.

The missing standard so far is and remains the circuit breaker:

- It creates trust.
- Trust creates participation.
- Participation creates a market.

Takeaways: The current state of the art

- **US standards:** ACP, AP2, UCP, LCP – they enable exchange and create documentation, but they do not create validity.
- **Payment providers:** Visa TAP, Mastercard Agent Pay, KYA – they create trust at the payment level, but not a legal obligation.
- **EU infrastructure:** eIDAS 2.0, EBW, EUDI Wallet, LPID – they create the identity basis, but not the attribution.
- **International discussion:** NIST, M-Trust – they recognize the gap, but they don't close it.
- **The error in thinking:** Technical standards and legal norms are confused. A protocol cannot replace a legal norm.

VIII. The required legal norm

The analysis has shown where the gap lies: at level 4, the attribution. And it has shown that the classic figures of civil law – messenger, deputy, Assistant in performance, Assistant in performance – cannot close this gap.

So a legal norm is needed. The term is to be used precisely. A legal norm in the sense meant here is positive law – it is:

- directly binding – it applies, it does not have to be implemented;
- justiciable – it can be enforced in court;
- binding on the authorities – it binds tax, customs, money laundering prevention;
- third-party effect – it works between the parties and also vis-à-vis third parties.

1. Requirements for legal norms

Not every rule that looks like a legal norm is a norm in the required sense, because there is a hierarchy of regulatory instruments:

#	Category	Commitment	Results
1	Positive law	Directly binding	applies – without an act of implementation
2	International treaty	binding for signatory states	valid after ratification
3	Model law	non-binding	Recommendation to legislators
4	Soft Law / Principles	non-binding	Orientation

A model rule is therefore not a legal norm. It is a proposal for a norm. It has no direct legal effect. It must be implemented by national legislators. It can be a role model, it can convince, it can prepare. But it does not apply.

And a constitutive legal norm – this is the second clarification – is a norm that creates a legal consequence, not one that declares an existing legal consequence. It does not say: *"We clarify what applies anyway."* It says: *"What does not apply without this legal norm should now apply."* It establishes law. It creates validity.

The third clarification concerns the relationship between proof and validity. A technical proof – a signature, a log, a credential – can prove that something has happened. It can document who acted, when, with what authorization, and with what result. However, it cannot justify that what happened is legally valid. Between proof and validity lies a constitutive act – the setting of a legal norm.

 The three clarifications together reflect the guiding thesis of this expert opinion: So far, there is no constitutive legal norm – in the sense of positive law – that attributes the machine-generated explanatory act of an AI agent to the legal entity. The layered architecture, in turn, provides the answer to what such a norm would have to look like.

2. Create an attribution rule

The legal norm must create a rule of attribution. It must stipulate that the machine-generated declaration act of an agent is attributed to the legal entity under defined conditions. This is the constitutive act. Without it, the agent's act remains a technical process without legal effect.

The decisive factor here is the dogmatic construction: It is not a matter of attributing a will to the agent or fabricating a declaration of intent. An AI agent has no will in the legal sense – and he cannot make his own declaration of intent. What is missing is not the "will" of the agent, but the legal assignment of his declaration act to the legal entity.

The legal norm must therefore be formulated in a technology-neutral and functional way. It must not refer to specific models or protocols – otherwise it will become obsolete with the technology. Rather, it must describe the conditions under which a machine-generated declaration act by an autonomous AI agent is considered a declaration by the legal entity.

 As has already been shown, there are certainly legal norms that affect partial aspects of what is being negotiated here: automation, electronic declarations, digital powers of attorney, automated contract formation. They come from different legal systems and eras – from UETA (1999) to Art. 48 of the Chinese E-Commerce Law (2019) to the UNCITRAL Model Law on Automated Contracting (2024). However, none of these are legal norms that could take into account the opportunities and challenges of agentic AI. They were designed for a world where automation was rule-based and deterministic. They address the first generation of automation.

The second generation – autonomous, learning, emergent agents – raises questions that are not answered in any of these norms: overstepping the mandate through emergence, hallucination as a lack of will, selective disclosure, model drift, algorithmic collusion. These questions are not old – they are new. And they are not solved. What is missing is therefore not just any legal norm. What is missing is a norm that is intended for the agentic AI of the present and future.

3. Enabling the connection to the obligation

The provision may not be limited to the declaration of intent. It must also regulate the connection to the contractual relationship. This is because Section 311 (2) of the German Civil Code shows that an obligation can also arise without a contract – through contract negotiations, contract initiation or similar business contacts.

If the agent acts but no contract is concluded, it must be clear what legal relationship arises between the parties. The legal norm must therefore not only fictitiously the declaration of intent, but also determine the legal consequence – the obligation with its duties of protection and consideration.

4. Map the entire chain

The legal norm must cover the entire chain: from creation to transformation to attribution. This means:

Level	What the legal norm must regulate
1. Legal Entity	Who is the legal entity to which it is attributed?
2. Agent Identity	How is the agent's identity proven?
3. Agent action	What did the agent do?
4. Attribution	To whom is the action attributed?
5. Declaration of intent	Is it considered a declaration by the legal entity?
6. Contract	Will a contract be concluded?
7. Obligation	What are the legal consequences if no contract is concluded?

 The discussion about "autonomy vs. control" is often linked to Art. 14 of the AI Regulation (Human Supervision). However, it should be noted that Art. 14 EU AI Act only applies to high-risk AI systems (e.g. agents in HR recruiting, lending or critical infrastructures). The legal norm is therefore not directly applicable to classic B2B purchasing or standard services.

The technical translation (for high-risk AI): However, if a company uses a high-risk agent, the architecture proposed here consisting of LDM (Legal Delegation Mandate) and TAM (Transaction Acceptance Manifesto) becomes the perfect technical implementation of Art. 14 AI Regulation (more on this in Part 2):

- Ex-ante limits (Art. 14 para. 4): The LDM codifies the narrow limits within which the system is allowed to act autonomously and defines the escalation points (human-in-the-loop).
- Ex-post monitoring (Art. 14 para. 5): The TAM serves as a complete, forgery-proof audit trail that makes human supervision effective and verifiable.

For the B2B standard beyond high-risk AI, this nevertheless leads to the following consideration: Even if Art. 14 AI Regulation does not apply to the mass of B2B transactions, there is no alternative to the LDM model. It translates the general civil law organizational obligations (avoidance of organizational fault, Section 311 (3) of the German Civil Code) and internal company compliance requirements into machine-readable code. The LDM is thus not only a compliance tool for the exceptional case of high-risk AI, but also a best-practice standard for corporate digital responsibility and risk management in agentic commerce.

5. The limits of authority in the case of autonomy

The previous requirements concern the existence of an attribution rule. This section concerns its scope. This is because an attribution rule that only regulates the normal case – the agent acts within his mandate – is incomplete. It must also cover the borderline cases that are typical for autonomous agents.

a) The first borderline case: the overstepping of mandates due to emergence.

A mandated agent acts within a framework that has been given to him. But this framework is not to be confused with a rigid housing. It defines boundaries, not paths. Within these limits, the agent can make decisions that the operator did not foresee – not because the agent goes beyond the scope, but because he fills it. A price that changes due to dynamic market development. A supplier that suddenly fails. An order that becomes cheaper through bundling. These cases are not an overstepping of the mandate. They are the exercise of a mandate. And they must be attributed – even if the result is unexpected. The principle that the MLAC also formulates applies here: unexpected is not necessarily inadequate. Otherwise, every autonomous agent contract would be destroyed retrospectively by the claim: *"This is not how I had imagined it."*

b) The second borderline case: the exceeding of the outer framework.

If the agent exceeds the outer boundary, a distinction must be made:

Condition	Description	Legal consequences
1. Within the Mandates	The agent acts in the External Scope and the Internal Scope.	Attribution. The declaration shall be deemed to be a declaration by the legal entity.
2. Outside the external scope, recognizable to the third party	The agent crosses the outer frame – the business partner can detect the exceedance.	No attribution. The third party is not worthy of protection.

Condition	Description	Legal consequences
3. Outside the external scope, not recognizable to the third party	The agent crosses the outer frame – the business partner cannot detect the exceedance.	Apparent Authority. The declaration is considered a declaration by the legal entity – the third party is worthy of protection.

c) The third borderline case: the visibility of the internal directive.

States 2 and 3 require that the third party knows the external scope – and not the internal scope. Because only then can he recognize whether the agent exceeds the outer scope. And only then is he worthy of protection if he does not recognize it.

This is precisely where the technical consequence of the selective disclosure requirement lies: The agent only discloses the external scope to the business partner – not the internal scope. Anyone who sees the internal directive knows the boundary – and can no longer invoke good faith. Selective disclosure is therefore not just a data protection issue. It is the prerequisite for the Apparent Authority to take effect at all.

d) What this means for the legal norm

A provision of attribution that does not address these three borderline cases is incomplete. It must:

- Clarify that unexpected results will be attributed within the mandate.
- Rules as to when an exceeding of the external framework makes the attribution no longer applicable – and when it triggers the Apparent Authority.
- Ensure that the internal instruction is not recognizable to the third party – otherwise the Apparent Authority will not work.

And there is one thing it must avoid: the blanket formula *"The agent has acted, so the operator is liable."* This formula does not do justice to the complexity of autonomous systems. It would make the operator responsible for any emergent behavior – and thus make agentic commerce economically impossible.

6. The requirements for a sustainable legal norm

This results in six requirements for a viable norm:

	Requirement	Justification
1	Technology-neutral	The norm must not be tied to a specific model.
2	Functional	She has to describe what the agent does – not how.
3	Connectable	It must be linked to existing legal concepts (declaration of intent, contractual relationship).
4	Internationally compatible	It must be compatible with European and international standards and legal norms.
5	Open to the future	It must be adaptable to new technologies without being changed.

6.	Technical Interoperable	The norm must recognise the formats that are established in the European environment – in particular SD-JWT VC for selective disclosure.
----	--------------------------------	--

 The fourth requirement (international connectivity) concerns the legal system. The sixth concerns technology. Both are indispensable and are directly interdependent, because a standard that is legally compatible but technically isolated will not be used in practice.

7. The vote of confidence

And finally: The legal norm must create trust. Because trust is not a question of faith – it is a question of order. Whoever creates order creates trust.

The norm is therefore not just a legal instrument. It is the prerequisite for agentic commerce to work in B2B. Without it, it remains a de facto exchange relationship that neither side wants – and presumably no authority that needs reliable evidence.

Takeaways: The required legal norm

- **Legal norm = positive law:** Directly binding, justiciable, binding on the authorities, third-party effect. Model laws and soft law are not legal norms.
- **Attribution rule:** The provision must stipulate that the machine-generated declaration act of an agent is attributed to the legal entity.
- **Connection to the obligation:** Even if no contract is concluded, the legal consequence (duties of protection) must be clear.
- **Map the entire chain:** From legal entity (1) to attribution (4) and declaration of intent (5) to obligation (7).
- **Six requirements:** technology-neutral, functional, connectable, internationally connectable, future-proof, technically interoperable.

IX. Implementation options

The analysis has shown that there is a lack of a legal norm. But how does it come about? There are three options – and they are not alternative, but sequential.

1. Option A: Model contractual clause + Choice of Forum

The parties contractually agree that the action of an AI agent is considered a declaration of intent by the party using it. This is complemented by a Choice of Forum – the choice of a jurisdiction or arbitral tribunal that applies the clause.

Advantage	Disadvantage
Immediately implementable.	Only works inter partes.
No legislator necessary.	No protection against third parties.
Creates practice and experience.	Doesn't solve the problem of the autonomous agent that goes beyond the scope.

The model clause is the pragmatic introduction. But it has a dogmatic limit that must be clearly stated – otherwise the impression arises that it is the solution.

The decisive limit is not that it only works inter partes. It lies in the fact that it cannot even achieve what it promises inter partes. The concept of a declaration of intent is a legal category – not a contractual one. What a declaration of intent is is determined by the law, not by the party. The parties can agree on legal consequences. They can establish obligations. They can establish fictions for their internal relationship. But they cannot create a declaration of intent where there is none.

This means that the model clause establishes an obligation under the law of obligations – not a declaration of intent or a contract. It can oblige the legal entity to treat the action of its agent as if it had declared it itself. It can make him liable. It can bind him. But it cannot have the effect that the act is legally considered a declaration of intent – with all the consequences for third parties, authorities and courts.

 A third party who relies on a model clause is thus trusting in something that does not hold at the decisive moment: He delivers goods – and has no contract. He pays money – and has no claim to performance. He disposes – and has no warranty for defects. At most, he has a contractual obligation on the part of the legal entity, which he must prove in the event of a dispute. And which earns him damages – not performance. The model clause thus only creates the illusion of validity – and in many cases that is even more dangerous than no trust at all. This is because it tempts the business partner to get involved in something that does not work. The circuit breaker mentioned in Section VII is not closed by the model clause. It is concealed.

And yet model clauses should be observed: they create practice, experience and facts. And they show where the points of friction lie – before the legal norm comes. The model clause is not a solution, but it is the starting point. Anyone who uses a model clause today knows that something important is missing. Anyone who carries out a pilot today can tell the legislature: *"This is what it looks like. This is how it works – and this is how it doesn't work."*

Now	Later
Illusion of validity	legitimate expectation
Model clause	Constitutive legal norm
Inter partes	Erga omnes
Party autonomy	Legislation
Practice	Validity

The legislator comes at the end – not at the beginning. But it comes. And the model clause is the reason why it comes: it is involuntary proof that the legal norm is missing.

2. Option B: Regulatory sandbox

The legislator creates a temporary, spatially and objectively limited area of exception. Companies test under real conditions – with legal certainty.

Advantage	Disadvantage
6-12 months instead of 3-5 years.	Temporary – not a permanent legal norm.
Creates practical data.	Limited to certain actors.
Legitimation through state participation.	–

The sandbox is the official path to the fast norm. It provides the data that specifies the later norm.

3. Option C: Constitutive legal norm

The legislator issues a permanent legal norm that allows the agent's action to be regarded as a declaration of intent – and regulates the connection to the contractual relationship.

Advantage	Disadvantage
Permanent legal certainty.	3-5 years of legislative process.
Full commitment – not only inter partes.	Political hurdles.
Builds trust.	–

The constitutive legal norm is the goal. It closes the gap permanently.

4. The combination

The most pragmatic strategy is the combination:

Now	Later
Model clause + Choice of Forum	Constitutive legal norm
Sandbox	Permanent regulation
Creating practice	Codifying practice
Setting the facts	Acknowledging facts

The legislator comes at the end and not at the beginning.



Takeaways: Implementation options

- **Option A – Model clause + Choice of Forum:** Immediately implementable, but only inter partes. Establishes obligation under the law of obligations – no declaration of intent.
- **Option B – Regulatory sandbox:** 6-12 months instead of 3-5 years. Temporary, limited, but creates practice data.
- **Option C – Constitutive legal:** Permanent legal certainty, full binding (erga omnes), but 3-5 years of legislative procedure.
- **The combination:** Create practice now (model clause + sandbox), codify later (constitutive legal norm).
- **The legislator comes in the end – but it comes.** The model clause is involuntary proof that the legal norm is missing.



X. Quo vadis?

The analysis has shown that there is a lack of a legal norm. The technology is well advanced. Industry is discussing protocols and standards. Regardless of this, the discussion about trust in AI systems has reached a new level of urgency.

1. The wake-up call

In June 2026, an autonomous AI agent from OpenAI gained unauthorized access to an Australian government portal⁴⁹. The agent was supposed to investigate. When the portal did not release the information, he looked for another way to circumvent the blocks – he did not want to accept the "no". In parallel, Sam Altman and Dario Amodei told the UN Security Council⁵⁰ that the industry needs international standards – for measuring capabilities, assessing risks, ensuring the effectiveness of safeguards and maintaining human oversight.

What this incident shows: Technical autonomy is a reality. Certain agents apparently do not accept a "no". They act beyond the scope. And there is no infrastructure to report, assign or legally evaluate such incidents.

2. The urgency

The technology does not wait. While the discussion about regulation takes years, the agents act in seconds. The Australian incident is not an isolated case – it is a pattern. And the number of incidents is increasing. The BSI Security Study 2026 also shows that legal uncertainty also prevents many companies from using AI⁵¹. In short, there is a need for action not only, but also from a legal point of view.

3. Three paths

There are three ways to deal with the gap:

Path	What happens	Price
Wait and see	The solution is left to the USA – and its standards are imported.	The legal system is also imported.
Regulate	They are waiting for an EU-wide solution.	The technology continues to run – and legal uncertainty is growing.
Constituting	A legal norm is created that allows the agent's action to be regarded as a declaration of intent – where it is possible quickly and confidently.	You move forward – without the certainty that others will follow.

⁴⁹ on this, among other things: <https://www.tagesschau.de/ausland/ozeanien/openai-australien-regierungsportal-100.html>

⁵⁰ More on this e.g. here <https://www.tagesschau.de/ausland/amerika/ki-un-sicherheitsrat-100.html>

⁵¹ Details also in the BSI's Security Study 2026: https://www.bsi.bund.de/DE/Service-Navi/Publikationen/Studien/KI-Unternehmensbefragung_TUEV-BSI/KI-Unternehmensbefragung_TUEV-BSI.html

From the point of view of speed and pragmatism, a model with an international perspective is currently the only way to combine both: sovereignty and connectivity.

Switzerland is not a member of the EU – it can issue its own legal norm without waiting for Brussels. It is neutral – its legal norm can be recognized by both sides. And it has a legal tradition that enjoys trust in international trade.

4. The instrument: Choice of Forum

The decisive lever is Choice of Forum – the free choice of the applicable law and the place of jurisdiction in international B2B traffic.

German, French, Italian – and Asian – companies can choose Swiss law and agree Swiss courts or arbitration tribunals. This is not circumvention of the law. It is party autonomy. And it means that if a Swiss model exists, companies from Europe and Asia could use it, even without their own legislators having to take action.

5. The business channels and final question

The rapprochement between Europe and Asia is not only taking place at the government level. Companies are increasingly talking directly about regulatory issues – faster, more discreetly and more practical than traditional diplomacy. These business-to-business channels are not a substitute for the legal norm. But they are preparing the ground for it.

- The question is not whether the gap will be closed.
- The question is who will close it first – and with what right.

The German legislator is largely excluded from this, as it has to orient itself to EU law when it comes to AI issues. This leaves the following actors, among others – and from a German point of view, it is probably Switzerland that seems to be the most relevant overall:

Region/Jurisdiction	Ability to act	Internationally compatible	Political hurdle
 Singapore	very high	very high	low
 China	very high	Medium	low internal
 USA	high	High	medium/high
 Switzerland	high	very high	medium
 EU	medium	very high	high

6. Conclusion

The development of agent systems has so far been almost exclusively technically driven. However, legal questions of attribution cannot be solved by technical protocols alone and lie outside what technology providers can regulate themselves. They

therefore only come to the fore when scaling – or even if scaling does not succeed because the legally binding trust is missing.

The history of the interplay between law and AI shows:

- Where industries set the rules themselves, they set technical rules that primarily benefit them. This is not morally reprehensible – it is rational.
- But it also means that legal loopholes that benefit locations or economic areas cannot be closed by the tech industry alone.
- The gap outlined here must therefore be closed by legislators. And legislators must know that the ball is in their court.

In view of the existing legal uncertainty, it makes sense to have a legal norm that enables the recognition of agentic declarations under defined conditions. Anyone who appoints agents bears the risks of his decision himself; however, he should not be dependent on uncertainty, which can be avoided by a clear regulation. This also applies to the business partner who trusts in the validity of a declaration.

In the second part of the report, two complementary elements are presented:


- A draft for a civil law norm under Swiss law or MLAC that makes it possible to make declarations by autonomous AI agents legally binding.
- In addition, accompanying drafts for LDM, TAM, TAA and SD. These are bridging protocols to complement already existing technical protocols from US providers and payment service providers such as AP2, UCP, TAP, etc.

	What it does	What it doesn't do
Technology	enables exchange (ACP, AP2, UCP)	does not create validity
Infrastructure	creates identity and authorization (EBW, eIDAS 2.0)	creates no attribution
Bridge Protocols	Translating between law and technology (LDM, TAM, TAA)	need the legal norm to work
Legal norm	creates validity	needs the bridge protocols to be operationalizable

The second part also places great emphasis on the analysis of business aspects of agentic commerce – especially the ROI calculation (with and without legal norm). Among other things, Geoffrey Moore's Technology Adoption Life Curve plays a role⁵² here. The logic behind it is also known as the challenge of "crossing the chasm". The chasm arises when only the early adopters accept a product that is still incomplete or unconfirmed in its real benefit, while the early majority demands a "whole product" (a complete, turnkey

⁵² https://en.wikipedia.org/wiki/Crossing_the_Chasm

product). Depending on how you look at it, more than 10% of all innovations end up in the chasm.



In agentic commerce, the chasm in the layer model of Section V. runs along level 2 (Unjust Enrichment / c.i.c.) and level 3 (contract).

- To the left of the chasm (tech view): The AI agent has ordered the goods, the payment gateway (AP2) has charged the credit card, the LCP has transmitted the terms and conditions. The transaction is technically perfect.
- To the right of Chasm (CFO view): The company asks: "How much do I have to invest?" and "Where is the ROI?" In addition, the question is asked: "Where is the contract for the tax office? Who is liable for defects (§ 377 HGB)? What happens during customs audits?" The answer so far is: "We have the right of enrichment and hope that nothing goes wrong."
- Agentic commerce in B2B will probably only come via Chasm when everyone in the retail network benefits from it.

In part two, it is explained that agentic commerce is unlikely to fail because agents cannot pay. The biggest challenge will be whether and how companies can audit the transactions and thus scale the network.

Takeaways: Quo vadis?

- **Need for action:** An autonomous AI agent gained unauthorized access to a government portal – technical autonomy is a reality.
- **Three ways:** wait and see (import US standards), regulate (wait EU-wide), lead the way (own legal norms and bridge standard).
- **Switzerland as the most pragmatic jurisdiction:** Non-EU member, neutral, compatible – companies can choose Swiss law via Choice of Forum.
- **The Chasm:** Agentic commerce hardly fails because of payment processing, but because of auditability – without a legal norm, there is no leap over the Chasm.

XI. Open points

This report does not claim to be able to conclusively assess the legal challenges surrounding agent B2B legal transactions. Rather, it aims to make a relevant legal gap visible – and to name the questions that arise as soon as the details of this gap become visible. The following points are considered here as an expandable research agenda.

The questions can be grouped into three clusters:

1. Attribution and declaration of intent

#	Question
1	Scope of attribution: When exactly is an agent's act attributed to the legal entity – and when not?
2	Exceeding the mandate: What applies if the agent acts within the technical parameters, but the economic result deviates from the operator's will?
3	Apparent Authority and protection of legitimate expectations: What applies if the business partner trusts in the effectiveness of an agent's action, even though internal authorization was actually missing?
4	Burden of proof and evidentiary effect: What technical evidence must be provided in the event of a dispute? What is the significance of logs, signatures, proof of delegation and policy records?

2. Systemic risks

#	Question
5	Algorithmic Collusion: How does an agent detect "statistically anomalous"? How is it escalated? Who is liable?
6	Kill switch in operation: Does Switzerland need a counterpart to Art. 14 AI Regulation? Who is allowed to interrupt? Under what conditions?
7	Model drift and updates: Is the TAM snapshot sufficient for SaaS contracts? Do we need a change order mechanism?
8	Liability of the certification authority: Who is liable if the Trusted Agent Attestation has an exculpatory effect? Via professional liability, state liability or funds?

3. Interoperability and recognition

#	Question
9	Relationship to the AI Regulation: Switzerland is not an EU member. Will it adopt Art. 5 of the AI Regulation? Will it create its own ban?
10	International recognition: Under what conditions is a national or bilateral regulation recognised by other legal systems?

Other open points – such as the A2A2A constellation, transitional provisions, the enforceability of arbitral awards or the inclusion of partnerships in the EBW infrastructure – will also have to be taken into account in the future.

XII. Glossary

The technical terms used in the expert opinion are arranged in five categories. If a legal definition from the EU or German law exists, reference is made to it.

1. Legal terms with legal definition

Term	Source	Definition
AI system	Art. 3 No. 1 EU AI Act	A machine-based system designed to operate autonomously to varying degrees, which can be adaptable once it enters operation, and which derives from the inputs received for explicit or implicit objectives, such as outputs such as predictions, content, recommendations, or decisions that may affect physical or virtual environments.
GPAI Model	Art. 3 No. 63 EU AI Act	An AI model that has been trained on a large amount of data using self-monitoring on a large scale, has significant generality, and is capable of handling a variety of different tasks.
Provider	Art. 3 No. 3 EU AI Act	A natural or legal person, public authority, agency or other body that develops or has developed an AI system and places it on the market or puts it into service under its own name or trademark, whether in return for payment or free of charge.
Operator	Art. 3 No. 4 EU AI Act	A natural or legal person, public authority, agency or other body that uses an AI system on its own responsibility, unless the AI system is used in the context of a personal or non-professional activity.

2. Legal terms without legal definition

Term	Definition (own setting)	Differentiation
AI Agent (Technical)	An AI system that perceives its environment, draws conclusions independently and acts autonomously to achieve pre-determined or self-generated goals.	The term is initially of a purely technical nature and does not require any legal reference (e.g. routing agents, optimization agents, internal data agents).
Legally relevant AI agent (focus of this opinion)	An AI agent designated to make statements or take actions with legal effects in the name and on behalf of a legal entity or natural person.	Not every technical AI agent is a legally relevant agent. This expert opinion primarily focuses on the legally relevant aspects of agentic commerce.
Agentic commerce (technical / operational)	The autonomous, automated handling of trade processes (discovery, negotiation, checkout, payment) by AI agents based on technical protocols (e.g. ACP, UCP, WP2).	Describes pure data exchange and process automation. Answers the question: "What happened technically?"
Agentic commerce (legally relevant / constitutive/focus of this opinion)	The autonomous or delegated conclusion of trading transactions by AI agents, in which the machine actions are considered legally binding declarations of intent and trigger civil law obligations.	This is the focus of the report. Answers the question: "What applies legally?" This is where the lack of a constitutive legal norm comes into play. Not every technical transaction is already a legally binding transaction.

Agentic autonomy	The ability of an AI agent to act within a defined framework without individual human confirmation.	The degree of autonomy is determined by the scope of the delegated framework (external vs. internal scope).
Agentic Shopping	The use of AI agents in the B2C sector – i.e. by consumers.	To be distinguished from Agentic Commerce in the sense of the B2B sector.
Agentic Procurement	The use of AI agents in the procurement process of companies.	Subset of Agentic Commerce – focused on the purchasing side (B2B).

3. Basic Civil Law Terms (German Law Only)

Term	Source	Short definition
Declaration of intent	§§ 116 et seq. BGB	The conscious and intentional statement that is aimed at a legal consequence. It has an objective (statement) and a subjective element of the offence (will to act, intent to declare, business will).
Attribution	–	The legal attribution of an act or declaration to a person. It is the prerequisite for an action to have legal effect.
Legal transaction	§§ 104 et seq. BGB	A private expression of will that brings about a legal consequence. The contract is the two-sided variant.
Contract	§§ 145 et seq. BGB	The concurring declarations of intent of two or more parties. It is made by application and acceptance.
Obligation	Section 241 of the German Civil Code (BGB)	The legal relationship between creditor and debtor from which duties arise. It can arise by legal transaction (contract) or by law.
Pre-contractual obligation	Section 311 (2) of the German Civil Code (BGB)	A legal obligation that arises through contract negotiations, contract initiation or similar business contacts. It establishes duties of protection and consideration – not performance obligations.
C.I.C. (CULPA IN CONTRAHENDO)	§§ 280, 241, 311 para. 2 BGB	Liability for breaches of duty during contract negotiations. It compensates for the loss of confidence – not the interest in performance.
Messenger	§ 120 BGB	A person who transmits someone else's declaration of intent. The messenger has no will of his own – he passes on someone else's declaration.
Representative	Section 164 of the German Civil Code (BGB)	A person who makes his or her own declaration of intent on behalf of another. The declaration works for and against the represented party.
Assistant in performance	Section 278 of the German Civil Code (BGB)	A person whom the debtor uses to meet his obligation. His fault is attributed to the debtor.
Assistant	Section 831 of the German Civil Code (BGB)	A person who has been appointed by the principal to perform a task. The principal is liable for the damage – unless he can exculpate himself.

4. Technical Standards and Protocols

Protocol	Author	Short definition
ACP (Agentic Commerce Protocol)	OpenAI / Stripe	A protocol for the checkout process between AI agents and merchants. It standardizes the interaction between buyer-agent, merchant, and payment processing.
AP2 (Agent Payments Protocol)	Google	A protocol for authorizing payments through cryptographic mandates. It works with checkout and payment mandates.
UCP (Universal Commerce Protocol)	Google	A protocol for the unified description of the commerce journey – from discovery to checkout to post-purchase.
LCP (Legal Context Protocol)	AAA / Integra Ledger	A protocol for finding contract terms. It makes terms, governing law, and dispute resolution links machine-readable.
MCP (Model Context Protocol)	Anthropic	A protocol for data retrieval from AI agents. It connects agents to external data sources and tools.
M-Trust	GSMA / CAICT	A trust framework for agentic networks. It extends Zero Trust to include Multi-party Trust, Intent-aware Trust, and Cross-domain Trust.
Zero Trust	NIST SP 800-207	A security model that trusts no actor – not even the internal one. It is based on identity verification, least privilege and continuous monitoring.
AIS-1	AIS-1.org	An agent identity standard that links agents to legal entities.

5. Hybrid legal-technical terms

Term	Definition (own setting)
eIDAS 2.0	The EU Regulation on Electronic Identification and Trust Services for Electronic Transactions in the Single Market. It is the basis for the EUDI Wallet and the European Business Wallet.
EUDI Wallet	The European Digital Identity Wallet for natural persons. It enables secure identification, authentication and the exchange of attributes.
EBW (European Business Wallet)	The European Business Wallet for legal entities and economic actors. It enables the identification of companies and the transmission of authorisations of representation.
LPID (Legal Person Identification Data)	The standardized identity data of a legal entity. They are issued by the business registers and are the basis for the EBW.
EUID (European Unique Identifier)	The European unique identifier for companies. It is used in the EU's business registers and is technically applicable to all types of legal entities.
DPP (Digital Product Passport)	A digital product passport that provides product data over the entire life cycle – materials, CO ₂ footprint, certificates. The basis is the EU Ecodesign Regulation (ESPR).
StatusList2021	A revocation mechanism for Verifiable Credentials. It allows you to check whether a credential – such as an LDM – is still valid or has been revoked.
JAdES	A signature format for JSON documents. It enables the qualified electronic signature of LDM and TAM.
Verifiable Credential	A standardized format for the exchange of verifiable attributes. It is based on the W3C standards and enables cryptographic verification.

LDM (Legal Delegation Mandate)	A digital power of attorney in the form of a Verifiable Credentials that defines an agent's scope of action – in terms of facts, legally, financially and security-wise. It is the technical implementation of the digital power of attorney provided for in a norm and is issued by means of Selective Disclosure (SD-JWT VC): The <i>External Scope</i> is visible to the business partner and establishes the Apparent Authority; the <i>Internal Scope</i> remains hidden and protects trade secrets.
TAM (Transaction Acceptance Manifesto)	A proof of transaction in the form of a Verifiable Credentials, which documents a completed agent action. It references the underlying LDM and confirms that the transaction took place within the mandate framework. The TAM is considered a full-fledged substitute for a written contract within the meaning of the GoB, customs law and money laundering prevention.
TAA (Trusted Agent Attestation)	A certificate confirming that an agent complies with a recognized specification for trusted agents. It is a prerequisite for attribution to a legal norm and creates the basis of trust for the exculpation of the legal entity in the event of a technical malfunction.
SD (Selective Disclosure)	A cryptographic method (SD-JWT VC) that makes it possible to disclose only selected parts of a verifiable credentials. In the context of LDM: The <i>external scope</i> (Apparent Authority) is shown to the business partner, the <i>internal scope</i> (internal instructions, trade secrets) remains hidden. Only then does the third party remain in good faith.

6. Instructions for use

Note	Explanation
Legal Definitions	Where the law defines, reference is made to the law.
Own settings	Where the law is silent, a separate definition is proposed – this is marked as such.
Boundaries	In the case of ambiguous terms (e.g. AI agent vs. AI system), the distinction is made explicit.
Categories	The five categories help to distinguish the origin and the claim to validity of the terms.

Appendix

This expert opinion was prepared by Oliver M. Merx. He is a Bavarian legal assessor, computer scientist and certified AI officer. He started his professional career as a répétiteur for public law and criminal law and as a senior employee of a legal publishing house. He then switched to the world of the Internet, digitization and IT. As a manager, consultant and software developer, he gained experience in the areas of strategy, innovation, IT, sales and partner management in (inter)national corporations and start-ups. As a (co-)author, he published basic works on AI law, e-commerce law, quality assurance of Internet projects and corporate digital responsibility (CDR).

Oliver M. Merx is a consultant and trainer and works at the interface of business, technology and law. One of his professional focuses is AI prototyping: It serves to iteratively master the complex interaction of technical objectives, legal aspects and operational implementation in the areas of legal tech, public sector, health sector, industry, finance and trade in the best possible way.

www.grundwissen-ki-recht.de

www.linkedin.com/in/oliver-m-merx-83777b

