

00:16

Welcome to a new episode of BFRR, Bitcoin, Fiat and Rock and Roll. The podcast that explores the intersection of traditional finance, digital assets and digital money and helps you understand how digital money and digital assets will evolve in the years ahead. This is part three of our series on agentic payments. And the premise of the series is this, agentic payments

00:46

is an entire ecosystem rather than a single topic. So we are not doing one episode and moving on, we are doing one layer per episode in this series and we are going vertical in this attempt. Part one of the series was Sonja Davidovic from the IMF on the shape of the phenomenon, what actually changes when software rather than a human initiates a payment.

01:14

Part two was Peter Großkopf, a CTO and COO at All Unity, one layer down into the money layer. Why stable coins for agentic payments? Where that argument for stable coin holds and where it might break. Today's third part of the series is a bit different. Today, there is no guest. Today, it's me taking stock because

01:42

After those two conversations, I've now got a problem. The two of them agree on the diagnosis and disagree on many things that follow from it. So this episode attempts to do three things. An interim summary, a primer, if you're joining the series here, and an account of how my own position has moved through these two episodes in the series before.

02:11

And I want to start with the question the whole series turns on because Peter answered it in one sentence in part two, and I've been doing on it ever since. Can the control layer, the IMF layer two covering the agent's mandate, its budget, its spending limits. So can this control layer live inside the form of digital money itself, or does it always end up

02:40

outside the token in a wallet, a facilitator, a PSP. There's already protocols and open source technology developed in those directions. For example, AP2 protocols, which pretty much is also a standard formed around mandates and using verifiable credentials. And there's also cryptographic control so that you, for example, can use in an X4-2 facilitator.

03:09

So that's definitely technology that we're currently looking into. I think it would, it makes perfectly sense where it's possible, where it doesn't change the standard, but where you lose kind of compatibility with most wallet infrastructure. that's definitely something where you, think, lose more than you gain. So two guests, and let me tell you properly who I've been arguing with because the sequence of their careers is the reason their answers.

03:39

differ and also the reason why we invited them. Sonja Davidovic has been at the front edge of this stuff for years. Bank for international settlements, the innovation hub in Singapore, the European Central Bank on the digital euro. And now back at the IMF, she was writing central bank digital currency papers when few or even nobody was actually talking about CBDCs.

04:05

She writes the framework before the industry has the vocabulary at hand. Her IMF note 2026-4 with RW Turp published April 2026 is, for my money, still the best framing out there when it comes to agentic payments. And then Peter Großkopf, he sat on other sides of the table.

04:33

Co-founder and CTO of Solaris, building banking as a service before that was a category.

Then CTO and managing director at Börse Stuttgart Digital Exchange, the first regulated digital asset MTF in Germany. Then Unstoppable Finance and Iron Jane in Berlin, an attempt at the three things that stop regulated institutions touching public blockchains. Privacy of account data, compliance governance.

05:02

Now CTO and COO at Allunity, the DWS Flow Traders Galaxy joint venture out of Frankfurt with a full BaFin e-money institutional license and live Micah compliant stable coins, plus the digital advisory board of the German finance ministry. So in my view, Sonja tends to write the model, Peter tends to then ship the product and

05:32

If I've misstated either position here, the comments are open and I'd really rather be corrected in public. just go ahead. Here's what I want to do in this episode's remainder. Everybody is talking about agentic payments and agentic commerce right now. Instant checkout in ChatGPT, Google's universal commerce protocol, Visa Intelligent Commerce.

06:00

MasterCard agent pay protocol after protocol. The slightly harder question few are already debating loudly is when your agent spends your money in a way you didn't specifically anticipate, was that particular payment authorized? And who eats it when it was not authorized?

06:28

The question, not a technology one, in my opinion, by the way, the question is the one I'm going to spend most of this episode investigating. So first let's get the foundation right. Because if you're joining our agentic payment series here, you need it. Payment infrastructure is deterministic by design. Same input, same output every time. Rules are fixed.

06:58

Outcomes are binary. Settlement is legally final. That's true of a real time gross settlement system. It's true of a card network. It's true of a blockchain. That determinism is the entire product. That's the value proposition. It's why you trust that money really moved and settled. It really is not a limitation someone forgot to fix. That's the feature.

07:28

Then in turn, agentic AI is the opposite by construction. Large language models generate outputs by sampling from probability distributions. The same prompt can give you a different answer. That's the property that makes agents useful. That's their value proposition. They adapt, they plan, they handle situations nobody specified in advance. They are probabilistic.

07:58

So the tension isn't that we're putting AI into payments. AI has been in payments for 40 years. American Express had an expert system doing credit authorization in the 80s. Falcon fraud manager was scoring transactions with neural networks in the 90s. None of that touched initiation. So far AI evaluated and humans decided. What's really new is

08:27

delegated authority at machine speed. Sonja's phrase for it in the IMF node is the shift from click to pay to decide to pay. The agent doesn't recommend the transaction. It initiates it under a mandate you gave earlier and possibly manages the whole cycle through dispute resolution. Which brings the obvious question. If you

08:56

can't make the agent deterministic and you can't make settlement probabilistic, where do you put the boundary? That's what the IMF three layer model is for. That's what I appreciate it for. That's why we invited Sonja to not only learn about IMF thinking and her views, but in particular about this three layer model.

09:21

Sonja's answer, and she was careful to call it a normative lens rather than a description of anything that exists. Her answer is to separate the functions into three layers. Layer one, intent and orchestration. This is where the agent lives. Goal formation, planning, search, comparison, negotiation, multi-agent coordination. It's messy, it's probabilistic, and that's fine.

09:50

because nothing here moves money. This is where MCP, agent to agent protocols and the universal commerce protocol operate. Then there is layer two, control and authorization. This is where the rules bite. This agent authorized to act. Is this inside the mandate, the budget, the compliance envelope of the agent? So these are deterministic

10:18

guardrails. This is where AP2 sits with its cryptographically signed mandates and where OAuth style attested agent identity does know your agent work in distinction to know your customer, of course. It's also where programmable wallet standards enforce spend limits. So ERC 6900 modular smart accounts, ERC 1812 for

10:48

off-chain verifiable claims or ERC 84 for agent identity registries. And then ultimately layer three, settlement. This remains unchanged, final, no room for probability. I'm thinking of RTGS, instant payment networks, card clearing, stablecoin rails, and in

11:14

DLT environments, ERC 4337, smart contract wallet execution. If you come from traditional payments, layer two is most likely not a new idea to you. It's the direct debit mandate. It's card on file credentialing. It's the corporate treasury approval metrics.

11:40

What is new in the world of agentic payments is that with a direct debit, the payee is known and the amount logic is rule bound. And with an agent, the payee might be one of the agent invented at runtime. Now, does the model survive contact with practice, with a practitioner building already?

12:08

stable coins and stable coins being used in a gigantic payments. I put it to Peter, Peter Grosskopf, who has spent his career at layer three in my reading. He said it matches, but with a caveat, I think is important and slightly uncomfortable for anyone doing research in the space, including me. He said the field is still forming and it's evolving at AI speed.

12:37

His estimate roughly 10 times the pace of the back then print to web transition. His concern was that people working in theory will have a hard time keeping up with what's actually happening in the market today. I'll take that one on the chin. It is also the reason we are doing this series.

13:02

one layer at a time rather than one grand framework or one grand episode. And it left one thing open. Sonya ended part one with a question rather than an answer. That is, can the agentic mandate, the caps, the parameters, the agent's constraints in her layer two, can this mandate be mapped onto the smart contracts underlying digital assets?

13:32

Is there real technical interoperability there? Peter's answer was technically yes in places. Allunity supports EIP 3009 style signed transfer authorizations. They are looking at AP2 with verifiable credentials. They run a permissioned X402 facilitator, but he wouldn't put the mandate inside a custom token standard because a non

14:02

standard token loses wallet support and a token nobody's wallet supports isn't money really. So the state of play is this. The control layer sits next to the money rather than inside it. And that turns out to matter enormously for the legal question. Before we get there, the money layer, because I made a strong claim in part one and Sonja declined to endorse it.

14:31

And I owe you the argument on both sides. My claim was agents need money that behaves like software, programmable by default, always on, always final. And stable coins are one of the very few forms of digital money doing that today. Yeah. I mean, you're entitled to your own opinion. I will stick by my opinion that.

14:59

It really depends on the use case. And here's the strongest version of that claim. And I formulated it as an economic argument rather than an ideological one. Take hard acceptance in Europe. Interchange is kept by the interchange fee regulation at 0.2 % for consumer debit and 0.3 % for consumer credit. But interchange is not what the merchant pays at the scheme fees. Acquire a markup.

15:29

gateway and a typical online merchant land somewhere around 1.5 % plus 20 to 30 cents fixed. Correct me if that's moved in your market, of course. Now, say you're willing to spend 3 % of transaction values on payment cost. Generous. With a 25 cent fixed fee counting just that fixed component against your budget.

15:58

Your smallest economically viable ticket is a bit over 8 euros. Count the ad valorem 2 and it's closer to 17. Either way, below that, the payment eats the product. A stablecoin transfer on a modern layer 2 or on Solana or on BASE, which is what AllUnity uses, is a fraction of a cent. Same 3 % budget.

16:27

and your minimum viable ticket drops to somewhere around three cents. That's roughly two and a half orders of magnitude collapse in the smallest transaction that can exist. And it matters because a gigantic payments is natively priced below the old line per article, per crawl, per data fetch, per API request or per inference call or Unity's own material puts publisher revenue

16:53

per AI crawl at zero euros today. And I don't think that's because publisher unwilling to charge. There has been no money primitive small enough to actually charge with. Then I have a second argument for why stable coins, which I care about even more. And it is about where the control lives.

17:18

On card and account rails, spending controls sit in somebody's back office, a rule book, a risk engine, a mandate table at a PSP. In a smart account, the controls travel with the money, spend caps, counterparty allow lists, time bounded session keys, per transaction limits, enforced in code verifiable by anyone.

17:45

That turns Sonja's IMF layer 2 from a policy into a cryptographically enforceable constraint. Now, the counter case and Peter, who sells stablecoins for a living, gave me a more structured version of it than I've heard so far. He splits the world three ways. Agent to service that's paying for an API call or a data point or an online article.

18:15

Then he sees agent to merchant. So that is the traditional shopping cart physical goods. And then third agent to agent payments, which he sees as the long run extension of agent to

service. In agent to service stable coins are close to unbeatable because the tickets are tiny and card economics simply don't reach. But then in agent to merchant,

18:44

Cards keep a real job. Baskets are large, goods sometimes don't arrive and cards have a 60-year-old judgeback framework built precisely because payers make mistakes. Now the payer is a large language model which will sometimes be confidently wrong. Stablecoin settlement is final and hard to serve. Peter's response in our second episode of the series was that reversibility could

19:14

be built. So escrow in a smart contract, a 14 day withdrawal window or a dispute period. So technically, of course, that's true, but it runs straight back into the standard problem. A specialized token standard needs critical mass. And without such a critical mass, no wallet will ever support you and your specialized token with its reversibility built.

19:43

in. And then I pushed him on something he conceded quickly. A charge back is a judgment rather than a transfer. Somebody has to be the arbiter. Does a MICA licensed e-money institution want that role? He said plainly that it makes sense to move a Chantig payment services into a separate entity once the regulatory duties are clear. That's a meaningful admission. Then there's an

20:13

Uncomfortable argument against stable coins for agentic payments. addition, Euro liquidity, total stable coin supply is a few hundred billion dollars and overwhelmingly dollar denominated. Euro tokens are a rounding error on that really, at least today in 2026, of course. If an agent wants deep, liquid 24-7 settlement today, it reaches for a dollar stable coin.

20:43

Allunity is shipping a dollar stablecoin too, tells you something. Peter's answer was short, mid, long term. The first wave was dollar tokens for crypto use cases. The second wave is real world companies who want their own currency and merchants don't want foreign exchange exposure. I think that's right directionally and it's also not true yet.

21:09

And the one argument against stablecoins in agentic payments, I'd add the transfer of funds regulation poses a challenge here. For crypto asset transfers, there is no de minimis threshold. Every transfer carries originator and beneficiary information requirements. Try applying that to a per crawl payment worth a tenth of a cent.

21:39

So my take survives, but way narrower than I stated it. The stable coins win where tickets are small and finality is a feature in a gigantic payments. Cards in turn hold agent to merchant until the reversibility gets built.

22:03

really. So which means that the money layer is roughly solvable and the thing that is not solvable with technology at least is now the next section. Here's the scenario that broke my thinking on this and it comes from a paper by Nidia Remolina at Singapore Management University titled, Agents Payments, When is a payment

22:32

unauthorized. It's on SSRN and it'll be in the show notes of course. Consider a fictional person called Elena. Elena tells her agent, register me for a payments conference in Amsterdam, arrange travel, keep the total under 1500 euros. The agent compares options, books the hotel, pays for the ticket. Strong customer authentication is satisfied. No password stolen, no card cloned. So all good so far.

23:01

Lower the package cost, however, the agent buys a 299 euro annual travel platform subscription. The merchant is legitimate. The price is correct. After the membership discount, Elena's total is still under 1500 and the agent explains it will save money on future trips. Elena objects. Naturally, she authorized a conference trip, not a year long membership.

23:28

The AI provider says it was within her budget instruction. The bank says it received a valid order through agreed credentials. Was Elena's 299 euro payment authorized? Sit with that for a second because European payment law actually has no clear answer. PSD2

23:56

Directive 2015-2366 says a transaction is authorized only if the payer, in that case Elena, consented. Its Article 64 lets you consent before execution and to a series of transactions. Fine. The PSR proposal keeps the architecture and swaps in the word permission to avoid confusion with GDPR consent.

24:27

also fine. But the whole regulatory framework assumes the payment service provider doesn't modify the amount or the payee. Strong custom authentication, dynamic linking, bind authentication to a specific amount and a specific payee. That's Article 97 and the SCA regulatory technical standards. That works when the user

24:54

can see those features at the moment they authenticate in an agentic payments flow. The amount and the pay may not exist yet when the user delegates. So the issue is actually not the authentication. Authentication worked perfectly in the example of Elena. The issue is the boundary of the delegation.

25:19

And the November 2025 provisional agreement on the payments package doesn't fix it. Payee name checks, spending limits, blocking measures, platform recourse is defined in defined fraud cases. All useful, but none of it answers Elena's questions. Payee verification is not

25:47

intent verification. It cannot tell you whether the agent, the PE, was allowed to create the annual transcription. So what could actually fix it? Remember the researcher Remolina at Singapore Management University? Remolina's proposal is what she calls a bounded delegated authorization test. Again, a bounded delegated authorization test.

26:16

And it has four cumulative elements. I'll give them to you because I think this is the most useful piece or publication on that topic so far. So the first piece, a legally cognizable mandate. Not every detail that would make agents pointless, of course, but enough detail to define a stable class of authorized transactions and

26:45

The paper draws a helpful distinction. A prompt-like, quote-unquote, choose good value products is a preference for good value products. In turn, a prompt-like, quote-unquote, purchase office stationery from approved suppliers up to 300 euros per order on this card and ask before any new supplier is authority. So preferences aren't payment law.

27:15

objects. Then a second element of this proposed bounded delegated authorization test. A payment constraint envelope. The mandate translated into machine readable, enforceable limits, amount caps, pay categories, white lists, new pay rules, recurrence limits, permitted rails, currency and geography, expiry,

27:43

or approval triggers and critically capability changes have to be covered. If the provider adds a new rail, a new tool that can initiate payment, a new merchant category that should trigger reauthorization, not unknowingly expand your authority into new LLM features. Then a third element of this test, runtime conformity.

28:12

Not just whether the final order matches the mandate on its face, but whether the whole path stayed inside the envelope. A 95 euro payment can be inside a 100 euro cap and still outside the mandate if it went to a new payee where approval was required. And then ultimately the fourth element of the proposed authorization test, auditable proof.

28:38

Here I'm thinking of mandate, constraints, approval events, tool calls, model and policy versions, the payee checks, and human overrides. This is not source codes. This is not the weights in the LLM. So just enough to test the authorization claim when it is disputed. And then there is a piece with teeth.

29:04

The proposal is that if a language model provider enables an autonomous payment under a delegated mandate and cannot produce the applicable constraints, the payment is presumed unauthorized and the payer facing PSP reimburses with recourse against the provider. The paper also proposes a new regulatory category to hang this on called

29:33

Digital Assistant Payment Services. Digital Assistant Payment Services. Light registration under PSD3 would be important here. Modeled on how account information service providers are registered rather than fully authorized under Article 33. Contact duties under the PSR plus a payment authority dashboard.

29:59

The PSR proposal already has a permission dashboard in Article 43, but that controls account access, not autonomous payment specification. These are really different objects. Now, if you work at a DREDFY financial institution like a bank, a PSP or a payment network, what I just said begins to slowly evolve into a build list. A build list that makes it really concrete.

30:29

how financial institutions can now, maybe should now actually act on agentic payments. And for this, I brought five things what I learned in the previous two episodes, what a financial institution should do. And I would order them by priority roughly this way. So number one is decide where you sit in the layer model explicitly and write it down.

30:59

Most financial institutions I talk to are implicitly assuming they are layer three and therefore safe. If you're the payer facing PSP, you're the first respondent when your customer disputes an agent initiated payment. You do not get to be layer three only, even if you wish. Then a second thing a financial institution should do or at least could do.

31:28

Build the evidence chain before you need it. Whichever way regulation lands, you'll be asked to show whether the agent stayed inside the mandate. If your logs showed a valid SCA event and a valid payment order and nothing about who set what constraint, you have nothing. Note that article 72 of PSD2 already says use of a recorded

31:55

payment instrument isn't by itself sufficient proof of authorization. That principle is going to get applied to agents. Then three, know your agent. The IMG node flags the shift from KYC

to know your agent, KYA, as a public sector measure and points to Singapore's IMDA model framework for agentic AI as an

32:25

early operationalization, which node is voluntary as I speak. Meanwhile, Visa and MasterCard are building registries and cryptographic agent credentials at network level. If you're a bank, you will be consuming somebody else's agent identity layer. So choose deliberately here really. Then a fourth action for a financial institution. Get your mandate dashboard on the roadmap.

32:54

Not the open banking permission dashboard you may already be building for the PSR, but a separate view, active mandates, constraints, budgets, permitted fees, recurring authority, expiry of a payment approval trigger rules, and a pause button that doesn't kill unrelated services. And five, ultimately what you can do as a financial institution, sort the recourse

33:24

contractually out now, while you still have negotiating leverage with the AI providers because the loss allocation question you reimburse first, then recover from whoever controlled the decision path is going to be settled either via regulation or by whoever writes the contract first. Which now brings me to where I've landed.

33:51

after the first two conversations of our BFRR, agentic payment series. I rather came into this series thinking the hard problem was the money. The agentic payments would be unlocked by a form of money, small enough, fast enough and programmable enough to carry them. This was basically a stable coin story. I no longer think the money is the hardest part.

34:20

Peter has a life license, life tokens, a life facilitator, a substance settlement and a discovery product. The money layer has a sufficiently credible answer today, at least to me. The rather hard part is that we can now build a mandate that is machine enforceable, but not yet legally cognizable. The cryptography is ahead of the law and payment law still runs on a binary

34:50

authorized or unauthorized anchored in human intent. Agents don't produce human intent. Agents produce conformity with the constraints given potentially through a prompt or the absence of such conformity with constraints. So my stake in the ground would be the

35:15

binding constraint on agentic payments, at least in Europe for the next couple of years is legal and evidential rather than technical. the institutions that win will be the ones who can prove after the fact exactly what their customers authorized. Institutions with just the cleverest agent in layer one will not win. So let's do a little 12 month test borrowing Peter's

35:44

three milestones he laid out at the end of episode two in our series. First milestone supply. We need more publishers, data providers and APIs actually accepting agentic payments. A second milestone would be cheap on and off ramps in euros for both businesses and retail. And a third milestone, he thinks this is the blocker for

36:13

adoption of a genetic payments, consumer site infrastructure for setting mandates and spending limits that doesn't require a command line. So essentially an architecture in which a user, a retail user is not demanded or asked to use a command line or a perfectly engineered prompt to inject his mandate and his spending limits. And then if

36:42

this time next year in 12 months, an ordinary person like our moms and dads can set a payment mandate in three clicks. Then a gigantic payments gets real. Then we would be

talking if they cannot, we'll be having the emergent conceptual conversation we do have today. So here we go to close the thread of today. Four things I would take away.

37:13

from those three parts of the series as an inter-remistic or preliminary result so far. The first, the deterministic probabilistic tension is real and structural. Payment rails are deterministic because that's what they are supposed to be. That's the product of payment rails. Agents in turn are probabilistic because that's the product for

37:42

agents. The design job is separating those two functions, this deterministic function and this probabilistic function. And the IMF's proposed three layer model has survived contact with a leading practitioner, namely Peter Großkopf. My second takeaway would be the strongest argument for stable coins in urgent payments is that

38:11

the minimum viable transaction size collapses by two to three orders of magnitude, which is the difference between agentic commerce existing and not existing. Third, what I learned, money was split by use case. So the form of money, the form of digital money to be precise, was split by the use case of an agentic payment.

38:39

I would say stable coins take agent to service. Cards in turn keep agent to merchant until reversibility gets built into the token layer. So into stable coins, for example. And that's a standards problem more than a technical one. And fourth, the takeaway that really changed my thinking. The challenging part is legal rather than technical when it comes to

39:08

making agentic payments real and globally adopted. Authentication can be perfect and the payment can still be outside the mandate. Bounded delegated authorization, which is the test I outlined above, and the presumption that a payment without producible constraints is unauthorized is the most useful proposal I've seen so far to mitigate this deterministic

39:37

probabilistic tension legally. And now a thank you to the two guests not in the recording studio today, Sonja for declining to endorse my stablecoin take on air and telling me it depends on the use case, which type of digital money is used by an agent. And she was right. And it made me do the work really. Thanks for that Sonja. And Peter.

40:03

came on without knowing which way I take it and answer the hard questions, the technical operational questions head on instead of simply steering it back to his products, mainly his stable coins, which so many stable coin issues do. So he conceded that cards have the advantage in e-commerce today. He also conceded that a judge back is a judgment his e-money institution, which all unity really is, may not want to make.

40:33

And I found that intellectual openness from a stable coin issuer is rarer than it should be. So yeah, kudos to Peter here for that openness. If you got something out of this, the single most useful thing you can do is send it to one colleague, be it in Treadfy, DeFi or CFi who's building in the space and drop us a rating wherever you're listening. That does more for us than anything else.

41:00

Come find us and keep the argument going. We are on YouTube and LinkedIn and our expert community lives on Telegram in both German and English since BFRR has German roots, right? All the links are in the show notes along with the IMF note, the Remolina paper and parts one and two of this agentic payment series. And do head to the website for the written deep dive on this one too.

41:26

Subscribe wherever you get your podcasts and be there for a fourth part of this series when we likely go after the question, we so far left open which form of digital money will agents really actually pay with? Until then, this is Bitcoin Featuring Rock and Roll bringing Europe's perspective on digital money and digital assets to the world.