ZETTL VI: Zlef

Typed, calibrated decisions over crypto-asset payments:
a legal ontology of asset types, a label schema, and a research methodology

Jeet Raut, ZETTL

Living document — Edition 1 (v0.5r5 draft), 15 October 2026

Edition 1 — 15 October 2026. A living document. This is a dissertation-style research document for Zlef, the payment decision model proposed in Paper IV. Zlef is not built, trained or evaluated; it has no data, accuracy, calibration, cost or latency results. Zlef is a working name for Zettl's proposed model; it is not Cloudflare's Clef, and Zettl has not implemented Cloudflare's RLCD objective. Legal and tax statements describe sources as of the status date and are not advice; each carries a status marker ([In force], [Enacted, not yet effective], [Proposed], [Unverified], [Sources conflict]). The edition log (Appendix B) records what changes between editions.

Series macros: [Product, off-protocol] proposed product; [Research hypothesis] research hypothesis; [External dependency] third-party dependency; [Unsupported, needs evidence] needs evidence.

Abstract

Zlef is a proposed decision model for crypto-asset payments. It answers a fixed set of typed questions about each payment and attaches a probability to every answer: what kind of asset moved, what happened to it, what the payment was for, whether beneficial ownership changed, who the counterparty was, how complete the evidence is, how urgently a person should look, and which compliance surfaces the transfer crossed. It may say "unknown" or decline to decide. It is inspired by typed decision models such as TypeSafe's Jev and Cloudflare's Clef, but it is not either of them, and nothing about it is built or measured.

This dissertation-style document lays the groundwork in four parts. First, a *legal ontology* of crypto-assets. We survey how the main crypto-asset types are classified, taxed and constrained in their movement across the ten jurisdictions with the largest on-chain flows, plus the EU and the main hub regimes. From this we argue that legal category is a function of asset facts, jurisdiction and date, so it belongs to a versioned rules layer and not to a model. Second, a *label schema*. It asks only jurisdiction-free economic and technical questions, uses a hierarchy to allow partial answers, separates the label "unknown" from the decision to abstain, and states for every value how an annotator should choose it and why the question exists. Third, a *modeling methodology*: losses per question type, calibration, three abstention mechanisms, cost matrices derived from downstream flags, cross-question consistency constraints, and explanations kept apart from scores. Fourth, a *data and evaluation* plan: annotation with anchoring controls, customer-, time- and asset-introduction splits, per-class calibration, risk–coverage analysis, flag-level metrics and preregistered decisions.

The document is living. Edition 1 records its status date (15 October 2026), its verification gaps and a research agenda, and later editions will log every change.

Contents

1 Introduction **4**

2 Background and Related Work **5**
2.1 Typed decision models: Clef and Jev . . . 6
2.2 Reinforcement learning for calibration . . . 10
2.3 Calibration . . . 12
2.4 Selective prediction, abstention and learning to defer . . . 14
2.5 Conformal prediction . . . 15
2.6 Ordinal and hierarchical classification . . . 16
2.7 Label noise, annotation and human factors . . . 17
2.8 Cost-sensitive learning . . . 18
2.9 Distribution shift and drift . . . 18
2.10 Leakage . . . 19
2.11 Structured outputs, constrained decoding and label smoothing . . . 19
2.12 Financial transaction and AML machine learning, and legal NLP . . . 20
2.13 Explanation faithfulness . . . 21
2.14 Privacy-preserving machine learning . . . 22
2.15 Prompt injection and adversarial text . . . 22
2.16 Summary: from literature to design commitments . . . 22

3 A Legal Ontology of Crypto-Assets **24**
3.1 Legal category as a function of facts, jurisdiction and date . . . 24
3.2 Reference taxonomies . . . 27
3.3 A crosswalk from Zlef asset classes to legal categories . . . 34
3.4 Dimensions that drive classification . . . 38
3.5 Where regimes disagree . . . 41
3.6 Consequences for the label schema . . . 43

4 Asset Types I **43**
4.1 NAT: native and protocol tokens . . . 45
4.2 STB: stablecoins . . . 51
4.3 TKN: tokenized traditional assets . . . 59
4.4 GOV: governance and utility tokens . . . 64
4.5 MEME: memecoins . . . 68
4.6 NFT: non-fungible tokens . . . 72
4.7 LST: liquid staking and restaking receipts . . . 76
4.8 Summary across the seven classes . . . 80

5 Asset Types II, Events and Movement **81**
5.1 Wrapped and Bridged Representations (WRP) . . . 81
5.2 DeFi Position Receipts (DFI) . . . 86
5.3 Privacy Coins (PRV) . . . 89
5.4 CBDCs and E-Money Outside MiCA Scope (CBDC) . . . 91
5.5 Other Claims: Points, In-Game Items, Unvested and Airdrop Claims (OTH) . . . 94
5.6 The Event Dimension . . . 97
5.7 How Assets Travel . . . 106

6 The Zlef Label Schema **110**
6.1 Design principles . . . 112

6.2 Q1: `asset_class` . . . 113
6.3 Q2: `event_type` . . . 115
6.4 Q3: `economic_purpose` . . . 116
6.5 Q4: `beneficial_owner_change` . . . 117
6.6 Q5: `counterparty_type` . . . 117
6.7 Q6: `evidence_completeness` (ordinal) . . . 118
6.8 Q7: `review_priority` (ordinal) . . . 118
6.9 Q8: `travel_surface` (multi-label) . . . 118
6.10 Consistency constraints . . . 119
6.11 From labels to rules . . . 120
6.12 What the schema deliberately leaves out . . . 121
6.13 Changes from Paper IV . . . 121
6.14 Versioning . . . 121

7 Modeling Methodology **121**
7.1 Problem statement . . . 122
7.2 Candidate architectures . . . 123
7.3 Losses per question type . . . 126
7.4 Post-hoc calibration . . . 129
7.5 Abstention and back-off . . . 130
7.6 Cost matrices from flag severity . . . 132
7.7 Consistency across questions . . . 134
7.8 Explanations kept apart from scores . . . 135
7.9 Adversarial robustness of text fields . . . 135
7.10 Training sequence and model registry . . . 136

8 Data, Labels and Evaluation **137**
8.1 Data sources . . . 137
8.2 Annotation protocol . . . 139
8.3 Anchoring control: blind and assisted labelling . . . 141
8.4 Candidate labels from reviewer corrections . . . 142
8.5 Splits . . . 142
8.6 Metrics . . . 143
8.7 Drift monitoring . . . 147
8.8 Preregistration template . . . 148

9 Risks, Ethics and Limits **148**

10 Research Agenda for Later Editions **150**

A Label Dictionary (Edition 1) **173**

B Edition Log **174**

1 Introduction

Paper IV proposed that ZETTL's first products would be a compliance and analytics dashboard and a payment decision model, and it fixed one boundary between them. The model interprets payments into typed facts with probabilities. Legal and tax treatment comes from versioned rules and human review. Paper V built the research base for the rules. This document does the same for the model, which Jeet named Zlef.

The problem. A finance team that handles crypto-asset payments across borders must answer, for each payment:

- what the asset was;
- what happened to it;
- what it was for;
- who was on the other side;
- what evidence exists;
- which rules might apply in each place the payment touched.

The first five are questions of fact, often underdetermined by the records. The last is a question of law that changes by jurisdiction and over time. General-purpose language models can produce fluent answers to all six, but fluency is not calibration. An answer that is confidently wrong about whether a token is a stablecoin or a tokenized security can suppress the very flag that should have fired.

The thesis. We defend three claims, each of which the rest of the document makes precise.

1. *Separation.* Legal category is a function of asset facts, jurisdiction and date (Section 3). A model should predict the facts, and a versioned rules layer should compute the category. This makes legal change a rule edit, not a retraining run, and it keeps the model's errors measurable.
2. *Typed, calibrated, abstaining outputs.* For the facts that records do not settle, a model with a fixed typed output space can be evaluated with proper scoring rules and risk–coverage analysis, and its errors can be priced by their downstream effect on flags (Sections 7 and 8). Free-text generation cannot be evaluated this way.
3. *Learning before building.* Whether a specialised model beats simple rules or a general model on a real buyer's workflow is an empirical question [Research hypothesis]. The document therefore specifies how to find out, and preregisters what would count as an answer.

What Zlef is and is not. Zlef is a working name for Zettl's proposed model. It is not Cloudflare's Clef or TypeSafe's Jev. Those are [External dependency] decision models whose published designs inspire the typed output space and calibration objective, and which Section 2 reviews in detail. Zettl has not implemented Cloudflare's RLCD objective. Zlef does not decide residency, sanctions status, legal category, tax treatment or compliance. It does not replace a reviewer. It has no data, training or results.

Organisation.

- Section 2: related work.
- Section 3: the legal ontology and the crosswalk between the Zlef asset classes and the regimes' categories.
- Sections 4 and 5: each asset type in turn (definition, legal classification by jurisdiction, taxation by event, how it travels and how Zlef labels it), then the event dimension and the travel surface.
- Section 6: the label schema.
- Section 7: the modeling methodology.
- Section 8: data, labels and evaluation.
- Section 9: risks, ethics and limits.
- Section 10: the research agenda for later editions.

Status and sources. Legal and tax statements describe sources as of 15 October 2026 and carry status markers. Some jurisdictions, notably Vietnam and Ukraine, and some tax cells, rest on a reduced verification base and are marked [Unverified]. Academic citations were checked where possible; entries not re-checked for page numbers or DOIs are noted in the bibliography. Nothing here is legal or tax advice.

2 Background and Related Work

This section reviews the work that Zlef depends on. Zlef is a proposed model. It answers a fixed set of typed questions about a normalized payment record, returns a probability distribution for each question, and may abstain. Legal and tax treatment comes from versioned rules and human review, never from the model. The review therefore asks one question of every body of work: what does it imply for a model with those commitments?

We cover sixteen topics. We begin with the two vendor decision models that motivated Paper IV, Cloudflare's Clef and TypeSafe's Jev (Section 2.1). We then turn to reinforcement learning for calibration (Section 2.2), calibration itself (Section 2.3), selective prediction and learning to defer (Section 2.4), and conformal prediction (Section 2.5). After that come ordinal and hierarchical classification (Section 2.6), label noise and annotation (Section 2.7), cost-sensitive learning (Section 2.8), distribution shift (Section 2.9) and leakage (Section 2.10). The last group covers structured outputs and label smoothing (Section 2.11), financial and legal machine learning (Section 2.12), explanation faithfulness (Section 2.13), privacy (Section 2.14) and prompt injection (Section 2.15). Each subsection ends with a paragraph headed *Implication for Zlef*. Section 2.16 collects these implications in one table.

Two conventions apply throughout. First, every vendor claim is a vendor self-report. The vendors ran the benchmarks themselves, and nobody has replicated them. We report the numbers so that readers can see what is claimed, not because we accept them as evidence for payments. Second, where we write "Zlef should", we state a design commitment or a hypothesis for the methodology in Section 7 and the evaluation in Section 8. Nothing here reports a result. Zlef has not been built, trained or measured.

2.1 Typed decision models: Clef and Jev

A *decision model*, in the vendors' usage, is a model that receives a state and a schema of typed questions and returns a bounded answer for each question. It does not return free text. Two such systems were announced in the weeks before this edition. TypeSafe launched Jev on 15 Sep 2026 [2]. Cloudflare released Clef and Clef-flash on 1 Oct 2026 [29], and then Clef-omni on 9 Oct 2026 [30]. Paper IV took three ideas from them and no more: typed, bounded outputs; a probability on every output; and no free-text decisions. ZETTL does not currently use either model; both are candidate [External dependency] baselines for the evaluation (Section 7.2). Zlef is a working name for a model that ZETTL proposes. It is not Cloudflare's Clef.

2.1.1 Clef

Sources and authorship. The launch post, "Introducing Clef: our open-source decision models, and new RL fine-tuning platform", is dated 1 Oct 2026 and lists Michelle Chen, Alex Reneau and Kevin Flansburg as authors [29]. A follow-up post on 9 Oct 2026 lists Michelle Chen, Alex Reneau and Kevin Jain [30]. We record each author list as shown and do not reconcile them. Hugging Face hosts model cards for `Cloudflare/clef` and `Cloudflare/clef-flash` [33].

Base models. Clef is built on a frozen backbone. In Cloudflare's words, the models were trained "By freezing Qwen3.8-27B for Clef and Qwen3.5-9B for Clef-flash" [29]. The model card for `Cloudflare/clef` says it was "post-trained from" `Qwen/Qwen3.8-27B` and keeps that model's vision encoder [33]. It reports 27B parameters stored in BF16. The card for `Cloudflare/clef-flash` names `Qwen/Qwen3.5-9B` as the base, also with a vision encoder, and 9B parameters. [Unverified] We report these checkpoint names as Cloudflare quotes them. We have not confirmed that they exist as public Qwen releases. Clef-omni is built on a "Qwen3-Omni-30B-A3B-Instruct mixture-of-experts (MoE) foundation" and accepts audio, video, images and text [30].

Routing head and adapters. Cloudflare trains two components on top of the frozen backbone. The team "jointly optimized the routing head alongside rank-256 low-rank adapters" [29]. The post does not define the routing head further. On our reading, not Cloudflare's, the design is therefore a frozen pretrained network, a set of low-rank adapters of rank 256, and a learned head that maps the adapted representation to scores over the valid options of each question.

Inference. Clef does not generate text. Cloudflare describes a single prefill pass that scores all valid schema choices in parallel, with the model constrained "to output only probabilities instead of text generation" [29]. The model card states the same design: the output is "one logit per allowed option" for each question, followed by a softmax per question, and "There is no free-form text generation and no output parsing" [33]. This is the design that the calibration literature in Section 2.3 studies most closely: a softmax over a fixed label set.

Question types and interface. A request carries `model`, `state` and `questions` [29]. Each question has a `type`. The type `noul` is a yes/no question. The type `choice` takes `criteria` as a map from label to description. The type `score` takes `criteria` as an ordered list of levels. The name `noul` is not a typo; TypeSafe's documentation uses the same name [216]. Cloudflare calls

Clef "fully Jev-API compatible" [29]. The model card describes the input as a `state` (a string or JSON), optional `images` or `videos` with `media_kwargs`, and `questions` keyed by question identifier [33]. Its `encode_record` step defaults to a maximum of 16,384 tokens. Answers differ by type. A choice answer returns `choice`, `confidence` and `probabilities`. A score answer returns an expected `score`, `confidence`, a `legend` and `probabilities`. A noul answer returns the probability of true.

Context and modality. The launch post says "our model has a 64k context window" and contrasts it with 32k for Jev [29]. Eight days later Cloudflare corrected this for one configuration: "The hosted version of Clef-flash now has a context window of 24k, instead of 64k as previously advertised" [30]. The same post says the open weights support 256k. Clef "has a vision encoder so it's able to take in images" [29]. The launch post describes Jev as text-only.

Training data and losses. Cloudflare trained on "internal synthetic datasets permutating field orders, prompts, and schema structures" [29]. It discloses no dataset size, compute or hardware. The supervised objective is "label-smoothed cross-entropy for valid schema outputs paired with a Brier loss to refine probability calibration". The 9 Oct post restates "label-smoothed cross-entropy loss with Brier score calibration" [30]. Both terms have a long history. Label smoothing comes from Szegedy et al. [204]. The Brier score comes from Brier [19]. We discuss each below.

RLCD. Cloudflare also describes a secondary optimisation target: "We also developed Reinforcement Learning for Calibrated Decisions (RLCD)" [29]. As described, the reward has three parts. It grants "partial credit to adjacent ordinal choices". It rewards "fully precise record outputs". It applies "a reference penalty to prevent distribution shift". The third part reads like a penalty on divergence from a reference policy, a common regulariser in reinforcement learning for language models. That reading is ours, not Cloudflare's. Cloudflare publishes no reward formula, so the mechanism cannot be reproduced from the post. The 9 Oct post restates the supervised losses and does not mention RLCD [30]. The TypeSafe launch post also names "Reinforcement Learning for Calibrated Decisions (RLCD)" but gives no mechanism [2]. In this paper "RLCD" always means the objective as Cloudflare describes it for Clef. ZETTL has not implemented it.

Benchmarks (vendor-reported). Cloudflare reports results on its own "Jev Decision Index" [29]. Table 1 reproduces the headline figures. On TypeSafe workflow evaluations, Cloudflare reports that Clef wins three of four areas. The exception is agent-trace observability, where Cloudflare reports Jev at 71.6, Clef at 68.5 and Clef-flash at 69.8. The model cards report a separate internal run, "Decision Index 0.2.1" [33]. On it, Clef scores 98.5 on BFCL, 90.3 on MMLU, 48.0 on GPQA Diamond and 99.0 on ARC-Easy. Clef-flash trails Clef on GSM8K (67.3 against 80.8) and on RAGTruth F1 (35.6 against 79.4). For Typesafe workflow evaluations, the card reports an invoice-processing exact-action score of 64.7 and a customer-service score of 76.3 for Clef. The cards display ForecastBench Brier figures of 13.9 for Clef and 10.6 for Clef-flash, with lower being better. [Unverified] The cards do not explain the scale, and we could not confirm whether these figures are Brier scores multiplied by 100.

Table 1: Clef benchmark figures as reported by Cloudflare on 1 Oct 2026 [29]. All figures are vendor-reported and unreplicated. None measures a payment-record task.

Benchmark (metric)	Clef	Clef-flash	Jev	Note
BFCL (case exact)	98.47	98.76	95.75	
BANKING77 (macro-F1)	94.20	90.93	79.74	intent queries, not transactions
CLINC150+OOS (macro-F1)	97.43	66.77	89.27	
BRIGHT (nDCG@10)	45.91	39.26	47.52	Jev higher
PhishNChips (accuracy)	79.60	75.05	62.55	"DiffusionGemma Jev" 85.35
Median latency (ms)	209.3	38.8	524.1	Laya 5.8
p95 latency (ms)	238.6	122.4	536.0	Laya 222.5

Price, licence and data policy. The launch post says Cloudflare is "fully open-sourcing these models on Hugging Face" under "an Apache 2.0 license" [29]. Both model cards list Apache-2.0 [33]. Clef-omni is released as "open-weights on HuggingFace", but its post does not state a licence [30]. Prices per million input tokens, as of 9 Oct 2026, are as follows. Clef-omni "launched today at $0.15 per M input tokens". Clef "remains at $0.24 per M input tokens". Clef-flash is "now $0.038 per M input tokens", down from $0.09. The same post reports that median Clef latency improved by roughly 1.7–2.0 times with no new weights. On data, Cloudflare states "we don't read, store, or train on your requests or responses" unless the customer opts into fine-tuning [29]. The post also describes an RL fine-tuning platform. It captures datasets through AI Gateway, runs rollouts on Workers AI, uses Containers as an RL sandbox, adds a new Trainer, and lets customers redeploy their own models. It starts as a service led by forward-deployed engineers.

What Cloudflare does not state. Neither the posts nor the cards describe an abstention or "unsure" output. Neither reports expected calibration error (ECE) or a reliability diagram. Neither gives the RLCD formula or the dataset size. The cards have no limitations section and do not name training data or losses [33]. The only calibration-related number is the ForecastBench Brier figure, whose scale is unexplained.

2.1.2 Jev

Launch claims. TypeSafe's founder, Diogo Almeida, announced "System One Models" and Jev on 15 Sep 2026 [2]. A System One model is "a new class of frontier models built to make fast, structured decisions that software can use directly". Intended uses are to "classify, route, score, extract, or branch". The post claims: "Calibrated: higher confidence means higher accuracy." It also says Jev "can't hallucinate" and that "The model never makes type errors". The post concedes that the zero hallucination figure "is not empirical". That figure follows from schema-constrained outputs, not from measured correctness. Reported end-to-end response time is "70ms-500ms", from evaluations "generally run from our laptops on the West Coast". Input pricing is "$0.042 / MTok". The reference answers for TypeSafe's evaluations are "the average of GPT-6 Astra and Fable 5.1", and the vendor acknowledges that this may introduce bias. Jev is

in early access. TypeSafe discloses no weights, parameter count or architecture details beyond "a new model architecture".

Question types. TypeSafe's documentation defines three question types [216]. A `noul` question is yes/no; its `criteria` are optional and use `true` and `false` keys. A `choice` question requires a `criteria` map with at most 255 options. A `score` question takes an ordered `criteria` array of 2 to 10 levels. Instructions may be a string, an object or an array. A noul answer is a single number in $[0, 1]$. A choice answer gives `choice`, `probabilities` and `confidence`. A score answer gives a probability-weighted `score`, which may fall between levels, plus `legend`, `probabilities` and `confidence`. Requests are limited to a 32 MiB body and 64 nesting levels, and numbers must be finite. The documentation advises sending identifiers and large exact values as strings.

Confidence semantics. The most important detail for Zlef is what Jev's `confidence` field means. It is not a calibrated probability of being right. It is a fixed formula applied to the probability vector [216]. For a choice question with n options and largest probability $p_{\max}$,

$$\mathrm{conf}_{\mathrm{choice}} = \frac{p_{\max} - 1/n}{1 - 1/n}. \tag{1}$$

For a score question with levels i, probabilities p_i and modal level m,

$$\mathrm{conf}_{\mathrm{score}} = \max\Big(0,\ 1 - \frac{\sum_i p_i \, |i - m|}{\mathrm{MAD}_{\mathrm{unif}}}\Big), \tag{2}$$

where $\mathrm{MAD}_{\mathrm{unif}}$ is the corresponding mean absolute deviation under a uniform distribution. A noul answer has no confidence field. The documentation suggests $|2p - 1|$ as a comparable number. TypeSafe itself calls this "one reasonable way to summarize a distribution, not the only one". Equation (1) rescales $p_{\max}$ so that a uniform vector maps to 0 and a one-hot vector maps to 1. If the probabilities are calibrated, $p_{\max}$ is the probability that the top choice is right, but $\mathrm{conf}_{\mathrm{choice}}$ is not. For example, with $n = 2$ and $p_{\max} = 0.75$, the confidence is 0.5. A system that thresholds Jev's confidence is therefore thresholding a monotone transform of $p_{\max}$ whose meaning depends on n.

Threshold guidance. The documentation gives example floors of 0.5 ("genuinely unsure") and 0.9 for destructive actions [216]. It also states: "The correct threshold values depend on your domain and the performance of the model for your use case." It notes that "0.5 is a weak answer among two options and a strong one among ten." It advises that noul answers in the middle range should go to a person. The primer says that "Calibration makes uncertainty usable by software" and that "These rates describe groups of predictions, not a guarantee about any single answer." The pages we reviewed show no ECE, reliability diagram or held-out calibration study for Jev.

Abstention. Neither vendor states an abstention mechanism. Clef's documentation does not mention one. TypeSafe's documentation offers threshold guidance for routing low-confidence answers to a person, which is a policy the caller applies. It is not a model output, and it comes with no coverage or risk guarantee.

Table 2: Clef and Jev compared on properties that matter for Zlef. Sources: [2, 29, 30, 33, 216]. All entries are vendor statements as of 15 Oct 2026.

Property	Clef / Clef-flash	Jev
Weights	Open, Apache-2.0 (omni licence not stated)	Not released
Architecture	Frozen Qwen backbone (names as quoted), routing head, rank-256 adapters	"a new model architecture"
Output	One logit per allowed option, softmax per question	Probabilities per option; noul in $[0, 1]$
Question types	noul, choice, score	noul, choice ($\leq$255), score (2–10 levels)
Losses	Label-smoothed CE plus Brier; RLCD secondary	Not disclosed; RLCD named
Confidence field	Returned; definition not given	Fixed formula of p (Equations (1) and (2))
Calibration evidence	ForecastBench Brier, scale unexplained	Claim only; no ECE shown
Abstention	Not stated	Not stated; threshold guidance only
Context	64k claimed; hosted Clef-flash 24k	32k (per Cloudflare)

Implication for Zlef. Clef is the closest public analogue to Zlef. It has a frozen backbone, a typed schema, probability-only outputs and a training loss that includes a proper scoring rule. It gives Zlef a precedent for two commitments: no free text from the scorer, and a proper-scoring term in training. It does not give Zlef any evidence on calibration, abstention or payments. Neither vendor reports class-wise calibration or calibration under shift, and neither offers formal abstention. Zlef must therefore establish its calibration and abstention claims through its own evaluation (Section 8). Four further points follow. First, Zlef should use the probability vector, after post-hoc calibration on ZETTL data, and not a vendor confidence statistic such as Equation (1). Second, the Clef-flash context correction shows that a serving configuration can change while weights stay fixed. A model registry should therefore pin the serving configuration, including the context limit and hosted or self-hosted mode, as well as the weights. A silent drop in context length could truncate a payment memo or an attached invoice. Third, the invoice-processing exact-action score of 64.7 is the most payment-adjacent number either vendor publishes. If it is anywhere near representative, a payment workflow needs an abstain-and-review path. Fourth, TypeSafe's advice to send large exact values as strings applies directly to token amounts and addresses, which must not be rounded as floating-point numbers. Clef and Jev remain candidate external baselines, used under the hosting and data-permission conditions in Paper IV.

2.2 Reinforcement learning for calibration

Supervised training with a proper scoring rule rewards calibrated probabilities directly. Reinforcement learning is different. A policy is rewarded for its sampled outputs, and a binary reward for being right gives no reason to report uncertainty honestly. Cloudflare describes an RL objective aimed at calibration, and TypeSafe names the same objective; neither publishes it.

RLCR. Damani et al. [51] give the nearest peer-reviewed analogue with a published formula and proof. Their method, Reinforcement Learning with Calibration Rewards (RLCR), asks a language model for an answer y and a verbalised confidence $q \in [0, 1]$. With y^* the correct answer, the reward is

$$R_{\text{RLCR}}(y, q, y^*) = \mathbb{1}[y \equiv y^*] - \big(q - \mathbb{1}[y \equiv y^*]\big)^2. \tag{3}$$

The second term is a Brier score on the event "the answer is correct". Their Theorem 1 shows that, for a given answer, expected reward is maximised when q equals the probability that the answer is right. That is a calibration incentive. Among calibrated reports, the answer most likely to be correct earns the most. That is a correctness incentive. The result holds for any *bounded* proper scoring rule. The unbounded log score does not meet the condition, and with it a degenerate policy can become preferable. Empirically, the authors report better calibration without loss of accuracy, in and out of domain. They also report that RL with a binary reward degrades calibration. The paper appeared on arXiv on 22 Jul 2025 and is marked as an ICLR 2026 conference paper.

RLCR and RLCD are different objects. RLCR has a formula and a theorem. RLCD, as Cloudflare describes it, has three named ingredients and no formula (Section 2.1.1). Two of those ingredients are not covered by RLCR's theorem. "Partial credit to adjacent ordinal choices" rewards near misses on ordered scales. Whether such a reward keeps the calibration incentive depends on its exact form, which is unpublished. The "reference penalty" is a regulariser. It trades reward against closeness to a reference model, so the optimum need not be the calibrated report. The names differ by one letter, and we keep them apart.

Calibration of language-model probabilities. Two further results bear on any design that reads probabilities out of a language model. Kadavath et al. [128] find that large language models are reasonably calibrated on multiple-choice and true/false formats when the options are shown explicitly. They also show that models can be trained to predict whether they know the answer. Calibration degrades under format changes and some distribution shifts. Tian et al. [210] find that, for models tuned with reinforcement learning from human feedback, verbalised confidences are often better calibrated than the model's token probabilities, which that tuning tends to distort. Zhao et al. [234] show that few-shot language-model classifiers are biased toward majority labels, recent labels and common tokens, and they correct this with "contextual calibration" using content-free inputs.

Implication for Zlef. If Zlef ever uses an RL objective, it should use a bounded proper score, such as the Brier score, as its calibration term, because that is the condition under which RLCR's incentive result holds. An RLCD-style objective, with partial credit for adjacent ordinal levels, stays a [Research hypothesis] for ordinal questions. It is not a method that ZETTL can reproduce from Cloudflare's description. Any reference penalty should be treated as a regulariser whose effect on calibration has to be measured. Kadavath et al. support Zlef's choice of closed questions with explicit options. Tian et al. and Zhao et al. warn that raw option probabilities from a tuned backbone may be miscalibrated or biased by prompt format. Post-hoc calibration and tests for invariance to option order are therefore mandatory, not optional.

2.3 Calibration

A probabilistic classifier is *calibrated* if, among the cases to which it assigns probability p to an outcome, that outcome occurs a fraction p of the time. Clef needs calibration because downstream thresholds, cost calculations and abstention rules all treat its outputs as probabilities.

2.3.1 Proper scoring rules

A scoring rule assigns a penalty $S(p, y)$ to a forecast distribution p when outcome y occurs. It is *proper* if the expected penalty under the true distribution is minimised by reporting that distribution. It is *strictly* proper if that minimiser is unique. Gneiting and Raftery [98] give the canonical treatment, covering the logarithmic, quadratic (Brier), spherical and continuous ranked probability scores. For a K-class forecast p and one-hot outcome e_y, the two rules used most in this paper are

$$S_{\log}(p, y) = -\log p_y, \tag{4}$$

$$S_{\text{Brier}}(p, y) = \sum_{k=1}^{K} \big(p_k - \mathbb{1}[y = k]\big)^2. \tag{5}$$

The Brier score was introduced for weather forecasts [19]. It can be decomposed into reliability, resolution and uncertainty terms, which separate miscalibration from lack of discrimination. For ordered outcomes, the ranked probability score (RPS) compares cumulative distributions:

$$S_{\text{RPS}}(p, y) = \sum_{j=1}^{K-1} \Big(\sum_{k \le j} p_k - \mathbb{1}[y \le j] \Big)^2. \tag{6}$$

It is a discretisation of the continuous ranked probability score and is proper [98]. Unlike Equation (5), it penalises probability mass less the closer it sits to the true level. That makes it a principled substitute for ad hoc "partial credit".

2.3.2 Calibration error and its estimators

Naeini et al. [151] gave expected calibration error (ECE) and maximum calibration error (MCE) the form now in standard use, alongside a Bayesian binning method. Guo et al. [102] popularised ECE and reliability diagrams for deep networks. Given N predictions with top-label confidence $\hat{c}_i$ and correctness $a_i \in \{0, 1\}$, partitioned into bins $B_1, \dots, B_M$,

$$\widehat{\text{ECE}} = \sum_{m=1}^{M} \frac{|B_m|}{N} \Big| \bar{a}(B_m) - \bar{c}(B_m) \Big|, \qquad \widehat{\text{MCE}} = \max_m \Big| \bar{a}(B_m) - \bar{c}(B_m) \Big|, \tag{7}$$

where $\bar{a}$ and $\bar{c}$ are the mean accuracy and mean confidence in a bin.

This estimator has known flaws. Nixon et al. [155] show that it is sensitive to the number and placement of bins. They also show that it looks only at the top label and so ignores calibration of the other classes. They propose Static Calibration Error (SCE), which averages a binned error over every class, and Adaptive Calibration Error (ACE), which uses equal-mass bins, plus thresholded variants. A further flaw follows from elementary probability. Even for a perfectly calibrated model, the within-bin difference $\bar{a}(B_m) - \bar{c}(B_m)$ is a noisy quantity with mean zero. Its absolute value has positive mean. The plug-in estimator Equation (7) is therefore biased upward,

and the bias grows as bins get smaller. With skewed label distributions, many bins are nearly empty and the bias is large. Estimators designed to remove this bias have been proposed in the literature. [Unverified] We do not cite a specific debiased estimator in this edition, and Section 8.6 states how Zlef handles the bias until one is added.

2.3.3 Top-label and class-wise calibration

Top-label calibration asks whether the confidence in the predicted class matches its accuracy. Class-wise calibration asks, for every class k, whether $\Pr(y = k \mid p_k = s) = s$. Kull et al. [134] formalise class-wise calibration and class-wise ECE and show they differ from top-label calibration. A model can be calibrated on its top label and still badly miscalibrated on a rare class, because the rare class is seldom the top label and contributes little to top-label ECE.

2.3.4 Post-hoc calibration methods

Post-hoc methods fit a map from model scores to probabilities on held-out data, leaving the model unchanged.

- **Platt scaling** fits a sigmoid $\sigma(as + b)$ to binary scores [181].
- **Histogram binning** replaces a score by the empirical frequency in its bin [232].
- **Isotonic regression** fits a non-decreasing step function. Zadrozny and Elkan [233] extend it to many classes by one-vs-rest or pairwise coupling. It is flexible but needs more data.
- **Bayesian binning into quantiles** (BBQ) averages over binning schemes [151].
- **Temperature scaling** divides all logits by one scalar $T > 0$ fitted on held-out negative log-likelihood. It does not change the predicted class. Guo et al. [102] find it a strong fix for modern networks, which are typically overconfident.
- **Dirichlet calibration** fits a multinomial logistic regression on log-probabilities with off-diagonal regularisation [134]. It generalises beta calibration to many classes and can correct class-specific miscalibration that one temperature cannot.

Minderer et al. [147] find that recent non-convolutional architectures are better calibrated than older convolutional networks. Model size and amount of pretraining do not fully explain the difference. For these models, calibration under shift tracks in-distribution calibration more closely.

2.3.5 Calibration under distribution shift

Ovadia et al. [169] compare uncertainty methods under dataset shift. Calibration fitted in-distribution, including temperature scaling, degrades as shift increases. Ensembles hold up best among the methods they compare. A calibration map is a fact about one data distribution. It does not carry over to a new one.

Implication for Zlef. Zlef should train with proper scores and report them per question: negative log-likelihood, Brier score and, for ordinal questions, RPS. The binned ECE in Equation (7) is not enough. Zlef should report adaptive-bin ECE and class-wise SCE, with uncertainty intervals and a correction for estimator bias, because payment labels are skewed and the plug-in

estimator overstates error in sparse bins. Class-wise calibration matters most for rare classes whose consequences are large, such as `TKN.EQ` under `asset_class`, `PRV`, or `vasp_unmatched` under `counterparty_type`. MCE, the worst bin, matters for the high-confidence region in which labels would be applied without review. Temperature scaling per question is the cheapest calibration baseline. Dirichlet calibration is the class-wise alternative. Isotonic regression suits high-volume binary questions. All calibration claims hold only in-distribution, so they must be re-measured after every adapter, backbone or serving change and on a schedule (Section 8.7). Minderer et al. caution against assuming that a large frozen backbone brings good calibration with it.

2.4 Selective prediction, abstention and learning to defer

A selective classifier may decline to answer. The literature treats abstention as a trade-off between the error on the cases it answers and the fraction it answers.

Chow's rule. Chow [31] derives the Bayes-optimal reject rule. With posterior $p(y \mid x)$, a cost of 1 for an error and a cost $c_r < 1$ for a rejection, the optimal rule rejects when

$$\max_k p(k \mid x) < 1 - c_r. \tag{8}$$

The paper also characterises the error–reject trade-off curve. With calibrated probabilities and stated costs, the threshold is a consequence of the cost policy, not a free hyperparameter. Section 2.8 extends this to unequal error costs.

Risk and coverage. El-Yaniv and Wiener [63] formalise selective classification as a trade-off between selective risk, the error rate on accepted cases, and coverage, the fraction accepted, with learning-theoretic guarantees. Geifman and El-Yaniv [95] make this practical for deep networks. Given a trained network and a target risk, their selection-with-guaranteed-risk (SGR) algorithm picks a confidence threshold with a high-probability bound on selective risk. SelectiveNet trains a selection head jointly with the predictor for a target coverage [96].

Deferral to a specific expert. Rejection sends a case nowhere in particular. Deferral sends it to a decision-maker whose own accuracy matters. Madras et al. [140] model the downstream decision-maker's accuracy and biases and show that learning to defer can improve both the accuracy and the fairness of the combined system. Mozannar and Sontag [149] give a consistent surrogate loss for learning a classifier and a deferral rule together from expert predictions. Raghu et al. [185] argue that optimising a model for overall accuracy differs from optimising it for triage, that is, for deciding which cases to automate and which to send to people.

Human–AI teams. Deferral assumes that the human adds something. Bansal et al. [13] find that AI explanations raised human acceptance of AI recommendations whether or not those recommendations were correct. They did not reliably make the team better than the better of the human or the AI alone.

Implication for Zlef. Zlef's abstention should rest on this literature, not on vendor confidence fields. Chow's rule makes the threshold a cost-policy decision once probabilities are calibrated. SGR gives a finite-sample bound on the error rate of labels applied without review. Evaluation should report risk–coverage curves, not a single accuracy figure. Zlef defers to reviewers with different skills, such as tax specialists and AML analysts, and to versioned rules. Once ZETTL has logs of reviewer decisions, a learned deferral rule with a consistent loss becomes possible. The goal is the performance of the whole system under a review-capacity limit, not the model's standalone accuracy. Because explanations can increase acceptance of wrong answers, the review interface should show calibrated probabilities first. Section 2.13 returns to this point.

2.5 Conformal prediction

Conformal prediction turns any score into a prediction set with a coverage guarantee. Under exchangeability of calibration and test data, a set $C(x)$ built at level α satisfies

$$\Pr\left(y_{\text{test}} \in C(x_{\text{test}})\right) \geq 1 - \alpha \tag{9}$$

in finite samples, whatever the model [226]. The guarantee is *marginal*: it holds on average over test points, not for each one.

Split conformal. Angelopoulos and Bates [3] give a practitioner's tutorial with code. In the split version, one computes a non-conformity score $s(x, y)$ on n held-out calibration points. One then takes the $\lceil (n+1)(1-\alpha) \rceil / n$ empirical quantile $\hat{q}$ and returns $C(x) = \{y : s(x, y) \leq \hat{q}\}$. The tutorial also covers checks of coverage and of conditional coverage.

Adaptive sets. With score $s(x, y) = 1 - p_y(x)$, sets have similar size for easy and hard inputs. Adaptive Prediction Sets (APS) use instead the cumulative probability mass of classes ranked above and including y [187]. Sets then grow for ambiguous inputs. Regularised APS (RAPS) adds a penalty on the rank of y beyond a set size, which yields smaller and more stable sets at the same coverage [4]. This helps when there are many classes.

Class-conditional coverage. Marginal coverage can hide under-coverage on rare classes. Mondrian conformal predictors calibrate separately within categories, for example per class, and so guarantee coverage within each category [225]. [Unverified] The Mondrian working paper is commonly cited as 2003, but we could not confirm its year. When some classes have few calibration examples, Ding et al. [55] pool calibration data across classes with similar score distributions (clustered conformal prediction) to approximate class-conditional coverage.

Risk control. Angelopoulos et al. [5] generalise the coverage guarantee to the expected value of any bounded loss that is monotone in a tuning parameter. If $L(\lambda)$ is such a loss, conformal risk control picks $\hat{\lambda}$ so that $\mathbb{E}[L(\hat{\lambda})] \leq r^*$ for a chosen target r^*.

Shift. Exchangeability fails when the test distribution differs from the calibration distribution. Tibshirani et al. [211] restore coverage under covariate shift by weighting calibration points with likelihood ratios between test and calibration feature distributions, where those ratios are known or can be estimated. Time-ordered streams also violate exchangeability, so weighted or online variants matter for them.

Implication for Zlef. Conformal sets give Zlef an abstention signal that does not depend on the model being calibrated. If the set for `asset_class` at the chosen level contains more than one class, the case can go to review. Set size can also drive routing. APS and RAPS are suitable scores for Zlef's choice questions, and RAPS matters most for `event_type`, which has twenty-five values. Marginal coverage is not enough for classes with large downstream consequences. Zlef should use Mondrian calibration for those classes, and clustered conformal calibration where data are thin. Conformal risk control links abstention to cost: it can bound the expected cost of labels applied without review, not just the miscoverage rate. Payment streams are time-ordered and new chains change feature distributions, so the exchangeability assumption must be checked and, where it fails, weighted conformal methods considered. Section 7.5 develops these uses.

2.6 Ordinal and hierarchical classification

Two of Zlef's question types have structure that a flat softmax ignores. Ordinal questions have ordered levels. Hierarchical questions have classes arranged in a tree.

Ordinal decomposition. Frank and Hall [91] decompose a K-level ordinal problem into $K-1$ binary problems of the form "is $y > j$?" and combine the answers. The approach respects order but can produce cumulative probabilities that are not monotone in j.

CORAL. Cao et al. [22] share weights across the $K-1$ binary heads and give each its own bias:

$$\Pr(y > j \mid x) = \sigma\big(w^\top h(x) + b_j\big), \qquad j = 1, \ldots, K-1. \tag{10}$$

If the biases are ordered, the cumulative probabilities are monotone, which they call rank consistency. A calibrated $\Pr(y \geq j)$ for every j supports rules such as "review if $\Pr(\text{priority} \geq \text{high}) > \tau$".

Hierarchical classification. Silla and Freitas [199] survey hierarchical classification across domains. They distinguish flat classifiers, local classifiers per node, per parent node and per level, and global classifiers. They also cover hierarchical evaluation metrics, including hierarchical precision, recall and F-measure. These count a prediction as partly right if it shares ancestors with the true class.

Making better mistakes. Bertinetto et al. [16] point out that top-1 error treats all mistakes alike, while some mistakes are worse than others. They propose a hierarchical cross-entropy (HXE) and "soft labels", a hierarchy-aware form of label smoothing in which target mass decays with tree distance from the true class. They report a trade-off between top-1 accuracy and mistake severity. We found no single canonical paper titled "hierarchical label smoothing". We use Bertinetto et al.'s soft labels as the citable representative and use the generic term only descriptively.

Implication for Zlef. The `asset_class` question is hierarchical: `STB.SYN` sits under `STB`, and `WRP.BRIDGE` under `WRP`. Zlef should be able to stop at an internal node, such as `STB`, when it cannot choose among the children. That is a principled partial abstention. It should be evaluated with hierarchical metrics and a measure of mistake severity, not just flat accuracy.

`evidence_completeness` (0–4) and `review_priority` (low to urgent) are ordinal. CORAL-style heads give monotone cumulative probabilities that map onto threshold rules, and RPS gives a proper score that credits near misses. Both are better grounded than an unpublished partial-credit reward. Soft labels are the hierarchical counterpart of the uniform label smoothing Clef uses.

2.7 Label noise, annotation and human factors

Noisy labels. Natarajan et al. [152] show that, under class-conditional random label noise, unbiased estimators and weighted surrogate losses give risk-consistent learning. Northcutt et al. [157] introduce confident learning. It estimates the joint distribution of noisy and true labels from out-of-sample predicted probabilities and flags likely label errors. A companion study finds label errors in widely used *test* sets [156]. When test labels are wrong, measured accuracy and calibration are wrong too.

Aggregating annotators. Dawid and Skene [52] estimate true labels and a confusion matrix for each annotator jointly, by maximum likelihood with the EM algorithm. With items i, annotators r, latent true class z_i with prior π, and annotator confusion matrices $\Theta^{(r)}$, the model is

$$\Pr\left(\{\ell_{ir}\}\right) = \prod_i \sum_k \pi_k \prod_r \Theta^{(r)}_{k,\ell_{ir}}. \tag{11}$$

Snow et al. [201] show that aggregating a few non-expert crowd labels can approach expert quality on several language tasks. Those tasks were perceptual or linguistic, not tasks that require expert judgement.

Automation bias. Parasuraman and Riley [170] classify how people misuse automation by over-relying on it, disuse it by under-relying on it, and abuse it by deploying it inappropriately. Goddard et al. [99] review automation bias in clinical decision support. They identify mediators, including task complexity, workload and trust, and mitigators, including training, presentation design and accountability. A reviewer shown a model's answer before forming their own may anchor on it.

Implication for Zlef. Zlef's training labels will be noisy. Historical records carry user-entered memos and inconsistent past accounting. Noise-robust losses are a baseline mitigation. The evaluation set must be double-annotated and adjudicated, or ECE and accuracy will be measured against noise. Confident learning can prioritise items for re-review. A Dawid–Skene model of the reviewer pool gives per-reviewer error estimates, which help build gold labels and later feed learning to defer. Snow et al. do not license crowd labelling of payment facts that need expert knowledge. As Paper IV states, reviewer corrections are *candidate* labels, not truth, and corrections made after seeing a model suggestion must be recorded as such. Automation-bias research supports measuring anchoring directly with blind and assisted labelling arms (Section 8.3).

2.8 Cost-sensitive learning

Elkan [64] shows that, given calibrated probabilities and a cost matrix C with C_{jk} the cost of predicting j when the truth is k, the optimal prediction is

$$\hat{y}(x) = \arg\min_j \sum_k p(k \mid x)\, C_{jk}. \tag{12}$$

He shows that shifting thresholds and re-weighting training examples are equivalent ways to account for class-dependent costs, and he states conditions for a "reasonable" cost matrix. MetaCost takes another route [60]. It estimates class probabilities by bagging, relabels each training example to the class with least expected cost, and retrains. This builds the costs into the model.

Implication for Zlef. Elkan's result is the bridge from Zlef's probabilities to actions. Calibrate once, then let a policy set decisions per question from explicit costs. When costs change, for example because a rule changes the severity of a flag, the policy changes and the model does not need retraining. Relabelling methods such as MetaCost bake costs into the weights. They are harder to audit and must be retrained whenever costs change, so Zlef prefers calibrated probabilities plus a separate decision layer. Chow's rule (Equation (8)) is the special case of Equation (12) with equal error costs and an added reject action. Section 7.6 derives Zlef's cost matrices from flag severities.

2.9 Distribution shift and drift

Quiñonero-Candela et al. [184] define the standard types of dataset shift. Under *covariate shift*, $p(x)$ changes but $p(y \mid x)$ does not. Under *prior* or *label shift*, $p(y)$ changes but $p(x \mid y)$ does not. Under *concept shift*, $p(y \mid x)$ changes. Gama et al. [93] survey concept drift: sudden, gradual, incremental and recurring. They cover detection methods and adaptation strategies.

Label shift. Lipton et al. [138] propose black-box shift estimation (BBSE). If $\hat{C}$ is the confusion matrix of a fixed classifier on source data, with $\hat{C}_{jk} = \Pr_{\text{src}}(\hat{y} = j, y = k)$, and $\hat{\mu}$ is the distribution of its predictions on target data, then the importance weights $w_k = p_{\text{tgt}}(y = k)/p_{\text{src}}(y = k)$ solve

$$\hat{\mu} = \hat{C}\, w. \tag{13}$$

The weights can then re-estimate class priors and correct posteriors without new labels.

Benchmarks of shift. WILDS collects real-world domain and subpopulation shifts [133]. Methods tuned on in-distribution data often fail to close the gap out of distribution. Results are reported per domain and for the worst group. Wild-Time is a benchmark for temporal shift [231]. In financial forensics, Weber et al. [227] split the Elliptic dataset by time step, not at random, and report that models failed to generalise after a dark-market shutdown late in the test period. [Unverified] The exact time steps should be re-checked in the paper before they are quoted.

Implication for Zlef. Crypto-asset payments shift suddenly, through token launches, depegs, exploits and rule changes, and recurrently, through tax-year cycles. Zlef needs drift detectors on inputs and also on its own outputs: the mix of predicted classes, confidence and abstention

rates. A change in the mix of payment types is label shift, and BBSE-style prior re-estimation is a candidate correction. A legal reclassification is concept shift in the legal sense. Zlef handles it in versioned rules, not in the model, because Zlef's labels are jurisdiction-free facts. Following WILDS, results should be reported by domain (chain, corridor, counterparty type) and for the worst group. Following Elliptic and Wild-Time, evaluation must use forward-in-time splits.

2.10 Leakage

Kaufman et al. [131] define leakage as the use, in training or evaluation, of information that would not be legitimately available at prediction time. They illustrate it with data-mining competition cases and recommend strict separation of learning and prediction data. Kapoor and Narayanan [129] survey leakage across scientific fields that use machine learning. They propose a taxonomy: no train/test separation, illegitimate features, temporal leakage, non-independence between training and test data, and sampling bias. They also propose "model info sheets" for reporting.

Implication for Zlef. Payment data offer many leakage paths. A record may include its later settlement status, compliance flags raised after the fact, an accountant's category that already encodes the label, or on-chain events after the payment. Zlef's feature store must time-stamp when each fact became available and must exclude facts produced by the rules layer or by review. Kapoor and Narayanan's taxonomy serves as a checklist for the evaluation protocol in Section 8.5. A model info sheet is a natural appendix for later editions.

2.11 Structured outputs, constrained decoding and label smoothing

Constrained decoding. Willard and Louf [229] recast regex- and grammar-constrained generation as transitions of a finite-state machine over the vocabulary, with precomputed indices. This is the basis of the open-source Outlines library. Geng et al. [97] use formal grammars to enforce output structure for information extraction, entity disambiguation and parsing without fine-tuning. Tam et al. [205] report that strict format constraints can degrade reasoning performance. Vendor "JSON modes" are documentation, not literature, and we do not rely on them.

Schema validity is not correctness. Constrained decoding guarantees that output matches a schema. A general language model with constrained decoding therefore also "never makes type errors". That property does not distinguish decision models. The difference lies in the probabilities. Constrained decoding yields token-level probabilities along one sampled path. These are not the same as a calibrated distribution over the classes. A decision model that scores every option in one pass (Section 2.1.1) returns a distribution over classes directly.

Biases of language models as classifiers. Zhao et al. [234] document majority-label, recency and common-token biases in few-shot classification. Zheng et al. [235] document position, verbosity and self-enhancement biases when language models judge other models' outputs. The second finding matters because TypeSafe's reference answers are averages of two language models' outputs [2]. Clef's training on "permutating field orders" [29] targets order sensitivity directly.

Label smoothing. Szegedy et al. [204] introduce label-smoothing regularisation, which replaces a one-hot target e_y with $(1-\varepsilon)e_y + \varepsilon u$ for a uniform vector u. Müller et al. [150] find that label smoothing often improves calibration, by an amount comparable to temperature scaling. It also pulls penultimate-layer representations into tight, equidistant clusters. That erases information about which wrong classes are similar, and it hurts knowledge distillation.

Implication for Zlef. Typed outputs are necessary but not sufficient. A general model with constrained decoding is a fair baseline for schema validity, but its sampled-path probabilities need separate calibration. Where possible, Zlef should derive option probabilities from a single scoring pass over all options, as Clef does, not from sampled JSON. Evaluation should test whether forcing a schema degrades accuracy and whether permuting option order changes answers. Vendor numbers built on model-generated references should be weighed with the biases above in mind. Label smoothing erases "which wrong class is close", which is the information that ordinal and hierarchical questions need. Zlef therefore prefers hierarchy-aware soft labels to uniform smoothing on structured questions. Because smoothing interacts with temperature scaling, Zlef calibrates *after* training in every case.

2.12 Financial transaction and AML machine learning, and legal NLP

Illicit-flow classification on Bitcoin. Weber et al. [227] release the Elliptic dataset, "a time series graph of over 200K Bitcoin transactions (nodes), 234K directed payment flows (edges)". A small fraction of nodes is labelled illicit, and most are unlabelled. They compare logistic regression, random forests and graph convolutional networks, and find random forests on the provided features strong, with graph features helping. The paper is on arXiv and was presented at a KDD 2019 workshop on anomaly detection in finance. [Unverified] We have not confirmed the workshop venue from the arXiv record. Bellei et al. [15] move from node to subgraph classification with Elliptic2: "122K labeled subgraphs of Bitcoin clusters" within "a background graph consisting of 49M node clusters and 196M edge transactions". Their aim is to classify the "shape" of laundering patterns. [Unverified] We found only the arXiv version and could not confirm a venue. Lo et al. [139] feed self-supervised graph embeddings (Deep Graph Infomax with a graph isomorphism network) into a random forest and report improvements over prior baselines on Elliptic. [Unverified] We could not confirm a journal version.

Transaction categorisation. Toran et al. [214] categorise bank transactions with weak supervision from heuristics and domain knowledge, plus text embeddings and discriminative networks, using few manual labels. They state that their method "outperforms market-leading solutions"; that is the authors' claim. García-Méndez et al. [94] classify terse bank descriptors with a Jaccard-based short-text similarity filter and a support vector machine. [Unverified] The journal is inferred from the DOI prefix and has not been confirmed. BANKING77, which appears in Clef's benchmarks, is an intent-classification dataset of banking customer queries [24]. It does not contain transaction records. Merchant category codes are standardised in ISO 18245. [Unverified] We could not confirm the current edition year. We found no peer-reviewed machine-learning paper on predicting merchant category codes.

Goods classification. Lee et al. [137] predict Harmonized System (HS) codes from text descriptions with a KoELECTRA pipeline. It predicts the 4-digit heading, retrieves supporting sentences

from the HS manual, then predicts the 6-digit subheading, using 129,084 cases from Chapter 85. The abstract reports top-3 accuracy of 95.5% over 265 subheadings. The paper's table gives 0.955 for the variant without retrieval and 0.937 for the full model. A field demonstration on eleven cases gave top-3 accuracy of 0.82. The authors note that experts apply General Rules of Interpretation and that the HS manual changes about every five years, so past cases can mislead.

Legal NLP. LexGLUE collects seven English legal tasks [26]. Legal-domain pretraining helps, but performance varies widely by task. Katz et al. [130] report that GPT-4 scored about 297 on a zero-shot run of the Uniform Bar Examination, above the passing threshold in every jurisdiction that uses it. Martínez [141] argues that the accompanying 90th-percentile claim is "overinflated". His estimates for the July exam fall below the 69th percentile overall, about the 48th on essays, and about the 15th on essays among those who passed. [Unverified] The DOI of the Martínez paper has not been confirmed.

Implication for Zlef. Laundering risk is a property of flows and subgraphs, not of one payment. A Zlef question about AML risk should take graph-risk features computed upstream, outside the decision model, as inputs. Zlef should not be expected to see laundering in a single record. For that reason, the Edition 1 schema has no AML-risk question. Payment memos resemble bank descriptors: short, abbreviated and noisy. Weak supervision is a practical way to bootstrap labels, subject to the noise caveats of Section 2.7. Vendor BANKING77 scores say little about Zlef's task. The HS-code work is the closest analogue to Zlef's design: the model *suggests* with top-k answers, confidence and evidence, and rules plus experts *decide*. Its stated limitation, that rule revisions invalidate historical labels, is why Zlef keeps legal and tax treatment in versioned rules. The legal-NLP results show that legal classification can be benchmarked but is far from solved, and that headline legal claims for AI depend on the comparison chosen. This supports keeping legal determinations out of the model and scrutinising vendor benchmarks.

2.13 Explanation faithfulness

Rudin [188] argues that post-hoc explanations of black-box models can be unfaithful and that high-stakes decisions should use models that are interpretable by design. Jacovi and Goldberg [125] separate *plausibility*, meaning that an explanation convinces people, from *faithfulness*, meaning that it reflects the model's actual reasoning. They argue that the two are often conflated. Turpin et al. [215] give direct evidence. Biasing features, such as reordering few-shot examples so that the answer is always "(A)", change model predictions, while chain-of-thought explanations systematically fail to mention them. Bansal et al. [13] find that explanations raised acceptance of AI recommendations whether or not they were correct (Section 2.4).

Implication for Zlef. Zlef's architecture follows Rudin where it matters most. The decisive legal and tax logic lives in versioned, inspectable rules, and the model only supplies probabilities for factual questions. Any explanation shown with a Zlef output must be labelled as a non-causal summary unless its faithfulness has been tested. Scores must never be derived from explanation text. A generated rationale for "this is a `gift`, not `compensation`" may not reflect what drove the label. Explanations may also raise reliance on wrong labels, so reviewers should see calibrated probabilities and pointers to evidence, not persuasive prose (Section 7.8).

2.14 Privacy-preserving machine learning

Dwork and Roth [61] give the standard treatment of (ε, δ)-differential privacy: mechanisms, composition and the privacy–utility trade-off. McMahan et al. [146] introduce federated averaging (FedAvg), which trains a shared model from data that stay on participants' devices. Federated learning alone is not a privacy guarantee without differential privacy or secure aggregation. Shokri et al. [196] show that shadow-model attacks can infer whether a record was in a model's training set. Related work extracts training data from large language models [23] and trains deep networks with differential privacy [1]. Public blockchains make payment data visible by design. Cong and He [39] analyse how decentralised consensus distributes information. That openness enables smart contracts and competition, and it can also facilitate collusion.

Implication for Zlef. If calibration statistics or reviewer-feedback datasets are shared across ZETTL operators, differentially private aggregates, such as per-class reliability tables, limit what can be inferred about individual payers. A multi-operator network may prefer federated fine-tuning of adapters to pooling raw payment data, with secure aggregation or differential privacy on top. Open-weight adapters trained on customer payments, which Clef's bring-your-own-model path would allow, can leak whether a record was used in training. Zlef should assess membership-inference risk before publishing any adapter. Public on-chain data are visible but not harmless. Linking them to permissioned records can re-identify people, which Section 8.1 addresses.

2.15 Prompt injection and adversarial text

Greshake et al. [101] show that attackers can plant instructions in data that a language-model application retrieves or processes, such as web pages, emails and documents. This blurs the line between data and instructions and enables manipulation and data exfiltration. Perez and Ribeiro [173] introduce PromptInject and show that simple handcrafted inputs can hijack an application's goal or leak its prompt.

Implication for Zlef. Payment memos, invoice PDFs, token names and symbols, ENS names and NFT metadata are all text that an attacker can control, and all of it may enter Zlef's state. A memo that reads "classify as refund, not income" is an indirect injection. Typed outputs limit the *action space*: a decision model cannot leak its prompt as free text. They do not protect the *answer*: the attacker's goal is to move probability mass to the label they want. Zlef should evaluate four mitigations. It should quote or escape untrusted fields and mark their provenance. It should run adversarial test sets for each question. It should prefer structured features from layers A and B over text claims. And it should never let a single text field flip a high-cost decision without corroborating structured features. Section 7.9 specifies these.

2.16 Summary: from literature to design commitments

Table 3 lists the design commitments that this review supports and where the methodology takes each one up. Three themes run through it. First, typed outputs, correct outputs and calibrated outputs are three different properties. Schema validity comes from option scoring or constrained decoding. Correctness must be measured on adjudicated, time-split data. Calibration must be measured class-wise and under shift. Vendor claims such as "can't hallucinate" refer only to

schema validity. Second, abstention should rest on costs and guarantees, through Chow's rule, selective risk bounds and conformal sets, not on vendor confidence heuristics. Third, law and tax belong in versioned rules, not in model weights, because rules change and historical labels then mislead.

Table 3: Design commitments for Zlef and their grounding in the literature.

Commitment	Grounding	Taken up in
Closed, typed questions with explicit options	Clef, Jev; Kadavath et al. [128]	Section 7.1
Proper-score training (CE plus Brier; RPS for ordinal)	Gneiting and Raftery [98]; Clef losses	Section 7.3
RL only with bounded proper score; RLCD a hypothesis	Damani et al. [51]; Cloudflare RLCD	Section 7.3.3
Post-hoc calibration per question and per class	Guo et al. [102], Kull et al. [134], Zadrozny and Elkan [233]	Section 7.4
Adaptive and class-wise calibration metrics	Kull et al. [134], Nixon et al. [155]	Section 8.6
Use p, not vendor confidence	TypeSafe confidence formula	Section 7.5
Cost-derived thresholds	Chow [31], Elkan [64]	Section 7.6
Selective risk guarantees; risk–coverage reporting	El-Yaniv and Wiener [63], Geifman and El-Yaniv [95]	Section 8.6
Conformal sets, class-conditional for high-stakes classes	Angelopoulos et al. [4], Ding et al. [55], Romano et al. [187]	Section 7.5
Cost-bounded automation	Angelopoulos et al. [5]	Section 7.5
Back-off to parent class	Bertinetto et al. [16], Silla and Freitas [199]	Section 7.3.2
Adjudicated evaluation labels; candidate labels from review	Dawid and Skene [52], Northcutt et al. [157]	Section 8.2
Measure anchoring of reviewers	Goddard et al. [99], Parasuraman and Riley [170]	Section 8.3
Customer, time, chain and asset splits	Kapoor and Narayanan [129], Kaufman et al. [131], Koh et al. [133]	Section 8.5
Drift monitoring of outputs and abstention	Gama et al. [93], Lipton et al. [138], Ovadia et al. [169]	Section 8.7
Explanations kept out of scoring	Jacovi and Goldberg [125], Rudin [188], Turpin et al. [215]	Section 7.8
Text fields treated as adversarial	Greshake et al. [101], Perez and Ribeiro [173]	Section 7.9

Commitment	Grounding	Taken up in
Membership-inference check before release	Dwork and Roth [61], Shokri et al. [196]	Section 8.1
Law and tax in versioned rules	Chalkidis et al. [26], Lee et al. [137], Martínez [141]	Section 7.1

3 A Legal Ontology of Crypto-Assets

A payment decision model must know what is being paid. For a card or a wire the answer is simple: a claim on a bank denominated in a currency. For a crypto-asset payment the answer is not simple. The same token can be a commodity in one country, a financial instrument in a second, an e-money substitute in a third and an object that the law does not name in a fourth. The answer can also change on a known date, because many regimes in our country set are enacted but not yet effective, or proposed. This section builds the ontology that the rest of the paper uses. It makes three claims.

1. A legal category is not a property of a token. It is the value of a function of the token's facts, the jurisdiction and the date. We formalise this in Section 3.1.
2. The reference taxonomies in force or proposed on 15 Oct 2026 can be mapped onto a small set of jurisdiction-free economic and technical facts. Zlef predicts those facts. A versioned rules layer computes the legal category (Sections 3.2 and 3.3).
3. Regimes disagree in predictable places: wrapped and bridged tokens, yield-bearing stablecoins, non-unique NFTs, staking receipts and payment use. The model should flag these places for review or abstain there, not hide them (Sections 3.4 and 3.5).

Every legal statement below carries a status marker and an as-of date. Unless stated otherwise the as-of date is 15 Oct 2026. Nothing in this section is legal advice. We describe rules, triggers and uncertainty. We do not reach conclusions about specific issuers or tokens beyond what a cited source states.

3.1 Legal category as a function of facts, jurisdiction and date

The problem with a single label. A common design for crypto compliance tools gives each token one label, such as "security", "stablecoin" or "utility". That design fails for three reasons that the sources make concrete.

First, the label depends on the jurisdiction. [In force] Under the joint SEC/CFTC interpretation effective 23 Mar 2026, BTC, ETH and SOL are named examples of a "digital commodity" and are not securities [172, 194]. [In force] Under MiCA the same assets are "other crypto-assets" within Title II [72]. [In force] In India they are "virtual digital assets", a concept used for tax and anti-money-laundering purposes only, with no regulatory classification law behind it [49, 178]. Three regimes, three categories, one token.

Second, the label depends on the date. [Enacted, not yet effective] The UK's FSMA 2000 (Cryptoassets) Regulations 2026 were made on 4 Feb 2026 but commence in full on 25 Oct 2027 [217]. Until then the UK category "qualifying cryptoasset" exists in the statute book but does not yet govern conduct. [Enacted, not yet effective] The GENIUS Act was enacted on 18 Jul 2025, but no issuer can be

"permitted" before its effective date, which secondary sources give as the earlier of 18 Jan 2027 or 120 days after final regulations [48, 218]. The statutory wording of the effective-date clause was not read directly in our sources. A label computed on 15 Oct 2026 and the label computed on 1 Feb 2027 can therefore differ for the same token and the same jurisdiction.

Third, the label often depends on facts that are not visible on chain. [In force] Under the US interpretation, a non-security crypto-asset can be sold *subject to* an investment contract, and that contract ends when the issuer has fulfilled, failed or abandoned its promised essential managerial efforts [172]. Whether an issuer has done so is a question about conduct and promises, not about token transfers.

Formal set-up. We separate what a model can learn from evidence from what a rule must decide. The definitions below fix notation for the rest of the paper.

Definition 3.1 (Fact tuple)**.** For a transfer or holding event τ, the *fact tuple* is

$$x(\tau) = (a, e, p, b, c, s, v, r),$$

where a is the asset class (Zlef question Q1, `asset_class`), e the event type (Q2), p the economic purpose (Q3), $b \in \{\text{yes}, \text{no}, \text{unknown}\}$ the beneficial-owner change (Q4), c the counterparty type (Q5), s the travel-surface set (Q8), v the protocol-authenticated quantities (amounts, timestamps, contract addresses; fact layer A), and r the residency, entity and authorisation facts (fact layers B, C or F). The components a, e, p, b, c, s are *jurisdiction-free*: their definitions refer to economic and technical properties, never to a statute.

Definition 3.2 (Legal category map)**.** Let J be the set of jurisdictions and $\mathcal{K}_j$ the set of legal categories that jurisdiction j recognises. Let $\Sigma = \{\text{in force}, \text{enacted}, \text{proposed}, \text{unverified}, \text{conflict}\}$ be the status vocabulary. For an event date t and an evaluation (as-of) date t', the rules layer computes

$$L(x, j, t; t') = (\kappa, \sigma, \rho), \qquad \kappa \in \mathcal{K}_j \cup \{\perp\},\ \sigma \in \Sigma,$$

where κ is the legal category, σ the status of the rule that produced it, ρ the source citation, and $\perp$ means "not determined by any rule in the current rule version".

The map has two time arguments on purpose. The event date t selects the law that applied to the event. The evaluation date t' records what was known when the computation ran. Rules can be retroactive or can be announced before they apply, so the two dates must be kept apart. The UK stablecoin tax outcome, for example, leaves "limited retrospection" as an open question [107]. A bitemporal map lets the rules layer answer both "what category applied on the payment date?" and "what did we believe on the reporting date?".

Three properties. The definitions make three properties explicit.

1. **Jurisdiction indexing.** In general $L(x, j, t; t') \neq L(x, j', t; t')$ for $j \neq j'$. The examples above for BTC show this. The categories are not even drawn from a shared set: $\mathcal{K}_{\text{US}}$ contains "digital collectible", $\mathcal{K}_{\text{EU}}$ contains "asset-referenced token", and $\mathcal{K}_{\text{IN}}$ contains "virtual digital asset". No bijection between them exists.

2. **Time indexing.** For fixed x and j, L is piecewise constant in t. Its breakpoints are the effective dates of rules. Table 4 lists the breakpoints that fall between 2026 and 2028 in our sources.

3. **Partiality.** L returns $\kappa = \perp$ in two situations. The rules may be silent (for example, bridged stablecoins in every regime we reviewed). Or the rules may turn on facts outside x that only a reviewer can establish (for example, whether essential managerial efforts are complete). In both cases the output is a review task, not a category.

Where Zlef sits. [Direction, 15 Oct 2026 (v0.5r4)] Zlef outputs a distribution $q_\theta(x \mid E)$, designed to be calibrated (a goal to be measured, not a property), over the jurisdiction-free components of the fact tuple, given evidence E. It does not output κ. The rules layer pushes the distribution forward:

$$\Pr[\kappa \mid E, j, t; t'] \;=\; \sum_{x} q_\theta(x \mid E)\, \mathbf{1}\big[L(x, j, t; t')_\kappa = \kappa\big].$$

The sum runs over the jurisdiction-free components. The residency and authorisation facts r are supplied from layers B, C or F and are never predicted.

Remark 3.3 (Calibration of the push-forward). If q_θ is calibrated as a joint distribution over the components that L reads, the push-forward probability of each legal category is calibrated as well, because each category is an event in the fact space. Calibration of each marginal is not enough. L usually reads several components together. An example is a stablecoin that is used as payment and crosses a border, where a rule reads asset class, event type and jurisdiction set at once. The evaluation chapter must therefore test joint calibration on the combinations that the rules actually read.

Remark 3.4 (Why legal labels are not model outputs). Predicting κ directly would have four costs. (i) Every rule change would invalidate training labels, and Table 4 shows about one material change per quarter. (ii) Labels could not be reused across jurisdictions, because $\mathcal{K}_j$ differ. (iii) The status σ of a rule cannot be learned from transactions. It is a fact about a legislature or regulator. (iv) An auditor must be able to trace a legal conclusion to a rule version and a source. The factorisation gives that trace for free: the model's output is a fact claim with a probability, and the rule's output is a category with a citation.

Assumption 3.5 (Observability gap)**.** Some inputs to L are not observable from transfers and must not be predicted. They include issuer authorisation status (e.g. entries on the ESMA interim MiCA register, last updated 7 Oct 2026 [65], or the HKMA licensee list [103]), reserve composition, contractual rights, and issuer conduct. They enter the rules layer as layer-B (source-attested) or layer-F (reviewed) facts. Where they are missing, the rules layer returns $\perp$ and raises a data flag in the DAT family.

Breakpoints in the period of interest. Table 4 lists dated rule changes that alter L for at least one asset class in our set. A model frozen in 2024 would miss most of them, which is the main reason for keeping legal categories in a versioned rules layer.

Table 4: Dated breakpoints of the legal category map, 2026–2028 (as of 15 Oct 2026).

Date	Juris.	Change	Status	Source
2 Feb 2026	BR	BCB Res. 519–521: VASP authorisation; fiat-referenced stablecoin trades and international crypto payments become FX operations	[In force]	[9, 54]

Date	Juris.	Change	Status	Source
23 Mar 2026	US	SEC/CFTC five-category taxonomy effective	[In force]	[194]
1 Jul 2026	EU	End of MiCA transitional period for CASPs across the EU	[In force]	[66]
1 Sep 2026	RU	New federal framework on digital currencies effective; transition to 1 Jul 2027	[In force]	[12]
1 Oct 2026	BR	BCB Res. 561: eFX providers may not settle cross-border remittances in crypto or stablecoins	[In force]	[34]
1 Jan 2027	KR	Virtual-asset income tax scheduled (22% above KRW 2.5m); deferral bill announced 15 Oct 2026 ([Proposed])	[Enacted, not yet effective]	[17]
by 18 Jan 2027	US	GENIUS Act general effective date (earlier of 18 Jan 2027 or 120 days after final regulations)	[Enacted, not yet effective]	[48]
2 Feb 2027	BR	Cross-border travel-rule phase under BCB Res. 520 (full compliance by 2 Feb 2028)	[Enacted, not yet effective]	[158]
4 Feb 2027	KR	Token-securities amendments to the capital markets and electronic securities acts	[Enacted, not yet effective]	[25, 28]
6 Apr 2027	UK	"No gain, no loss" rules for crypto lending and liquidity pools (draft legislation)	[Proposed]	[113]
Apr 2027	UK	CGT exemption for eligible stablecoins for individuals (draft)	[Proposed]	[107]
10 Jul 2027	EU	AMLR applies, including the Art. 79 ban on CASP accounts for anonymity-enhancing coins	[Enacted, not yet effective]	[36, 73]
by about Jul 2027	JP	Crypto assets move from the PSA to the FIEA (cabinet order within one year of promulgation on 23 Jul 2026)	[Enacted, not yet effective]	[202, 228]
25 Oct 2027	UK	Full commencement of SI 2026/102 (qualifying cryptoasset, qualifying stablecoin, staking as a regulated activity)	[Enacted, not yet effective]	[217]
1 Jan 2028 (expected)	JP	20.315% separate taxation for specified crypto assets, from 1 Jan of the year after FIEA effectiveness; expected from 2028	[Sources conflict]	[206, 230]
18 Jul 2028	US	Digital asset service providers barred from offering non-permitted payment stablecoins	[Enacted, not yet effective]	[218]

3.2 Reference taxonomies

This subsection surveys the taxonomies that the rules layer must encode. We group them into supranational and industry references, the EU, the United States, the other jurisdictions in the country set, and the reference hubs. Table 5 at the end summarises them.

FINMA (Switzerland), 2018. [In force] FINMA's ICO Guidelines of 16 Feb 2018 divide tokens by function into payment tokens, utility tokens and asset tokens, and allow hybrids [89]. Payment tokens "are synonymous with cryptocurrencies". Utility tokens are "intended to provide digital access to an application or service" and are treated as securities if they also have an investment function. Asset tokens "represent assets such as participations in real physical underlyings, companies, or earnings streams" and are treated as securities. [In force] Later Swiss layers include ledger-based securities under the DLT Act (in force 1 Aug 2021) and FINMA Guidance 06/2024 of 26 Jul 2024, under which a stablecoin issuer is generally conducting deposit-taking business unless a compliant bank default guarantee is in place [177]. [Proposed] A Federal Council

consultation (opened 22 Oct 2025, closed 6 Feb 2026) proposes FinIA licences for a "Payment Instrument Institution" and a "Crypto Institution" [177].

FINMA's triad is the oldest regulator taxonomy still in use, and it is useful as a spine. Most later regimes split one of its three boxes. Payment tokens split into e-money tokens, asset-referenced tokens and digital payment tokens. Utility tokens split into utility tokens, digital tools and digital commodities. Asset tokens split into financial instruments, digital securities and electronically recorded rights.

EU: MiCA Titles II, III and IV. [In force] MiCA (Regulation (EU) 2023/1114) defines a crypto-asset as "a digital representation of a value or of a right that is able to be transferred and stored electronically, using distributed ledger technology or similar technology" (Art. 3(1)(5)) [72]. It then sorts crypto-assets into three titles:

- **Title IV, e-money tokens (EMTs):** crypto-assets that purport to maintain a stable value "by referencing the value of one official currency" (Art. 3(1)(7)).
- **Title III, asset-referenced tokens (ARTs):** crypto-assets that are not EMTs and purport to maintain a stable value "by referencing another value or right or a combination thereof, including one or more official currencies" (Art. 3(1)(6)).
- **Title II, other crypto-assets:** everything else in scope, including BTC-type assets, governance tokens, memecoins and utility tokens. A utility token is a crypto-asset "only intended to provide access to a good or a service supplied by its issuer" (Art. 3(1)(9)).

[In force] Titles III and IV have applied since 30 Jun 2024 and the rest of MiCA since 30 Dec 2024 [65, 72]. [In force] The CASP transitional period "will officially expire across the EU on 1 July 2026", and it has now done so [66].

The exclusions matter as much as the titles. [In force] Art. 2(3) excludes crypto-assets "that are unique and not fungible with other crypto-assets". Recital 11 limits that exclusion: fractional parts of a unique asset are not unique, and issuance "in a large series or collection should be considered an indicator of their fungibility" [72]. [In force] Art. 2(4) excludes financial instruments, deposits including structured deposits, funds, securitisation positions and certain insurance and pension products [72]. [In force] ARTs (Art. 40) and EMTs (Art. 50) may not carry interest; the Commission's 2026 consultation restates that "MiCA currently prohibits the granting of interest or any interest-equivalent remuneration of stablecoins" [68].

EU: ESMA guidelines on qualification as financial instruments. [In force] ESMA's guidelines under MiCA Art. 2(5) were finalised on 17 Dec 2024, translated on 19 Mar 2025, and apply from 18 May 2025 on a comply-or-explain basis [183, 186]. They set cumulative tests for a transferable security (not an instrument of payment; within a "class of securities"; negotiable on the capital market), require pooled capital and a defined investment policy for a collective investment undertaking, and assess derivatives by reference to the underlying, with perpetual futures in view. Where a token qualifies as a financial instrument, that qualification "should take precedence". NFTs are generally outside MiCA, subject to an assessment of genuine uniqueness. The guidelines are technology-neutral: "the form of the technology used does not affect the qualification of a crypto-asset as a financial instrument" [62].

EU: the 2026 MiCA review. [Proposed] The Commission's targeted consultation of 20 May 2026 (closed 31 Aug 2026) names the hardest assets to classify: "hybrid tokens", "tokenised fund interests", "wrapped assets" and "assets marketed as NFTs but issued in series" (Q7) [68, 200]. It asks whether to keep the interest ban (Q20), notes that "MiCA does not currently prohibit multi-issuer models" (Q30), records that staking is "currently not addressed in MiCA regulation" (Q66), and asks about lending and borrowing (Q67), decentralisation criteria for DeFi (Q61–65) and tokenised deposits. For the ontology, Q7 is a useful signal: the regulator itself names four of our confusion zones.

United States: the five-category taxonomy. [In force] The joint SEC/CFTC interpretation (Release Nos. 33-11412 and 34-105020) was issued on 17 Mar 2026 and took effect on 23 Mar 2026 [192, 194]. It supersedes the 2019 staff framework for investment-contract analysis [53]. It defines five categories:

- **Digital commodity:** "a crypto asset that is intrinsically linked to and derives its value from the programmatic operation" of a crypto system. Named examples include APT, AVAX, BTC, BCH, ADA, LINK, DOGE, ETH, HBAR, LTC, DOT, SHIB, SOL, XLM, XTZ and XRP. Not a security.
- **Digital collectible:** "designed to be collected and/or used". Examples include CryptoPunks, Chromie Squiggles, fan tokens and WIF. Memecoins sit here. Fractionalised collectibles "could constitute the offer or sale of a security".
- **Digital tool:** "performs a practical function", for example ENS names and tickets.
- **Stablecoin:** "designed to maintain a stable value relative to a reference asset". GENIUS payment stablecoins and 2025 "Covered Stablecoins" are not securities; others are fact-dependent.
- **Digital security:** a tokenized security "is a financial instrument".

The interpretation adds an overlay: a non-security asset can be sold subject to an investment contract, which ends when the issuer's essential managerial efforts are complete, failed or abandoned [172]. It treats protocol mining, protocol staking, one-to-one "Staking Receipt Tokens", one-to-one "Redeemable Wrapped Tokens" (including bridge-generated ones) and nominal-consideration airdrops of non-security assets as non-securities transactions. The CFTC will administer the Commodity Exchange Act consistently with the release, but "commodity" is used in an economic sense and the release takes no position on whether every digital commodity meets the statutory definition [172].

United States: proposed Regulation Crypto Assets. [Proposed] On 18 Aug 2026 the SEC proposed Regulation Crypto Assets [193]. It would add two offering exemptions ($5m once over four years, and $75m per 12 months), a safe harbour deeming a crypto-asset no longer subject to an investment contract once the issuer has "completed or permanently ceased all essential managerial efforts", and state pre-emption, with a 60-day comment period. If adopted, the safe harbour would turn part of the investment-contract question into a dated, source-attested fact. That would move some cells of L from $\perp$ to a determinate category.

United States: CLARITY Act. [Proposed] The CLARITY Act passed the House on 17 Jul 2025 (294–134). A Senate cloture vote failed 49–50 on 15 Sep 2026, and a motion to reconsider

preserves the possibility of revival [59, 168]. Its statutory "digital commodity" and "mature blockchain" paths are therefore not law. The rules layer must not encode them as in force.

United States: GENIUS Act payment stablecoins. [Enacted, not yet effective] The GENIUS Act (Public Law 119-27, approved 18 Jul 2025) defines a "payment stablecoin" and excludes a permitted issuer's payment stablecoin from "security" and "commodity" status [218]. Issuers may not "pay the holder of any payment stablecoin any form of interest or yield" (§4(a)(11)). Issuers must have the capability to "block, freeze, and reject" (§4(a)(5)(A)(iv)) and to comply with any lawful order (§4(a)(6)(B)). From three years after enactment, about 18 Jul 2028, digital asset service providers may not offer non-permitted payment stablecoins (§3(b)(1)). [Proposed] Implementing rules from the OCC (2 Mar 2026), Treasury (Aug 2026) and the Federal Reserve (29 Sep 2026) are proposals [78, 92, 165].

United Kingdom: qualifying cryptoassets and stablecoins. [Enacted, not yet effective] SI 2026/102, made on 4 Feb 2026 with full commencement on 25 Oct 2027, creates the "qualifying cryptoasset" [217]. Our reading of art. 88F is that a qualifying cryptoasset must be fungible and transferable, and must not be "solely a record of … rights in another cryptoasset". Electronic money, fiat currency "including a central bank digital currency", closed-loop tokens and "specified investment cryptoassets" (security tokens) are excluded (art. 88F(4)). A "qualifying stablecoin" (art. 88G) is a qualifying cryptoasset that "seeks or purports to maintain a stable value in relation to a particular fiat currency" and is backed by fiat or other assets. Staking, "the use of a qualifying cryptoasset in blockchain validation", becomes a regulated activity (art. 9Z6). [Proposed] The Bank of England's policy statement and draft rules for systemic sterling stablecoins (22 Jun 2026) propose up to 70% backing in short-term gilts and no holding limits [10]. [In force] Separately, digital assets can be personal property under the Property (Digital Assets etc) Act 2025, which received Royal Assent in Dec 2025 [135]; the exact day was not confirmed.

Japan: crypto assets, electronic payment instruments and security tokens. [In force] Japan's Payment Services Act separates "crypto assets" (Art. 2(14)) from "electronic payment instruments" (EPIs, Art. 2(5)), the category for fiat stablecoins whose issuers are limited to banks, funds-transfer providers and trust companies [202]. [In force] Tokenized shares, bonds, trust beneficiary interests and fund interests are "electronically recorded transferable rights" under the FIEA [202]. [Enacted, not yet effective] The FIEA amendment (Act No. 64 of 2026, passed 15 Jul and promulgated 23 Jul 2026) moves crypto assets into the FIEA as a financial-instrument category distinct from securities, by cabinet order within one year [202, 228]. Stablecoins stay under the PSA. The FIEA also brings insider-trading rules to crypto assets.

South Korea: VAUPA. [In force] The Virtual Asset User Protection Act has applied since 19 Jul 2024 and uses the "virtual asset" concept of the Specific Financial Information Act [85]. CBDCs, CBDC deposit tokens and in-game currency are excluded. [In force] The FSC's NFT guideline of 10 Jun 2024 excludes NFTs unless they are issued in mass or series, are divisible, are usable as payment, or are exchangeable for virtual assets among the public [87]. [In force] There is "currently no statutory definition of stablecoins" [28]. [Proposed] The Digital Asset Basic Act and its won-stablecoin phase remain proposals, with an FSC parliamentary review targeted for Nov 2026 [25, 47].

India: virtual digital assets. [In force] India defines a "virtual digital asset" in s. 2(47A) of the Income Tax Act, carried into the Income-tax Act 2025 from 1 Apr 2026, and uses it for a 30% tax and 1% TDS [21, 49]. [In force] VDA service providers have been PMLA reporting entities since the March 2023 notification [161]. India has no regulatory classification law for crypto-assets. The Reserve Bank of India continues to urge prohibition (testimony of 2 Jul 2026) [49]. The VDA concept is a tax and AML perimeter, not a market-structure taxonomy.

Brazil: Law 14,478 and the BCB resolutions. [In force] Law 14,478/2022 defines the "ativo virtual" [54]. [In force] BCB Resolutions 519, 520 and 521 (10 Nov 2025, effective 2 Feb 2026) set VASP authorisation, conduct and travel-rule duties, and FX treatment [9, 54]. Under Res. 521, both "pagamento ou a transferência internacional usando ativos virtuais" and "compra, venda ou troca de ativos virtuais referenciados em moeda fiduciária" are FX operations. Brazil thus attaches a category to the *transaction* as well as to the asset. Tokens structured as securities fall to the CVM; the CVM guidance for this boundary was not verified in our sources.

Russia: digital currency and digital financial assets. [In force] Federal Law 259-FZ distinguishes "digital currency" from "digital financial assets" (DFAs), which are tokenised monetary claims, securities rights and equity in non-public companies issued on Bank of Russia-registered platforms [179]. [In force] According to the Bank of Russia, a new framework took effect on 1 Sep 2026 with a transition to 1 Jul 2027. It allows exporters and importers to use crypto in cross-border payments "without limitations", keeps domestic crypto payments banned, reportedly caps non-qualified investors at RUB 300,000 a year per intermediary [Sources conflict], and applies the same requirements to foreign stablecoins [12]. [Unverified] Secondary press reports the law as Federal Law 282-FZ "On Digital Currencies and Digital Rights", signed 4 Aug 2026. An industry tracker dated 20 May 2026 still described the draft as not enacted [179]; the Bank of Russia's July 2026 statement post-dates it.

Indonesia: digital financial assets. [In force] Indonesia places crypto-assets within "digital financial assets", described as "digital commodities intended for investment", under POJK 27/2024, with OJK supervision since January 2025 [27]. A crypto-asset is tradable only if it is on the OJK list. It is not a means of payment, because the rupiah is the sole legal tender. The rules "do not separate fiat-backed from algorithmic stablecoins".

Vietnam: digital asset law and pilot market. [In force] The Law on Digital Technology Industry, passed in June 2025, recognises digital assets as property and has applied since 1 Jan 2026, according to a secondary source [189]. [In force] Government Resolution 05/2025/NQ-CP (Sep 2025) sets a five-year pilot crypto-asset market [189]. [In force] No pilot operator had been licensed as of 30 Aug 2026; five applicants had passed initial review [46]. [Unverified] The statutory definitions of "digital asset", "virtual asset" and "crypto-asset" in the 2025 law were not read in our sources. Vietnam is therefore a low-confidence jurisdiction for every asset class.

Ukraine: draft virtual-assets law. [Proposed] A comprehensive virtual-assets bill passed first reading on 3 Sep 2025 and was awaiting second reading in 2026 [203, 223]. It would involve the National Bank and a designated regulator, admit EU or regulator-approved foreign exchanges through simplified authorisation, and tax gains at 18% PIT plus a 5% military levy, with a 5%

preferential rate in 2026 [100]. The Verkhovna Rada page refers to Bill No. 10225-d [223] (also referred to in some trackers as 3637; unverified). [Unverified] Whether the 2022 law "On Virtual Assets" has commenced, and whether the bill was adopted at second reading by Oct 2026, could not be confirmed. Ukraine is therefore an abstention zone for legal categories.

Reference hubs: Singapore, Hong Kong and the UAE. [In force] Singapore uses the "digital payment token" for payment-type crypto-assets and the "capital markets product" for investment-type tokens; MAS's stablecoin framework of 15 Aug 2023 creates a "MAS-regulated stablecoin" label for compliant single-currency SGD or G10 stablecoins issued in Singapore [148, 180]. [Unverified] The in-force status of statutory backing for that label could not be confirmed. [In force] Hong Kong's Stablecoins Ordinance (Cap. 656) has applied since 1 Aug 2025; issuing fiat-referenced stablecoins "is a regulated activity in Hong Kong and a licence is required", and the first licences were granted on 10 Apr 2026 [103, 104]. [In force] The UAE distinguishes the CBUAE "Payment Token" (dirham tokens for mainland payments), registered "Foreign Payment Tokens", Dubai VARA "virtual assets" and fiat-referenced virtual assets, and ADGM fiat-referenced tokens [175].

Standard setters: FSB, IOSCO and BIS. [In force] The FSB's high-level recommendations for global stablecoin arrangements (final, 17 Jul 2023) are non-binding. They apply requirements "on a functional basis and proportionate to their risks" (Rec. 2) and state that for single-fiat stablecoins "redemption should be at par into fiat" (Rec. 9) [88]. [In force] The Basel Committee's prudential standard for cryptoasset exposures (SCO60, 16 Dec 2022, implementation by 1 Jan 2025) divides exposures into tokenised traditional assets, stablecoins with effective stabilisation, and unbacked cryptoassets [14]. [Unverified] The Group 1a/1b/2a/2b labels and any 2025–26 review of SCO60 were not confirmed in our sources. [Unverified] IOSCO's Nov 2023 policy recommendations for crypto and digital asset markets were not verified in our sources.

OECD CARF: the "Relevant Crypto-Asset". [In force] The OECD published the Crypto-Asset Reporting Framework and amended CRS on 8 Jun 2023 [167]. The amended CRS was changed "to bring certain electronic money products and central bank digital currencies in scope" [167]. [Unverified] The CARF definition of "Relevant Crypto-Asset" excludes central bank digital currencies, Specified Electronic Money Products, and crypto-assets that the reporting provider has adequately determined "cannot be used for payment or investment purposes"; the 2023 text was not re-read in our sources [166]. This is the only taxonomy in our set that is defined for tax reporting. It cuts across the market-structure taxonomies. A CBDC and a Specified Electronic Money Product are excluded from CARF but brought into the amended CRS; the CARF treatment of tokenized deposits is [Unverified] in our sources, so the rules layer needs a reporting-perimeter output that is separate from the market-structure category.

Industry taxonomies: TTF and ITC. The InterWork Alliance Token Taxonomy Framework is a technical, not legal, taxonomy [118]. It types tokens as fungible (τ_F) or non-fungible (τ_N); by units (fractional, whole, singleton); by value type (intrinsic or reference); by representation (common or unique); by supply (fixed, capped-variable, gated, infinite); and by behaviours such as mintable, burnable, transferable, divisible and delegable. It is a good source of features for Zlef because each property is checkable against a contract. [Unverified] The International Token Classification maintained by ITSA could not be fetched [124]; we do not rely on it.

Table 5: Reference taxonomies encoded by the rules layer (as of 15 Oct 2026).

Framework	Instrument	Top-level categories	Status	Source
FINMA (CH)	ICO Guidelines, 16 Feb 2018	payment / utility / asset token; hybrids	[In force]	[89]
MiCA (EU)	Reg. 2023/1114	EMT (Title IV), ART (Title III), other incl. utility (Title II); exclusions Art. 2(3), 2(4)	[In force]	[72]
ESMA (EU)	Guidelines under Art. 2(5), apply 18 May 2025	transferable security, MMI, CIU unit, derivative; precedence over MiCA	[In force]	[183]
MiCA review (EU)	Consultation, 20 May 2026	hard cases: hybrids, tokenised funds, wrapped, serial NFTs	[Proposed]	[68]
SEC/CFTC (US)	Rel. 33-11412, eff. 23 Mar 2026	digital commodity / collectible / tool / stablecoin / security; IC overlay	[In force]	[194]
SEC (US)	Reg. Crypto Assets, 18 Aug 2026	exemptions; safe harbour for end of IC	[Proposed]	[193]
Congress (US)	CLARITY Act	statutory digital commodity; cloture failed 15 Sep 2026	[Proposed]	[59]
Congress (US)	GENIUS Act, 18 Jul 2025	payment stablecoin (permitted issuer)	[Enacted, not yet effective]	[218]
HMT (UK)	SI 2026/102	qualifying cryptoasset; qualifying stablecoin; specified investment cryptoasset	[Enacted, not yet effective]	[217]
FSA (JP)	PSA; FIEA amendment Act No. 64 of 2026	crypto asset; EPI; ERTR	PSA [In force]; FIEA [Enacted, not yet effective]	[202]
FSC (KR)	VAUPA; NFT guideline	virtual asset; NFT exclusion; security (FSCMA)	[In force]	[85, 87]
CBDT (IN)	Income-tax Act s. 2(47A)	virtual digital asset (tax/AML only)	[In force]	[49]
BCB (BR)	Law 14,478; Res. 519–521, 561	ativo virtual; fiat-referenced VA in FX market	[In force]	[54]
OJK (ID)	POJK 27/2024	digital financial asset incl. crypto-asset	[In force]	[27]
Gov't (VN)	Law on Digital Technology Industry; Res. 05/2025	digital asset; pilot crypto-asset	[In force] (definitions [Unverified])	[189]
Rada (UA)	Virtual-assets bill	virtual asset (draft)	[Proposed]	[223]
CBR (RU)	259-FZ; 2026 framework	digital currency; DFA	[In force]	[12, 179]
FSB	GSC recommendations, 17 Jul 2023	functional; par redemption	[In force] (non-binding)	[88]
BCBS	SCO60, 16 Dec 2022	tokenised traditional / stablecoin / unbacked	[In force]	[14]
OECD	CARF and CRS 2.0, 8 Jun 2023	Relevant Crypto-Asset; CBDC and e-money to CRS	[In force] (RCA wording [Unverified])	[167]
IWA	Token Taxonomy Framework	technical types and behaviours	n/a (industry)	[118]

3.3 A crosswalk from Zlef asset classes to legal categories

The crosswalk below maps each Zlef `asset_class` code to the category that each framework most often assigns. It is a table of *priors* for the rules layer, not a set of conclusions. Three cautions apply.

1. A cell gives the usual category for the typical token in that class. Individual tokens can differ because of rights, issuer conduct or structure. The rules layer must apply the full rule, not the cell.
2. Cells marked "U" are unverified in our sources. The rules layer must return $\perp$ there and must not borrow a neighbouring jurisdiction's answer.
3. Some cells depend on facts outside the asset class, such as issuer authorisation. Those facts come from registers (layer B), not from the model.

Notation. DC = digital commodity; COL = digital collectible; TOOL = digital tool; PSC = GENIUS payment stablecoin; DSEC = digital security or security; IC? = investment-contract overlay possible; CA-II = MiCA Title II crypto-asset; UT = MiCA utility token; EMT/ART = MiCA e-money/asset-referenced token; FI = MiFID II financial instrument; OUT = out of scope; QCA/QSC/SIC = UK qualifying cryptoasset/qualifying stablecoin/specified investment cryptoasset; CA = Japanese crypto asset; EPI = electronic payment instrument; ERTR = electronically recorded transferable right; VA = virtual asset; RCA = CARF Relevant Crypto-Asset; DEP = deposit. A suffix (e) marks an enacted category not yet effective. A suffix (p) marks a proposal. A dagger (†) marks a cell resting on background knowledge that was not re-verified, which has status [Unverified]. Cells without a suffix are [In force] as of 15 Oct 2026. Sources are given in Section 3.2 and, per asset class, in the next section.

Table 6: Crosswalk I: Zlef `asset_class` codes to reference frameworks (as of 15 Oct 2026). Priors only; see notation above.

Zlef code	FINMA (CH)	MiCA / ESMA (EU)	SEC/CFTC 2026 (US)	SI 2026/102 (UK)	PSA/FIEA (JP)	VAUPA (KR)	CARF (OECD)	TTF (IWA)
NAT.L1 / NAT.L2	payment token	CA-II	DC (named list); IC? at initial sale	QCA (e)	CA; FIEA CA (e)	VA	RCA†	τ_F, intrinsic
STB.FIAT. SINGLE	issuer: deposit unless bank guarantee	EMT	PSC (e) once permitted; not a security	QSC (e)	EPI	VA; no stablecoin statute	RCA unless e-money†	τ_F, reference
STB.FIAT. BASKET	U	ART	stablecoin; fact-dependent	U (QSC needs "a particular fiat currency")	U	VA	RCA†	τ_F, reference
STB.COMM	asset token (claim on gold)	ART	no clean category; CEA commodity; IC?	QCA, not QSC	U	VA	RCA†	τ_F, reference
STB.SYN	U	EMT/ART if it references; algorithmic EMT not authorisable†	not PSC; fact-dependent	QCA, not QSC (no backing)	CA, not EPI†	VA	RCA†	τ_F, reference
STB.YLD	U	interest ban (Arts. 40/50); possibly FI	not PSC; may be a security	QCA	U	VA	RCA†	τ_F, reference
TKN.MMF	asset token; ledger-based security	FI (CIU unit); out of MiCA	DSEC	SIC	ERTR	security; token form (e) 4 Feb 2027	U	τ_F, roles
TKN.EQ	asset token	FI (share or derivative)	DSEC (issuer, custodial, synthetic, SBS)	SIC	ERTR	security	U	τ_F, roles
TKN.DEP	deposit	deposit; out of MiCA	deposit†	U	U	U (CBDC deposit tokens excluded)	U	τ_F
TKN.DEBT	asset token	FI (transferable security)	DSEC	SIC	ERTR	security	U	τ_F, roles
TKN.RE	asset token	FI or CA-II by rights	DSEC if a security	SIC or QCA	ERTR (trust interests)	U	U	τ_F or τ_N
GOV	utility token; hybrid → security	UT or CA-II; FI if profit rights	TOOL / DC / IC?; Reg. CA (p)	QCA; SIC if share- or debt-like	CA; specified CA (e)	VA unless security	RCA†	τ_F, delegable
MEME	payment token	CA-II	COL (DOGE, SHIB are DC)	QCA	CA	VA	RCA†	τ_F, intrinsic
NFT.UNIQUE	case by case†	OUT (Art. 2(3))	COL or TOOL	not QCA (non-fungible)	not CA (indicators)	excluded	in scope if payment or investment†	τ_N, singleton
NFT.FRAC	U	CA-II or FI (Recital 11)	DSEC possible	QCA if fungible	CA possible	VA (divisible or mass)	RCA†	τ_F over τ_N (hybrid)
LST.STAKE	U	CA-II or FI; staking unaddressed	SRT, not a security if 1:1 on DC	QCA; staking regulated (e)	U	U	RCA†	τ_F, reference

Zlef code	FINMA (CH)	MiCA / ESMA (EU)	SEC/CFTC 2026 (US)	SI 2026/102 (UK)	PSA/FIEA (JP)	VAUPA (KR)	CARF (OECD)	TTF (IWA)
LST.RESTAKE	U	CA-II or FI	fact-specific (excluded from LST guidance)	QCA	U	U	RCA†	τ_F, reference
WRP.CUST / WRP.BRIDGE	U	CA; open question (Q7)	redeemable wrapped token; non-security if underlying is	possibly not QCA (record of rights)	U	U	RCA†	τ_F, reference
WRP. NATIVEBURN	U	as underlying (analysis)	issuer liability (analysis)	U	U	U	RCA†	τ_F, reference
DFI.LP / LEND / DEBT	U	CA-II or FI; lending unregulated (Q67)	no direct treatment (U)	QCA or excluded record; lending regulated (p)	U	U	RCA†	τ_F, reference
PRV	U	CA-II; CASP ban from 10 Jul 2027 (e)	DC or COL; no ban	QCA	exchange listing barred†	exchange handling barred†	RCA†	τ_F
CBDC	n/a	OUT (e-money; EMT is e-money)	n/a	excluded	prepaid / EPI†	excluded	excluded; CRS 2.0	τ_F
OTH (points, in-game)	U	not CA if non-transferable	airdrop not a security; in-game items COL	closed-loop excluded	prepaid†	in-game excluded	U	varies

Table 7: Crosswalk II: Zlef codes to umbrella categories in IN, BR, ID, VN, UA, RU and the hubs (as of 15 Oct 2026). AV = ativo virtual; AV-FX = fiat-referenced VA inside the FX market; DFA = Indonesian digital financial asset; DCur = Russian digital currency; RU-DFA = Russian digital financial asset; DPT = digital payment token; CMP = capital markets product; PT = UAE payment token.

Zlef code	IN	BR	ID	VN	UA	RU	SG	HK	AE
NAT	VDA	AV	DFA (crypto-asset); not payment	crypto-asset (pilot); definition U	VA (p)	DCur; barred from domestic payment	DPT	VA†	VA
STB.FIAT.*	VDA	AV-FX	DFA	U	U	not defined; same rules as DCur under 2026 law	DPT; SCS label if SGD/G10 and SG-issued	licensed fiat-referenced stablecoin	PT; foreign PT; fiat-referenced VA
STB.COMM	VDA	AV	DFA	U	U	RU-DFA if domestic†	DPT or CMP	U	zone-dependent
STB.SYN / STB.YLD	VDA	AV; FX status ambiguous	DFA (no split by backing)	U	U	DCur	DPT or CMP	U	U
TKN.*	security if a share etc. (U)	CVM security†	U	U	U	RU-DFA	CMP	securities†	securities (U)
GOV	VDA	AV	DFA	U	U	DCur	DPT or CMP	VA	VA
MEME	VDA	AV	DFA	U	U	DCur	DPT	VA	VA
NFT.*	VDA†	U	U	U	U	U	U	U	U
LST.*	VDA	U	U	U	U	U	U	U	U
WRP.*	VDA	AV-FX? (bridged stablecoin)	U	U	U	U	U	U	U
DFI.*	VDA	U	U	U	U	U	U	U	U
PRV	VDA; "unacceptably high-risk" for VASPs	U	U	U	U	U	U	U	banned by VARA†
CBDC	e-Rupee is legal tender, not VDA	U	U	U	U	digital ruble†	U	U	digital dirham (U)

Reading the crosswalk. Three patterns stand out. First, the umbrella regimes of India, Brazil, Indonesia and Russia assign almost every class to one bucket (VDA, ativo virtual, digital financial asset, digital currency). For these countries the asset class matters less for the legal category than for tax mechanics and transaction-level rules, such as Brazil's FX characterisation. Second, the fine-grained regimes (US, EU, UK, JP) split stablecoins, tokenized assets and NFTs, and they split them along different lines. Third, most cells for staking receipts, wrapped tokens and DeFi positions are "U" outside the US and the UK. That is where Zlef must route to review.

3.4 Dimensions that drive classification

The crosswalk is a summary. The rules that generate it read a small set of underlying dimensions. If Zlef predicts these dimensions well, the rules layer can compute any of the categories above. If it does not, no amount of rule engineering can recover them. This subsection names the dimensions, shows where each is decisive, and maps each to Zlef fields and to observable features.

1. Issuer existence and redemption claim. Whether there is an identifiable issuer who owes the holder something is the first cut in most regimes. [In force] MiCA EMTs must be issued by an EU credit institution or e-money institution [65]. [Enacted, not yet effective] GENIUS payment stablecoins must come from a permitted issuer [218]. [In force] Japanese EPIs can be issued only by banks, funds-transfer providers and trust companies [202]. [In force] The FSB recommends par redemption for single-fiat stablecoins (non-binding) [88]. Native tokens have no issuer at all, which is why MiCA places the white-paper duty for BTC on the person seeking admission to trading [72]. Wrapped tokens have a custodian or bridge that owes the underlying, which is why the US treats them as receipts [194].

2. Reference asset. Given an issuer, the next cut is what the token references. One official currency gives an EMT in the EU; anything else that is stabilised gives an ART [72]. The UK qualifying stablecoin requires "a particular fiat currency" and backing [217]. Brazil's fiat-referenced definition requires reserve assets "created for the purpose of keeping its value pegged" [54]. A gold token therefore lands in different boxes: an ART in the EU, a qualifying cryptoasset but not a qualifying stablecoin in the UK, and outside the five clean categories in the US [35].

3. Rights conferred. Profit, dividend, participation or debt rights pull a token into securities law in every fine-grained regime: an asset token under FINMA, a financial instrument under the ESMA guidelines, a digital security in the US, an ERTR in Japan, a security under Korea's FSCMA. Access rights alone point to a utility token or digital tool. Governance rights alone are not named as a category in the US taxonomy [194]. ESMA notes that utility tokens are not CIU units "merely because they grant participatory rights" [183].

4. Decentralisation and control. [In force] The US investment-contract overlay turns on whether essential managerial efforts are complete [172]. [Proposed] The proposed Regulation Crypto Assets would add a safe harbour for that state [193]. [Proposed] The MiCA review asks how to define "fully decentralised" services [68]. [Proposed] The UK would catch DeFi where there is an "identifiable controlling entity" (FCA consultation CP25/40) [198]. Control is the hardest

dimension to observe. Admin keys and upgradeability are visible on chain, but promises, roadmaps and team conduct are not.

5. Fungibility and uniqueness. Fungibility separates NFTs from other crypto-assets, but each regime tests it differently. The EU uses uniqueness plus a series indicator (Art. 2(3), Recital 11). Korea uses mass or serial issuance, divisibility and payment use [87]. [In force] Japan's FSA, in a comment of 24 Mar 2023, used price and quantity indicators: at least JPY 1,000 per unit is a "high" price and at most 1,000,000 units a "limited" quantity, as indicators and not a safe harbour [202]. The UK requires fungibility for a qualifying cryptoasset. The US separates collectibles from fractionalised collectibles.

6. Transferability restrictions. [In force] MiCA requires that a crypto-asset be "able to be transferred" [72]. [Enacted, not yet effective] The UK excludes closed-loop tokens that are transferable only by redemption within "a limited network of service providers" [217]. Non-transferable points therefore fall outside both definitions on our reading. At the other end, permissioned security-token standards such as ERC-3643 and ERC-1400 restrict transfers to whitelisted identities. They carry no legal classification of their own, but they are strongly predictive of a security or financial instrument.

7. Yield. [Enacted, not yet effective] GENIUS bars permitted issuers from paying "any form of interest or yield" [218]. [Proposed] The OCC's proposed rule would add a rebuttable presumption that issuer arrangements with an affiliate or "related third party" to pay yield breach the ban [165, 174]. [In force] MiCA Arts. 40 and 50 bar interest on ARTs and EMTs [68]. [In force] The US staking-receipt treatment requires that rewards be protocol-driven and not guaranteed or influenced by the provider [79]. Yield is the dimension that most often moves a stablecoin out of the payment category and toward securities or fund law.

8. Use as payment. Some regimes define a category by use. Korea's NFT guideline asks whether a token is usable as payment [87]. [Unverified] CARF excludes assets that cannot be used for payment or investment [166]. [In force] Indonesia states that crypto-assets are not a means of payment [27]. [In force] Russia bars digital currency from domestic payments [12]. [In force] Brazil treats a cross-border payment with virtual assets as an FX operation [9]. [In force] The UAE allows only licensed dirham Payment Tokens for mainland retail payments [175]. Use as payment is a property of the *event*, not of the asset. Zlef captures it in `event_type` and `economic_purpose`, not in `asset_class`.

Two further dimensions. Representation depth (native, custodial wrap, lock-and-mint bridge, issuer burn-and-mint) and privacy enhancement are not classification dimensions in most statutes, but they decide the travel and evidence surface. [Enacted, not yet effective] The EU AMLR bars CASPs from keeping accounts that allow anonymisation "including through anonymity-enhancing coins" from 10 Jul 2027 [36, 73]. Representation depth decides whether a token is an issuer liability or a claim on a bridge contract.

Feature table. Table 8 maps each dimension to Zlef fields and to features that can be observed. The feature names are proposals for the method chapter and are [Research hypothesis] until tested.

The fact-layer column says where the feature comes from: A (protocol-authenticated), B (source-attested), C (absent) or F (reviewed conclusion).

Table 8: Classification dimensions, the regimes where they are decisive, and the Zlef fields and observable features that carry them.

Dimension	Decisive in	Observable features (proposed)	Layer	Zlef field(s)
Issuer and redemption claim	MiCA EMT/ART; GENIUS; JP EPI; FINMA deposit view; US receipts	mint/burn role holders; redemption events; issuer-controlled admin roles; blacklist functions; issuer attestation documents	A, B	`asset_class` (NAT vs STB vs WRP); `issuer_freezable`
Reference asset	EU EMT vs ART; UK QSC; BR fiat-referenced	price stability against a fiat or commodity reference; reserve attestations; declared peg in issuer documents	A, B	STB subcodes
Rights conferred	FINMA asset token; ESMA FI tests; US DSEC; JP ERTR; KR FSCMA	distribution or dividend functions; fee-switch; prospectus or offering documents; transfer-agent links	A, B, F	TKN vs GOV; STB.YLD
Decentralisation and control	US IC overlay; Reg. CA (p); MiCA Recital 22; UK controlling entity	admin keys; upgradeability; holder concentration; governance activity	A, F	none (rules layer, review); feeds `review_priority`
Fungibility and uniqueness	EU Art. 2(3); KR NFT guideline; JP indicators; UK fungibility; US fractionalisation	token standard (ERC-20, 721, 1155); edition size; fractionalisation vaults	A	NFT.UNIQUE vs NFT.FRAC; MEME
Transferability restrictions	MiCA "able to be transferred"; UK closed-loop; points	transfer-hook checks; identity registry; revert patterns; lockups	A	`permissioned_transfer`; TKN; OTH
Yield	GENIUS §4(a)(11); OCC (p); MiCA Arts. 40/50; US SRT conditions	rebasing balances; exchange-rate accrual; vault share tokens; reward distribution events	A	STB.YLD; LST; DFI
Use as payment	KR NFT; CARF; ID; RU; BR FX; UAE PT	merchant counterparty; invoice; memo fields	A, B	`event_type`; `economic_purpose`
Representation depth	US RWT; UK record of rights; EU Q7	bridge lock and mint events; custodian addresses; issuer burn-and-mint messages	A	WRP subcodes; `crosses_chain`; `passes_bridge`
Privacy	EU AMLR Art. 79 (e); exchange listing rules	shielded pool use; ring signatures; mixer contracts	A	PRV; `privacy_enhanced`; `touches_mixer`

What the model must not learn. Two dimensions in the table are not for the model to decide. Decentralisation and control is a reviewed conclusion (layer F), because it turns on conduct and promises. Rights conferred are partly source-attested (layer B), because they live in documents. Zlef can raise `review_priority` when features suggest that one of these dimensions is in play. For example, admin keys on a governance token or a distribution function on a stablecoin. It does not output a decentralisation label or a security label.

3.5 Where regimes disagree

The following cases show the same jurisdiction-free facts mapped to different legal categories. Each case states what Zlef outputs and what the rules layer must do.

Case 1: wrapped tokens. [In force] In the US a one-to-one redeemable wrapped token is a receipt, and is not a security if the underlying is not [194]. [Enacted, not yet effective] In the UK, on our reading of art. 88F(2), a wrapped token may fall outside "qualifying cryptoasset" as "solely a record of … rights in another cryptoasset"; FCA guidance confirming this was not found [217]. [Proposed] In the EU, "wrapped assets" are an open question in the MiCA review [68]. Zlef outputs `WRP.CUST` or `WRP.BRIDGE` plus the travel-surface flags. The rules layer returns three different categories, one of them $\perp$.

Case 2: bridged stablecoins. [Enacted, not yet effective] Under GENIUS only the permitted issuer's token is a payment stablecoin [218]. A third-party-bridged USDC.e is a claim on a bridge contract that holds USDC, not a liability of the issuer. [Unverified] No regime we reviewed addresses whether such a token inherits payment-stablecoin status [194]. A native USDC transfer via issuer burn-and-mint (CCTP) remains an issuer liability on our analysis [32]. Zlef separates these cases through `WRP.BRIDGE` versus `WRP.NATIVEBURN` (or the native `STB` code), which is exactly the fact that the law would need.

Case 3: yield-bearing stablecoins. Issuer yield is banned in the US ([Enacted, not yet effective] GENIUS §4(a)(11)) and the EU ([In force] MiCA Arts. 40/50) [68, 218]. [Proposed] The OCC proposes extending the US ban to affiliates and related third parties [165]. [Proposed] The Commission is consulting on relaxing the EU ban [68]. [In force] Under the US taxonomy, stablecoins outside GENIUS and the covered-stablecoin statement "may be securities depending on the facts" [194]. The regimes agree that yield matters but may move in opposite directions. Zlef outputs `STB.YLD`. The rules layer reads the date.

Case 4: a single offshore fiat stablecoin. For a large USD stablecoin issued outside the EU and the US, the categories in our sources differ by jurisdiction. [In force] In the EU an EMT must be issued by an EU credit institution or e-money institution [65]. [Unverified] Our sources record that USDT has no MiCA-authorised issuer and that CASPs had to restrict it for EEA clients by end-Q1 2025; the restriction deadline is background knowledge, not re-verified. [Enacted, not yet effective] In the US it is not yet GENIUS-permitted, and the foreign-issuer path depends on a Treasury comparability determination that is still proposed [92]. [In force] In Japan, Korea, Brazil and Indonesia it is tradable as a crypto-asset or virtual asset subject to local listing rules, and in Brazil its purchase through a VASP is an FX operation [9]. Zlef outputs `STB.FIAT.SINGLE` in every case. Authorisation status is a layer-B fact from issuer registers, not a model output.

Case 5: the NFT boundary. For an edition of 10,000 identical-looking image tokens with a floor price below JPY 1,000, the regimes ask different questions. The EU asks whether the series indicates fungibility (Recital 11). Korea asks whether issuance is mass or serial [87]. Japan asks about price and quantity indicators [202]. The US asks whether the item is a collectible or a fractionalised collectible [194]. Zlef outputs `NFT.UNIQUE` or `NFT.FRAC` and exposes edition size and token standard as features. The rules layer applies each test.

Case 6: staking receipts. [In force] The US treats one-to-one staking receipt tokens on non-security assets as non-securities, but restaking tokens are outside the staff guidance [79]. [Enacted, not yet effective] The UK makes staking itself a regulated activity from 25 Oct 2027 [217]. [Proposed] The EU records staking as "currently not addressed" [68]. Zlef separates `LST.STAKE` from `LST.RESTAKE` because the US line runs exactly there.

Case 7: memecoins. "Memecoin" is not a legal category anywhere in our sources. [In force] In the US, WIF is named as a digital collectible while DOGE and SHIB are named as digital commodities [194]. A rules layer that mapped `MEME` to "digital collectible" would therefore be wrong for the two largest memecoins. The rules layer must check named-asset lists before it applies a class default.

Case 8: payment use. [In force] In Brazil a cross-border payment with any virtual asset, and any purchase or sale of a fiat-referenced one, is an FX operation [9]. [In force] In the UAE only licensed dirham Payment Tokens may be used for mainland retail payments [175]. [In force] In Russia domestic crypto payments remain banned while foreign-trade use is allowed [12]. The asset is the same; the event differs. Zlef carries this through `event_type` and `economic_purpose`, and the rules layer raises FXC or STB family flags.

Summary of disagreements. Table 9 condenses the eight cases.

Table 9: Where regimes disagree, and what Zlef outputs (status as in the text, as of 15 Oct 2026).

Case	US	EU	UK	Other	Zlef output
Wrapped token	receipt; non-security if underlying is	open (Q7)	possibly not QCA	U	WRP.CUST / WRP.BRIDGE
Bridged stablecoin	unaddressed	unaddressed	U	BR: FX status unclear	WRP.BRIDGE vs WRP. NATIVEBURN
Yield stablecoin	issuer ban; OCC presumption (p)	interest ban; relaxation consulted (p)	U	—	STB.YLD
Offshore USD stablecoin	not yet permitted	no authorised issuer	U	JP/KR/ID: tradable VA; BR: FX	STB.FIAT.SINGLE + register fact
Serial NFT	collectible vs fractionalised	series indicates fungibility	fungibility test	KR mass issuance; JP price/quantity	NFT.UNIQUE / NFT.FRAC
Staking receipt	SRT non-security; LRT fact-specific	unaddressed	staking regulated (e)	U	LST.STAKE / LST.RESTAKE
Memecoin	COL or DC by name	CA-II	QCA	umbrella category	MEME + named-asset lookup
Payment use	—	—	—	BR FX; AE PT only; RU domestic ban	`event_type`, `economic_purpose`

3.6 Consequences for the label schema

The ontology leads to six design rules for Zlef. The schema chapter turns them into questions and codes.

1. **Facts, not categories.** Zlef predicts jurisdiction-free facts. Legal categories, tax characters and reporting duties are outputs of L and carry a status and a source.
2. **Codes follow the dimensions.** Each `asset_class` subcode exists because at least one regime draws a line there: single versus basket fiat (EMT versus ART), commodity versus fiat (UK QSC), yield (GENIUS, MiCA), stake versus restake (US staff guidance), custodial versus bridge versus burn-and-mint (issuer liability), unique versus fractional (EU, KR, US).
3. **Abstention is a feature.** Where L returns $\perp$ because rules are silent (VN, UA, most staking and DeFi cells outside the US and UK), the system routes to review. Zlef itself abstains with `UNKNOWN` only when the *facts* are uncertain, not when the *law* is.
4. **Registers, not predictions.** Issuer authorisation, licence and listing status come from layer-B registers with a dated snapshot.
5. **Bitemporal rules.** Every category carries an event date and an as-of date, so a rule change in 2027 does not silently rewrite a 2026 conclusion.
6. **Named-asset overrides.** Where a regulator names assets (for example the US list of digital commodities), the named list overrides the class default in the rules layer.

The next section applies this ontology to seven asset classes in turn.

4 Asset Types I

This section applies the ontology of Section 3 to seven asset classes: native tokens (NAT), stablecoins (STB), tokenized traditional assets (TKN), governance and utility tokens (GOV), memecoins (MEME), NFTs (NFT) and liquid staking and restaking receipts (LST). The remaining classes (WRP, DFI, PRV, CBDC, OTH) are treated in the next section.

Each subsection follows the same five steps.

1. **Definition and technical form.** What the asset is, in jurisdiction-free terms, and how it is implemented.
2. **Legal classification across jurisdictions.** A table with the category, regulator, status and source in each jurisdiction of the country set and in the hubs where our sources cover them.
3. **Taxation by event.** A table of the tax treatment of each Zlef `event_type` in the ten corridor countries.
4. **How it travels.** Transfer mechanics, travel-rule reach, issuer freeze powers, transfer restrictions, bridges, the split between on-chain and off-chain evidence, and the Paper V flag families that the asset tends to raise.
5. **How Zlef labels it, and why.** The `asset_class` code and subcode, the observable features, the usual confusions and how the schema separates them, when Zlef must abstain, which questions are ordinal and which nominal, and an illustrative label.

Conventions for the tax tables. The tax tables are written from the perspective of an individual resident taxpayer unless stated otherwise. Table 10 gives the cell codes. Each tax table ends with a status row that gives the status of the jurisdiction's regime as a whole. A cell marker overrides the status row for that cell. The NAT table is the baseline. Later tables show only the events where an asset class changes the baseline; all other events follow Table 13.

Table 10: Cell codes and markers used in the tax tables of this section.

Code	Meaning
D	Disposal: a gain or loss is realised on the asset given up.
I	Income on receipt, measured at market value.
N	Not a taxable event.
FT	Final or transaction tax on the gross value of the transfer.
unsettled	No authoritative guidance was found; treatment is ambiguous.
U	Not covered or not verified in our sources; status [Unverified].
†	Background knowledge, with the primary source named but not re-fetched on 15 Oct 2026; status [Unverified].
(p)	Draft legislation, bill or proposal; status [Proposed].
(e)	Enacted but not yet effective; status [Enacted, not yet effective].

Conventions for flags. We name flags by Paper V family. Identifiers of the form `FAM-XX-nn` in this section are generic placeholders, for example `STB-XX-01` or `AST-XX-01`. Paper V holds the definitive identifiers, triggers and severities. A flag raised here is a pointer to a rule in layer E, not a conclusion.

Conventions for labels. The illustrative labels use the field names and codes of the Zlef schema exactly. They are invented examples, labelled as such, and do not describe real transactions. Calibrated probabilities are omitted to keep the examples short.

Travel-rule baseline. Several subsections refer to the travel-rule thresholds of the corridor countries. Table 11 collects them once.

Table 11: Travel-rule thresholds for virtual-asset transfers (as of 15 Oct 2026).

Juris.	Rule and threshold	Self-hosted wallets	Status	Source
US	31 CFR 1010.410(f); transmittal orders of USD 3,000 or more	no BSA travel-rule duty; risk-based	[In force]	[219]
EU	TFR (EU) 2023/1113; no de minimis CASP-to-CASP	ownership or control check above EUR 1,000	[In force] since 30 Dec 2024	[71]
UK	MLRs Part 7A; basic data below EUR 1,000, full data above	collection not mandatory; risk assessment	[In force] since 1 Sep 2023	[144]
JP	APTCP; all transfers, no threshold; covers stablecoin providers	collect, retain, assess risk	[In force] since 1 Jun 2023	[83]
KR	KRW 1,000,000	whitelisted wallets in practice†	[In force]	[207]
KR	threshold removed for all VASP-to-VASP transfers, from about 20 Feb 2027	same-person rule for foreign exchanges and personal wallets	[Enacted, not yet effective]	[207]

Juris.	Rule and threshold	Self-hosted wallets	Status	Source
IN	no de minimis (FIU-IND guidelines, updated 8 Jan 2026)	higher risk; enhanced CDD	[In force]	[161]
BR	no threshold; domestic from 2 Feb 2026; cross-border from 2 Feb 2027; full by 2 Feb 2028	VASP must identify the owner; FX framework	domestic [In force]; cross-border [Enacted, not yet effective] (2 Feb 2027)	[158]
ID	all amounts; extended data at USD 1,000 or more	registered, verified withdrawal wallets	[In force]	[160]
VN	no VASP travel rule (FATF self-report); pilot licensees must meet FATF standards	U	[In force] (pilot)	[76, 189]
UA	legislation "Yes" per FATF self-report; threshold U	U	[Unverified]	[76]
RU	"in progress" per FATF self-report	only foreign-trade companies may send to self-custody (reported)	[Unverified]	[76]

The asymmetry between the US threshold and the zero thresholds elsewhere is the most common travel-rule friction in our corridor set. A US VASP may send below USD 3,000 without data, and the receiving VASP in Japan, India, Brazil, Indonesia or the UK must then treat the transfer as missing information [76]. This is a TRV-family flag, for example `TRV-XX-01`, and it applies to every asset class below that moves through VASPs.

4.1 NAT: native and protocol tokens

Definition and technical form. A native token is the unit of account of a blockchain's own ledger. It pays transaction fees and, in proof-of-stake systems, secures consensus through staking. It has no issuer in the legal sense. New units come from the protocol's issuance schedule (block rewards or staking rewards), not from a party that owes the holder anything. Transfers are protocol-level state changes signed with the holder's private key. On account-based chains the balance is a field of the account; on UTXO chains it is a set of unspent outputs. Zlef distinguishes two subcodes. `NAT.L1` is the native token of a base-layer chain, such as BTC, ETH or SOL. `NAT.L2` is the native token of a layer-2 or application chain whose own ledger is secured by, or settles to, another chain.

The defining technical property is the absence of a token contract. On an EVM chain, a native transfer is a value field on a transaction, not a call to an ERC-20 contract. This is observable and is the single strongest feature for the NAT class.

Legal classification across jurisdictions. Table 12 gives the categories. The pattern is simple. Fine-grained regimes give native tokens their residual or commodity category. Umbrella regimes give them the umbrella category. No jurisdiction in our set treats a native token as a security by default, but the US investment-contract overlay can attach to an initial sale.

Table 12: NAT: legal classification (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
US	Digital commodity; BTC, ETH and SOL are named examples; investment-contract overlay possible at initial sale	CFTC (derivatives, spot anti-fraud); SEC only under an investment contract	[In force] since 23 Mar 2026	[172, 194]
EU	Crypto-asset, Title II ("other"); white-paper duty on the person seeking admission when there is no issuer	NCAs, ESMA	[In force] since 30 Dec 2024	[72]
UK	Qualifying cryptoasset (art. 88F)	FCA	[Enacted, not yet effective]; commences 25 Oct 2027	[217]
JP	Crypto asset (PSA Art. 2(14)); moving to a distinct FIEA category	FSA, JVCEA	PSA [In force]; FIEA [Enacted, not yet effective]	[202, 228]
KR	Virtual asset (SFTA, VAUPA)	FSC, FSS, KoFIU	[In force] since 19 Jul 2024	[85]
IN	Virtual digital asset (tax and AML only)	CBDT, FIU-IND	[In force]	[49]
BR	Ativo virtual (Law 14,478/2022)	BCB (VASPs), Receita Federal	[In force]	[54]
ID	Crypto-asset within "digital financial assets"; tradable only if OJK-listed; not a means of payment	OJK	[In force]	[27]
VN	Crypto-asset in the pilot market; digital assets are property	MoF (pilot)	[In force]; statutory definition [Unverified]	[189]
UA	Virtual asset under the pending bill	NBU and designated regulator (bill)	[Proposed]	[223]
RU	Digital currency (259-FZ); barred from domestic payment; foreign-trade use allowed from 1 Sep 2026	Bank of Russia, MinFin	[In force]	[12, 179]
SG	Digital payment token	MAS	[In force]	[180]
HK	Virtual asset (AMLO)	SFC	[Unverified]	—
AE	Virtual asset (VARA; ADGM accepted virtual asset); not usable for mainland retail payment	VARA, FSRA, SCA	[In force]	[175]
CH	Payment token (AML rules only for pure payment tokens)	FINMA	[In force]	[89]

Three edge cases matter for the rules layer. First, the US list of named digital commodities is a dated, source-attested fact. It should override any class default. Second, [In force] the CFTC permitted listed spot crypto trading on designated contract markets from Dec 2025, which changes where a native-token trade can be executed but not its category [6]. Third, [In force] in Japan only assets screened by the self-regulatory organisation trade on registered domestic exchanges, so "listed in Japan" is a layer-B fact that matters for the planned separate tax regime [202].

Taxation by event. Table 13 is the baseline for the whole section. Across the ten countries the event, not the asset, decides the outcome [111, 123, 153]. The main split is between countries that tax gains (US, UK, JP, KR from 2027, IN, BR) and countries that levy a final tax on gross transfer value (ID, VN).

Table 13: NAT: taxation by event (baseline for all classes). Codes as in Table 10. Sources: US [119–121, 123]; UK [109–112]; JP [153]; KR [17]; IN [21, 117, 171]; BR [142, 212]; ID [56, 115]; VN [7]; UA [100].

Event	US	UK	JP	KR	IN	BR	ID	VN	UA	RU
`acquire_fiat`	N; basis = cost incl. fees†	N; cost enters s. 104 pool	N	N†	N; buyer deducts 1% TDS if seller is not an exchange	N; reportable	N (tax on seller side)	N	U	U
`dispose_fiat`	D (property)†	D	D (miscellaneous income)	D; 22% above KRW 2.5m	D at 30%; seller suffers 1% TDS	D; exempt if monthly sales $\leq$ BRL 35,000; 15–22.5%	FT 0.21% domestic platform; 1% foreign	FT 0.1% via licensed provider	U; bill: 18% + 5% levy (p)	U; property, PIT 13/15%†
`swap_crypto`	D on asset given up†	D	D (FAQ 1-3)	D†	D; TDS on each leg†	D†	FT on platform trade	FT	U	U
`pay_goods_services` (payer)	D (IRS FAQ 14)	D	D (FAQ 1-2)	D†	D; payee deducts 1% TDS	D†	off-platform: seller self-remits 1%	off-platform U	U	domestic payment barred [In force]
`receive_payment_goods_services` (business)	ordinary income at FMV†	trading receipt†	business income†	business income†	business income†	reportable	no VAT on crypto; income U	U	U	U
`compensation`	wages at FMV; withholding†	employment income; PAYE if readily convertible†	salary income†	earned income†	salary at slab rates†	salary†	salary†	U	U	U
`staking_reward`	I on dominion and control (Rev. Rul. 2023-14)	I (misc. or trade) at GBP value on receipt	I on acquisition (FAQ 1-7)	U	I (other sources); later sale at 30%	reportable; character unsettled	U	U	U	U
`mining_reward`	I at FMV†	I†	I (FAQ 1-7)	U	I	reportable; ordinary income [Sources conflict]	VAT 12% $\times$ 11/12 on mining services; income U	U	U	U
`airdrop`	I on dominion and control (Rev. Rul. 2019-24)	I if for a service; else N, CGT on disposal	U	U	I if FMV $>$ INR 50,000 from non-relatives	reportable	U	U	U	U
`fork`	N without new units; I if received	N; cost apportioned†	N at fork; basis JPY 0 (FAQ 1-6)	U	U	U	U	U	U	U

Event	US	UK	JP	KR	IN	BR	ID	VN	UA	RU
`internal_transfer`	N; per-wallet basis moves with units	N if beneficial ownership retained	N†	U	N; TDS can misfire on P2P self-transfers†	N; reported under DeCripto	U	U	U	U
`gift` (donor)	N; gift-tax regime†	D at market value; spouse no gain, no loss	donor includes market value in income (FAQ 2-10)	U	N for donor†; recipient I above INR 50,000	state gift tax (ITCMD)†	U	U	U	U
`loss_theft`	limited for personal-use property†; unsettled	not a disposal; negligible-value claim†	misc. losses do not offset other income (FAQ 2-11)	U	no set-off of VDA losses	U	n/a (gross tax)	n/a (gross tax)	U	U
`fee`	D of fee units; fees adjust basis†	allowable cost†	included in cost†	U	only cost of acquisition deductible	unsettled	U	U	U	U
Status of regime	[In force]	[In force]	[In force]	[Enacted, not yet effective] from 1 Jan 2027; deferral bill announced 15 Oct 2026 (p)	[In force] (Income-tax Act 2025 from 1 Apr 2026)	[In force]; MP 1.303 lapsed 8 Oct 2025	[In force] since 1 Aug 2025 (PMK 50/2025)	[In force] (Circulars 32/2026, 41/2026); no licensed provider at 30 Aug 2026	[Unverified]	[Unverified]

Four points from the table need emphasis. First, [In force] US basis has been tracked per wallet or account since 1 Jan 2025 under Rev. Proc. 2024-28, so an internal transfer is non-taxable but must carry its basis records with it [90, 121]. Second, [In force] UK pooling runs same-day, then 30-day, then the s. 104 pool per token type [112]. Third, [In force] Japan's default cost method for individuals is the total-average method, with moving average by election [153]. Fourth, [In force] Brazil's provisional measure MP 1.303/2025, which would have set a 17.5% flat rate, lost validity on 8 Oct 2025 and never took effect [182]. The pre-existing progressive scale applies.

Two further points are proposals and must not be encoded as law. [Proposed] US de minimis bills (S. 2207, H.R. 9178 and the ADAPT Act announced 30 Sep 2026) had not passed a chamber in our sources [37, 45]. [Proposed] The ADAPT Act would also extend the wash-sale rule to digital assets [45].

How it travels. A native transfer is signed by the holder and validated by the protocol. No issuer can freeze it. At protocol level, nobody can stop it except through censorship by block producers. Control points sit at the edges: the VASPs where value enters and leaves, and the address-screening tools they use [75]. The chain records sender and receiver addresses, amounts, timestamps and fees. It does not record who owns either address, why the transfer was made, the cost basis of the units, or their fiat value at the time.

This split defines the compliance surface. The travel rule applies when a VASP sits on either side (Table 11). Sanctions screening uses address lists. [In force] Treasury removed Tornado Cash from the sanctions list on 21 Mar 2025, which removes a sanctions flag for exposure to that protocol but does not remove AML risk scoring for mixer exposure [222]. The Paper V flag families that native tokens tend to raise are:

- **TRV**, for threshold asymmetry and missing originator data (`TRV-XX-01`), and transfers to self-hosted wallets (`TRV-XX-02`).
- **SAN**, for address-list matches (`SAN-XX-01`), including DPRK-attributed addresses after large thefts [75].
- **TAX**, for a disposal on every spend or swap (`TAX-XX-01`) and income on staking or mining receipt.
- **DAT**, for missing evidence that both wallets in an internal transfer belong to the same owner (`DAT-XX-01`). This is the main source of disposal mislabelling.
- **FXC**, in Brazil, where a cross-border payment with any virtual asset is an FX operation [9], and in Russia, where domestic payment is barred [12].

How Zlef labels it, and why. *Code.* `NAT.L1` or `NAT.L2`. The split follows the ledger, not the law. No regime in our set draws a legal line between L1 and L2 tokens, but the split matters for travel: an L2 native token often reaches the user through a bridge, and the L2's settlement path affects which events are observable.

Observable features. The transfer carries value in the protocol's native field rather than through a token contract. The asset identifier is the chain's native asset, not a contract address. Issuance comes only from protocol rewards. There are no admin roles, no blacklist functions and no mint authority held by a party.

Typical confusions. (i) A wrapped version of a native token on its own chain (for example a wrapped ether contract) is an ERC-20 contract, so Zlef labels it `WRP`, not `NAT`. (ii) A native token

bridged to another chain by a third-party bridge is `WRP.BRIDGE`. (iii) A token that pays fees on an application chain but was minted by a foundation contract on a different chain may be `GOV` on the origin chain and `NAT.L2` on the application chain. [Research hypothesis] We propose labelling by the ledger on which the observed transfer occurs, with `crosses_chain` and `passes_bridge` recording the path. (iv) A memecoin that is itself a chain's native asset, such as DOGE, is `NAT.L1`, not `MEME`. This matches the US source, which names DOGE as a digital commodity rather than a collectible [194].

When to abstain. Zlef outputs `UNKNOWN` for `asset_class` when the chain cannot be resolved (for example an unrecognised chain identifier) or when a transfer record lacks the asset field. It does not abstain because the law is unclear. A native token on a chain unknown to the rules layer is still `NAT`; the rules layer then returns $\perp$ for the legal category.

Ordinal and nominal questions. For native tokens the decisive nominal questions are `event_type` and `beneficial_owner_change`, because the tax outcome turns on whether a disposal occurred. `evidence_completeness` (ordinal 0–4) is usually 1 for a bare chain record. It rises to 2 with counterparty identity and to 3 with an invoice. `review_priority` (ordinal) rises when the counterparty is `self_hosted` and `beneficial_owner_change` is `unknown`, because an undetected self-transfer becomes a false disposal.

Illustrative label (invented example: a business pays a contractor's invoice in a base-layer native token to a self-hosted wallet).

```
{
  "asset_class": "NAT.L1",
  "event_type": "pay_goods_services",
  "economic_purpose": "service_purchase",
  "beneficial_owner_change": "yes",
  "counterparty_type": "self_hosted",
  "evidence_completeness": 3,
  "review_priority": "medium",
  "travel_surface": {
    "crosses_chain": "no", "passes_bridge": "no",
    "issuer_freezable": "no", "permissioned_transfer": "no",
    "privacy_enhanced": "no", "touches_mixer": "no",
    "touches_unhosted": "yes"
  }
}
```

The rules layer reads this tuple with the payer's residence and the date. For a US payer it would raise a disposal on the units given up. For a UK payer the same. For an Indian payer it would also raise the 1% TDS duty on the payee side. None of those outcomes is in the label.

4.2 STB: stablecoins

Definition and technical form. A stablecoin is a token designed to hold a stable value against a reference asset. Zlef defines its subcodes by the economic mechanism that holds the value, not by any statute.

- `STB.FIAT.SINGLE`: an issuer holds reserves in one fiat currency and redeems at par in that currency. Examples in our sources include USD, EUR and JPY coins.
- `STB.FIAT.BASKET`: as above, but against a basket of currencies.
- `STB.COMM`: an issuer holds a commodity, typically gold, and the token is a claim on it.

- `STB.SYN`: no fiat or commodity reserve of the token's own; the value is held by an algorithm, over-collateralised crypto, or a delta-neutral hedge.
- `STB.YLD`: a wrapper or design that passes yield to the holder, for example a staked or savings version of a stablecoin.

Technically, most stablecoins are token contracts (ERC-20 on EVM chains, SPL tokens on Solana). The issuer holds privileged roles: mint, burn, pause, and in most fiat designs a blacklist or freeze function. Supply changes when the issuer mints against deposits and burns on redemption. Cross-chain movement happens in one of two ways. The issuer may burn on one chain and mint on another, as in Circle's Cross-Chain Transfer Protocol (CCTP), which "natively transfers USDC across blockchains using burn-and-mint" with attestation by Circle's Iris service [32]. Or a third-party bridge may lock the original and mint a representation, which Zlef labels `WRP.BRIDGE`, not `STB`.

Legal classification across jurisdictions. Stablecoins are the asset class where regimes have moved fastest. Table 14 covers fiat-backed coins. Table 15 covers commodity-backed, synthetic and yield-bearing coins.

Table 14: STB.FIAT.*: legal classification of fiat-backed stablecoins (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
US	Payment stablecoin if issued by a permitted issuer; excluded from "security" and "commodity". No issuer is permitted before the effective date	OCC, Fed, FDIC, NCUA, states; Treasury	[Enacted, not yet effective]; effective by 18 Jan 2027	[48, 218]
US	Stablecoin category of the 2026 taxonomy; GENIUS and covered stablecoins not securities, others fact-dependent	SEC, CFTC	[In force]	[194]
US	Implementing rules (OCC 2 Mar 2026; Treasury Aug 2026; Fed 29 Sep 2026)	OCC, Treasury, Fed	[Proposed]	[78, 92, 165]
EU	E-money token if it references one official currency (Title IV); asset-referenced token if it references several	NCA of issuer; EBA if significant	[In force] since 30 Jun 2024	[72]
UK	Qualifying stablecoin (art. 88G): stable value against "a particular fiat currency", backed by fiat or other assets; e-money excluded	FCA; BoE for systemic sterling coins	[Enacted, not yet effective]; 25 Oct 2027	[217]
UK	Systemic sterling stablecoins: up to 70% short-term gilts; GBP 40bn issuance guardrail; no holding limits	BoE	[Proposed] (draft rules, 22 Jun 2026)	[10]
JP	Electronic payment instrument (PSA Art. 2(5)); issuers limited to banks, funds-transfer providers and trust companies; stays under the PSA after the FIEA move	FSA	[In force]	[202, 209]
KR	Virtual asset; "no statutory definition of stablecoins"	FSC; BoK on issuer eligibility	[In force]	[28]

Juris.	Category	Regulator	Status	Source
KR	Won-stablecoin regime (Digital Asset Basic Act, phase 2); bank-led consortia debated	FSC, BoK	[Proposed]	[25, 47]
IN	Virtual digital asset (stablecoins included); not a payment instrument; RBI opposed	CBDT, FIU-IND	[In force]	[49, 178]
BR	Fiat-referenced virtual asset inside the FX market when bought, sold or exchanged via a VASP; self-custody transfers are FX	BCB	[In force] since 2 Feb 2026	[9, 54]
BR	eFX providers barred from settling cross-border remittances in crypto or stablecoins	BCB	[In force] since 1 Oct 2026	[34, 136]
ID	Crypto-asset; no stablecoin category; tradable if OJK-listed; no payment use	OJK	[In force]	[27]
VN	U	—	[Unverified]	—
UA	U	—	[Unverified]	—
RU	Not separately defined; foreign stablecoins subject to the same requirements as digital currency under the 2026 framework; domestic payment barred	Bank of Russia, MinFin	[In force]	[12, 179]
SG	Digital payment token; "MAS-regulated stablecoin" label for compliant SGD or G10 single-currency coins issued in Singapore	MAS	[In force] (framework); statutory label [Unverified]	[148, 180]
HK	Fiat-referenced stablecoin; issuance is a licensed activity (Cap. 656); first licences 10 Apr 2026	HKMA	[In force] since 1 Aug 2025	[103, 104]
AE	Payment Token (dirham, mainland payments); Foreign Payment Token; VARA fiat-referenced VA; ADGM fiat-referenced token	CBUAE, VARA, FSRA	[In force]	[175]
CH	Issuer usually conducts deposit-taking unless a compliant bank guarantee exists	FINMA	[In force] (Guidance 06/2024)	[177]
CH	Payment Instrument Institution licence	FINMA	[Proposed]	[177]

Several details matter for the rules layer.

- **Permitted versus not.** [Enacted, not yet effective] Before the GENIUS effective date there is no "permitted" status in the US, so no current stablecoin is GENIUS-permitted. From about 18 Jul 2028, digital asset service providers may not offer non-permitted payment stablecoins (§3(b)(1)) [218]. [Proposed] Foreign issuers need a Treasury comparability determination, and Treasury's proposal sets a path [92]. [In force] A CFTC staff no-action position of 8 Dec 2025 lets futures commission merchants accept payment stablecoins as margin [132]; it is a non-binding staff view.
- **Multi-issuance.** [Proposed] The Commission notes that "MiCA does not currently prohibit multi-issuer models" and asks whether it should (Q30–40) [68].

- **Japanese issuance.** [In force] A yen stablecoin was first issued under an EPI licence in Oct 2025, and a trust-bank yen stablecoin launched on 24 Jun 2026 [202, 209].
- **Brazil's cross-border cap.** [Sources conflict] Our sources disagree on the limit for international payments with virtual assets where the counterparty is not an authorised FX institution: a law-firm summary gives USD 50,000 per transfer [54], while the BCB release gives "US$100.000,00" [9]. The rules layer must hold both values until the resolution text is read.
- **Brazil's transition.** [In force] The VASP transition ends on 30 Oct 2026 according to an industry tracker [176].
- **Russia.** [In force] The EU's 19th sanctions package of 23 Oct 2025 prohibited transactions in the rouble-pegged A7A5 token across the EU [40], and [In force] the 20th package of 23 Apr 2026 banned transactions in RUBx [41]. [In force] OFAC designated the issuer of A7A5 on 14 Aug 2025; the token itself is not listed as an SDN [221].

Table 15: STB.COMM, STB.SYN and STB.YLD: legal classification (as of 15 Oct 2026).

Juris.	Subcode	Category	Regulator	Status	Source
US	COMM	Not one of the five clean categories; likely a CEA commodity; *Howey* still applies; GENIUS does not apply	CFTC	[In force]	[35, 194]
US	SYN, YLD	Not a payment stablecoin if it pays yield or is not fully reserved in eligible assets; may be a security depending on facts	SEC, CFTC	[In force] (taxonomy); GENIUS [Enacted, not yet effective]	[194, 218]
US	YLD	Rebuttable presumption that issuer arrangements with an affiliate or related third party to pay yield breach the ban	OCC	[Proposed]	[165, 174]
EU	COMM	Asset-referenced token (Title III); interest ban (Art. 40)	NCA; EBA if significant	[In force]	[35, 72]
EU	SYN, YLD	EMT or ART if it references a value; interest ban (Arts. 40, 50); may be a financial instrument	NCA, EBA, ESMA	[In force]	[68, 183]
EU	YLD	Whether to relax the interest ban (Q20)	Commission	[Proposed]	[68]
UK	COMM, SYN	Qualifying cryptoasset, not a qualifying stablecoin (which needs fiat reference and backing)	FCA	[Enacted, not yet effective]	[217]
JP	SYN	Not an EPI (no redemption at face value); a crypto asset if tradable	FSA	[Unverified]	[202]
JP	COMM	Crypto asset or trust beneficiary right by structure	FSA	[Unverified]	—
KR	all	Virtual asset; a 2025 Seoul court decision distinguished fiat-collateralised from algorithmic coins	FSC	[In force]	[28]
IN	all	Virtual digital asset	CBDT	[In force]	[178]

Juris.	Subcode	Category	Regulator	Status	Source
BR	SYN	Ativo virtual; whether it is "fiat-referenced" (and so inside FX) is ambiguous	BCB	[In force]	[54]
ID	all	Crypto-asset; rules "do not separate fiat-backed from algorithmic stablecoins"	OJK	[In force]	[27]
RU	COMM	Digital financial asset if issued on a Russian DFA platform	Bank of Russia	[Unverified]	[179]
SG	COMM	Digital payment token if a medium of exchange; capital markets product if it gives investment returns	MAS	[In force]	[35]
CH	COMM	Asset token (security) if a claim on gold	FINMA	[In force]	[89]
VN, UA, HK, AE	all	U	—	[Unverified]	—

Taxation by event. For tax, a stablecoin is usually just a crypto-asset. Gains are small because the price is stable, but they are still gains. Table 16 shows where the stablecoin class changes the baseline.

Table 16: STB: taxation by event, deltas from the NAT baseline. Sources: US [37, 45, 122]; UK [107, 113]; JP [206]; BR [9, 142]; ID [56]; VN [7].

Event	US	UK	JP	KR	IN	BR	ID	VN	UA	RU
`swap_stable`	D; usually de minimis gain; H.R. 9178 would use redemption value (p)	D today; eligible stablecoins CGT-exempt from Apr 2027 (p); no rollover between exempt and non-exempt	if both are EPIs, outside crypto-asset rules; tax U	D†	D†	D; also FX if international	FT	FT	U	U
`dispose_fiat`, `pay_goods_services`	D; ADAPT Act: no gain on regulated USD stablecoin purchases (p)	D; exempt from Apr 2027 (p); no relief for de-pegging	EPI tax treatment U; excluded from planned 20.315% regime	D†	D; 1% TDS	D; international payment is FX; eFX settlement barred	FT	FT	U	domestic payment barred [In force]
`lend_out`, `lend_return`	unsettled	individuals: returns taxed like savings interest (p); companies: money debt (p)	lending income (FAQ 1-7)	U	U	crypto-backed loans reportable	U	U	U	U
issuer yield (STB.YLD)	unsettled; event code open	savings-interest treatment for lending returns (p); other yield U	U	U	U	U	U	U	U	U
Reporting	1099-DA: no report if qualifying-stablecoin sales $\leq$ USD 10,000 a year	CARF data from 1 Jan 2026	U	CARF 2027	s. 509 statements	DeCripto monthly; BCB 521 FX reports	U	U	U	U
Status	[In force] except (p)	[In force] except (p)	[In force]	[Enacted, not yet effective]	[In force]	[In force]	[In force]	[In force]	[Unverified]	[Unverified]

The UK proposal is the largest asset-specific delta in our set. [Proposed] HM Treasury's outcome of 13 Jul 2026 gives individuals a full CGT exemption for eligible stablecoins, including non-sterling ones, through Finance Bill 2026–27 from Apr 2027; it rejects "no gain, no loss" because of avoidance risk, gives no relief for de-pegging, and treats eligible stablecoins as a money debt for corporation tax [107]. A possible withholding tax on stablecoin lending and limited retrospection remain open. The rules layer must store this with its draft status and its start date.

[In force] In the US, the GENIUS Act contains no tax provisions in the portions read; stablecoins remain property for tax [218]. [In force] The 2025 Form 1099-DA instructions exempt qualifying-stablecoin sales from reporting when proceeds "do not exceed $10,000 for the year", with optional aggregate reporting above that [122]. This is a reporting rule, not a tax exemption.

[Unverified] For CARF, stablecoins are in scope unless they qualify as Specified Electronic Money Products, which move to the amended CRS [166]. [In force] DAC8 covers stablecoins and EMTs from 1 Jan 2026 [67].

How it travels. Fiat stablecoins move like any token, but with an issuer in the loop. The issuer can stop them. [Enacted, not yet effective] GENIUS requires issuers to have "technical capabilities, policies, and procedures to block, freeze, and reject" (§4(a)(5)(A)(iv)) and the capability to comply with any lawful order (§4(a)(6)(B)); a lawful order includes orders to "seize, freeze, burn, or prevent the transfer of payment stablecoins" (§2(16)(A)); foreign issuers' market access is conditioned on the same capability (§3(b)(2)) [218]. In practice, major fiat stablecoins already carry contract-level blacklist functions. A freeze is therefore a contract event that the chain records. For a stablecoin, "this address is frozen" is an on-chain, authoritative fact. For a native token there is no equivalent.

Two FATF findings shape the risk picture. [In force] FATF reports that "most identified on-chain illicit activity now involves stablecoins" and that VASPs "may be unable to rely on issuer-level asset freeze/burn mechanisms" [76]. Stablecoins are virtual assets for travel-rule purposes, and Japan's travel rule expressly covers electronic-payment-instrument service providers [83].

Cross-chain movement is the main travel risk. A burn-and-mint transfer by the issuer keeps the claim on the same issuer. [In force] CCTP's fast transfer takes about 8–20 seconds and its standard transfer 15–19 minutes on Ethereum and layer-2 chains, according to the issuer [32]. A lock-and-mint bridge creates a new liability on the destination chain. Each leg of a bridge transfer is visible on its own chain, but linking the legs requires decoding bridge messages, such as CCTP's message events, or a commercial attribution graph. Intent or solver bridges that settle from the solver's own inventory leave no common message, so the legs are harder to link.

Brazil adds a transaction-level rule. [In force] Under BCB Res. 521, a purchase, sale or exchange of a fiat-referenced virtual asset through a VASP is an FX operation, and transfers to or from self-custodied wallets require the VASP to identify the owner [9].

Typical flags:

- **STB**: non-permitted or non-authorised issuer for the jurisdiction (`STB-XX-01`); issuer yield or interest (`STB-XX-02`); eFX settlement in stablecoins in Brazil (`STB-XX-03`).
- **FXC**: Brazil's FX characterisation (`FXC-XX-01`).
- **SAN**: issuer freeze event or listed issuer, for example A7A5 in the EU (`SAN-XX-02`).
- **TRV**: as for native tokens.

- **REP**: US 1099-DA de minimis treatment; CARF versus CRS perimeter (`REP-XX-01`).

How Zlef labels it, and why. *Code.* One of the five `STB` subcodes. The subcodes are not chosen for elegance. Each one tracks a line that at least one regime draws: single versus basket (EMT versus ART in the EU), fiat versus commodity (qualifying stablecoin versus qualifying cryptoasset in the UK), reserved versus synthetic (GENIUS reserve rules; Japan's face-value redemption test), and yield (GENIUS §4(a)(11), MiCA Arts. 40 and 50).

Observable features. A token contract with an issuer-held mint and burn role; blacklist or freeze functions; supply that tracks issuer mint and burn events; a price series that stays near a fiat or commodity reference; issuer attestations or reserve reports (layer B). For `STB.YLD`, a rebasing balance, an exchange rate that grows against the base stablecoin, or a vault share token. For `STB.SYN`, collateral held in other crypto-assets or hedging positions rather than fiat reserves.

Typical confusions.

1. `STB.FIAT.SINGLE` versus `WRP.BRIDGE`. A bridged representation of a fiat stablecoin is not an issuer liability. It is a claim on a bridge contract. Zlef separates the two by the contract's mint authority: if the issuer mints, the class is `STB`; if a bridge mints, it is `WRP.BRIDGE`. Native coins at rest keep the `STB` code. The issuer's own burn-and-mint transfer event (CCTP) is labelled `WRP.NATIVEBURN` with `event_type` = `bridge`, `crosses_chain` = `yes` and `passes_bridge` = `no`, following the convention in Section 6.2.

2. `STB.FIAT.SINGLE` versus `TKN.DEP`. A tokenized bank deposit is a deposit, not a stablecoin. [In force] The EU excludes deposits from MiCA (Art. 2(4)) [72]. The separating feature is the issuer type (a bank issuing against its own deposit book), which is a layer-B fact.

3. `STB.FIAT.SINGLE` versus `CBDC`. A central bank liability is `CBDC`. The issuer identity separates them.

4. `STB.YLD` versus `TKN.MMF`. A tokenized money-market fund share also holds a near-stable value and pays a return. The difference is legal form (a fund unit versus a stablecoin wrapper). The separating features are transfer restrictions (fund shares are usually permissioned) and issuer documents. When neither is available, Zlef should stop at the parent level or abstain.

5. `STB.SYN` versus `STB.FIAT.SINGLE`. A synthetic dollar quotes like a fiat coin. Its reserves are not fiat. The separating feature is the collateral, which is visible on chain for most synthetic designs.

When to abstain. Zlef outputs `UNKNOWN` when it cannot tell whether a stable-priced token is issuer-backed, synthetic or a fund share, for example a new token with no attestations and opaque collateral. When the token is known to be a stablecoin but the subtype is not, Zlef emits the internal node `STB`, as the schema allows (Section 6.1, P4). Issuer authorisation ("is this issuer permitted in the US or authorised in the EU?") is never a Zlef output. It is a layer-B register fact.

Ordinal and nominal questions. The subcode is nominal. `evidence_completeness` matters more than for native tokens, because redemption and mint records at the issuer can raise the level to 3 or 4. The multi-label `travel_surface` carries two stablecoin-specific bits: `issuer_freezable` (true for most fiat designs) and `crosses_chain`. `review_priority` is corridor-free: it may rise for record facts such as a yield-bearing subcode or a missing reserve attestation. Escalation for a Brazil corridor (FX characterisation) or for yield-bearing coins in US or EU corridors is applied by the decision policy or the rules layer, which see the jurisdiction context.

Illustrative label (invented example: a business pays a supplier's invoice in a single-currency fiat stablecoin, on one chain, to the supplier's account at a regulated VASP).

```
{
  "asset_class": "STB.FIAT.SINGLE",
  "event_type": "pay_goods_services",
  "economic_purpose": "goods_purchase",
  "beneficial_owner_change": "yes",
  "counterparty_type": "regulated_vasp",
  "evidence_completeness": 3,
  "review_priority": "low",
  "travel_surface": {
    "crosses_chain": "no", "passes_bridge": "no",
    "issuer_freezable": "yes", "permissioned_transfer": "no",
    "privacy_enhanced": "no", "touches_mixer": "no",
    "touches_unhosted": "no"
  }
}
```

`issuer_freezable` is read from the issuer contract, not predicted (C7 in Section 6.10). `review_priority` is low because nothing in the record calls for urgency. If the payment ends in Brazil, the rules layer tests FX characterisation and the decision policy may escalate the item; the label does not change. Had the coin first moved across chains by the issuer's burn-and-mint protocol, that movement would be a separate event labelled `WRP.NATIVEBURN` (Section 5.1).

4.3 TKN: tokenized traditional assets

Definition and technical form. A tokenized traditional asset is a token that represents a claim that exists in traditional law independent of the token. Zlef has five subcodes.

- `TKN.MMF`: shares or units of a money-market or treasury fund, or tokenized treasury bills.
- `TKN.EQ`: equity, or a tracker or derivative that references equity.
- `TKN.DEP`: a commercial bank deposit recorded on a blockchain.
- `TKN.DEBT`: bonds, notes and other debt.
- `TKN.RE`: real estate and other real-world assets.

Technically these tokens are usually permissioned. Security-token standards such as ERC-3643 (T-REX) and ERC-1400 add an identity registry and a compliance module. Under ERC-3643 the token contract checks each transfer against an identity registry that maps wallets to identity contracts holding signed claims (KYC, accreditation, jurisdiction) from trusted claim issuers. A compliance contract applies modular rules such as country restrictions, holder caps and lock-ups. Agent roles can freeze addresses or partial balances, force transfers, recover lost wallets and pause the token. The official register may be held off chain by a transfer agent. These standards carry no legal classification. A token on ERC-3643 is a security only because of the rights it represents, but the standard is a strong feature.

Legal classification across jurisdictions. Table 17 shows near-uniform treatment for fund, equity and debt tokens: they are securities or financial instruments wherever our sources cover them. Tokenized deposits are deposits. The open questions are about trackers and synthetic structures, not about plain tokenized securities.

Table 17: TKN: legal classification (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
US	Digital security: "a security is a security regardless of whether it is issued … offchain or onchain"	SEC, FINRA; CFTC for swaps	[In force]	[172, 194]
US	Staff taxonomy of tokenization models: issuer-sponsored; third-party custodial; third-party synthetic (linked securities; security-based swaps)	SEC staff	[In force] (staff statement, 28 Jan 2026; non-binding)	[195, 197]
US	DTC no-action relief for tokenization (11 Dec 2025); tokenized MMF shares as collateral via CFTC guidance (8 Dec 2025)	SEC staff, CFTC	[In force] (staff positions)	[132]
US (DEP)	Bank deposit, not a payment stablecoin	bank regulators	[Unverified]	—
EU	Financial instrument under MiFID II, excluded from MiCA (Art. 2(4)); fund tokens are CIU units; T-bill notes transferable securities	NCAs, ESMA	[In force]	[72, 183]
EU	DLT Pilot Regime (Reg. 2022/858) for market infrastructures	NCAs, ESMA	[In force]	[62]
EU (DEP)	Deposit, excluded from MiCA; tokenised deposits addressed in the review	NCAs, ECB	[In force]; review [Proposed]	[68, 72]
UK	Specified investment cryptoasset, excluded from "qualifying cryptoasset" and regulated as the underlying	FCA	[Enacted, not yet effective]	[217]
JP	Electronically recorded transferable rights (FIEA securities); "no change" under the 2026 amendment	FSA	[In force]	[202]
KR	Security under the FSCMA; token securities are a new form of issuance, not a new category	FSC, FSS	[Enacted, not yet effective]; effective 4 Feb 2027	[25, 28]
KR (DEP)	CBDC deposit tokens excluded from "virtual asset"	FSC, BoK	[In force]	[85]
IN	Security under securities law if a share etc.; no tokenization framework found	SEBI	[Unverified]	—
BR	Security (valor mobiliário) if it meets the securities test	CVM	[Unverified]	—
ID	Boundary between digital financial asset and security; OJK consultation draft (Sep 2025)	OJK	[Proposed]	[27]
VN	U; possibly outside the crypto pilot	—	[Unverified]	—
UA	U	—	[Unverified]	—
RU	Digital financial assets (259-FZ): tokenized monetary claims, securities rights, equity in non-public companies; issued only on registered platforms	Bank of Russia	[In force]	[179]
SG	Capital markets product	MAS	[In force]	[35]
HK	Securities (SFO)	SFC	[Unverified]	—

Juris.	Category	Regulator	Status	Source
AE	Securities (ADGM FSRA, DFSA, SCA)	—	[Unverified]	—
CH	Asset token; ledger-based security (DLT Act, in force 1 Aug 2021)	FINMA	[In force]	[35, 89]
CH (DEP)	Deposit (banking)	FINMA	[In force]	[177]

The US staff taxonomy is the most useful source for Zlef. It shows that three identical-looking "stock tokens" can carry three different legal labels: a tokenized share issued by the issuer or its transfer agent; a custodial security entitlement; or a synthetic token that "may represent a security-based swap" because it gives executory payments tied to a security's value without ownership [197]. A law-firm gloss puts it as "tokenization changes the 'plumbing,' not the regulatory perimeter" [197]. [Unverified] The specific EU classification of retail stock-token products offered as trackers or derivatives was not verified in our sources. [Proposed] The MiCA review lists "tokenised fund interests" among the hardest assets to classify [68].

Taxation by event. Tokenized securities are taxed as securities. The token wrapper does not change the character. Table 18 shows the deltas. Most cells rest on background knowledge, which is itself a finding: our sources verify very little asset-specific tax guidance for tokenized securities.

Table 18: TKN: taxation by event, deltas from the NAT baseline. Sources: UK [113, 190]; other cells background (†).

Event	US	UK	JP	KR	IN	BR	ID	VN	UA	RU
`dispose_fiat`, `swap_crypto`	D under securities rules; wash-sale rule applies to stock and securities†	D; share-identification rules†	separate taxation as securities†	securities rules†	shares or securities, not VDA where excluded†; scope U	securities (CVM)†	securities†	U	U	U
distribution (dividend, coupon)	income†	income†	income†	U	U	U	U	U	U	U
`lend_out`	unsettled; securities-lending non-recognition covers securities only†	excluded from draft no-gain, no-loss relief (p)	U	U	U	U	U	U	U	U
`internal_transfer`	N†	N	N†	U	U	U	U	U	U	U
Reporting	digital-asset and securities reporting; dual reporting on 1099-B in some cases†	U	U	U	U	U	U	U	U	U
Status	[Unverified] except as noted	[In force] except (p)	[Unverified]	[Unverified]	[Unverified]	[Unverified]	[Unverified]	[Unverified]	[Unverified]	[Unverified]

[Proposed] The UK draft no-gain, no-loss rules for crypto lending and liquidity pools (effective 6 Apr 2027) exclude securities and tokenised assets from "qualifying cryptoassets" [190]. This is the one verified asset-specific tax rule for this class in our sources. [Unverified] Whether tokenized securities report under the CRS rather than CARF needs checking; F-layer review should treat this as open.

How it travels. Tokenized securities move only between whitelisted addresses. A transfer to a non-whitelisted counterparty reverts; it does not complete and is not flagged after the fact. "Transfer succeeded" is therefore evidence that the counterparty passed the issuer's onboarding gate. It is not evidence that the gate was adequate. The issuer or its agent can freeze, force-transfer, recover and pause. On chain, the observer sees transfers between whitelisted addresses and identity claims that are hashed or attested. The cap table, investor identity, accreditation and corporate actions sit off chain.

[Unverified] For FATF purposes, financial assets covered by other regulation are excluded from the virtual-asset definition, so the travel rule may not apply in the same way; securities transfer-agent records take its place. Tokenized deposits travel inside a bank's system and carry the bank's own payment controls.

Typical flags:

- **AST**: tokenized security with transfer restrictions (`AST-XX-01`); synthetic stock token that may be a security-based swap (`AST-XX-02`).
- **LIC**: offering a security token to residents of a jurisdiction without the required licence or prospectus (`LIC-XX-01`).
- **REP**: securities reporting rather than crypto reporting (`REP-XX-02`).
- **DAT**: missing evidence of the token's legal structure (`DAT-XX-02`).

How Zlef labels it, and why. *Code.* One of the five `TKN` subcodes. The subcodes follow the reference asset because tax and reporting follow it: a fund unit, a share, a deposit, a bond and a property claim each have their own rules in every jurisdiction.

Observable features. A permissioned transfer hook; an identity registry contract; agent roles with freeze and forced-transfer powers; transfers only among a small set of whitelisted addresses; a net asset value or coupon schedule published off chain; issuer or transfer-agent documents (layer B). For `TKN.DEP`, a bank issuer and redemption into the bank's own deposits.

Typical confusions. (i) `TKN.MMF` versus `STB.YLD`, discussed above; permissioning and fund documents separate them. (ii) `TKN.EQ` issued by the company versus a synthetic tracker. Both are `TKN.EQ` in the schema, because the schema records the reference asset, not the legal wrapper. The legal wrapper is a layer-B fact that the rules layer needs, and it should be stored alongside the label. (iii) `TKN.DEP` versus `STB.FIAT.SINGLE`, discussed above. (iv) `TKN.RE` versus `NFT.UNIQUE`. A single-property token may be a non-fungible token in form. If it conveys a claim on the property, Zlef labels it `TKN.RE`; the NFT form is a feature, not the class.

When to abstain. Zlef outputs `UNKNOWN` when a permissioned token's reference asset cannot be established from documents or registries. A permissioned transfer hook alone says "probably a security", but it does not say which kind.

Ordinal and nominal questions. The subcode is nominal. `permissioned_transfer` in `travel_surface` is almost always true and is the main binary feature. `evidence_completeness`

tends to be high (3 or 4), because the issuer's register and the investor's onboarding documents exist. `counterparty_type` is usually `regulated_vasp`, `bank_or_psp` or `business`, rarely `self_hosted`.

Illustrative label (invented example: a treasury firm redeems tokenized money-market fund shares into a fiat stablecoin with the fund's transfer agent).

```
{
  "asset_class": "TKN.MMF",
  "event_type": "swap_crypto",
  "economic_purpose": "investment",
  "beneficial_owner_change": "yes",
  "counterparty_type": "business",
  "evidence_completeness": 4,
  "review_priority": "low",
  "travel_surface": {
    "crosses_chain": "no", "passes_bridge": "no",
    "issuer_freezable": "yes", "permissioned_transfer": "yes",
    "privacy_enhanced": "no", "touches_mixer": "no",
    "touches_unhosted": "no"
  }
}
```

The schema labels a redemption into a stablecoin as `swap_crypto`. The rules layer, reading `TKN.MMF`, applies securities disposal rules rather than crypto rules.

4.4 GOV: governance and utility tokens

Definition and technical form. A governance or utility token is a fungible token, issued by a project or protocol through a token contract, that gives voting rights over a protocol, access to a service, or both. It is not the native fee token of a chain (that is `NAT`). It does not represent a traditional claim (that is `TKN`). Governance tokens often support delegation of voting power, and many are distributed with vesting schedules, lock-ups and airdrops. Some carry admin keys that can pause or upgrade the contract. Some protocols route fees to token holders through a "fee switch" or buy-back, which moves the token toward profit rights.

Legal classification across jurisdictions. GOV is the class where the classification is most fact-specific. Table 19 gives the categories.

Table 19: GOV: legal classification (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
US	Fact-specific: digital tool, digital commodity, or a token sold under an investment contract until essential efforts are complete; governance rights alone are not a named category	SEC, CFTC	[In force]	[172, 194]
US	Offering exemptions and a safe harbour for the end of the investment contract (Regulation Crypto Assets)	SEC	[Proposed] (18 Aug 2026)	[193]
US	Statutory "mature blockchain" and "digital commodity" paths (CLARITY Act)	Congress	[Proposed]; cloture failed 15 Sep 2026	[59, 168]

Juris.	Category	Regulator	Status	Source
EU	Utility token if "only" access to the issuer's good or service; otherwise Title II crypto-asset; financial instrument if it carries profit or dividend rights	NCAs, ESMA	[In force]	[72, 183]
UK	Qualifying cryptoasset; specified investment if it gives share- or debt-like rights	FCA	[Enacted, not yet effective]	[217]
JP	Crypto asset if tradable with unspecified persons; "specified crypto asset" with disclosure duties where a specified issuer has issuance authority; FIEA insider-trading rules	FSA	PSA [In force]; FIEA [Enacted, not yet effective]	[202, 228]
KR	Virtual asset unless an investment-contract security under the FSCMA	FSC	[In force]	[28]
IN	Virtual digital asset	CBDT	[In force]	[49]
BR	Ativo virtual	BCB	[In force]	[54]
ID	Crypto-asset (OJK list)	OJK	[In force]	[27]
VN	U	—	[Unverified]	—
UA	U	—	[Unverified]	—
RU	Digital currency	Bank of Russia	[In force]	[179]
SG	Digital payment token or capital markets product	MAS	[In force]	[180]
HK	Virtual asset	SFC	[Unverified]	—
AE	Virtual asset	VARA, FSRA	[Unverified]	—
CH	Utility token, not a security if functional at issuance; hybrid with investment function → security	FINMA	[In force]	[89]

Two US staff signals inform the rules layer. [In force] No-action letters for network and utility tokens (29 Sep 2025 and 24 Nov 2025) show the staff approach [132]; they are non-binding. [In force] The 2026 interpretation codified several 2025 staff statements; of these, the protocol-staking and covered-stablecoin statements are confirmed in secondary sources, while the dates of the memecoin and mining statements rest on background knowledge [172]. [Unverified] Whether EU utility tokens for existing services benefit from the white-paper exemption in MiCA Art. 4(3) was not re-verified in our sources.

The legal analysis for GOV turns on two dimensions from Section 3.4: rights conferred and decentralisation and control. Fee-switch or revenue-sharing designs look more like profit rights, which pull the token toward securities law in the US, the EU and Switzerland [89, 183]. Control by an identifiable team keeps the US investment-contract overlay alive.

Taxation by event. Governance tokens follow the NAT baseline for most events. Table 20 shows the deltas.

Table 20: GOV: taxation by event, deltas from the NAT baseline. Sources: JP [206, 230]; UK [110]; IN [171]; other cells background (†).

Event	US	UK	JP	KR	IN	BR	ID	VN	UA	RU
`compensation` (vesting tokens)	compensation on vesting†	employment income†	salary income†	U	salary†	salary†	U	U	U	U
`airdrop` (governance distribution)	I on dominion and control	I only if for or in expectation of a service	U	U	I above INR 50,000 from non-relatives	reportable	U	U	U	U
`dispose_fiat`	as NAT	as NAT	as NAT; 20.315% only if on the specified list, expected from 2028 [Sources conflict]	as NAT	as NAT	as NAT	as NAT	as NAT	U	U
Status	[In force]	[In force]	[In force]; reform [Sources conflict]	[Enacted, not yet effective]	[In force]	[In force]	[In force]	[In force]	[Unverified]	[Unverified]

The Japanese reform matters most for this class. [Sources conflict] The FY2026 tax reform outline provides 20.315% separate taxation for "specified crypto assets" on the FIBO register transferred through a registered operator, starting on 1 Jan of the year after the FIEA amendment takes effect; secondary press reports 105 specified assets and an expected start in January 2028 [206, 230]. Our sources conflict on whether the implementing tax provisions were enacted and on the start year (2027 or 2028); we write "expected from 2028". A governance token's tax rate in Japan may therefore depend on a list, which is a layer-B fact.

How it travels. Governance tokens move as ordinary ERC-20 or equivalent transfers. Most have no issuer freeze. Some have an admin pause or upgrade key. Vesting contracts and lock-ups restrict transfer for insiders, and delegation moves voting power without moving the tokens. VASPs apply the travel rule at the edges, and DEX front-ends may geo-block. On chain, the observer sees transfers, delegations and votes. Vesting and compensation records, lock-up terms and the issuer's promises sit off chain.

Typical flags:

- **AST**: possible investment-contract or security status, especially where fee-sharing or admin control is present (`AST-XX-03`).
- **TAX**: compensation on vesting (`TAX-XX-02`); airdrop income (`TAX-XX-03`).
- **LIC**: marketing to residents of a jurisdiction that requires an authorised offeror or white paper (`LIC-XX-02`).
- **TRV**: as for native tokens.

How Zlef labels it, and why. *Code.* `GOV`. The schema has no subcodes for this class. Governance and utility are combined because no regime in our set gives governance tokens a category of their own, and the line between governance and utility is a matter of rights that Zlef cannot see.

Observable features. A token contract deployed by a project, not by the chain; a delegation function; a governor or voting contract that reads token balances; a vesting or lock-up contract holding a large share of supply; a distribution history that starts with an airdrop or team allocation. Admin keys and upgradeability are separate features. They feed `review_priority`, not the class.

Typical confusions. (i) `GOV` versus `MEME`. Both are project tokens. The separating features are a governance contract or a service the token unlocks (GOV) against neither (MEME). Launch venue helps: tokens launched on a bonding-curve launchpad with no governance contract lean MEME. (ii) `GOV` versus `NAT.L2`. A token that pays fees on its own chain is `NAT.L2` on that chain, even if it is `GOV` on its origin chain. (iii) `GOV` versus `TKN.EQ`. A token with explicit revenue rights and a legal wrapper is closer to a security. Zlef still labels it `GOV` unless there is a traditional legal claim, but it raises `review_priority` so a reviewer can apply the rights test.

When to abstain. Zlef outputs `UNKNOWN` when a project token has neither a governance contract nor an identifiable service, and when its launch history is unknown. It does not abstain merely because the token's securities status is unclear; that is a layer-F question.

Ordinal and nominal questions. `event_type` is often `airdrop` or `compensation` for this class, both nominal. `economic_purpose` distinguishes `payroll` from `distribution`. `review_priority` is the main ordinal output. It should rise with admin control, fee-sharing features and insider allocations, because these are the facts that a reviewer needs for the investment-contract and profit-rights tests.

Illustrative label (invented example: a protocol foundation pays a contributor in vested governance tokens released from a vesting contract).

```
{
  "asset_class": "GOV",
  "event_type": "compensation",
  "economic_purpose": "payroll",
  "beneficial_owner_change": "yes",
  "counterparty_type": "smart_contract_protocol",
  "evidence_completeness": 3,
  "review_priority": "medium",
  "travel_surface": {
    "crosses_chain": "no", "passes_bridge": "no",
    "issuer_freezable": "no", "permissioned_transfer": "no",
    "privacy_enhanced": "no", "touches_mixer": "no",
    "touches_unhosted": "yes"
  }
}
```

The counterparty is the vesting contract that releases the tokens. The foundation behind it is a layer-B fact that a reviewer attaches from the grant agreement.

4.5 MEME: memecoins

Definition and technical form. A memecoin is a fungible token whose value comes from attention and community, not from a protocol function, a claim or a governance right. Memecoins are usually ordinary token contracts. Many launch on bonding-curve launchpads and then trade on decentralised exchanges. Some contracts include functions that matter for compliance: blacklists, transfer taxes, or "honeypot" logic that lets buyers in but stops them from selling. Supply is often fixed at launch, with large allocations to the deployer or insiders.

Legal classification across jurisdictions. "Memecoin" is not a legal category in any regime we reviewed. Table 21 shows that each regime puts memecoins into its residual or collectible bucket.

Table 21: MEME: legal classification (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
US	Digital collectible (WIF named); not a security generally; may become a digital commodity if it becomes functional; DOGE and SHIB are named digital commodities	SEC (fraud via other laws), CFTC anti-fraud	[In force]	[194]
US	2025 staff statement on meme coins (27 Feb 2025)	SEC staff	[Unverified] (date)	—
EU	Title II crypto-asset; admission to trading needs a white paper; with no issuer, the CASP seeking admission bears the duty	NCAs	[In force]	[72]
UK	Qualifying cryptoasset	FCA	[Enacted, not yet effective]	[217]
JP	Crypto asset; domestic listing only after self-regulatory screening	FSA, JVCEA	[In force]	[202]

Juris.	Category	Regulator	Status	Source
KR	Virtual asset; listing and delisting under VAUPA unfair-trading rules	FSC	[In force]	[85]
IN	Virtual digital asset	CBDT	[In force]	[49]
BR	Ativo virtual	BCB	[In force]	[54]
ID	Crypto-asset if OJK-listed	OJK	[In force]	[27]
VN	U	—	[Unverified]	—
UA	U	—	[Unverified]	—
RU	Digital currency	Bank of Russia	[In force]	[179]
SG	Digital payment token	MAS	[In force]	[180]
HK, AE	Virtual asset	SFC; VARA, FSRA	[Unverified]	—
CH	Payment token	FINMA	[In force]	[89]

The US entries show why the rules layer needs named-asset overrides. A class default of "digital collectible" would mislabel DOGE and SHIB, which the interpretation names as digital commodities [194]. A memecoin can also be sold subject to an investment contract if promoters promise yield or a roadmap.

Taxation by event. Memecoins follow the NAT baseline in every jurisdiction in our set. Table 22 shows the few deltas, which are about loss events and Japan's planned list.

Table 22: MEME: taxation by event, deltas from the NAT baseline. Sources: JP [153, 206]; IN [171]; other cells background (†).

Event	US	UK	JP	KR	IN	BR	ID	VN	UA	RU
`dispose_fiat`	as NAT; no special rule†	as NAT†	comprehensive taxation; likely outside the specified-asset list [Unverified]	as NAT	as NAT	as NAT	as NAT	as NAT	U	U
`loss_theft` (rug pull, honeypot)	limited for personal-use property†; unsettled	negligible-value claim†	no offset against other income	U	no set-off	U	n/a	n/a	U	U
Status	[In force]	[In force]	[In force]	[Enacted, not yet effective]	[In force]	[In force]	[In force]	[In force]	[Unverified]	[Unverified]

The practical tax issue for memecoins is evidence, not rate. High rates of theft and rug pulls make loss claims frequent, and loss rules are among the least settled cells of the baseline.

How it travels. Memecoins move through ordinary token transfers and through DEX swaps. Usually nobody can freeze them, but some contracts include blacklist, honeypot or transfer-tax functions. Everything is on chain, including clusters of deployer and insider wallets. Off-chain promotion and insider allocation agreements are not. Travel-rule duties apply only where a VASP is involved. Launchpad and DEX front-ends are the main screening points.

Typical flags:

- **AST**: asset-class status, including the named-asset override and a possible investment-contract overlay where promoters promise returns (`AST-XX-04`).
- **AML**: insider clusters, rapid wash-like trading and honeypot contracts (`AML-XX-01`).
- **TAX**: loss and theft claims (`TAX-XX-04`).
- **DAT**: missing price evidence for illiquid tokens (`DAT-XX-03`).

How Zlef labels it, and why. *Code.* `MEME`. The class exists in the schema although no law uses the word. It exists because the economic and evidence profile differs from GOV: there is no governance contract or service, valuation is noisy, insider concentration is high, and contract-level traps are common. These facts change how the rules layer and the reviewer should treat a payment.

Observable features. No governance, staking or service function; a launchpad origin; a fixed supply with large early allocations; thin and volatile DEX liquidity; transfer-tax or blacklist functions in the contract; very high holder turnover in the first days.

Typical confusions. (i) `MEME` versus `NAT.L1`: a chain's native asset that began as a joke is `NAT`, because it has no token contract. (ii) `MEME` versus `GOV`: a memecoin that later adds a governance contract moves to `GOV` from that date. Labels are dated. (iii) `MEME` versus `NFT.FRAC`: fungible fan or community tokens are `MEME`, while fractions of a single NFT are `NFT.FRAC`.

When to abstain. Zlef outputs `UNKNOWN` when a token's contract cannot be inspected (for example, an unverified contract on a chain without source verification) and its launch path is unknown. It should not default every unknown fungible token to `MEME`; that would make MEME a waste-basket class and spoil its calibration.

Ordinal and nominal questions. `review_priority` (ordinal) matters most. It should rise for honeypot features, insider-cluster exposure and illiquid pricing. `evidence_completeness` is usually 1. `event_type` is commonly `swap_crypto`.

Illustrative label (invented example: a user swaps a launchpad memecoin for a fiat stablecoin on a DEX).

```
{
  "asset_class": "MEME",
  "event_type": "swap_crypto",
  "economic_purpose": "investment",
  "beneficial_owner_change": "yes",
  "counterparty_type": "smart_contract_protocol",
  "evidence_completeness": 1,
  "review_priority": "medium",
  "travel_surface": {
    "crosses_chain": "no", "passes_bridge": "no",
```

```
    "issuer_freezable": "no", "permissioned_transfer": "no",
    "privacy_enhanced": "no", "touches_mixer": "no",
    "touches_unhosted": "yes"
  }
}
```

4.6 NFT: non-fungible tokens

Definition and technical form. An NFT is a token with a unique identifier within its contract. On EVM chains it follows ERC-721 (one owner per token identifier) or ERC-1155 (which allows several copies of the same identifier). Ownership is on chain. The referenced content (image, ticket, name, membership) usually sits off chain, linked by a URI. Sales are matched by marketplaces. Royalties are optional and enforced only by some marketplaces. Zlef has two subcodes. `NFT.UNIQUE` is a genuinely one-of-one item or an item in a small set of distinct items. `NFT.FRAC` covers fractionalised NFTs (a vault holds the NFT and issues fungible shares) and large series of interchangeable editions.

Legal classification across jurisdictions. The NFT boundary is drawn differently in every fine-grained regime (Table 23). That is why the schema splits the class.

Table 23: NFT: legal classification (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
US	Digital collectible (CryptoPunks, Chromie Squiggles) or digital tool (ENS names, tickets); not securities; fractionalised collectibles "could constitute the offer or sale of a security"; NFT defined as "a non-interchangeable crypto asset with a unique digital identifier"	SEC	[In force]	[194]
EU	Out of scope if "unique and not fungible" (Art. 2(3)); fractions not unique; large series indicate fungibility (Recital 11); may still be a financial instrument	NCAs, ESMA	[In force]	[72, 183]
EU	"Assets marketed as NFTs but issued in series" (Q7); regulating NFT service providers (Q68)	Commission	[Proposed]	[68]
UK	Not a qualifying cryptoasset if not fungible; NFTs not named; digital assets can be personal property	FCA	[Enacted, not yet effective]	[135, 217]
JP	Not a crypto asset without a payment function; indicators: JPY 1,000 or more per unit ("high" price), 1,000,000 units or fewer ("limited" quantity), not a safe harbour; dividend-bearing NFTs may be securities	FSA	[In force] (FSA comment, 24 Mar 2023)	[202]

Juris.	Category	Regulator	Status	Source
KR	Excluded from virtual assets unless mass or serial issuance, divisible into fractional units, usable as payment, or exchangeable for virtual assets among the public; securities check first	FSC	[In force] (guideline, 10 Jun 2024)	[87]
IN	Virtual digital asset definition covers NFTs; notification excludes some	CBDT	[Unverified]	—
BR	U	—	[Unverified]	—
ID	U	—	[Unverified]	—
VN	U	—	[Unverified]	—
UA	U	—	[Unverified]	—
RU	U	—	[Unverified]	—
CH	Utility- or asset-token analysis; typically not a security	FINMA	[Unverified]	—
SG, HK, AE	U	—	[Unverified]	—

Taxation by event. NFTs change the NAT baseline in the US (collectibles rate and reporting thresholds), the UK (pooling), Japan (exclusion from the planned regime) and Korea (outside the virtual-asset perimeter). Table 24 shows the deltas.

Table 24: NFT: taxation by event, deltas from the NAT baseline. Sources: US [122]; JP [206]; other cells background (†).

Event	US	UK	JP	KR	IN	BR	ID	VN	UA	RU
`dispose_fiat`, `swap_crypto`	D; possible collectible rate up to 28% via look-through†	D; not pooled where non-fungible†	not covered in FAQ excerpt; excluded from planned 20.315% regime	outside VA perimeter unless conditions met†; tax U	VDA; some NFTs excluded by notification†	DeCripto scope for NFTs U	U	U	U	U
creator primary sale (`receive_payment_goods_services`)	ordinary income†	U	U	U	U	U	U	U	U	U
Reporting	1099-DA: no report if specified-NFT proceeds $\leq$ USD 600 a year	U	U	U	U	U	U	U	U	U
Status	[In force] except †	[Unverified]	[Unverified] (exclusion reported; reform [Sources conflict])	[Unverified]	[Unverified]	[Unverified]	[Unverified]	[Unverified]	[Unverified]	[Unverified]

[In force] The 2025 Form 1099-DA instructions exempt specified-NFT sales from reporting when proceeds "do not exceed $600 for the year" [122]. [Unverified] For CARF, NFTs are in scope where they can be used for payment or investment, on a case-by-case basis [166].

How it travels. NFTs move through ERC-721 or ERC-1155 transfers and marketplace order matching. Marketplaces and collection operators can delist or block transfers through operator filters. There is rarely an issuer freeze. On chain, the observer sees the token identifier, the transfer and the sale price when the sale settles on chain. Bundled trades, off-chain buyer identity, wash-trading intent and fair value are off chain. [Unverified] FATF treats NFTs as virtual assets when they are used for payment or investment in practice, so the travel rule applies case by case. Marketplace KYC is often absent, so the payment leg, not the NFT leg, is usually where screening happens.

Typical flags:

- **AST**: NFT boundary test differs by jurisdiction; fractional or serial issuance may bring the token into the crypto-asset or securities perimeter (`AST-XX-05`).
- **AML**: wash trading and self-dealing at inflated prices (`AML-XX-02`).
- **TAX**: possible US collectible rate; creator income (`TAX-XX-05`).
- **REP**: US specified-NFT de minimis (`REP-XX-03`).
- **DAT**: missing fair value for bundled or illiquid trades (`DAT-XX-04`).

How Zlef labels it, and why. *Code.* `NFT.UNIQUE` or `NFT.FRAC`. The split exists because every fine-grained regime moves fractional and serial NFTs toward the crypto-asset or securities perimeter: the EU through Recital 11, Korea through its mass-issuance and divisibility tests, Japan through its quantity indicator, and the US through its fractionalised-collectible remark.

Observable features. The token standard (ERC-721 or ERC-1155); the number of token identifiers in the collection; the number of copies per identifier under ERC-1155; whether the NFT is held by a fractionalisation vault that issues fungible shares; floor price and sale prices; whether the token is used as a ticket, name or membership (digital-tool cues).

Typical confusions. (i) `NFT.UNIQUE` versus `NFT.FRAC` for large collections of distinct items. [Research hypothesis] We propose that Zlef record collection size and copies per identifier as features and use `NFT.FRAC` only for fractional shares and editions of interchangeable copies; collection size then feeds the rules layer's series tests, which differ by jurisdiction. (ii) `NFT.UNIQUE` versus `TKN.RE`: an NFT that conveys a claim on a property is `TKN.RE`. (iii) `NFT` versus `OTH`: in-game items that cannot leave the game are `OTH`; transferable in-game NFTs are `NFT`. (iv) The shares issued by a fractionalisation vault are fungible tokens; the schema labels a transfer of those shares as `NFT.FRAC`, not `MEME`.

When to abstain. Zlef outputs `UNKNOWN` when the collection's structure cannot be read (for example, an unverified contract with custom transfer logic) or when it cannot tell whether a fungible token is a vault share or an independent token.

Ordinal and nominal questions. The subcode is nominal. `evidence_completeness` is often 1, because marketplace identity is weak. `review_priority` should rise for high-value sales between related wallets (wash-trading cues) and for fractional shares, which are record facts; any jurisdiction-specific escalation is left to the decision policy.

Illustrative label (invented example: a business buys a one-of-one domain-name NFT from an individual through a marketplace, paying in fiat; the label is for the NFT leg, from the buyer's side).

```
{
  "asset_class": "NFT.UNIQUE",
  "event_type": "acquire_fiat",
  "economic_purpose": "goods_purchase",
  "beneficial_owner_change": "yes",
  "counterparty_type": "individual",
  "evidence_completeness": 2,
  "review_priority": "low",
  "travel_surface": {
    "crosses_chain": "no", "passes_bridge": "no",
    "issuer_freezable": "no", "permissioned_transfer": "no",
    "privacy_enhanced": "no", "touches_mixer": "no",
    "touches_unhosted": "yes"
  }
}
```

The label describes the acquisition of the NFT by the buyer. Had the buyer paid in crypto, the payment leg would be a separate event with its own label (`pay_goods_services` on the asset paid), and the rules layer would link the two.

4.7 LST: liquid staking and restaking receipts

Definition and technical form. A liquid staking token is a transferable receipt for a native token deposited with a staking protocol or provider. The protocol stakes the deposit with validators and passes staking rewards to receipt holders. Zlef has two subcodes. `LST.STAKE` is a receipt for protocol staking of a native token. `LST.RESTAKE` is a receipt for restaking, where staked assets or liquid staking tokens are deposited again to secure additional services, with operators chosen by the restaking protocol.

Rewards reach holders in one of two ways. A rebasing token increases each holder's balance as rewards accrue. A value-accruing token keeps the balance fixed and raises its exchange rate against the native token. Minting happens on deposit. Redemption usually passes through a withdrawal queue. Many protocols have governance or admin roles that can pause contracts, and some restaking tokens sit on upgradeable contracts.

Legal classification across jurisdictions. Table 25 shows that the US and the UK have addressed staking directly, the EU has not, and most other cells are unverified.

Table 25: LST: legal classification (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
US	Staking Receipt Token on a non-security asset is not a security; receipts "evidence Depositors' ownership of the deposited digital commodities"; a receipt for a security is a security	SEC	[In force]	[194]

Juris.	Category	Regulator	Status	Source
US	Conditions: one-to-one issuance; protocol-driven rewards; no guaranteed or provider-influenced rewards; administrative or ministerial activity; restaking tokens excluded	SEC staff	[In force] (staff statement, 5 Aug 2025; non-binding)	[79]
EU	Title II crypto-asset, or a financial instrument if it pays a managed return; staking "currently not addressed"; custodial staking needs custody authorisation	NCAs, ESMA	[In force]; review [Proposed]	[68]
UK	Qualifying cryptoasset; staking ("the use of a qualifying cryptoasset in blockchain validation") is a regulated activity	FCA	[Enacted, not yet effective]; 25 Oct 2027	[198, 217]
UK	Whether issuing a liquid staking token is "making arrangements for staking"	FCA	[Unverified]	—
JP	Crypto asset if tradable	FSA	[Unverified]	—
KR	Virtual asset	FSC	[Unverified]	—
IN	Virtual digital asset	CBDT	[In force]	[49]
BR	U (ativo virtual by default)	BCB	[Unverified]	—
ID	U	—	[Unverified]	—
VN	U	—	[Unverified]	—
UA	U	—	[Unverified]	—
RU	U	—	[Unverified]	—
SG, HK, AE, CH	U	—	[Unverified]	—

The US line between `LST.STAKE` and `LST.RESTAKE` is explicit. Restaking tokens are outside the staff guidance, so they remain fact-specific; discretionary selection of the services to secure points toward managerial efforts [79].

Taxation by event. LSTs raise two tax questions that the baseline does not answer. Is minting or redeeming the receipt a disposal of the native token? And when does staking income arise for the holder of a receipt? Table 26 shows that our sources answer neither question with authority in most countries.

Table 26: LST: taxation by event, deltas from the NAT baseline. Sources: US [120, 127]; UK [109, 111, 113]; JP [153]; IN [171]; BR [142]; ID [56].

Event	US	UK	JP	KR	IN	BR	ID	VN	UA	RU
`wrap_unwrap` (mint or redeem receipt)	unsettled; receipt exchange may be a disposal†	unsettled; turns on beneficial ownership; staking scope of draft rules U	unsettled; may be a swap (FAQ 1-3 by analogy)	U	U	reportable	U	U	U	U
`staking_reward` (rebasing)	I on dominion and control; timing per receipt unsettled	I on receipt at GBP value	I on acquisition (FAQ 1-7)	U	I (other sources)	reportable; character unsettled	U	U	U	U
`staking_reward` (value-accruing)	unsettled: income vs gain on redemption	unsettled	unsettled	U	U	U	U	U	U	U
`dispose_fiat`	as NAT	as NAT	as NAT	as NAT	as NAT	as NAT	FT on platform sale	as NAT	U	U
Status	[In force]	[In force]; draft (p)	[In force]	[Enacted, not yet effective]	[In force]	[In force]	[In force]	[In force]	[Unverified]	[Unverified]

[In force] US staking rewards are "included in the taxpayer's gross income in the taxable year in which the taxpayer gains dominion and control" (Rev. Rul. 2023-14) [120]. A refund suit challenging this was dismissed as moot without a ruling on the merits [127]. Neither source addresses receipt tokens. Our own reading is that rebasing and value-accruing designs may change the timing of income, because one adds units and the other raises a price. That reading is [Unsupported, needs evidence] and is an analytical point, not guidance. [Proposed] The UK draft no-gain, no-loss rules cover lending and automated market making; whether they cover staking is not confirmed in our sources [113].

How it travels. An LST is minted by depositing the native token into the protocol, moves like an ERC-20 token, and is redeemed through a withdrawal queue. Protocol governance or an admin may pause it. Restaking tokens add operator delegation and, often, upgradeable contracts. On chain, the observer sees mint and burn events, rebasing events and operator delegation. The attribution of rewards to each user, which income tax needs, is off chain or must be reconstructed. Travel-rule duties apply through VASPs at the edges. Restaking operator sets are an additional screening point.

Typical flags:

- **TAX**: ambiguity on mint and redeem (`TAX-XX-06`); income timing for rebasing versus value-accruing receipts (`TAX-XX-07`).
- **AST**: restaking receipt outside the US staff guidance (`AST-XX-06`).
- **LIC**: staking as a regulated activity in the UK from 25 Oct 2027 (`LIC-XX-03`).
- **DAT**: missing per-user reward attribution (`DAT-XX-05`).

How Zlef labels it, and why. *Code.* `LST.STAKE` or `LST.RESTAKE`. The split tracks the US staff line, which is the only verified legal line in our sources for this class, and it also tracks a real difference in risk: restaking adds discretionary operator selection and slashing exposure beyond the base protocol.

Observable features. A token minted only against deposits of a native token (or of another LST for restaking); a staking deposit contract or validator registry interaction; a rebasing balance or a growing exchange rate; a withdrawal queue; operator delegation events for restaking.

Typical confusions. (i) `LST.STAKE` versus `WRP.CUST`. Both are one-to-one receipts for a native token. The separating feature is whether the deposit is staked: an LST's backing sits in validators, a custodial wrap's backing sits in a custody account. (ii) `LST.STAKE` versus `DFI.LEND`. A lending receipt also accrues yield, but from borrowers, not from consensus rewards. The source of yield (staking deposit contract versus lending pool) separates them. (iii) `LST.RESTAKE` versus `LST.STAKE` when an LST is deposited into a restaking protocol. The receipt issued by the restaking protocol is `LST.RESTAKE`; the underlying LST, while it sits in that protocol, is not transferred to a new owner.

When to abstain. Zlef outputs `UNKNOWN` when it cannot tell whether a yield-bearing receipt draws on staking, lending or a strategy vault. That happens with aggregator vaults that hold several strategies.

Ordinal and nominal questions. The mint and redeem events are labelled `wrap_unwrap` in `event_type`; the tax consequence is left to the rules layer. Under the per-leg rule of Section 6.5, the deposit leg of a mint moves the native token into a pooled protocol contract in exchange for a receipt, so `beneficial_owner_change` is `yes`; whether the UK or any other regime treats that

leg as a disposal is a rules question. `review_priority` may rise for restaking receipts as a record fact (operator delegation, upgradeable contracts); escalation for particular corridors is applied by the decision policy.

Illustrative label (invented example: an individual deposits a native token into a liquid staking protocol and receives a value-accruing receipt).

```
{
  "asset_class": "LST.STAKE",
  "event_type": "wrap_unwrap",
  "economic_purpose": "investment",
  "beneficial_owner_change": "yes",
  "counterparty_type": "smart_contract_protocol",
  "evidence_completeness": 1,
  "review_priority": "medium",
  "travel_surface": {
    "crosses_chain": "no", "passes_bridge": "no",
    "issuer_freezable": "no", "permissioned_transfer": "no",
    "privacy_enhanced": "no", "touches_mixer": "no",
    "touches_unhosted": "no"
  }
}
```

`beneficial_owner_change` is `yes` under the schema's economic rule, and the event satisfies C3 of Section 6.10 because one side is an `LST` receipt. The label is not a disposal finding: the rules layer, which reads the event and the jurisdiction, raises an ambiguity flag where the law is unsettled.

4.8 Summary across the seven classes

Table 27 condenses the section into the facts the rules layer and the reviewer need most.

Table 27: Summary of the seven asset classes.

Class	Strongest observable feature	Main confusion	Abstain when	Most frequent flag families
NAT	no token contract; protocol issuance	wrapped native (WRP); L2 gas token	chain unresolved	TRV, SAN, TAX, DAT
STB	issuer mint/burn and freeze roles; stable price	bridged copy (WRP.BRIDGE); deposit token; fund share	mechanism unclear	STB, FXC, SAN, TRV, REP
TKN	permissioned transfer; identity registry	yield stablecoin; NFT-form RWA	reference asset unknown	AST, LIC, REP, DAT
GOV	governance or delegation contract	MEME; NAT.L2; security-like rights	no function and unknown launch	AST, TAX, LIC
MEME	launchpad origin; no function	GOV; native memecoin (NAT)	contract unreadable	AST, AML, TAX, DAT
NFT	ERC-721/1155; collection structure	series vs unique; RWA; in-game	structure unreadable	AST, AML, TAX, REP, DAT
LST	mint against staking deposit; reward accrual	custodial wrap; lending receipt	yield source unclear	TAX, AST, LIC, DAT

Three conclusions carry into the schema chapter. First, every class has at least one confusion that only an issuer or document fact can resolve. Those facts belong in layer B, and the label

should carry a pointer to them. Second, the cells marked U in the legal and tax tables are concentrated in VN, UA and RU and in the LST and NFT classes. The rules layer must return $\perp$ there and route to review. Third, the subcodes that look fine-grained (single versus basket fiat, stake versus restake, unique versus fractional) each match a legal line drawn by at least one regime in our set. That is the test for adding any further subcode in later editions.

5 Asset Types II, Events and Movement

This section completes the asset ontology and then turns from *what* an asset is to *what happens* to it and *how it moves*. It has three parts.

- **Part 1** covers the five remaining Zlef asset classes: wrapped and bridged representations (`WRP`), DeFi position receipts (`DFI`), privacy coins (`PRV`), CBDCs and e-money outside MiCA scope (`CBDC`), and the residual class (`OTH`). Each class follows the same five headings as the classes in the preceding section.
- **Part 2** (Section 5.6) covers the event dimension. It gives one consolidated table of tax character for every Zlef `event_type` across the ten jurisdictions plus an EU reference column, and then analyses the hard cases.
- **Part 3** (Section 5.7) covers the `travel_surface` label set: what each label means, what the chain shows, and which regimes it activates.

Reading conventions. Every legal or regulatory statement carries one status marker, as of 15 Oct 2026. Some sources were known to us but could not be re-checked against the primary text on that date. We mark those [Unverified] even where the rule is long-standing, because Zlef must not inherit stale law. In the tax tables, *unsettled* means that we located no statute, regulation, ruling or official guidance that answers the question. It is a finding about the sources, not a legal claim, so it carries no marker. Nothing here is legal or tax advice, and nothing here describes the position of any particular person or firm.

Why these five classes are hard. The classes in the preceding section are mostly *things*: a coin, a stablecoin, a share, a collectible. The five classes here are mostly *relationships*. A wrapped token is a claim on someone who holds another token. An LP token is a claim on a pool. A lending receipt is a claim on a protocol. A privacy coin hides the relationship between sender and receiver. A CBDC is a claim on a central bank that most crypto regimes exclude by definition. A points balance or an unvested allocation is a claim that may not yet exist on a chain at all. For each of these, the legal outcome depends less on the token and more on the event and the relationship it creates. That is why Part 2 follows directly.

5.1 Wrapped and Bridged Representations (WRP)

Definition and technical form. A wrapped or bridged representation is a token whose value depends on a claim to another token, usually on another chain. Zlef separates three technical forms, because they differ in who holds the underlying and who can stop a transfer.

- `WRP.CUST`**, custodial wrap.** A custodian holds the underlying asset and mints a token that represents it. Bitcoin represented on Ethereum (for example WBTC or cbBTC) is the

common case. Redemption runs through the custodian. The custodian, or an issuer acting for it, can freeze the representation [Unverified].

- `WRP.BRIDGE`**, lock-mint bridge.** A bridge contract or a validator set locks the original token on the source chain and mints a representation on the destination chain. The representation is a liability of the bridge, not of the original issuer. A third-party-bridged stablecoin such as USDC.e is a claim on a bridge contract that holds USDC. It is not a direct liability of the stablecoin issuer.
- `WRP.NATIVEBURN`**, issuer-native burn-mint.** The issuer burns the token on the source chain and mints the same issuer's token on the destination chain. Circle's Cross-Chain Transfer Protocol (CCTP) "natively transfers USDC across blockchains using burn-and-mint", with attestation by Circle's Iris service. Circle documents a V2 Fast Transfer of about 8–20 seconds and a Standard Transfer of 15–19 minutes on Ethereum and its layer-2 networks [32].

Two further cases need a convention. First, a same-chain contract wrap of a native gas token (ETH to WETH) crosses no chain and has no human custodian; the wrapping contract holds the underlying and redeems one-for-one. Edition 1 places it under `WRP.CUST`, treating the contract as the custodian. We record this as a schema question for Edition 2. Second, a CCTP transfer leaves the holder with *native* USDC on the destination chain. At rest, that token is a fiat stablecoin (`STB.FIAT.SINGLE`). Edition 1 uses `WRP.NATIVEBURN` only for the asset as it moves through the burn-mint event, with `event_type` = `bridge`, `crosses_chain` = `yes` and, because CCTP is the issuer's own protocol, `passes_bridge` = `no` (Section 6.2). The rules layer then treats `WRP.NATIVEBURN` as the same issuer claim on both sides.

Legal classification across jurisdictions. No jurisdiction in our set has a statute written for wrapped tokens. The US has the most explicit treatment, through the joint SEC–CFTC interpretation of March 2026. The EU names wrapped assets as an open question. The UK statutory definition may exclude some wrapped tokens altogether. Table 28 summarises.

Three points matter for Zlef. First, the US securities answer does not settle the US tax answer. The interpretation tells us that wrapping is not a securities transaction; it says nothing about whether wrapping is a realisation event for income tax. Second, the bridged-stablecoin gap is commercially large. A token that looks like USDC on a block explorer may be a claim on a bridge, with a different issuer, a different freeze authority and possibly a different regulatory status. Third, the UK "solely a record of rights" exclusion means a UK rule could treat the wrapper and the underlying as one thing for regulatory purposes even while tax law asks whether a disposal occurred.

Taxation by event. Table 29 gives the tax position for the three events that matter: wrapping or unwrapping, bridging, and a later disposal of the wrapped unit. The pattern is plain. Disposal of a wrapped unit for fiat or another asset follows the general rule in each jurisdiction. Wrapping and bridging themselves are unsettled everywhere.

The UK case shows why Zlef needs an explicit `beneficial_owner_change` label. HMRC's manual says there is no disposal where the individual keeps beneficial ownership [111]. So under current UK law the answer for a wrap depends on a fact about the relationship, not on the token. A one-for-one contract wrap that the holder can always reverse looks different from a deposit to a custodian that can refuse redemption. Zlef cannot decide the legal test. It can,

Table 28: Legal classification of wrapped and bridged representations (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
US	"Redeemable Wrapped Token": a receipt, not a security if the underlying is not a security and the token is issued one-for-one "without directly or indirectly offering any return, yield, profit opportunity". Bridges that "programmatically generate and redeem" are covered. A wrapper of a security is a security.	SEC, CFTC	[In force] (effective 23 Mar 2026)	[220]
US	Bridged stablecoin: under GENIUS only the permitted issuer's token is a payment stablecoin. Whether a third-party-bridged representation inherits that status is not addressed.	OCC, Fed, Treasury	[Unverified] (gap; GENIUS itself [Enacted, not yet effective])	[218]
EU	Crypto-asset under MiCA. A wrapped EMT may be an EMT, an ART or another crypto-asset depending on who issues and redeems it. The Commission's review names "wrapped assets" as hard to classify (Q7).	NCAs, ESMA	[Proposed] (consultation closed 31 Aug 2026)	[68]
EU	Custodial wrapping as custody by a CASP, needing authorisation.	NCAs	[Unverified]	[72]
UK	Qualifying cryptoasset, unless "solely a record of … rights in another cryptoasset" (art. 88F(2)). A wrapped token may therefore fall outside the definition; not confirmed in FCA guidance.	FCA	[Enacted, not yet effective] (full commencement 25 Oct 2027)	[217]
JP	Treated as the underlying (crypto asset) in practice; no specific guidance located.	FSA	[Unverified]	—
KR	As the underlying (virtual asset); no specific guidance located.	FSC	[Unverified]	—
IN	As the underlying (virtual digital asset).	CBDT, FIU-IND	[Unverified]	—
BR	As the underlying (ativo virtual); a bridged fiat-referenced token may sit in the FX framework.	BCB	[Unverified]	—
ID, VN, UA, RU	No classification located.	—	[Unverified]	—

however, report the facts that the test uses: who controls the underlying, whether redemption is unconditional, and whether the same holder controls both legs.

How it travels and the flags it raises. Wrapped and bridged assets move in two legs on two ledgers. Each leg is visible on its own chain. Linking the legs requires decoding the bridge's messages or using a commercial attribution graph. For CCTP, the link runs through the `MessageSent` and `MessageReceived` events, a nonce and the source and destination domain identifiers [32]. Linking is weaker where an intent or solver bridge settles the destination leg from the solver's own inventory, because no common message exists (our inference).

Control points differ by form. A custodian or bridge operator can usually freeze or refuse redemption, and a bridge's validator or guardian set controls minting [Unverified]. Lock-mint

Table 29: Tax treatment of wrapped and bridged representations by event (as of 15 Oct 2026).

Juris.	Wrap / unwrap (`wrap_unwrap`)	Bridge (`bridge`)	Later disposal of the wrapped unit
US	Unsettled. No IRS guidance. An exchange of one property for another may be a realisation; many practitioners treat one-for-one redeemable wraps as non-realisation [Unverified].	Unsettled. Burn-mint by the same issuer is arguably the same asset (our inference).	Disposal of property [Unverified]; spending is a disposal [In force] [123].
UK	Unsettled. Current law turns on whether beneficial ownership is retained [In force] [111]. The draft 2027 relief does not mention wrapping [Proposed] [105].	Unsettled; the same beneficial-ownership test (our inference).	Disposal for CGT [In force] [111].
JP	Unsettled; the NTA FAQ is silent.	Unsettled.	Disposal; gain is miscellaneous income [In force] [154].
KR	Unsettled.	Unsettled.	Disposal under the virtual-asset income tax from 1 Jan 2027, 22% on annual gains above KRW 2.5m; [Enacted, not yet effective]; a deferral bill was announced on 15 Oct 2026 [Proposed] [18].
IN	Unsettled.	Unsettled.	Transfer taxed at 30% with 1% TDS [In force] [20, 116].
BR	Unsettled.	Unsettled.	Capital gain above the BRL 35,000 monthly sales exemption [In force] [213].
ID	Unsettled.	Unsettled.	Final tax of 0.21% (domestic licensed platform) or 1% (foreign designated or self-remitted) on transaction value [In force] [57].
VN	Unsettled.	Unsettled.	0.1% PIT on gross transfer value via licensed providers (pilot) [In force] [7].
UA	Unsettled.	Unsettled.	General PIT on investment income [Unverified].
RU	Unsettled.	Unsettled.	Property treatment for tax [Unverified].

bridges concentrate risk in custodian keys or validator sets; large bridge exploits in 2022 (Ronin, Wormhole, Nomad) are the standard examples [Unverified]. Whether bridge operators are VASPs is unsettled [Unverified].

Bridges are also a laundering route. The FBI attributed the theft from Bybit of about USD 1.5bn on about 21 Feb 2025 to the DPRK group it calls "TraderTraitor" [77]. FATF puts the theft at USD 1.46bn, the largest virtual-asset theft on record, with "only 3.8%" recovered [80]. Laundering reportedly involved swaps and bridging from Ethereum assets into bitcoin [Unverified]. FATF's 2020 red-flag indicators list rapid sequential bridging (chain-hopping) as a typology [Unverified].

In Paper V's flag taxonomy, a wrapped or bridged payment can raise flags in five families: `AST` (the classification gap for wrapped and bridged stablecoins), `TAX` (disposal ambiguity), `TRV` (a travel-rule message may cover only one leg), `SAN` (screening must run on both legs), and `AML`

(chain-hopping patterns).

How Zlef labels it, and why. Zlef assigns `WRP.*` from features it can observe:

- a mapping from the token contract to a known underlying, held in a versioned registry of wrappers and bridges;
- mint events on the destination chain paired with lock or burn events on the source chain;
- issuer-native message events (for CCTP, the message and attestation pattern);
- mint and freeze roles held by a single custodian address, which points to `WRP.CUST`;
- contract code that redeems one-for-one on demand with no custodian role, which points to a contract wrap.

Token names and symbols (a "w" prefix, an ".e" suffix) are weak evidence, because anyone can choose them.

Four confusions recur. (i) `WRP.BRIDGE` against `WRP.NATIVEBURN` for the same brand of stablecoin. This error matters most, because it changes the issuer, the freeze authority and possibly the GENIUS status. (ii) `WRP.*` against `STB.*`: a native stablecoin at rest is `STB`; a third-party-bridged one is `WRP.BRIDGE`. (iii) `WRP` against `LST`: both are receipts, but a liquid staking token accrues staking rewards and a wrap does not. (iv) `WRP` against `DFI.LEND`: an exchange-rate lending receipt also redeems for the underlying, but at a rising rate.

Zlef should abstain (`UNKNOWN`) when the bridge contract is not in the registry, when a solver fill leaves no common message, or when a multi-hop route passes through more than one unverified bridge. For `beneficial_owner_change`, Zlef predicts `no` when the same key controls both legs (for example, the CCTP mint recipient equals the burning address and the address is known to be the holder's), and `unknown` when it cannot link the legs to one holder. It does not predict whether the event is a disposal. That is a rules-layer output (Section 5.6).

Illustrative label (invented example, not a model output). A holder moves USDC from Ethereum to an L2 network via CCTP, from and to the same self-hosted address, with no invoice or accounting record.

```
{
  "asset_class": "WRP.NATIVEBURN",
  "event_type": "bridge",
  "economic_purpose": "internal_transfer",
  "beneficial_owner_change": "no",
  "counterparty_type": "smart_contract_protocol",
  "evidence_completeness": 1,
  "review_priority": "low",
  "travel_surface": {
    "crosses_chain": "yes", "passes_bridge": "no",
    "issuer_freezable": "yes", "permissioned_transfer": "no",
    "privacy_enhanced": "no", "touches_mixer": "no",
    "touches_unhosted": "yes"
  }
}
```

The purpose `internal_transfer` with a `bridge` event is allowed by C2 of Section 6.10, because the owner moves its own funds. If the same transfer had used a third-party lock-mint bridge,

the label would change to `WRP.BRIDGE` with `passes_bridge = yes`, and the issuer-freeze flag would refer to the bridge operator rather than to the stablecoin issuer. The rules layer would then add an `AST`-family review flag for the bridged-stablecoin status gap.

5.2 DeFi Position Receipts (DFI)

Definition and technical form. A DeFi position receipt is a token, or a contract-recorded balance, that represents a holder's position in a protocol. Zlef separates three forms.

- `DFI.LP`**, liquidity-pool token.** The holder deposits two or more assets into an automated market maker (AMM) and receives a token for a pro-rata share of the pool. The pool trades the deposited assets against other users, so the quantities the holder can withdraw change over time. Withdrawal burns the LP token.
- `DFI.LEND`**, lending receipt.** The holder supplies an asset to a lending pool and receives a receipt. Interest accrues in one of two ways: the receipt balance grows (rebasing, as with aTokens) or the receipt's exchange rate against the underlying rises (as with cTokens). Redemption burns the receipt.
- `DFI.DEBT`**, debt position.** The holder borrows against collateral. Many protocols record the debt as a balance or a non-transferable debt token. If collateral value falls below a threshold, an oracle-driven liquidation sells the collateral [Unverified].

Legal classification across jurisdictions. Regulators have said little about these receipts directly. Table 30 records what we found.

The gap is structural. Securities and crypto regimes classify *instruments*. A position receipt is closer to a contract claim whose value depends on what the pool or protocol does. The EU tests (pooling, investment policy) and the US tests (investment contract, note) both look through the token to that activity. Zlef therefore cannot read a legal category from the receipt alone, and Edition 1 does not try.

Taxation by event. Tax law is more developed than regulatory law here, mainly because of the UK draft legislation published on 13 Jul 2026. Table 31 sets out the position by event.

Two details in the UK draft are directly observable on chain and so useful to Zlef. First, relief for AMM contributions depends on "same-type, same-quantity" treatment, and a gain arises if the quantity differs. The quantities deposited and withdrawn are visible in the pool's events. Second, the draft excludes companies and tokenised securities. Those are entity and asset facts, which the rules layer takes from layer B or F records and from `asset_class`. The draft is not law. Its status is draft legislation, published 13 Jul 2026 and last updated 30 Sep 2026 [Proposed] [105]. HMRC's consultation outcome of 26 Nov 2025 preceded it [108].

How it travels and the flags it raises. DeFi receipts rarely move through VASPs. They are minted and burned by contracts, usually from self-hosted wallets. The control points are protocol governance (pause, asset freeze), oracle-driven liquidation and front-end geo-blocking; some front-ends also screen connected addresses against OFAC lists [Unverified]. Pools are fungible, so funds do not become "tainted" in a legal sense when a sanctioned address has used the same pool, although risk vendors may score that exposure [Unverified].

Table 30: Legal classification of DeFi position receipts (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
US	No direct treatment in the 2026 interpretation. Pooled-return receipts may be investment contracts; interest-paying lending receipts resemble notes (our analysis).	SEC	[Unverified]	[220]
US	CLARITY Act DeFi provisions: Senate cloture failed 49–50 on 15 Sep 2026.	—	[Proposed] (stalled)	[58]
EU	Crypto-asset, or a financial instrument (for example a collective investment unit if pooled under an investment policy) under ESMA's qualification guidelines, applying from 18 May 2025.	NCAs, ESMA	[In force]	[183]
EU	DeFi services outside MiCA only if "fully decentralised"; lending and borrowing not regulated; both under review (Q61–67).	EC, ESMA	[Proposed] (review)	[68]
UK	Lending and borrowing as regulated activities; DeFi caught where there is an "identifiable controlling entity".	FCA	[Proposed] (CP25/40)	[198]
UK	An LP token may be "solely a record of rights in another cryptoasset" and so outside the definition (our reading).	FCA	[Unverified]	[217]
IN	Virtual digital asset (broad tax definition).	CBDT	[Unverified]	—
JP, KR, BR, ID, VN, UA, RU	No classification located.	—	[Unverified]	—

The US reporting gap matters. A 2024 rule that would have made DeFi front-ends brokers was nullified under the Congressional Review Act on 10 Apr 2025 [In force] [114]. DeFi trades therefore carry no third-party Form 1099-DA in the US, and the taxpayer must build the record. FATF approved a targeted report on DeFi regulatory challenges at its June 2026 plenary, published in July 2026 [81].

Flags in Paper V's families are mainly `TAX` (disposal ambiguity on deposit and withdrawal; character of returns), `REP` (no broker reporting in the US; DeCripto self-reporting in Brazil) and `DAT` (allocation of pool returns to tax categories needs records the chain does not hold).

How Zlef labels it, and why. Observable features are strong for this class. An LP deposit shows two or more assets moving into a pool contract and an LP token minted to the depositor in the same transaction. A lending deposit shows a supply event and a receipt minted. A borrow shows a debt balance or debt token minted and the borrowed asset sent out. Liquidations emit their own events. The protocol's interest index or exchange rate is readable from contract state.

The confusions are among receipts. (i) `DFI.LEND` against `STB.YLD`: a savings wrapper over a stablecoin and a lending receipt for a stablecoin can be technically similar; the difference is whether the token is issued as a stablecoin product or as a pool share. (ii) `DFI.LP` against `TKN.MMF` or other fund units: both are pro-rata shares, but one is a regulated fund. (iii) `DFI.LEND` against `LST`: both are receipts that accrue. (iv) Vault tokens that route deposits through several strategies may fit `DFI.LEND`, `STB.YLD` or `LST`. Zlef should abstain on the class when the vault's

Table 31: Tax treatment of DeFi positions by event (as of 15 Oct 2026).

Juris.	Deposit (`lp_deposit`, `lend_out`)	Returns (fees, interest)	Withdrawal (`lp_withdraw`, `lend_return`)
US	Unsettled. No guidance; receiving an LP token for two assets is argued to be a disposal on general principles [Unverified]. Securities-lending non-recognition (s.1058) covers securities only [Unverified].	Unsettled.	Unsettled.
UK	Today: a disposal where beneficial ownership passes [Unverified]. From 6 Apr 2027: no gain, no loss (NGNL) for lending and for same-type, same-quantity AMM contributions; individuals and trustees only; tokenised securities excluded [Proposed] [105, 191].	Returns remain taxable [Proposed] [191]. Eligible-stablecoin lending returns taxed like savings interest [Proposed] [106].	NGNL from 6 Apr 2027; a gain arises if the quantity returned differs [Proposed] [105].
JP	Unsettled.	Lending gains are income [In force] [154].	Unsettled.
KR	Unsettled.	Unsettled.	Unsettled.
IN	Unsettled.	Income from other sources [In force] [20].	Unsettled.
BR	Unsettled; crypto-backed loans are reportable under DeCripto [In force] [143].	Unsettled.	Unsettled.
ID	Unsettled.	Unsettled.	Unsettled.
VN	Unsettled.	Unsettled.	Unsettled.
UA	Unsettled.	Unsettled.	Unsettled.
RU	Unsettled.	Unsettled.	Unsettled.

strategy is not in the protocol registry, and on the event when a single transaction loops deposits and borrows (leverage loops).

For protocol deposits, `beneficial_owner_change` follows the per-leg rule of Section 6.5. A leg that moves an asset into a pooled protocol contract, or to a custodian, in exchange for a receipt or claim is `yes`; a leg between addresses that records show the same owner controls, including the owner's own smart-contract wallet, is `no`; otherwise `unknown`. This is an economic fact, not a legal conclusion, and it is not the same as the UK legal test. Whether annotators can apply it consistently is an open question [Research hypothesis].

Illustrative label (invented example, not a model output). A self-hosted wallet adds two assets to an AMM pool and receives an LP token.

```
{
  "asset_class": "DFI.LP",
  "event_type": "lp_deposit",
  "economic_purpose": "investment",
  "beneficial_owner_change": "yes",
  "counterparty_type": "smart_contract_protocol",
  "evidence_completeness": 1,
  "review_priority": "medium",
```

```
  "travel_surface": {
    "crosses_chain": "no", "passes_bridge": "no",
    "issuer_freezable": "no", "permissioned_transfer": "no",
    "privacy_enhanced": "no", "touches_mixer": "no",
    "touches_unhosted": "yes"
  }
}
```

For a UK individual, the rules layer would apply current law before 6 Apr 2027 and, if the draft is enacted unchanged, NGNL for contributions on or after that date. It would mark the second branch [Proposed] until enactment. Zlef's label is the same in both cases. Only the rule version changes.

5.3 Privacy Coins (PRV)

Definition and technical form. A privacy coin is a native asset whose protocol hides some or all of the sender, receiver and amount. Two designs dominate. Monero (XMR) hides all three by default, using ring signatures, stealth addresses and confidential amounts (RingCT). Zcash (ZEC) has a transparent pool, which behaves like a public chain, and shielded pools that use zero-knowledge proofs (zk-SNARKs) to hide parties and amounts [Unverified]. The asset class `PRV` describes the *coin*. It is distinct from the travel label `privacy_enhanced`, which describes a *movement*. Ether sent through a shielded pool on Ethereum is a native token (`NAT`) with `privacy_enhanced = yes`; it is not `PRV`. A transparent-to-transparent ZEC transfer is `PRV` with `privacy_enhanced = no`.

Legal classification across jurisdictions. No jurisdiction in our set creates a separate legal category for privacy coins. Restrictions work on *intermediaries*: who may list, hold or transfer them. The EU is the only jurisdiction in which we found an enacted, dated rule (Table 32).

The AMLR rule is the most important dated change for this class. From 10 Jul 2027, an EU CASP that receives XMR from a customer faces a prohibition, not a risk score. That turns a Paper V `AML`-family review flag into a candidate `block` for EU-regulated parties from that date. The exact scope depends on the article text and on implementing guidance, neither of which we have read.

Taxation by event. Privacy coins are taxed like other crypto assets in every jurisdiction where we found a rule. The asset type changes the *evidence*, not the tax character. Table 33 applies the general rules.

The tax consequence of privacy is indirect but real. Basis, holding period and provenance must be shown from records the chain does not expose. In the US, where basis must be tracked per wallet from 1 Jan 2025 [121], and where the IRS FAQ says that without documentation the basis of a gifted asset is zero [123], missing evidence can change the tax computed even when the tax character is clear.

How it travels and the flags it raises. Privacy coins have no issuer and no protocol-level control point. They can be stopped only at the edges: an exchange deposit or withdrawal. Shielded transfers are practically incompatible with the travel rule, because a VASP cannot see its counterparty on chain [Unverified]. Many exchanges have delisted or restricted XMR and similar

Table 32: Legal classification and restrictions for privacy coins (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
EU	Crypto-asset under MiCA (Title II).	NCAs	[In force]	[72]
EU	AMLR Art. 79: credit institutions, financial institutions and CASPs may not keep accounts that allow anonymisation, "including through anonymity-enhancing coins". Applies from 10 Jul 2027. Wording taken from a secondary source; the article text was not read directly. Whether ZEC with optional shielding is caught is open.	AMLA (from 2027), NCAs	[Enacted, not yet effective]	[36, 73]
US	No asset-level ban; delistings are exchange-driven.	—	[Unverified]	—
UK	No asset-level ban located; travel-rule friction drives delistings.	FCA	[Unverified]	—
JP	De facto barred from registered exchanges since 2018 (self-regulatory practice).	FSA, JVCEA	[Unverified]	—
KR	Exchanges barred from handling "dark coins" under the 2021 enforcement decree.	FSC, KoFIU	[Unverified]	—
IN	Virtual digital asset; FIU-IND guidelines (updated 8 Jan 2026) treat anonymity-enhanced coins as "unacceptably high-risk".	FIU-IND	[In force]	[162]
BR, ID, VN, UA, RU	No specific rule located.	—	[Unverified]	—

Table 33: Tax treatment of privacy coins by event (general crypto rules; as of 15 Oct 2026).

Juris.	Disposal for fiat or swap	Spending for goods or services	Evidence note
US	Disposal of property [Unverified].	Disposal [In force] [123].	Per-wallet basis applies from 1 Jan 2025 [In force] [121]; shielded flows make lots hard to prove.
UK	Disposal; pooled per token type [In force] [111].	Disposal [In force] [111].	Pool records must come from off-chain sources.
JP	Disposal; miscellaneous income [In force] [154].	Disposal [In force] [154].	Not listed by registered exchanges [Unverified].
KR	Disposal from 1 Jan 2027 [Enacted, not yet effective] [18].	Disposal from 1 Jan 2027 [Unverified].	Domestic listing restricted [Unverified].
IN	30% on transfer; 1% TDS [In force] [116].	Disposal; TDS by the person acquiring the VDA (our reading of the TDS rule).	No loss set-off [In force] [20].
BR	Capital gain above BRL 35,000 monthly sales [In force] [213].	Disposal [Unverified].	DeCripto self-reporting for foreign or decentralised platforms [In force] [143].
ID	Final tax 0.21% via domestic platform; 1% self-remitted off platform [In force] [57].	Seller self-remits 1% off platform [In force] [57].	Platform listing depends on the OJK list.
VN	0.1% via licensed providers (pilot) [In force] [7].	Off-platform treatment unsettled.	No licensed operator as of 30 Aug 2026.
UA	General PIT [Unverified].	Unsettled.	—
RU	Property treatment [Unverified].	Domestic crypto payments banned [In force] [11].	—

coins (for example in Korea in 2021 and on one large global exchange in several jurisdictions in 2024) [Unverified].

What the chain shows depends on the design. For XMR, almost nothing beyond the existence of a transaction. For ZEC, transparent-pool transfers show addresses and amounts, and shielded transfers do not. A movement into or out of a shielded pool is itself visible, which gives a weak signal even when the shielded leg is not.

Flags fall in the `AML` family (privacy coins are named there in Paper V), `AST` (the EU prohibition from 10 Jul 2027), `TRV` (originator and beneficiary data cannot be checked against the chain) and `DAT` (basis and provenance need off-chain attestation).

How Zlef labels it, and why. The asset class itself is easy: the chain or the contract identifies it. A wrapped privacy coin on another chain is `WRP.*` with an underlying of `PRV`; the registry records the link. The hard labels are the others. For a shielded transfer, Zlef cannot see the counterparty, the purpose or whether ownership changed. Its correct output is `UNKNOWN` or `unknown` on those fields, with low `evidence_completeness`, and a raised `review_priority`. Forcing a guess here would produce confident errors exactly where a reviewer most needs honest uncertainty.

One schema issue follows. For a shielded transfer, `touches_unhosted` cannot be observed. The draft specification defined the `travel_surface` labels as yes/no; Edition 1 adopts `unknown` as a third annotation value for each travel label, so that "not observed" is not recorded as "no" (Section 6.9). The example below uses it.

Illustrative label (invented example, not a model output). A VASP receives an XMR deposit to a customer account from an unknown source.

```
{
  "asset_class": "PRV",
  "event_type": "UNKNOWN",
  "economic_purpose": "UNKNOWN",
  "beneficial_owner_change": "unknown",
  "counterparty_type": "UNKNOWN",
  "evidence_completeness": 1,
  "review_priority": "high",
  "travel_surface": {
    "crosses_chain": "no", "passes_bridge": "no",
    "issuer_freezable": "no", "permissioned_transfer": "no",
    "privacy_enhanced": "yes", "touches_mixer": "no",
    "touches_unhosted": "unknown"
  }
}
```

Here `evidence_completeness = 1` means only that on-chain data exists; for XMR that data says very little. The rules layer, not Zlef, decides whether an EU CASP may accept the deposit after 10 Jul 2027.

5.4 CBDCs and E-Money Outside MiCA Scope (CBDC)

Definition and technical form. The `CBDC` class covers two kinds of fiat-denominated digital money that most crypto regimes exclude by definition:

- **central bank digital currencies**: a direct liability of a central bank in digital form, such as India's e-Rupee or Russia's digital ruble;
- **e-money products recorded on a ledger that sit outside the MiCA crypto-asset perimeter**, which the OECD's amended Common Reporting Standard calls Specified Electronic Money Products.

Three neighbours must be kept apart. An EU e-money token (EMT) under MiCA is a private issuer's crypto-asset and belongs in `STB.FIAT.SINGLE`. A tokenised commercial-bank deposit belongs in `TKN.DEP`. A privately issued token that references a national currency, such as the ruble-pegged A7A5, is a stablecoin (`STB`), not a CBDC, whatever its marketing says. CBDC ledgers are usually permissioned and run by the central bank or its intermediaries. Holders typically reach them through a bank or payment provider, not through a public block explorer.

Legal classification across jurisdictions. The consistent pattern is exclusion: crypto statutes define CBDCs out of scope (Table 34).

Table 34: Legal classification of CBDCs and out-of-scope e-money (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
UK	Excluded from "qualifying cryptoasset": "currency of the United Kingdom or any other country … including a central bank digital currency"; electronic money also excluded (art. 88F(4)).	BoE, FCA	[Enacted, not yet effective] (full commencement 25 Oct 2027)	[217]
KR	CBDC and CBDC deposit tokens excluded from "virtual asset".	FSC, BoK	[In force] (since 19 Jul 2024)	[86]
EU	CBDCs outside MiCA; EMTs deemed e-money; e-money otherwise under the E-Money Directive. Status of digital euro legislation not confirmed.	ECB, NCAs	[Unverified]	[72]
EU	EU support for the digital rouble prohibited (20th sanctions package, 23 Apr 2026).	Council	[In force]	[43]
IN	e-Rupee is RBI-issued legal tender and is promoted by the RBI; not a virtual digital asset.	RBI	[Unverified]	[50]
RU	Digital ruble under a separate law.	Bank of Russia	[Unverified]	—
US	No US CBDC.	—	[Unverified]	—
JP	Prepaid or electronic payment instrument rules may apply to e-money-like tokens.	FSA	[Unverified]	—
BR, ID, VN, UA	No classification located.	—	[Unverified]	—

Taxation by event. We found no crypto tax rule in the ten jurisdictions that addresses CBDCs. That is consistent with the exclusions above: a CBDC is a form of the currency itself. Our inference is that general currency rules are the starting point. A domestic CBDC used domestically would raise no crypto disposal. A foreign CBDC might raise foreign-currency questions. Neither inference is sourced, so the table records the position as unsettled.

What *is* sourced is the reporting split. The OECD's 2023 package amended the CRS to bring "certain electronic money products and central bank digital currencies in scope" [In force] (as

an approved OECD standard; domestic adoption varies) [167]. The Crypto-Asset Reporting Framework (CARF) excludes CBDCs and Specified Electronic Money Products from "Relevant Crypto-Assets" [Unverified]. In the EU, DAC8 covers stablecoins and EMTs [In force] [69]; its treatment of CBDCs through e-money reporting is [Unverified]. The practical point for Zlef is that a CBDC holding is reported, if at all, as a financial account under CRS, not as a crypto-asset under CARF.

Table 35: Tax and reporting treatment of CBDCs by event (as of 15 Oct 2026).

Juris.	Conversion to or from fiat	Payment for goods or services	Reporting regime
US	Unsettled (no US CBDC).	Unsettled.	No CRS participation located.
UK	Unsettled; outside crypto CGT guidance.	Unsettled.	CRS [Unverified].
JP	Unsettled.	Unsettled.	CRS [Unverified].
KR	Unsettled; excluded from "virtual asset", so outside the 2027 virtual-asset tax on our reading.	Unsettled.	CRS [Unverified].
IN	Unsettled; e-Rupee is legal tender [Unverified].	Unsettled.	CRS [Unverified].
BR	Unsettled.	Unsettled.	CRS [Unverified].
ID	Unsettled.	Unsettled.	CRS [Unverified].
VN	Unsettled.	Unsettled.	Not CARF-committed.
UA	Unsettled.	Unsettled.	CRS participant [Unverified].
RU	Unsettled; digital ruble [Unverified].	Unsettled.	Most Western partners suspended exchange [Unverified].

How it travels and the flags it raises. A CBDC moves on a ledger the issuer controls. Freezing and permissioning are design features, so `issuer_freezable` and `permissioned_transfer` are normally both `yes`. Whether FATF's virtual-asset definition, which excludes digital representations of fiat currencies, keeps CBDCs out of the VASP travel rule is [Unverified] in our sources. Identity is held by the intermediary that onboarded the user, and Zlef sees it only as a layer B record.

Sanctions create the sharpest flags. The EU prohibited support for the digital rouble in its 20th package [In force] [43]. The 19th package banned transactions in the A7A5 ruble token across the EU [In force] [42], and the 20th banned transactions in RUBx [In force] [43]. A model that labels a ruble-referenced private token as `CBDC`, or a digital ruble transfer as an ordinary stablecoin, would send the rules layer to the wrong sanctions rule. Flags fall in `SAN`, `REP` (CRS rather than CARF) and `FXC` (cross-border use of a foreign CBDC).

How Zlef labels it, and why. The decisive feature is issuer identity: a central bank or a regulated e-money issuer whose product is outside MiCA. On permissioned ledgers this comes from the network operator's records (layer B), not from public chain data. Secondary features are the ledger itself (a known CBDC network) and the absence of a smart-contract token standard on a public chain.

Confusions: `CBDC` against `STB.FIAT.SINGLE` (a private EMT); `CBDC` against `TKN.DEP` (a commercial-bank liability); and `CBDC` against a currency-referenced private token such as A7A5. Zlef should abstain when the issuer cannot be confirmed, for example a pilot token on a shared

ledger where commercial-bank and central-bank tokens coexist.
Illustrative label (invented example, not a model output). A business receives a retail CBDC payment for an invoiced service, through its bank, which supplies the invoice reference.


```
{
  "asset_class": "CBDC",
  "event_type": "receive_payment_goods_services",
  "economic_purpose": "service_purchase",
  "beneficial_owner_change": "yes",
  "counterparty_type": "individual",
  "evidence_completeness": 3,
  "review_priority": "low",
  "travel_surface": {
    "crosses_chain": "no", "passes_bridge": "no",
    "issuer_freezable": "yes", "permissioned_transfer": "yes",
    "privacy_enhanced": "no", "touches_mixer": "no",
    "touches_unhosted": "no"
  }
}
```


5.5 Other Claims: Points, In-Game Items, Unvested and Airdrop Claims (OTH)

Definition and technical form. `OTH` collects claims that are not yet, or not fully, crypto-assets in the sense of the other classes:

- **points**: balances that a protocol or platform records, usually off chain and non-transferable, often in anticipation of a later token distribution;
- **in-game tokens and items**: units usable inside a game, which may be closed-loop (redeemable only in the game) or freely transferable;
- **unvested claims**: token allocations held in a vesting or lock-up contract, released on a schedule;
- **airdrop claims**: the right to claim tokens, often recorded as an entry in a claim contract (for example a Merkle-proof list) before the holder calls the claim function.

Once claimed, vested or tokenised and transferable, the unit usually becomes another class (`GOV`, `MEME`, `NFT`). `OTH` labels the claim before that point, and the event that turns it into something else.

Legal classification across jurisdictions. Table 36 shows that transferability is the hinge. Where a unit cannot be transferred, the crypto definitions in the EU and UK do not reach it on our reading. Where it can, it is classified like the underlying.

Taxation by event. Airdrops are the best-documented event in this class. Points and in-game units are largely unaddressed. Vesting is treated as compensation where the recipient is an employee or contractor (Table 37).

The timing question is the same as for staking (Section 5.6). For an airdrop claim, the claim right exists before the claim transaction. The US test (dominion and control) and the UK test

Table 36: Legal classification of points, in-game items, unvested and airdrop claims (as of 15 Oct 2026).

Juris.	Category	Regulator	Status	Source
US	Airdrops of non-security crypto for no or nominal consideration involve no "investment of money" and are not securities transactions.	SEC, CFTC	[In force]	[220]
US	"Digital collectibles" include in-game items; fan tokens named as collectibles.	SEC, CFTC	[In force]	[220]
US	Unvested or locked tokens: same category as the underlying; insider holdings relevant to when an investment contract ends.	SEC	[In force]	[220]
US	Points programmes that require deposits may involve "consideration".	SEC	[Unverified]	—
EU	Non-transferable points fail "able to be transferred" (Art. 3(1)(5)), so are not crypto-assets (our reading); transferable in-game tokens are Title II crypto-assets.	NCAs	[Unverified]	[72]
UK	Closed-loop exclusion (art. 88F(4)): not transferable except by redemption within "a limited network of service providers". Non-transferable points are not qualifying cryptoassets (our reading).	FCA	[Enacted, not yet effective] (full commencement 25 Oct 2027)	[217]
KR	In-game currency excluded from "virtual asset".	FSC	[In force]	[86]
JP	In-game units may be prepaid payment instruments if not exchangeable with unspecified persons.	FSA	[Unverified]	—
IN, BR, ID, VN, UA, RU	No classification located.	—	[Unverified]	—

(whether the airdrop was for a service) both need facts the claim contract does not hold: whether the holder could claim, and why the allocation was made.

How it travels and the flags it raises. Points usually do not travel on chain at all. In-game units travel inside the game's ledger or, if tokenised, as ordinary tokens. Unvested claims are held in a vesting contract and released on a schedule; the release is visible. Airdrop claims are visible when claimed, and the claim contract may be paused or may exclude addresses (for example on sanctions screening) [Unverified]. Flags are mainly `TAX` (income timing and character), `DAT` (the reason for the allocation is off chain) and `AML` (sybil farming of airdrops by address clusters, which is a fraud pattern rather than a legal category).

How Zlef labels it, and why. Features: transfers out of a known vesting contract; claim-function calls against a known distribution contract; mints from a game contract; for points, nothing on chain, only a layer B statement from the platform. Confusions: (i) `OTH` against `GOV` when a claimed governance token arrives (the claim event is `airdrop`; the token at rest is `GOV`); (ii) `OTH` against `NFT` for in-game items issued as non-fungible tokens (they are `NFT`; the US treats them as collectibles); (iii) points earned for deposits against `DFI` receipts. Zlef should abstain on `economic_purpose` when a release could be either `payroll` or `distribution` and no contract or

Table 37: Tax treatment of other claims by event (as of 15 Oct 2026).

Juris.	Airdrop received (`airdrop`)	Vesting or release to a worker (`compensation`)	Points or in-game units
US	Income when the taxpayer gains dominion and control [In force] [119].	Compensation (property transferred for services) [Unverified].	Unsettled.
UK	Income if received "in return for, or in expectation of, a service"; otherwise no income and CGT on later disposal [In force] [110].	Employment income; PAYE and NIC if a readily convertible asset [Unverified].	Unsettled.
JP	Unsettled (not addressed in the FAQ sections reviewed).	Salary income at market value [Unverified].	Unsettled.
KR	Unsettled.	Earned income [Unverified].	Unsettled.
IN	Income if fair market value exceeds INR 50,000 from non-relatives [In force] [20].	Salary at slab rates [Unverified].	Unsettled.
BR	Reportable under DeCripto; character unsettled [In force] [143].	Salary [Unverified].	Unsettled.
ID	Unsettled.	Salary [Unverified].	Unsettled.
VN	Unsettled.	Unsettled.	Unsettled.
UA	Unsettled.	Unsettled.	Unsettled.
RU	Unsettled.	Unsettled.	Unsettled.

employment record is supplied.

Illustrative label (invented example, not a model output). A protocol credits non-transferable points to an address that had deposited into it, recorded as an entry in its claim contract; no documentation is supplied.

```
{
  "asset_class": "OTH",
  "event_type": "airdrop",
  "economic_purpose": "UNKNOWN",
  "beneficial_owner_change": "yes",
  "counterparty_type": "smart_contract_protocol",
  "evidence_completeness": 1,
  "review_priority": "medium",
  "travel_surface": {
    "crosses_chain": "no", "passes_bridge": "no",
    "issuer_freezable": "no", "permissioned_transfer": "no",
    "privacy_enhanced": "no", "touches_mixer": "no",
    "touches_unhosted": "yes"
  }
}
```

`economic_purpose` is `UNKNOWN` because no document says why the points were allocated. Whether this is income on receipt is a rules-layer question that differs by jurisdiction. In the UK it turns on whether the distribution was for a service; the deposit history is relevant evidence, which is why `review_priority` is not `low`.

5.6 The Event Dimension

Across all twelve top-level asset classes, one finding holds. *The event decides the tax outcome, and the asset type changes the event.* In every jurisdiction we studied, the first question is whether a disposal or an income receipt occurred. The asset type matters mainly where it changes that answer: stablecoins (a UK exemption in draft), wrapped and DeFi positions (whether a disposal occurred at all), tokenised securities (securities rules), and NFTs (special US rates and reporting thresholds). This subsection therefore turns the axis. It takes each Zlef `event_type` and asks how each jurisdiction characterises it.

How to read the consolidated table. Table 38 gives one row per `event_type`, from the payer's or holder's perspective, across the ten jurisdictions and an EU reference column. Because the EU does not harmonise direct taxes, the EU column uses Germany as the member-state reference for income tax and notes EU-level VAT where relevant. Each cell holds a code for the tax *character*, and a superscript for the *status* of the source.

- **Character codes.** **D** disposal or realisation (gain or loss); **I** income on receipt at market value; **N** not a taxable event; **FT** final transaction tax on gross value; **X** the activity is prohibited; **U** unsettled (no authoritative guidance located). A slash gives two outcomes that depend on facts (for example **I/N**).
- **Status superscripts.** No superscript: [In force], verified against a primary or labelled secondary source on 15 Oct 2026. ^{u}: [Unverified]. ^{p}: [Proposed] (draft legislation, bill or consultation). ^{e}: [Enacted, not yet effective] (Korea's virtual-asset income tax from 1 Jan 2027; a deferral bill announced on 15 Oct 2026 is [Proposed]). ^{c}: [Sources conflict].

The table records character only. It does not record rates, thresholds, holding periods or reporting duties; those sit in the rules layer and in Paper V. Sources are given in the notes column and in the per-class tables of Part 1 and the preceding section.

Table 38: Tax character of each Zlef `event_type` by jurisdiction (as of 15 Oct 2026). Codes and superscripts are defined in the text. EU column: Germany for income tax; EU VAT noted.

`event_type`	US	UK	JP	KR	IN	BR	ID	VN	UA	RU	EU/DE	Notes and sources
`acquire_fiat`	N^{u}	N	N	N^{u}	N	N	N	N^{u}	U	U	N^{u}	IN: a buyer from a non-exchange seller deducts 1% TDS [116]. UK: cost enters the s.104 pool [111]. BR: reportable [143].
`dispose_fiat`	D^{u}	D	D	D^{e}	D	D	FT	FT	D^{u}	D^{u}	D^{u}	IN 30% plus 1% TDS [20]; ID 0.21%/1% [57]; VN 0.1% [7]; KR from 1 Jan 2027 [18]; DE exempt after one-year holding^{u}.
`swap_crypto`	D^{u}	D	D	D^{u}	D^{u}	D^{u}	FT	FT	U	U	D^{u}	UK lists swapping as a disposal [111]; JP FAQ 1-3 treats A-for-B as buying B with A [154].
`swap_stable`	D	D; N^{p}	U^{u}	D^{u}	D^{u}	D	FT	FT	U	U	D^{u}	US 1099-DA de minimis USD 10,000 for qualifying stablecoins [122]. UK: eligible stablecoins CGT-exempt from Apr 2027 (draft) [106]. JP: many fiat stablecoins are electronic payment instruments. BR: FX operation (Res. 521) [8].
`pay_goods_services`	D	D	D	D^{u}	D	D^{u}	FT	U	U	X	D^{u}	IRS FAQ 14 [123]; JP FAQ 1-2 [154]. IN: TDS by the person acquiring the VDA. ID: off-platform seller self-remits 1%. RU: domestic crypto payments banned [11]. EU VAT: see text.
`receive_payment_goods_services`	I^{u}	I^{u}	I^{u}	I^{u}	I^{u}	U	I^{u}	U	U	X	I^{u}	Payee income at market value; basis equals that value^{u}. ID: no VAT on the crypto transfer itself [115].
`compensation`	I^{u}	I^{u}	I^{u}	I^{u}	I^{u}	I^{u}	I^{u}	U	U	U	I^{u}	Wages at market value; withholding and payroll taxes follow general rules^{u}.
`staking_reward`	I	I	I	U	I	U	U	U	U	U	I^{u}	US Rev. Rul. 2023-14 [120]; UK CRYPTO21200 [109]; JP FAQ 1-7 [154]; IN other sources [20]; BR reportable, character unsettled [143].
`mining_reward`	I^{u}	I^{u}	I	U	I	I^{c}	U	U	U	U^{u}	I^{u}	JP FAQ 1-7. BR: reportable; secondary sources conflict on character [213]. ID: 12% VAT on 11/12 of the base for mining services [115].

Continued on next page

Continued from previous page

event_type	US	UK	JP	KR	IN	BR	ID	VN	UA	RU	EU/DE	Notes and sources
airdrop	I	I/N	U	U	I	U	U	U	U	U	N[u]	US Rev. Rul. 2019-24 [119]; UK service test [110]; IN above INR 50,000 from non-relatives.
fork	N/I	N[u]	N	U	U	U	U	U	U	U	N[u]	US: income only if new units are received [119]; JP FAQ 1-6: no income at the fork, basis zero [154].
wrap_unwrap	U	U	U	U	U	U	U	U	U	U	U	UK beneficial-ownership test applies by analogy; draft 2027 relief silent [105].
bridge	U	U	U	U	U	U	U	U	U	U	U	Issuer-native burn-mint is the strongest "same asset" case (our inference).
lp_deposit	U	D[u]; N[p]	U	U	U	U	U	U	U	U	U	UK NGNL from 6 Apr 2027 for same-type, same-quantity contributions (draft) [105].
lp_withdraw	U	D[u]; N[p]	U	U	U	U	U	U	U	U	U	UK draft: gain if the quantity differs. JP: returns are income [154].
lend_out	U	D[u]; N[p]	U	U	U	U	U	U	U	U	U	UK draft NGNL; individuals and trustees only [191]. DE: lending income[u].
lend_return	U	D[u]; N[p]	U	U	U	U	U	U	U	U	U	Returns: JP income; IN other sources; UK taxable (draft).
borrow	U	U	U	U	U	U	U	U	U	U	U	BR: crypto-backed loans reportable [143].
repay	U	U	U	U	U	U	U	U	U	U	U	No source located for any jurisdiction.
gift (donor)	N[u]	D	D	U	N[u]	U[u]	U	U	U	U	U[u]	UK: spouse or civil partner at no gain, no loss [111]. JP: donor includes market value in income (FAQ 2-10). US, BR, DE: gift-tax regimes[u]. IN recipient taxed above INR 50,000.
internal_transfer	N	N	N[u]	U	N[u]	N	U	U	U	U	N[u]	IRS FAQ 38 [123]; per-wallet basis [121]; BR reports wallet-to-wallet transfers [143].
refund_reversal	U	U	U	U	U	U	U	U	U	U	U	US general rescission doctrine may apply within one tax year[u].
fee	D[u]	D[u]	U	U	U	U	U	U	U	U	U	US H.R. 9178 would exclude gains on fees up to USD 10 (bill) [38]. IN: only cost of acquisition deductible [20].
loss_theft	U[u]	N[u]	U	U	U	U	U	U	U	U	U	JP: miscellaneous-income losses cannot offset other income (FAQ 2-11). IN: no set-off. UK: theft not a disposal; negligible-value claim[u].

Continued on next page

Continued from previous page

event_type	US	UK	JP	KR	IN	BR	ID	VN	UA	RU	EU/DE	Notes and sources
UNKNOWN	—	—	—	—	—	—	—	—	—	—	—	No character assigned; the rules layer routes to human review.

What the table shows. Four patterns stand out.

1. **Agreement at the core.** Disposals for fiat, swaps and spending are disposals in every jurisdiction with an income-based regime, and receipts from staking, mining and work are income where a rule exists. A model that only had to label these cases would not need jurisdiction-specific outputs.
2. **Gross-value regimes cut across the core.** Indonesia and Vietnam tax the *transaction value* on licensed platforms, not the gain. For those two, the D/I distinction matters less than whether the transaction ran through a licensed or designated platform. That is a counterparty fact (`counterparty_type`), not an asset fact.
3. **The DeFi and representation-change rows are almost entirely U.** Wrapping, bridging, LP and lending events have no authoritative answer in nine of ten jurisdictions. The UK is the exception, and only through draft legislation effective from 6 Apr 2027.
4. **Ukraine and Russia are almost entirely unverified.** The parliament's own page records first-reading passage of Ukraine's virtual-asset Bill No. 10225-d on 3 Sep 2025 [Proposed] [224] (also referred to in some trackers as 3637; unverified). Its second reading was not confirmed. A rule set built on either bill would encode a proposal as law.

We now take the hard cases in turn.

Wrapping and bridging: disposal or not? The question is whether exchanging token A for a token that represents a claim to A is a realisation of A. Three sources bear on it, and none answers it.

- **US.** There is no IRS guidance. A general realisation rule applies to exchanges of property, and many practitioners treat one-for-one, contract-redeemable wraps as non-realisation [Unverified]. Separately, the March 2026 SEC–CFTC interpretation treats wrapping into a one-for-one redeemable token, including via a bridge that "programmatically generates and redeems" it, as a non-securities transaction [In force] [220]. That is a securities ruling. It does not bind the IRS.
- **UK.** HMRC's manual frames the test as beneficial ownership: there is no disposal where the individual retains it [In force] [111]. The draft DeFi legislation for April 2027 covers lending and liquidity pools but does not mention wrapping [Proposed] [105].
- **Elsewhere.** Japan's NTA FAQ is silent. We found no guidance in Korea, India, Brazil, Indonesia, Vietnam, Ukraine or Russia.

Our analysis distinguishes the three `WRP` forms. Issuer-native burn-mint (`WRP.NATIVEBURN`) is the strongest case for "same asset, no disposal", because the issuer and the claim are identical before and after. A contract wrap with unconditional one-for-one redemption is the next strongest. A lock-mint bridge or custodial wrap gives the holder a claim on a new obligor; this is the weakest case. This ranking is our inference, not law. Zlef encodes the form in `asset_class`, the event in `event_type`, and the ownership fact in `beneficial_owner_change`. The rules layer maps these to "ambiguous; human review" in every jurisdiction until guidance exists.

LP deposits. An LP deposit exchanges two or more assets for one pool token. In the US, there is no guidance; on general principles the receipt of a new token for two assets looks like a disposal [Unverified]. Under current UK law, the deposit is a disposal where beneficial ownership

passes [Unverified]. The UK draft would change that for individuals and trustees from 6 Apr 2027. Contributions to and withdrawals from automated market makers of the same type and quantity of tokens would be treated on a no-gain, no-loss basis, and "gains arise if quantity differs" [Proposed] [105]. Companies are outside the draft; tokenised securities are excluded; returns are still taxed; historic transactions are not reopened [Proposed] [191].

For Zlef, the useful fact is that the UK draft test is quantitative and the quantities are on chain. An AMM's impermanent loss shows up exactly as a different quantity mix on withdrawal. Zlef labels the deposit `lp_deposit` and the withdrawal `lp_withdraw`; the ledger records the quantities; the rules layer compares them. No model inference about value is needed for that comparison.

Lending. Lending raises the same disposal question as LP deposits, plus a question about the character of returns.

- **US.** The non-recognition rule for securities lending covers securities only, so it does not obviously reach crypto loans [Unverified]. There is no other guidance.
- **UK.** Under the draft, lending by individuals and trustees "will be treated as being on a NGNL (no gain, no loss) basis" from 6 Apr 2027 [Proposed] [105]. Returns remain taxable. For eligible stablecoins, the separate stablecoin outcome treats them as a money debt for Corporation Tax, so lending them is lending money under the loan-relationship rules; individuals' stablecoin lending returns are taxed like savings interest; a withholding tax on stablecoin lending is still under consideration [Proposed] [106].
- **Japan.** Lending gains are income [In force] [154].
- **Brazil.** Crypto-backed loans are reportable under DeCripto [In force] [143]; their tax character is unsettled.

Zlef separates `lend_out`, `lend_return`, `borrow` and `repay` because the regimes may treat the two sides of a loan differently. The UK draft speaks of lending; we found no source on the borrower side in any jurisdiction.

Staking rewards: timing. Every jurisdiction with a staking rule treats rewards as income. The hard part is *when*.

- **US.** Rev. Rul. 2023-14 holds that rewards are "included in the taxpayer's gross income in the taxable year in which the taxpayer gains dominion and control", whether the taxpayer stakes directly or through an exchange [In force] [120]. The *Jarrett* litigation challenged the taxation of staking rewards on receipt. The government paid the refund (a check for USD 4,001.83), and the Sixth Circuit dismissed the appeal as moot in 2023 without reaching the merits [126]. The legal question it raised therefore remains unanswered by a court; the revenue ruling is the operative guidance.
- **UK.** Staking returns are miscellaneous income or trading income, valued at "the pound sterling value (at the time of receipt)" [In force] [109].
- **Japan.** Rewards are income on acquisition; corporations mark locked-up assets to market at year end (FAQ 3-1-6) [In force] [154].
- **India.** Income from other sources at fair market value; a later sale is taxed at 30% with that value as cost [In force] [20].

- **Proposals.** The US ADAPT Act, announced on 30 Sep 2026, includes staking and mining sourcing rules [Proposed] [45].

"Dominion and control" is a legal concept. The chain shows something narrower: when a reward was credited, and whether the credited units could be moved. Rebasing liquid staking tokens credit rewards continuously as balance changes; value-accruing tokens credit nothing until redemption (our inference). Zlef therefore labels the *event* (`staking_reward`) and its observable timestamp and transferability, and leaves the legal timing to the rules layer. It does not predict a tax year.

Internal transfers: `beneficial_owner_change = no`**.** A transfer between two wallets of the same owner is not taxable in the jurisdictions with a rule. The IRS FAQ states that "the transfer is a non-taxable event, even if you receive an information return" [In force] [123]. HMRC treats moving tokens between one's own wallets as no disposal [In force] [111]. Brazil does not tax the transfer but DeCripto reports wallet-to-wallet transfers [In force] [143].

Three later rules make the label more important, not less.

- **US per-wallet basis.** Since 1 Jan 2025, basis "must be applied on an account-by-account basis" [In force] [121]. Basis now moves with the units, so a self-transfer must carry its lot records. Notice 2025-7, extended by Notice 2026-20, lets taxpayers identify units on their own books through 31 Dec 2026 [In force] [90].
- **India TDS.** Peer-to-peer self-transfers can trigger withholding mechanics that assume a buyer and a seller [Unverified].
- **Travel rules.** Several regimes ask the same question for AML purposes: is the self-hosted address the customer's own? Brazil requires VASPs to identify the owner of self-custodied wallets [In force] [8]. Korea's amendment, effective about 20 Feb 2027, generally allows transfers to foreign exchanges and personal wallets only where sender and recipient are the same person [Enacted, not yet effective] [208]. (Part 3 covers these rules.)

The chain cannot prove that two addresses share an owner. Heuristics (common funding, co-spending, timing) give probabilities, not proof. This is the largest single source of disposal mislabelling. Zlef predicts `beneficial_owner_change` with a probability that is designed to be calibrated (calibration is to be measured), and the rules layer treats `no` as conclusive only when it is supported by a layer A or B fact, such as a signed ownership attestation or the VASP's own wallet-verification record.

Refunds. A refund in crypto reverses a payment, but not exactly. Between payment and refund the asset's price may move, and the refund may arrive in a different asset or quantity. We found no crypto-specific guidance in any jurisdiction. In the US, a general rescission doctrine (*Penn v. Robertson*) may treat a reversal within the same tax year as unwinding the original transaction [Unverified]. Otherwise, a refund looks like two transactions: the original disposal by the payer, and a new acquisition when the refund arrives.

Zlef labels `refund_reversal` with `economic_purpose = refund`, and should link the refund to the original payment where it can. That link is often off chain (an order number, a merchant reference), so `evidence_completeness` usually needs level 3 (an underlying document) before the rules layer can apply any rescission logic.

Fees paid in crypto. Every on-chain transaction pays a network fee in some token, so "fee" events are everywhere. The tax analysis is simple in principle and tedious in practice.

- **US.** Using crypto to pay a network fee disposes of the fee units; fees reduce proceeds or increase basis [Unverified]. H.R. 9178 (8 Jun 2026) would exclude gains on fees of USD 10 or less [Proposed] [38].
- **UK.** Fees are allowable costs, and the fee token is disposed of [Unverified].
- **Japan.** Fees are included in cost [Unverified].
- **India.** Only the cost of acquisition is deductible from transfer income [In force] [20].
- **Indirect taxes.** Indonesia levies VAT on platform services, mining and verification, at 12% on 11/12 of the base [In force] [115]. India's GST on exchange and platform services is [Unverified].

Zlef emits a `fee` event only for an explicit transfer paid as a fee (a protocol fee, an exchange fee, a relayer fee). The network fee of a transaction is recorded as an attribute of the main event, not as a separate event. This keeps the event count tractable and lets the rules layer apply each jurisdiction's fee rule in one place.

Stablecoin-to-stablecoin swaps. Swapping USDC for USDT is a disposal in the US today, with gains or losses that are usually tiny but reportable. Brokers need not report qualifying-stablecoin sales on Form 1099-DA if designated-sale proceeds "do not exceed $10,000 for the year" [In force] [122]. Two bills would go further: H.R. 9178 would value qualifying stablecoins at redemption value with a 99.5–100.5% safe harbour, and the ADAPT Act would give no gain or loss on regulated USD stablecoin purchases [Proposed] [38, 45].

The UK has chosen a full exemption. The outcome published on 13 Jul 2026 gives individuals a full CGT exemption for eligible stablecoins, including non-sterling ones, from April 2027 via Finance Bill 2026-27. HM Treasury rejected a no-gain, no-loss approach because of avoidance risk. There is no relief for de-pegging, and no rollover between exempt and non-exempt tokens [Proposed] [106]. On our reading of the outcome, a UK swap from an eligible stablecoin into a non-eligible token would be an exempt disposal of the eligible leg and the acquisition of a taxable asset. The reverse swap would be a taxable disposal of the non-eligible leg. Because there is no rollover, the boundary between the two kinds of token is itself a tax event.

Japan is different again. Many fiat-backed stablecoins are electronic payment instruments, not crypto assets, and press reports say they are excluded from the planned 20.315% regime for specified crypto assets [206]. Their tax treatment as electronic payment instruments is [Unverified]. In Brazil, buying, selling or exchanging fiat-referenced virtual assets is an FX operation [In force] (since 2 Feb 2026) [8].

Zlef labels `swap_stable` when both legs are `STB.*`. It does not decide eligibility. "Eligible stablecoin" in the UK is a statutory category that the rules layer checks against a versioned list. Zlef's job is to get the sub-class right (`STB.FIAT.SINGLE` against `STB.SYN` or `STB.YLD`), because a synthetic or yield-bearing leg is unlikely to qualify.

Payments for goods and services: payer disposal plus payee income. One on-chain transfer creates two tax events. For the payer, it is a disposal: the IRS says "you have exchanged a capital asset for that service" [In force] [123]; HMRC lists "paying for goods or services" as a disposal [In force] [111]; Japan's FAQ 1-2 treats it as a transfer of the crypto asset held [In force] [154]. For

a business payee, it is ordinary or trading income at market value, and the payee's basis is that value [Unverified]. India adds withholding: the TDS rule makes the person acquiring a VDA for consideration deduct 1% [116], which on our reading places the duty on a merchant who receives crypto as payment. Russia bans domestic crypto payments outright [In force] [11]. In Indonesia, crypto is not a means of payment, because the rupiah is the sole legal tender [In force] [27].

Indirect taxes add a third layer. In the EU, the CJEU held in *Hedqvist* (C-264/14, 22 Oct 2015) that exchanging bitcoin for traditional currency is a VAT-exempt financial transaction [44]; the scope of the holding as applied to other tokens is [Unverified]. VAT on the goods themselves follows barter rules [Unverified]. Vietnam states that crypto transfers "are not subject to VAT" [In force] [7], and Indonesia removed VAT on crypto transfers from 1 Aug 2025 [In force] [115].

The US de minimis proposals for small purchases (S. 2207: USD 300 per transaction and USD 5,000 per year) are bills, not law [Proposed] [38]. Zlef's schema handles the two-sided event through perspective: the same transfer is labelled `pay_goods_services` for the payer and `receive_payment_goods_services` for the payee, with `economic_purpose` of `goods_purchase` or `service_purchase`.

Why Zlef predicts `event_type` and `beneficial_owner_change`, not "taxable: yes/no". An obvious alternative design would train Zlef to output "taxable" directly. We reject it for seven reasons, each visible in Table 38.

1. **The same facts give different answers in different places.** A gift is a disposal for a UK donor and not for a US donor. A stablecoin swap is a disposal in the US and, from April 2027 if the draft is enacted, exempt for a UK individual. A swap in Indonesia is taxed on gross value whatever the gain. A single "taxable" output would have to be indexed by jurisdiction, and the model would have to learn ten tax codes.

2. **The same facts give different answers at different times.** The UK DeFi and stablecoin drafts (6 Apr 2027) [Proposed], Korea's tax (1 Jan 2027) [Enacted, not yet effective], with a deferral bill announced on 15 Oct 2026 [Proposed], Japan's 20.315% regime (expected from 2028) [Sources conflict] and the end of US standing-order relief (31 Dec 2026) [In force] all fall within a plausible deployment window. A fact label stays valid across these dates; a legal label expires.

3. **Legal answers depend on facts the model must not guess.** The UK draft relief covers individuals and trustees, not companies. India's airdrop rule depends on whether the sender is a relative. Residency, entity type and relationships are layer B, C or F facts. A model that predicted "taxable" would be implicitly predicting them.

4. **Ambiguity is the true state for many rows.** Wrapping, bridging, LP and lending events are unsettled in most jurisdictions. A binary "taxable" label would force a guess where the honest output is "the law does not say". With fact labels, the rules layer can return "ambiguous; review" while the facts themselves are well calibrated.

5. **Calibration needs a verifiable target.** Whether two addresses share an owner, or whether a transfer was an LP deposit, has a fact of the matter that annotators can establish. Whether a transaction is "taxable" may have no answer until a tax authority or court gives one. A probability is only meaningful against a target that can be checked.

6. **Errors must be attributable.** With layered outputs, a wrong result can be traced to a wrong

fact (layer D, the model), a wrong rule (layer E) or a wrong review (layer F). A single legal label mixes all three.

7. **A model should not give legal conclusions.** Legal and tax treatment comes from versioned rules plus human review. Keeping Zlef's outputs jurisdiction-free keeps that boundary clean.

`beneficial_owner_change` is the single most useful fact because it drives disposal analysis in most income-tax regimes and drives self-hosted-wallet analysis in most travel-rule regimes. `event_type` is the second, because it fixes which rule family applies. Together with `asset_class`, `counterparty_type` and dated jurisdiction facts, they are enough for the rules layer to derive the character in every non-U cell of Table 38. Whether they are enough in practice is an empirical question that Edition 1 cannot answer [Research hypothesis].

5.7 How Assets Travel

Tax follows the event. Compliance follows the *route*. Two payments with the same asset, amount and purpose can raise different obligations because one crossed a bridge, touched a mixer or ended in a self-hosted wallet. Zlef's `travel_surface` label captures the route as a set of seven facts, each annotated `yes`, `no` or `unknown`. Each fact is chosen because it switches on a distinct body of rules. Table 39 summarises them; the paragraphs below define each one, say what the chain shows, and list the regimes it activates.

The travel-rule baseline. Most of these labels matter because of the FATF travel rule. FATF applies Recommendation 16 to virtual-asset transfers through the interpretive note to Recommendation 15, with a de minimis of USD/EUR 1,000 [In force] (as a standard) [82]. FATF's 2025 revision of Recommendation 16 widens payment-transparency duties and requires name, address and date of birth for cross-border peer-to-peer transfers above USD/EUR 1,000, to be implemented "by the end of 2030" [Enacted, not yet effective] (as a standard) [74]. FATF has said that it is not applying the revised Recommendation 16 directly to VASPs; VASP obligations continue through the interpretive note to Recommendation 15 [74]. FATF's 2026 update reports that 91 of 109 responding jurisdictions (83%) have passed travel-rule legislation, and that 55 of those 91 have taken no travel-rule supervisory or enforcement action [82]. Thresholds differ sharply across our ten jurisdictions (Table 40), which is why the route facts must be recorded rather than assumed.

`crosses_chain`**.** *Definition.* `yes` when the labelled movement starts on one ledger and ends on another. A movement between two L2 networks that settle to the same L1 still counts, because the user-facing ledgers and the token contracts differ.

Observability. Each leg is fully visible on its own chain: addresses, amounts and timestamps. What is *not* visible is the link between legs. Establishing it needs either the bridge's message format (for CCTP, the message events, nonce and domain identifiers [32]) or a commercial attribution graph. Without a link, the two legs look like two unrelated transfers, and a chain-hop looks like a payment to a stranger followed by a receipt from another.

Regimes. FATF's 2020 red-flag indicators name chain-hopping: rapid sequential bridging, conversion into privacy coins, use of non-KYC swap services, peel chains, and structuring below travel-rule thresholds [Unverified]. Sanctions screening must run on both legs, because a

Table 39: The `travel_surface` labels: definition, on-chain observability and main regimes (as of 15 Oct 2026).

Label	Definition (`yes` when ...)	Observability	Main regimes activated	Paper V families
`crosses_chain`	value leaves one ledger and arrives on another within the labelled movement	Each leg visible; the link needs message decoding or attribution	FATF chain-hopping red flags; per-leg sanctions screening	AML, SAN, TRV
`passes_bridge`	a bridge contract, validator set or custodian intermediates the cross-ledger movement	Bridge contracts and events visible; solver fills weakly linked	Bridge-laundering typologies (DPRK); unsettled VASP status of bridges	AML, SAN, AST
`issuer_freezable`	an issuer, custodian or agent can freeze, block, burn or seize the units	Contract roles and freeze events visible	GENIUS lawful-order capability; MiCA issuer duties; issuer blacklists	SAN, STB
`permissioned_transfer`	the token checks an identity or allow-list before each transfer	Revert behaviour and registry calls visible; identities hashed	Securities transfer rules; ERC-3643 identity registry	AST, LIC
`privacy_enhanced`	protocol or contract features hide sender, receiver or amount	Partial or none	AMLR anonymity-enhancing coins prohibition (from 10 Jul 2027); exchange delistings	AML, AST, TRV
`touches_mixer`	the funds pass through, or come directly from, a mixing service	Visible for contract mixers; custodial mixers inferred	OFAC listing history (Tornado Cash); AML exposure scoring	SAN, AML
`touches_unhosted`	one end of the movement is an address not controlled by a VASP	Inferred: absence of a known VASP attribution	Travel-rule self-hosted wallet rules; CARF transfer reporting	TRV, REP

destination-chain address can be listed when the source-chain address is not. A travel-rule message, where one exists, usually covers only the VASP-to-VASP or VASP-to-wallet leg, not the hop in between.

`passes_bridge`. *Definition.* `yes` when a bridge contract, validator or guardian set, or custodian intermediates the cross-ledger movement. It differs from `crosses_chain` in two cases, both of which are `crosses_chain = yes` with `passes_bridge = no`: an issuer-native burn-mint (CCTP) crosses chains through the issuer's own protocol, not a bridge, and is labelled `WRP.NATIVEBURN` (Section 6.2); and an exchange deposit on one chain with a withdrawal on another crosses chains through a VASP's internal books.

Observability. Bridge contracts are public, and their deposit and mint events are decodable. Intent or solver bridges, which settle the destination leg from a solver's inventory, leave no common message, so the link is weaker (our inference).

Regimes. Bridges concentrate two risks. The first is theft. Lock-mint bridges hold large pools of locked assets, and the major exploits of 2022 hit bridges [Unverified]. The second is laundering. The FBI attributed the Bybit theft of about 21 Feb 2025 to DPRK actors it calls "TraderTraitor" [77]; FATF puts the loss at USD 1.46bn, with "only 3.8%" recovered [80]. Laundering of the proceeds

reportedly involved swaps and bridging from Ethereum assets to bitcoin at high velocity [Unverified]. FATF's 2026 update links DPRK cyber-theft laundering to scam-centre infrastructure through a Cambodia-based conglomerate that "laundered at least USD 4 billion" between 2021 and 2025 [82]. Whether bridge operators are themselves VASPs is unsettled [Unverified]. For wrapped stablecoins, the bridge also changes the obligor (Section 5.1).

`issuer_freezable`. *Definition.* `yes` when an issuer, custodian or agent can freeze, block, burn or seize the units being moved. This is a property of the token contract and its governance, evaluated at the time of the movement.

Observability. Strong. Freeze and blacklist powers are contract roles, and their use emits contract events. A freeze or blacklist event is therefore an on-chain, authoritative signal that something happened, though not why. For native assets there is no equivalent; sanctions exposure must be inferred from off-chain address attribution.

Regimes. The GENIUS Act (Public Law 119-27, approved 18 Jul 2025) makes freeze capability a condition of issuance [Enacted, not yet effective] [218]:

- a permitted issuer must have "the technological capability to comply, and will comply, with the terms of any lawful order" (§4(a)(6)(B));
- it must have "technical capabilities, policies, and procedures to block, freeze, and reject" transactions (§4(a)(5)(A)(iv));
- a "lawful order" includes an order to "seize, freeze, burn, or prevent the transfer of payment stablecoins" (§2(16)(A));
- foreign issuers' market access is conditioned on the same capability (§3(b)(2));
- the prohibition on digital asset service providers offering non-permitted payment stablecoins begins "3 years after the date of enactment" (§3(b)(1)), that is about 18 Jul 2028.

The Act's general effective date is reported as the earlier of 18 months after enactment (about 18 Jan 2027) or 120 days after final regulations [Unverified] (the text was not confirmed in our reading). Issuer blacklists already operate in practice: Circle froze Tornado Cash-associated USDC in August 2022, and Tether can freeze and destroy blacklisted funds [Unverified]. In the EU, EMT and ART issuers must be authorised under MiCA [In force] [72]. FATF warns that VASPs "may be unable to rely on issuer-level asset freeze/burn mechanisms" [82]. In Zlef, `issuer_freezable = yes` tells the rules layer that a freeze is possible and that freeze events should be checked; it does not say that a freeze is likely.

`permissioned_transfer`. *Definition.* `yes` when the token contract checks an identity registry or allow-list before it lets a transfer complete.

Observability. The ERC-3643 (T-REX) design illustrates the pattern [Unverified] (technical standard, not re-checked). A token contract calls a `canTransfer` check; an identity registry maps each wallet to an identity contract (ONCHAINID) holding signed claims from trusted claim issuers (KYC, accreditation, jurisdiction); a compliance contract applies modular rules such as country restrictions, holder caps and lock-ups; and agent roles can freeze addresses or partial balances, force transfers, recover lost wallets and pause the token. On chain, one sees the registry calls and the outcome, and the claims appear only as hashes or attestations.

Regimes. Permissioned tokens are mostly tokenised securities, which are regulated as securities in every jurisdiction where we found a rule. A transfer to a non-registered address

reverts; it does not complete and then trigger a flag. So, for Zlef, a successful transfer is evidence that the counterparty passed the issuer's onboarding gate. It is not evidence that the gate was adequate. The UK draft DeFi relief excludes tokenised securities [Proposed] [191], so `permissioned_transfer` also feeds the tax rules. CBDCs on permissioned ledgers also take `yes` (Section 5.4).

`privacy_enhanced`**.** *Definition.* `yes` when protocol or contract features hide the sender, receiver or amount of the movement. This includes privacy-coin transfers (Section 5.3) and shielded pools on otherwise transparent chains. It does not include mixers, which have their own label.

Observability. For Monero, almost nothing beyond the existence of a transaction. For Zcash, transparent transfers are visible and shielded transfers are not; entering and leaving the shielded pool is visible. For shielded pools on transparent chains, deposits and withdrawals are visible but the link between them is hidden.

Regimes. The EU AMLR prohibits credit institutions, financial institutions and CASPs from keeping accounts that allow anonymisation, including through anonymity-enhancing coins, from 10 Jul 2027 [Enacted, not yet effective] [36, 73]. India's FIU-IND guidelines treat anonymity-enhanced coins as "unacceptably high-risk" [In force] [162]. Exchange practice in Japan and Korea excludes privacy coins [Unverified]. Where the label is `yes`, travel-rule data cannot be checked against the chain, and basis and provenance need off-chain attestation.

`touches_mixer`**.** *Definition.* `yes` when the funds pass through a mixing service, or arrive directly from one, within the window the rules layer specifies. Edition 1 uses direct exposure only; indirect exposure (several hops away) is a risk score, not a label.

Observability. Contract-based mixers are public contracts, so deposits and withdrawals are visible, though the mixer is designed to break the link between them. Custodial mixers are identified only through attribution.

Regimes. The legal history matters because it shows why sanctions flags and AML flags must be separate. OFAC designated Tornado Cash in August 2022 [Unverified]. In *Van Loon v. Department of the Treasury* (5th Cir., 26 Nov 2024), the court held that immutable smart contracts are not "property" that IEEPA lets OFAC block [Unverified]. On 21 Mar 2025, Treasury removed the sanctions: "we have exercised our discretion to remove the economic sanctions against Tornado Cash" [In force] [222]. The delisting removed a *sanctions* flag. It did not remove *AML* risk from mixer exposure, and it did not affect other listings, such as DPRK-linked addresses, or separate criminal cases [Unverified]. FATF's 2020 red-flag indicators list conversion through mixers as a chain-hopping typology [Unverified]. Zlef therefore labels mixer contact as a fact. The rules layer applies a dated sanctions-list snapshot (binary, list-driven) and, separately, an AML exposure policy (risk-scored).

`touches_unhosted`**.** *Definition.* `yes` when at least one end of the movement is an address not controlled by a VASP: a self-hosted wallet, a smart contract acting for a user, or an address of unknown control.

Observability. Indirect. The chain does not say who controls an address. Zlef infers "unhosted" from the absence of a known VASP attribution and from wallet behaviour. Errors in both directions are common: a VASP's fresh deposit address can look unhosted, and an individual's address can look like an exchange's.

Regimes. This label activates the self-hosted-wallet rules, which differ more across our jurisdictions than any other travel-rule feature (Table 40). It also feeds tax reporting: CARF treats transfers to wallets not known to be associated with a VASP as reportable transfers [Unverified], and Brazil's DeCripto reports wallet-to-wallet transfers [In force] [143]. FATF's 2026 update reports that 88% of responding jurisdictions (58 of 66) rate peer-to-peer activity as high risk, while only 23% collect peer-to-peer market metrics [82].

The table shows a structural asymmetry. A US VASP may send a transfer below USD 3,000 with no travel-rule data, while a Japanese, Indian, Brazilian, Indonesian or (from 2027) Korean VASP must treat the same transfer on receipt as missing information. And three jurisdictions (Brazil, Korea after the amendment, and the EU above EUR 1,000) make the question "is this self-hosted wallet the customer's own?" a legal test. That is the same question as `beneficial_owner_change = no` for an internal transfer. One Zlef prediction, confirmed by a layer B record, could therefore serve both tax and AML rules, provided its calibration is shown in evaluation.

What is visible on-chain and what needs off-chain records. Table 41 sets out, data element by data element, what the chain shows across the route types, and what must come from off-chain records. It is the evidential basis for `evidence_completeness`: level 1 means only the left-hand columns are available.

Implications for Zlef. Four design consequences follow from this part.

1. **The route is a set, not a class.** One movement can cross a chain, pass a bridge, use a freezable stablecoin and end in a self-hosted wallet. The labels are therefore separate outputs, each a marginal probability with its own calibration, not a single "route type"; their dependencies are handled by the consistency constraints (Section 6.10).
2. **Unobservable is not `no`.** For shielded transfers and unlinked bridge legs, several labels cannot be observed. Edition 1 adopts `unknown` as an annotation value for each travel label (Section 6.9), so that a reviewer can tell "checked and absent" from "not checkable".
3. **Facts, then dated rules.** Sanctions lists, travel-rule thresholds and prohibitions change on known dates: Tornado Cash was delisted on 21 Mar 2025; Korea's threshold falls to zero around 20 Feb 2027; the EU AMLR prohibition applies from 10 Jul 2027; the GENIUS DASP prohibition starts around 18 Jul 2028. Zlef's travel labels do not change on those dates. The rules that read them do.
4. **The strongest on-chain signals are issuer events.** Freeze, blacklist and registry-revert events are authoritative facts recorded by the contract. They belong in layer A or B, not in model prediction. Zlef should consume them as inputs where available, and predict only where they are absent.

6 The Zlef Label Schema

[Product, off-protocol] This chapter fixes the Edition 1 schema: the eight questions Zlef answers, their typed domains, the meaning of every value, the rules an annotator follows to choose one, and the reason each question exists. Chapters 4 and 5 applied the schema asset by asset. Section 7

Table 40: Travel-rule thresholds and self-hosted wallet treatment (as of 15 Oct 2026).

Juris.	Threshold	Self-hosted wallet treatment	Status	Source
US	USD 3,000 (31 CFR 1010.410(f))	No specific travel-rule duty; a self-hosted wallet is not a financial institution. The December 2020 unhosted-wallet proposal was never finalised.	[In force] (threshold); proposal [Unverified]	[219]
EU	No de minimis for CASP-to-CASP	TFR applicable since 30 Dec 2024; above EUR 1,000 the CASP must assess whether the self-hosted address is owned or controlled by its customer. EUR-Lex metadata shows the TFR consolidated into the AML package from 10 Jul 2027.	[In force] (application date); content [Unverified]	[70]
UK	Effectively zero; full data above EUR 1,000	Information collection for unhosted transfers "not mandatory"; a risk assessment is required.	[In force] (since 1 Sep 2023)	[145]
JP	None	Transfers to individuals and unregistered VASPs: collect and retain data and assess wallet risk.	[In force] (since 1 Jun 2023)	[84, 164]
KR	KRW 1m, falling to zero from about 20 Feb 2027	Today exchanges whitelist external wallets [Unverified]. After the amendment, transfers to non-low-risk foreign exchanges and personal wallets generally only where sender and recipient are the same person; high-risk counterparties prohibited.	[Enacted, not yet effective] (amendment)	[208]
IN	None	Unhosted wallets and P2P treated as higher risk: data collection, enhanced CDD, possible restriction.	[In force] (guidelines updated 8 Jan 2026)	[162]
BR	None; cross-border phase from 2 Feb 2027	VASP must identify the owner of self-custodied wallets; such transfers fall in the FX framework (Res. 521).	[In force] (domestic, since 2 Feb 2026); cross-border [Enacted, not yet effective] (2 Feb 2027)	[8, 159]
ID	None; wider data at or above USD 1,000	Withdrawal wallets must be "clear, registered, and previously verified".	[In force]	[163]
VN	No VASP travel rule (FATF self-report)	Pilot licensees must apply FATF standards; no licensee as of 30 Aug 2026.	[Unverified]	[82]
UA	Legislation reported; threshold not located	Not located.	[Unverified]	[82]
RU	"In progress" (FATF self-report)	Reported: only foreign-trade companies may send directly to self-custodied wallets; individuals route via foreign custodial wallets.	[Unverified]	—

Table 41: On-chain visibility by route type, and the off-chain record needed.

Data element	Public chain, no issuer	Issuer-freezable or permis-sioned	Bridged (per leg)	Privacy-enhanced	Off-chain record needed
Sender and receiver address	Yes	Yes	Yes	No (XMR); partly (ZEC)	Address is not a person: identity from VASP or issuer KYC
Amount	Yes	Yes	Yes	No (XMR)	—
Time and block	Yes	Yes	Yes	Yes	Dated price source for fiat value
Asset identity	Contract or chain	Contract plus registry	Must map wrapped to underlying	Partial	Versioned registry of wrappers, bridges and issuers
Link between legs	n/a	n/a	Only by message decoding	No	Bridge message index or attribution graph
Freeze or blacklist status	n/a	Yes (events, state)	Custodian tokens: yes	n/a	Reason for the freeze (lawful order, sanctions, fraud)
Identity and KYC	No	Hashed claims at most	No	No	VASP or issuer records; CARF self-certification
Same owner on both ends	No (heuristics only)	Registry may show same identity	No	No	Ownership attestation or VASP wallet verification
Purpose (payment, gift, loan, refund)	No	Sometimes encoded	No	No	Invoice, contract, order or payroll record
Cost basis and lot	No	No	No	No	Taxpayer ledger (per wallet in the US)
Fiat value	No (needs a price feed)	About par for stablecoins (de-peg risk)	Underlying price	No	Dated, cited price source
Tax residence	No	Claim may encode country	No	No	Self-certification; layer B/F fact
Sanctions status	Address-list match	List match plus issuer freeze	List match per leg	Exchange edges only	Dated list snapshot

gives losses and decision policies for it, and Section 8 describes how its labels would be collected and evaluated. The label dictionary in Appendix A restates every value in one place.

6.1 Design principles

Five principles shape the schema. Each answers a failure mode that the preceding chapters made concrete.

P1. Predict facts, not law. Every question asks for a jurisdiction-free economic or technical fact. None asks for a legal category, a tax treatment or a flag. Section 3 showed that legal category is a function $L(\text{facts}, \text{jurisdiction}, \text{date})$ that changes with the jurisdiction and over time.

A tokenized money-market fund share can be a security in one regime, a crypto-asset in another and unregulated in a third. A stablecoin can be a payment stablecoin, an e-money token, an electronic payment instrument or an FX operation, depending on where it is held and moved. If the model predicted these categories directly, it would have to learn every regime and relearn it whenever a statute changed. It would also turn a statutory change into a silent model error. A rule change should be a versioned edit to the rules layer (Paper V, Section 6.11 below), not a retraining run.

P2. Read before predicting. A question whose answer is fixed by a protocol-authenticated (A) or source-attested (B) fact is answered from that fact, not by the model. For example:

- a token contract in a maintained registry fixes `asset_class`;
- two chain identifiers fix `crosses_chain`;
- the count of linked documents fixes most of `evidence_completeness`.

The model is for the residue that records do not settle. This keeps model predictions (layer D) apart from the facts they would otherwise duplicate, and it makes the model's measured error rates apply only to inputs that genuinely need inference.

P3. UNKNOWN is a label, abstention is a decision. Every nominal question includes `UNKNOWN`. An annotator chooses it when the available facts do not determine the answer, not when the annotator is unsure. Abstention is a separate system decision taken on the model's output (Section 7.5). Keeping the two apart lets us measure two different things: how often the world is underdetermined, and how often the model is unreliable.

P4. Hierarchy allows partial answers. `asset_class` is a tree. A confident parent with an uncertain child ("a stablecoin, sub-type unclear") is a useful answer, because many rules depend only on the parent. The decision policy may therefore return an internal node (Section 7.3.2). The tree is built so that rules which depend only on a parent never need the child.

P5. Ordinal where order is real, and nowhere else. Two questions have a natural order (`evidence_completeness`, `review_priority`). There a near miss is less wrong than a far miss, so partial-credit objectives in the spirit of Cloudflare's RLCD are candidates (Section 2.2). The others are nominal, and an RLCD-style objective has no obvious meaning there. Paper IV left this as a research question, and Section 10 keeps it open.

6.2 Q1: `asset_class`

Question. What kind of asset moved, in economic and technical terms?

Domain. The tree in Table 42, plus `UNKNOWN`.

Annotation rules.

1. **Representation before economics.** Label a wrapped or bridged token `WRP.*`, not the class of its underlying asset. The underlying is recorded as a separate fact, a link to its own record. A bridged USDC is `WRP.BRIDGE` with an underlying of `STB.FIAT.SINGLE`. Native

Parent	Child	Defining facts (annotator test)
NAT	NAT.L1, NAT.L2	The protocol's own fee or staking unit; no issuer redemption claim. L2 if the token's native ledger settles to another chain.
STB	.FIAT.SINGLE	Issuer promises redemption at par in one fiat currency, backed by reserves.
	.FIAT.BASKET	Par claim against a basket of fiat currencies.
	.COMM	Claim on, or reference to, a commodity such as gold.
	.SYN	Price target held by collateral, hedging or algorithm, with no issuer par claim on reserves.
	.YLD	A stablecoin wrapper that passes yield to the holder.
TKN	.MMF, .EQ, .DEP, .DEBT, .RE	On-chain representation of an off-chain financial asset or claim: fund shares, equity or trackers, bank deposits, debt instruments, real estate or other RWAs.
GOV	—	Confers governance or protocol-use rights; no redemption claim; not mainly a meme.
MEME	—	No function beyond trading and community; value by attention.
NFT	.UNIQUE, .FRAC	Non-fungible token; FRAC where a series or fractionalisation makes units fungible in practice.
LST	.STAKE, .RESTAKE	Receipt for staked (or restaked) assets that accrues rewards.
WRP	.CUST, .BRIDGE, .NATIVEBURN	Representation of another asset: custodial wrap; lock-and-mint bridge; burn-and-mint native transfer (e.g. CCTP).
DFI	.LP, .LEND, .DEBT	Receipt for a DeFi position: liquidity share; lending deposit; debt position.
PRV	—	Asset whose protocol hides sender, receiver or amount by default or by option.
CBDC	—	Central-bank money, or e-money tokens outside MiCA's crypto-asset scope.
OTH	—	Points, in-game items, unvested or claim rights that are not yet transferable tokens.

Table 42: The asset_class tree. Full definitions, examples and confusions are in Sections 4 and 5.

USDC moved by the issuer's burn-and-mint protocol (CCTP) is STB.FIAT.SINGLE at rest. The burn-and-mint transfer event itself is labelled WRP.NATIVEBURN, with event_type = bridge and crosses_chain = yes. Because CCTP is the issuer's own protocol, passes_bridge = no for CCTP itself; it is yes only if a third-party bridge is also used (Section 5). The reason is that issuer freeze powers, bridge risk and the disposal analysis attach to the representation (Section 6.9).

2. **Receipts before underlyings.** An LST, LP token or lending receipt is labelled as the receipt (LST, DFI), for the same reason: tax events and transfer mechanics attach to the receipt.

3. **GOV versus MEME.** Choose GOV only if the token confers a documented governance or protocol-use right. "Has a DAO" is not enough when the DAO has no on-chain powers. When in doubt, choose UNKNOWN rather than guess, because some regimes treat the distinction differently (Section 3).

4. **STB.SYN versus STB.FIAT.** The test is whether a named issuer promises redemption at par from reserves. Marketing language ("fully backed") does not settle it; the issuer's terms do. A missing issuer document is a layer C fact, and the label should then be the parent STB.

5. **TKN versus everything else.** `TKN` requires an identifiable off-chain asset or claim and an issuer or custodian of record. A token that tracks a price without a claim is not `TKN`; it is `STB.SYN` if it targets a currency, otherwise `UNKNOWN` or `OTH`.

Why this question exists. Almost every rule family in Paper V branches on it:

- `AST` flags: securities status, privacy coins, transfer restrictions;
- `STB` flags: recognised versus non-recognised stablecoins, yield bans;
- `TAX` rules: the UK's draft CGT exemption covers only eligible stablecoins, and the UK draft no-gain, no-loss relief excludes tokenised securities [Proposed]; the US collectibles rate can reach NFTs [Unverified];
- `REP` scope: CARF's "Relevant Crypto-Asset" and its exclusions.

Why jurisdiction-free. The tree describes what the asset *is*, not what any law calls it. The crosswalk in Section 3 maps each node to each regime's category, and that crosswalk belongs to the rules layer.

6.3 Q2: `event_type`

Question. What happened to the asset, from the perspective of the record holder (payer or holder)?

Domain. The values are:

- acquisitions and disposals: `acquire_fiat`, `dispose_fiat`, `swap_crypto`, `swap_stable`;
- payments and income: `pay_goods_services`, `receive_payment_goods_services`, `compensation`, `staking_reward`, `mining_reward`, `airdrop`, `fork`;
- representation and DeFi: `wrap_unwrap`, `bridge`, `lp_deposit`, `lp_withdraw`, `lend_out`, `lend_return`, `borrow`, `repay`;
- other: `gift`, `internal_transfer`, `refund_reversal`, `fee`, `loss_theft`;
- `UNKNOWN`.

Annotation rules.

1. **Label the economic event, not the transaction type.** One on-chain transaction can contain several events. A swap with a fee is `swap_crypto` plus `fee`. The record is split into event records before labelling, and each event record gets one `event_type`.
2. **Perspective is fixed.** The same transfer is `pay_goods_services` in the payer's books and `receive_payment_goods_services` in the payee's. The schema always labels from the record holder's side, so a dashboard user who holds both sides gets two event records.
3. **Do not decide taxability.** `wrap_unwrap` is the label whether or not a jurisdiction treats wrapping as a disposal. Wrapping is unsettled in the US and treated differently elsewhere (Section 5.6). The tax consequence comes from the rules.

4. `internal_transfer` **needs evidence of common ownership.** Use it only when the records show that both endpoints belong to the same beneficial owner, for example two addresses registered to the same customer. Otherwise use `UNKNOWN`.

Why this question exists. Tax rules are keyed on events, not on transfers. Disposal, income, non-event and unsettled outcomes differ by event and by jurisdiction (Section 5.6). The event also selects which reporting and travel-rule rules apply. Why not "taxable: yes/no"? Because the answer differs across the ten jurisdictions and over time:

- UK lending and LP treatment changes from 6 April 2027 under draft legislation [Proposed];
- Japan's 20.315% separate rate is expected from 2028 [Sources conflict];
- Korea's tax begins on 1 Jan 2027 [Enacted, not yet effective], although a deferral bill was announced on 15 Oct 2026 [Proposed].

A model trained to answer "taxable" would be wrong somewhere at every moment.

6.4 Q3: `economic_purpose`

Question. What was the payment for? (This extends Paper IV's purpose question; see Section 6.13.)

Domain. `service_purchase`, `goods_purchase`, `refund`, `internal_transfer`, `fee`, `payroll`, `loan`, `distribution`, `investment`, `remittance_family`, `UNKNOWN`.

Annotation rules.

1. **Purpose comes from documents, not from amounts or memos alone.** A memo that says "refund" is evidence, but it is attacker-controllable text (Section 7.9). Without a linked document or relationship, the annotator should prefer `UNKNOWN`.
2. `distribution` **versus** `payroll`**.** Distributions go to owners or members as owners (dividends, DAO distributions). Payroll goes to workers for work. Contractor payments are `service_purchase` from the payer's side.
3. `remittance_family` covers person-to-person transfers to family without consideration. It is kept separate because its evidence differs (there is no invoice or contract). Whether any regime treats it differently from a gift is for the rules layer to state [Unsupported, needs evidence].

Why this question exists. Purpose drives which documents a reviewer needs and which reporting regime applies. For example, payroll in crypto raises employment-tax questions in every jurisdiction where our sources address it. It also feeds Brazil's FX purpose codes for international crypto payments (Paper V).

Why separate from Q2. Event and purpose are orthogonal. A `dispose_fiat` event can serve any purpose. A `service_purchase` can be settled by a `pay_goods_services` event, or by a `swap_stable` routed to the seller. Merging them would multiply the label space and blur two kinds of evidence.

6.5 Q4: `beneficial_owner_change`

Question. Does beneficial ownership of the asset change in this event? **Domain:** `yes`, `no`, `unknown`.

Why this question exists. Across the ten jurisdictions, the disposal analysis turns on a change of beneficial ownership more than on any on-chain feature.

- Internal transfers between one's own wallets are non-events in every jurisdiction where our sources give a rule (the US, UK, Japan, India, Brazil and Germany); elsewhere the position is unverified.
- Wrapping, bridging and LP deposits are hard legal cases. The per-leg rule below fixes the economic label; whether a given regime treats the leg as a disposal stays with the rules layer.
- The travel rule's self-hosted-wallet rules (for example Korea's "same person" rule for transfers to foreign exchanges and personal wallets from about February 2027 [Enacted, not yet effective]) ask the same question.

Making it a separate yes/no/unknown question lets the rules layer combine it with `event_type` under each jurisdiction's own test.

Annotation rule. The label is given per leg, in economic terms.

- A leg that moves an asset into a pooled protocol contract, or to a custodian, in exchange for a receipt or claim is `yes`. This covers LP deposits, lending deposits, liquid-staking and restaking mints, and custodial or lock-and-mint wraps.
- A leg between addresses that records show are controlled by the same owner, including the owner's own smart-contract wallet, is `no`. Choose `no` only on such positive evidence of common ownership.
- A leg to a different owner, shown by records or implied by the event type (a purchase from a merchant), is `yes`.
- Otherwise choose `unknown`.

Whether a jurisdiction treats the event as a disposal is a rules question, not part of this label. The asymmetry is deliberate: a false `no` can suppress a tax or travel-rule flag, which is the costlier error (Section 7.6).

6.6 Q5: `counterparty_type`

Domain. `regulated_vasp`, `vasp_unmatched`, `self_hosted`, `smart_contract_protocol`, `bank_or_psp`, `individual`, `business`, `UNKNOWN`.

`vasp_unmatched`**.** A VASP-like service that is not matched in any register consulted. The label records only the failed match. Whether registration is required anywhere, and so whether the provider is "unregistered" in a legal sense, is decided by the rules layer for each jurisdiction.

Read before predicting. Whether a VASP is registered is a fact about a register, and Paper V's flags read it from registers and travel-rule messages (layer B). The model is used only where no register match exists. Even then its output never establishes that a provider is registered or unregistered somewhere. A model prediction of `vasp_unmatched` can raise a `LIC` or `TRV` flag for review, capped at `review` under Paper V's rule for prediction-driven flags. It cannot clear one.

Why this question exists. Travel-rule obligations, sunrise handling, self-hosted-wallet rules and the market-access flags all depend on who is on the other side. The value set mixes institutional type (VASP, bank, protocol) with natural-versus-legal person (individual, business). The annotation rule is that the institutional type is chosen when it is known; this is a schema convention, not a claim about any regime.

6.7 Q6: `evidence_completeness` (ordinal)

Domain. The levels are:

- 0: no evidence beyond the record itself;
- 1: on-chain data only;
- 2: plus counterparty identity;
- 3: plus an underlying document (invoice, contract, payslip);
- 4: plus reconciliation to accounting records.

Read before predicting. Most of this question is counting, so the record supplies it. The model is needed only to judge whether a linked document actually relates to the payment, for example whether the attached invoice matches the amount and the counterparty.

Why ordinal. The levels are cumulative. Mistaking 3 for 4 is less harmful than mistaking 0 for 4. This is the question on which an RLCD-style partial-credit objective most clearly makes sense [Research hypothesis].

6.8 Q7: `review_priority` (ordinal)

Domain. `low` < `medium` < `high` < `urgent`.

Semantics. Priority orders a reviewer's queue. It never replaces a flag's severity, and a `block` flag is shown first whatever priority the model assigns. Priority is learned from reviewer behaviour (what reviewers chose to handle first and what turned out to matter), which makes it the question most exposed to feedback loops. Section 8.3 treats it as the last question to be trained and the first to be audited. Priority is corridor-free: it reflects facts in the record, not the jurisdictions involved. Escalation for a particular corridor or jurisdiction belongs to the decision policy or the rules layer, which see the jurisdiction context; Zlef does not.

6.9 Q8: `travel_surface` (multi-label)

Domain. Seven indicators:

- `crosses_chain`
- `passes_bridge`
- `issuer_freezable`
- `permissioned_transfer`
- `privacy_enhanced`
- `touches_mixer`
- `touches_unhosted`

Annotation labels are `yes`, `no` or `unknown`; this is adopted in Edition 1. The model outputs a Bernoulli marginal probability of `yes` for each indicator, and `unknown` annotations are masked out of the loss (Section 7.3.4). The draft specification had `yes/no` only. The change follows the privacy-coin case in Section 5, where some indicators cannot be observed at all. Examples in this paper write the values as the JSON strings `"yes"`, `"no"` and `"unknown"`.

Why this question exists. These indicators decide which compliance surfaces a payment crosses (Section 5):

- issuer freeze powers, relevant to sanctions response and to the GENIUS Act's lawful-order capability requirement for foreign issuers [Enacted, not yet effective];
- identity-registry transfer restrictions on permissioned tokens;
- bridge and mixer exposure, which are recurring laundering typologies;
- self-hosted endpoints, which trigger travel-rule and FX rules in several of the ten jurisdictions.

Most indicators are directly observable on-chain and should be read, not predicted (P2). The model's role is to propagate them through multi-hop paths and to infer them where the records are partial.

6.10 Consistency constraints

Answers to different questions must be jointly coherent. This list is the canonical Edition 1 constraint set; Section 7.7 refers to it by these numbers (C1–C11). Each constraint is stated over label values and applies in two places:

- to *annotation labels*, where `yes`, `no` and `unknown` (or `UNKNOWN`) are all allowed values, and a constraint whose antecedent or consequent involves `unknown/UNKNOWN` is satisfied; a violation sends the item back to adjudication;
- to *decisions*, where each question's decision is a value, an internal node or an abstention $\perp$, and each `travel_surface` indicator is decided as `yes`, `no` or $\perp$ from its marginal; abstentions satisfy every constraint.

Violations are never silently repaired. A violation in the raw decided outputs always emits a `DAT`-family review flag. Joint decoding may be used to *propose* a consistent alternative, which is shown to the reviewer next to the raw outputs (Section 7.7).

C1. `event_type` $=$ `internal_transfer` $\Rightarrow$ `beneficial_owner_change` $=$ `no`.

C2. `event_type` = `internal_transfer` ⇒ `economic_purpose` = `internal_transfer`. The converse does not hold: purpose `internal_transfer` is allowed with a `bridge` or `wrap_unwrap` event that moves the owner's own funds.

C3. If `event_type` is `wrap_unwrap` or `bridge`, then on one side of the event the `asset_class` is a `WRP` node or a receipt class (`LST`, `DFI.LEND`, `DFI.LP`), and on the other side it is a `WRP` node, a receipt class or the class of the underlying asset.

C4. `event_type` ∈ {`lp_deposit`, `lp_withdraw`} ⇒ one side of the event is `DFI.LP`.

C5. `event_type` = `staking_reward` ⇒ `asset_class` ∈ {`NAT`, `LST`}, or `UNKNOWN`.

C6. `asset_class` = `PRV` ⇒ `privacy_enhanced` ≠ `no`.

C7. `asset_class` ∈ `STB.FIAT` ∪ `TKN` ⇒ `issuer_freezable` is read from the issuer contract, not predicted. (This is a reader rule; a predicted value for these classes is itself a violation.)

C8. `counterparty_type` = `self_hosted` ⇒ `touches_unhosted` = `yes`.

C9. `event_type` = `gift` ⇒ `beneficial_owner_change` = `yes`.

C10. `event_type` = `bridge` ⇒ `crosses_chain` = `yes`.

C11. `passes_bridge` = `yes` ⇒ `crosses_chain` = `yes`, which follows from the definition of `passes_bridge` in Section 5.7.

6.11 From labels to rules

Table 43 shows which Paper V flag families consume each question. A question that feeds `block`-capable families (`SAN`, some `LIC` and `AST`) carries the highest error cost. For those families, Paper V requires that a `block` never rests on a model prediction alone.

Question	Consumed by (Paper V families)	Typical use
Q1 `asset_class`	AST, STB, TAX, REP, TRV	securities and privacy-coin flags; stablecoin recognition; transaction-tax scope; CARF scope
Q2 `event_type`	TAX, REP, TRV	disposal versus income versus non-event; reporting triggers
Q3 `economic_purpose`	TAX, FXC, DAT	payroll and distribution rules; FX purpose codes; missing-document flags
Q4 `beneficial_owner_change`	TAX, TRV, FXC	disposal analysis; same-person rules; self-custody as FX
Q5 `counterparty_type`	TRV, LIC, AML	travel-rule and sunrise handling; review of unmatched providers
Q6 `evidence_completeness`	DAT	missing-evidence flags; review queue
Q7 `review_priority`	— (queue only)	ordering of human review
Q8 `travel_surface`	SAN, AML, STB, TRV, AST	freeze capability; mixer and bridge exposure; self-hosted endpoints; transfer restrictions

Table 43: Which rule families consume each Zlef question. Rules also take jurisdiction context (residency, provider location, dates, amounts) as layer A, B, C or F facts, never as predictions.

6.12 What the schema deliberately leaves out

The following are *not* questions in Edition 1, and the reasons matter:

- **Residency, nationality, entity status and jurisdiction.** These are facts about people and firms that must be attested or reviewed. A model that inferred them from wallet behaviour, IP or language would be making exactly the inference Paper IV forbids.
- **Sanctions status.** A sanctions hit is a list match against a versioned list (layer B), not a probability. Zlef may help with *exposure* (Q8 indicators for mixers or bridges), but it never decides whether a party is designated.
- **Legal category, tax treatment, flags and severity.** These are rules-layer outputs (P1).
- **Fair value.** Valuation is a pricing-data problem with its own sources and rules (for example basis methods), not a classification.

6.13 Changes from Paper IV

Paper IV gave an illustrative schema to be replaced by the taxonomy chosen with the first buyer. Edition 1 changes it in four ways.

- **Relationship.** Paper IV's "relationship" question (stand-alone, refund of x, reversal of x, instalment of x) is not a Zlef question in Edition 1. A link to an earlier payment is recorded as a linked-record fact, read from references in the records, and `refund_reversal` plus purpose `refund` carry the event and purpose.
- **Evidence sufficient for purpose.** Paper IV's yes/no question is subsumed by `evidence_completeness` (Q6); sufficiency for a given rule is a threshold the rules layer applies to Q6.
- **Purpose values.** Edition 1 adds three values to Paper IV's purpose list: `goods_purchase`, `investment` and `remittance_family`.
- **Mechanics.** Paper IV's rule that transaction mechanics are read from records, not predicted, is generalised as principle P2 (read before predicting). Mechanics questions (`asset_class`, `event_type`, `travel_surface`) are answered by readers wherever records settle them, and the model is used only for the residue that records leave undetermined.

The edition log (Appendix B) records these changes.

6.14 Versioning

The schema is versioned with the document. A change to a domain (a new `asset_class` child, a new event) is a schema version. Every label and prediction stores the schema version under which it was made, and the edition log (Appendix B) records the change and its reason. Labels under an old schema are mapped to the new one by an explicit, reviewed mapping, or retired. They are never silently relabelled.

7 Modeling Methodology

[Product, off-protocol] This section specifies how Zlef would be modelled. It states the problem formally (Section 7.1), compares candidate architectures (Section 7.2), and gives a loss for each

question type (Section 7.3). It then describes post-hoc calibration (Section 7.4), three abstention mechanisms plus back-off to a coarser class (Section 7.5), and how to derive cost matrices from flag severities (Section 7.6). The final subsections cover consistency across questions (Section 7.7), explanations (Section 7.8) and adversarial text (Section 7.9). Nothing in this section has been implemented. Every choice between alternatives is an open comparison, to be settled by the preregistered evaluation in Section 8. We give no numerical thresholds, weights or costs. Where an equation needs a constant, the constant is a placeholder to be agreed with the first buyer and fixed before the experiment that uses it.

7.1 Problem statement

Inputs. A payment record is represented as a set of normalized facts

$$x = \{(n_i, v_i, \lambda_i, t_i)\}_{i=1}^{m}, \tag{14}$$

where n_i is a field name from a fixed vocabulary, v_i its value, $\lambda_i \in \{\mathrm{A}, \mathrm{B}, \mathrm{C}\}$ its fact layer, and t_i the time at which the fact became available to ZETTL. Layer A facts are protocol-authenticated. Layer B facts are attested by a source. Layer C entries are explicit absences: the field is named, and its value is the marker "absent". The model never sees layer D, E or F facts as inputs. Layer D is the model's own output. Layer E is a rule match. Layer F is a reviewed conclusion. Feeding any of them back would leak the label or the rules into the model (Section 2.10). The layer tag is part of the input, so a source-attested value is never presented as protocol-authenticated, as Paper IV requires. For prediction at time t^*, only facts with $t_i \leq t^*$ are admissible. Free-text fields, such as memos, invoice text and token names, are facts like any other. Each carries its layer, plus a provenance mark that says who could have written it (Section 7.9).

Questions. The questions form a fixed set $\mathcal{Q} = \{q_1, \ldots, q_8\}$, the Edition 1 Zlef schema. Each question q has a typed domain $\mathcal{Y}_q$ and a structure type κ_q:

- q_1 `asset_class`: nominal and hierarchical. $\mathcal{Y}_{q_1}$ is the set of nodes of a tree $\mathcal{T}$ whose root has children `NAT`, `STB`, `TKN`, `GOV`, `MEME`, `NFT`, `LST`, `WRP`, `DFI`, `PRV`, `CBDC`, `OTH` and `UNKNOWN`. Some of these have children, such as `STB.FIAT.SINGLE` or `WRP.NATIVEBURN`.
- q_2 `event_type`: nominal, 24 values plus `UNKNOWN`, from the payer's or holder's perspective.
- q_3 `economic_purpose`: nominal, 10 values plus `UNKNOWN`.
- q_4 `beneficial_owner_change`: `yes`, `no` or `unknown`.
- q_5 `counterparty_type`: nominal, 7 values plus `UNKNOWN`.
- q_6 `evidence_completeness`: ordinal, levels 0 to 4.
- q_7 `review_priority`: ordinal, `low` $<$ `medium` $<$ `high` $<$ `urgent`.
- q_8 `travel_surface`: multi-label over seven indicators, `crosses_chain`, `passes_bridge`, `issuer_freezable`, `permissioned_transfer`, `privacy_enhanced`, `touches_mixer` and `touches_unhosted`. Annotation labels take values in $\{\texttt{yes}, \texttt{no}, \texttt{unknown}\}^7$. The model outputs seven Bernoulli marginals $\Pr(\texttt{yes})$, written $y_s \in \{0, 1\}$ with $1 = \texttt{yes}$ where a label is known; `unknown` annotations are masked out of the loss (Section 7.3.4).

Every label is a jurisdiction-free economic or technical fact. No question asks for a legal category, a tax treatment or a flag. That is the central design decision of the schema.

Outputs. For each question q, the model returns a distribution and the system returns a decision:

$$p_\theta(\cdot \mid x, q) \in \Delta(\mathcal{Y}_q), \qquad \delta_q(x) \in \mathcal{Y}_q \cup \mathcal{V}_q^{\text{int}} \cup \{\perp\}. \tag{15}$$

Here $\Delta(\mathcal{Y}_q)$ is the probability simplex over the domain. For q_8 it is replaced by seven marginal probabilities. $\mathcal{V}_q^{\text{int}}$ is the set of internal nodes of the hierarchy, which is non-empty only for q_1. $\perp$ means abstain. The model and the decision are kept apart. The model produces p_θ, which is a layer D fact. A separate, versioned decision policy maps p_θ, a cost matrix and a review budget to δ_q. The policy can change without retraining (Section 2.8).

`UNKNOWN` is not abstention. The schema value `UNKNOWN` (or `unknown` for q_4) is a *label*. It asserts that the facts in x do not determine the answer, and a careful annotator who saw the same facts would choose it. Abstention $\perp$ is a *decision*. It says that the model is not reliable enough on this input for its answer to be used without review. The two can come apart. A model can be confident that the answer is `UNKNOWN`, for example when there is no memo, no invoice and an unlabelled counterparty. A model can also be unsure between two definite labels. Collapsing the two notions would hide the difference between missing evidence, which calls for a request for documents and a `DAT`-family flag, and model weakness, which calls for review. The evaluation reports them separately (Section 8.6).

Read before predicting. Some answers are determined by layer A or B facts. A token contract found in a maintained registry fixes `asset_class`. A record that shows two chain identifiers fixes `crosses_chain`. The count of attached documents fixes most of `evidence_completeness`. Paper IV states that transaction mechanics are read from records, not predicted. We adopt a "read before predicting" rule. For each question and record, a deterministic reader $r_q(x)$ returns either a value, tagged with the layer of the facts it used, or "undetermined". The model is queried only when the reader returns "undetermined". The evaluation reports how much of each question is answered by readers and how much by the model, and it scores the model only on the residual. This keeps model error from being diluted by easy cases that need no model.

The rules layer. Let J be the jurisdiction context: the jurisdiction set, dates, amounts and residency facts, all of which are layer A, B, C or F facts and never predictions. Let $y = (y_{q_1}, \ldots, y_{q_8})$ be the vector of answers. The rules layer is a versioned function

$$\mathcal{R}_\rho(y, J) = \{(f, \text{sev}(f))\}, \qquad \text{sev}(f) \in \{\text{block}, \text{review}, \text{inform}\}, \tag{16}$$

which returns the set of flags f that fire, with their severities, under rule version ρ. The flags follow the Paper V taxonomy: `LIC`, `TRV`, `SAN`, `AML`, `TAX`, `REP`, `FXC`, `STB`, `AST` and `DAT`. Zlef never evaluates $\mathcal{R}$. The cost analysis in Section 7.6 uses $\mathcal{R}$ only to say how much each model error matters.

7.2 Candidate architectures

We compare five architectures, A0 to A4. All share the input representation of Equation (14), the output contract of Equation (15) and the downstream decision policy. They differ in how they produce p_θ.

Table 44: Edition 1 questions and the methodology proposed for each. All entries are proposals to be compared, not decisions.

Question	Structure	Training loss	Calibrator	Abstention
`asset_class`	nominal, tree	marginal CE, soft labels, Brier	temperature; Dirichlet per class	cost threshold, conformal, back-off to parent
`event_type`	nominal (25)	label-smoothed CE plus Brier	temperature; Dirichlet	cost threshold; RAPS sets
`economic_purpose`	nominal (11)	label-smoothed CE plus Brier	temperature; Dirichlet	cost threshold; APS sets
`beneficial_owner_change`	ternary	CE plus Brier	temperature	cost threshold; consistency check
`counterparty_type`	nominal (8)	label-smoothed CE plus Brier	temperature; Mondrian checks	cost threshold; class-conditional sets
`evidence_completeness`	ordinal (5)	CORAL or RPS; RLCD-style a hypothesis	isotonic on cumulative	reader first; threshold on $\Pr(y \geq j)$
`review_priority`	ordinal (4)	CORAL or RPS; RLCD-style a hypothesis	isotonic on cumulative	threshold on $\Pr(y \geq j)$
`travel_surface`	multi-label (7)	per-label BCE plus Brier	Platt or isotonic per label	reader first; per-label threshold

A0: rules only. Hand-written rules map facts to answers. An example is "if the counterparty address is in the registry of regulated providers, then `regulated_vasp`". A0 has no probabilities. It either answers or returns `UNKNOWN`. We include it as the floor against which every learned model must justify its cost, and as the deterministic reader r_q of Section 7.1. If A0 answers a question well, no model is needed for it.

A1: a general language model with constrained output. A general instruction-tuned model receives a serialized x, the question and the allowed values. It produces an answer constrained to the schema [97, 229]. There are two ways to obtain probabilities. In the *sampling* variant, the model generates a constrained JSON answer, and probabilities are estimated from the frequencies of repeated samples or from token probabilities along the path. These are not class probabilities and need calibration (Section 2.11). In the *option-scoring* variant, each allowed value k is written as a string s_k, and

$$p(k \mid x, q) = \frac{\exp\left(\ell_k / \tau\right)}{\sum_{k'} \exp\left(\ell_{k'} / \tau\right)}, \qquad \ell_k = \log P_{\mathrm{LM}}(s_k \mid \mathrm{prompt}(x, q)), \tag{17}$$

with τ a temperature fitted afterwards. Option strings of different token lengths get different log-probabilities for reasons unrelated to the answer. Length normalisation, or single-token option codes, must therefore be fixed in advance and reported. Option order must be permuted in tests (Section 2.11). A1 is the baseline that tests whether a specialised model adds anything beyond schema validity.

A2: an external decision model. Clef, Clef-flash or Jev is called with Zlef's questions as `choice`, `noul` and `score` questions (Section 2.1). Hierarchical `asset_class` is posed as a flat `choice` over

Table 45: Candidate architectures. "Class probabilities" means a distribution over the allowed values obtained in one pass, before calibration.

Architecture	Class probabilities	Data leaves customer	Main use	Main risk
A0 rules	none	no	floor; deterministic reader	brittle; no uncertainty
A1 general model, constrained	option scoring only	depends on host	baseline for schema validity	length and order bias
A2 Clef or Jev	yes (vendor softmax)	yes for API; no if Clef self-hosted	external baseline	vendor change; no abstention
A3 open base plus heads	yes	no if self-hosted	primary candidate	training cost; label demand
A4 shared encoder	yes	no if self-hosted	candidate if A3 is too costly	coupled calibration

leaves plus internal nodes, or as several questions, one per level. Multi-label `travel_surface` is posed as seven `noul` questions. A2 is an [External dependency] baseline. It is used under the hosting and data-permission conditions in Paper IV, and we use only its `probabilities`, never its `confidence` field. Its serving configuration, including model version, hosted or self-hosted mode and context limit, is pinned and logged with every run, because the vendor changed one such limit after launch.

A3: an open-weight base with a head per question. A frozen open-weight backbone ϕ encodes the serialized record and question. Low-rank adapters ψ are trained, as in Clef, and each question has its own head:

$$h = \phi_\psi\big(\mathrm{enc}(x), q\big), \qquad z_q = W_q h + b_q, \qquad p_\theta(\cdot \mid x, q) = \mathrm{softmax}(z_q). \tag{18}$$

For ordinal and multi-label questions, the softmax is replaced by the heads of Section 7.3.3 and Section 7.3.4. A3 can run entirely on infrastructure the customer controls. Its heads can be retrained when the schema changes without touching the backbone. Each head can be calibrated and audited separately. Clef's weights under Apache-2.0 are one possible backbone. Any other open-weight model is another.

A4: a shared encoder for several questions at once. One encoder pass produces h, and all heads read it together. The questions are therefore answered from a single representation. A4 can also add explicit interaction between answers, for example by feeding the logits of one question into the head of another, or by the joint decoding of Section 7.7. The gain is cost and coherence: one pass per record, and answers that can share evidence. The risk is coupling. A change to one question's training data can shift another question's calibration, so every head must be re-calibrated whenever any head is retrained.

Hypotheses. [Research hypothesis] A3 or A4 gives better class-wise calibration and a better risk–coverage trade-off than A1 and A2 on ZETTL data, after each is post-hoc calibrated. [Research hypothesis] The gain, if any, is large enough to justify the cost of training. Paper IV's research question RQ4 asks the same thing. Neither hypothesis is supported by evidence yet.

7.3 Losses per question type

All losses are computed over a training set $\mathcal{D} = \{(x_i, Y_i)\}$, where Y_i holds the labels that exist for record i. Many records will have labels for only some questions. The total loss is a masked, weighted sum:

$$\mathcal{L}(\theta) = \sum_{q \in \mathcal{Q}} w_q \, \frac{1}{|\mathcal{D}_q|} \sum_{i \in \mathcal{D}_q} \mathcal{L}_q\big(p_\theta(\cdot \mid x_i, q),\, y_{iq}\big), \tag{19}$$

where $\mathcal{D}_q$ is the set of records with a label for q and w_q is a weight per question, to be chosen on validation data. Labels come with provenance (Section 8.2). Records labelled only by a reviewer who saw the model's suggestion may be down-weighted or excluded. That choice is fixed before training.

7.3.1 Nominal questions

For `event_type`, `economic_purpose`, `beneficial_owner_change` and `counterparty_type`, we follow Clef's published objective, label-smoothed cross-entropy plus a Brier term [29]. With K classes, smoothing ε and uniform vector u, the smoothed target is $\tilde{e}_y = (1 - \varepsilon)e_y + \varepsilon u$ [204], and

$$\mathcal{L}_{\text{nom}}(p, y) = -\sum_{k=1}^{K} \tilde{e}_{y,k} \log p_k \;+\; \beta \sum_{k=1}^{K} \big(p_k - \mathbb{1}[y = k]\big)^2. \tag{20}$$

The Brier term uses the unsmoothed target, so it remains a proper score [19, 98]. The smoothing ε and weight β are hyperparameters. Cloudflare does not publish its values, and we do not invent any. Smoothing can improve calibration but erases information about which wrong classes are close [150]. Post-hoc calibration (Section 7.4) is applied regardless, and the ablation without smoothing is part of the evaluation.

Two refinements apply to nominal questions. First, `UNKNOWN` is an ordinary class in training, but its labels need care. An annotator chooses it only when the guidelines say the facts do not determine the answer (Section 8.2). It is not a dumping ground for hard cases. Second, rare classes with large downstream cost may be up-weighted in training. Re-weighting changes the learned class priors, so the model's probabilities are no longer calibrated to the true class frequencies. If we re-weight, post-hoc calibration must undo the effect, and evaluation must confirm that it did. Our default is not to re-weight, and to put costs in the decision layer instead (Section 2.8).

7.3.2 Hierarchical question: `asset_class`

Let Λ be the set of leaves of $\mathcal{T}$. A top-level class with no children, such as `GOV`, is its own leaf. The model outputs a distribution p over leaves. Every node v then gets a probability by marginalisation:

$$p(v \mid x) = \sum_{\ell \in \text{leaves}(v)} p(\ell \mid x). \tag{21}$$

For example, $p(\texttt{STB} \mid x)$ is the sum of the probabilities of `STB.FIAT.SINGLE`, `STB.FIAT.BASKET`, `STB.COMM`, `STB.SYN` and `STB.YLD`. Marginal probabilities are automatically consistent: a parent is at least as probable as any of its children.

Partial labels. Annotators will sometimes be able to say "stablecoin" but not which kind, because the reserve disclosure is missing. Such a label sits at an internal node v. Marginalisation gives it a natural loss:

$$\mathcal{L}_{\text{part}}(p, v) = -\log p(v \mid x) = -\log \sum_{\ell \in \text{leaves}(v)} p(\ell \mid x). \tag{22}$$

This rewards any distribution of mass among the children of v and penalises only mass outside v. A label at a leaf reduces to ordinary cross-entropy. The model therefore learns from coarse labels without being forced to guess the fine one.

Hierarchy-aware soft labels. In place of uniform smoothing, the target for a leaf label y spreads mass according to tree distance, following the idea of Bertinetto et al. [16]:

$$\tilde{e}_{y,\ell} = \frac{\exp\big(-\gamma\, d_{\mathcal{T}}(y, \ell)\big)}{\sum_{\ell'} \exp\big(-\gamma\, d_{\mathcal{T}}(y, \ell')\big)}, \tag{23}$$

where $d_{\mathcal{T}}$ is a distance on the tree, such as the height of the lowest common ancestor, and $\gamma > 0$ controls how fast mass decays. With large γ, this approaches a one-hot target. With smaller γ, `STB.FIAT.SINGLE` shares more mass with `STB.FIAT.BASKET` than with `MEME`. The tree distance is a proxy. What matters for Zlef is the downstream consequence of confusing two classes, which Section 7.6 measures directly. We therefore treat $d_{\mathcal{T}}$ as a starting point and compare it with a distance derived from the cost matrix.

Level weighting. The loss can be written as a sum of conditional terms along the path from the root to y. If $v_0, v_1, \ldots, v_L = y$ is that path, then

$$\mathcal{L}_{\text{HXE}}(p, y) = -\sum_{l=1}^{L} \omega_l \log \frac{p(v_l \mid x)}{p(v_{l-1} \mid x)}, \tag{24}$$

with weights ω_l per level [16]. With all $\omega_l = 1$, this equals ordinary cross-entropy on the leaf. Giving more weight to upper levels penalises confusions between top-level classes, such as `TKN` against `GOV`, more than confusions within a class, such as `STB.FIAT.SINGLE` against `STB.FIAT.BASKET`. That is the direction a cost analysis is likely to favour, but it is to be tested, not assumed. The full hierarchical loss is Equation (24) or Equation (23) plus the Brier term of Equation (20) on leaves.

Back-off. Marginalisation makes partial answers possible. If no leaf is confident enough to act on, but a parent is, the decision policy can return the parent. Answering `STB` instead of `STB.SYN` is the canonical example. Section 7.5 gives the rule. For the rules layer to use a partial answer, each rule that reads `asset_class` must say what it does for an internal node. We propose a conservative convention. A rule fires on internal node v at the highest severity it would have for any leaf under v, downgraded to `review` where that severity is `block`. If the leaves under v disagree, a `DAT`-family flag asks for the missing fact. This convention belongs to the rules layer, and the Paper V rule owners must agree to it.

7.3.3 Ordinal questions

For `evidence_completeness` ($K = 5$) and `review_priority` ($K = 4$), we compare three objectives.

CORAL. Following Cao et al. [22], the head outputs $K - 1$ cumulative probabilities $\pi_j = \Pr(y > j \mid x) = \sigma(w^\top h + b_j)$ with shared w. The loss is a sum of binary cross-entropies:

$$\mathcal{L}_{\text{CORAL}}(\pi, y) = -\sum_{j=1}^{K-1} \Big[\mathbb{1}[y > j] \log \pi_j + \big(1 - \mathbb{1}[y > j]\big) \log(1 - \pi_j) \Big]. \tag{25}$$

Level probabilities are $p_1 = 1 - \pi_1$, $p_j = \pi_{j-1} - \pi_j$ for $1 < j < K$, and $p_K = \pi_{K-1}$. These are non-negative when the π_j are monotone, which the shared weights guarantee if the biases are ordered.

RPS. A softmax head over the K levels is trained with the ranked probability score of Equation (6) [98]. RPS is proper and penalises mass more the further it sits from the true level. It is the principled form of "partial credit".

An RLCD-style objective, as a hypothesis only. Cloudflare describes RLCD as granting "partial credit to adjacent ordinal choices", together with a reward for "fully precise record outputs" and "a reference penalty to prevent distribution shift" [29]. It publishes no formula. ZETTL has not implemented RLCD and cannot reproduce it. To make the hypothesis testable, we write down one instantiation of our own. It is *not* Cloudflare's objective. A policy π_θ outputs a level $\hat{y}$ and a probability vector p. With a partial-credit function $c(\hat{y}, y) \in [0, 1]$ that equals 1 when $\hat{y} = y$ and decreases with $|\hat{y} - y|$, a reference policy π_{ref}, and coefficients $\alpha, \eta \geq 0$, the objective is

$$J(\theta) = \mathbb{E}_{x, \hat{y} \sim \pi_\theta} \Big[c(\hat{y}, y) - \alpha \, S_{\text{RPS}}(p, y) \Big] - \eta \, \mathbb{E}_x \Big[\text{KL}\big(\pi_\theta(\cdot \mid x) \,\big\|\, \pi_{\text{ref}}(\cdot \mid x)\big) \Big]. \tag{26}$$

The first term is partial credit. The second is a bounded proper score. RLCR's calibration incentive is proved for a bounded proper score on a scalar confidence in correctness [51]; our use of RPS on a level distribution alongside a sampled level is an analogy, not covered by that theorem, and its incentive properties must be checked. RPS on K levels is bounded by $K - 1$. The third is a reference penalty. Three cautions apply. First, the partial-credit term alone does not reward calibration. A policy maximising it reports the level that minimises expected distance, not its true belief. The calibration incentive comes only from the proper-score term. Second, the reference penalty pulls π_θ toward π_{ref} and can move the optimum away from calibration. Its effect must be measured, not assumed. Third, Paper IV's research question RQ3 asks whether an analogous objective helps nominal questions, and what "adjacent" would mean there. For `asset_class`, tree distance or cost-matrix distance is a candidate. For flat nominal questions, there is no obvious adjacency, and we do not propose one. Paper IV's ordering of options holds: an RLCD-style objective is considered only after rules, prompts, calibration of existing models and supervised fine-tuning, and only if a measured gap justifies the spending.

Thresholds on ordinal outputs. Ordinal outputs are used through cumulative probabilities. A policy such as "send to senior review if $\Pr(\texttt{review_priority} \geq \texttt{high}) > \tau$" needs $\Pr(y \geq j)$ to be calibrated for each j. That is a set of binary calibration problems, handled in Section 7.4. Note that `review_priority` is partly circular. Priority depends on which flags fire, and flags depend on other answers. In Zlef it should capture only what the rules layer cannot compute, such as signs of urgency in the record. The final queue priority is a policy that combines the model's `review_priority` with the severities of the fired flags.

7.3.4 Multi-label question: `travel_surface`

Each of the seven labels $s \in \mathcal{S}$ gets a binary head with probability $p_s = \sigma(z_s)$. The loss is per-label binary cross-entropy plus a Brier term:

$$\mathcal{L}_{\text{ML}}(p, y) = \sum_{s \in \mathcal{S}_x} \Big[-y_s \log p_s - (1 - y_s) \log(1 - p_s) + \beta \, (p_s - y_s)^2 \Big], \tag{27}$$

where $\mathcal{S}_x \subseteq \mathcal{S}$ is the set of labels that the deterministic reader left undetermined for record x *and* that the annotator labelled `yes` or `no`. Indicators annotated `unknown` are masked out: they contribute no loss, so "not observable" is never trained as "no". Several labels are often determined by records. `crosses_chain` follows from chain identifiers, and `issuer_freezable` or `permissioned_transfer` may follow from contract metadata. Others, such as `touches_mixer`, depend on upstream graph analysis, which runs outside the model (Section 2.12). The model is then a fallback, and its label must say so. Labels are not independent. `passes_bridge` and `crosses_chain` are related, for instance. Per-label heads ignore this, and the dependency is handled by the consistency layer (Section 7.7). Each label is calibrated and abstains separately.

7.4 Post-hoc calibration

Calibration is fitted after training, on a calibration split that is disjoint from training, from conformal calibration and from test data (Section 8.5). Calibration is fitted per question, and where data allow, per class. We compare the following maps.

Per question. Temperature scaling fits one scalar T_q per question by minimising negative log-likelihood on the calibration split [102]:

$$p^{\text{cal}}(k \mid x, q) = \text{softmax}\big(z_q / T_q\big)_k. \tag{28}$$

It keeps the predicted class and is the cheapest baseline.

Per class. Dirichlet calibration fits a linear map on log-probabilities with off-diagonal regularisation [134]:

$$p^{\text{cal}}(\cdot \mid x, q) = \text{softmax}\big(A_q \log p(\cdot \mid x, q) + c_q\big). \tag{29}$$

It can correct a rare class that the model systematically over- or under-predicts. The number of parameters grows with the square of the number of classes. For `event_type`, that requires either strong regularisation or restriction to a diagonal A_q, which is equivalent to vector scaling. For binary labels, including each `travel_surface` label and each cumulative threshold of an ordinal question, we compare Platt scaling [181] with isotonic regression [233]. Isotonic regression is reserved for labels with enough calibration data.

Ordinal outputs. The $K-1$ cumulative probabilities are calibrated one by one. Separately calibrated thresholds can then cross, so that $\Pr(y \geq j+1) > \Pr(y \geq j)$. We restore monotonicity by isotonic projection across j and report how often it was needed.

Hierarchy. We calibrate the leaf distribution and then marginalise by Equation (21). Marginals of calibrated leaves need not be calibrated at the parent level, because calibration is a statement about conditional frequencies, and summing does not preserve it in general. We therefore measure calibration at every level that the decision policy uses. If a parent is miscalibrated, we fit a separate map for that node's marginal and use it for back-off decisions only.

Thin classes. Some classes, such as `CBDC` or `LST.RESTAKE`, may have too few calibration examples for a per-class map. For these, we pool classes with similar score distributions, as clustered conformal prediction does [55], or fall back to the per-question temperature. The calibration report must state which classes were calibrated individually and which were pooled.

When to recalibrate. Calibration is a property of a model, a serving configuration and a data distribution together. It is re-fitted and re-measured after any change to the backbone, adapters, heads, prompt template, option strings or serving configuration, and on the drift triggers of Section 8.7. Label shift alone can be corrected without new labels by BBSE-style prior re-estimation [138]. The new priors $\hat{w}$ from Equation (13) re-weight the posteriors: $p'(k \mid x) \propto \hat{w}_k \, p(k \mid x)$. This correction assumes that $p(x \mid y)$ is unchanged, which a new asset class or a new scam pattern violates.

7.5 Abstention and back-off

We compare three abstention mechanisms, M1 to M3, plus back-off to a coarser hierarchy level, which applies to `asset_class`. All three share one interface: for each question they return a decision $\delta_q \in \mathcal{Y}_q \cup \mathcal{V}_q^{\text{int}} \cup \{\perp\}$, and $\perp$ sends the question to review. None of them uses a vendor confidence statistic such as Equation (1).

M1: `UNKNOWN` as an explicit class. The decision is the most probable class, with `UNKNOWN` among the candidates:

$$\delta_q^{\text{M1}}(x) = \arg\max_{k \in \mathcal{Y}_q} p^{\text{cal}}(k \mid x, q), \tag{30}$$

and the item goes to review when the decision is `UNKNOWN`. M1 is the simplest mechanism and was one of the two mechanisms in Paper IV. Its weakness is the one explained in Section 7.1. `UNKNOWN` is a label about the evidence, so M1 abstains when evidence is missing but not when the model is torn between two definite labels. M1 does not respond to costs either.

M2: a cost-based threshold. Let C_q be the cost matrix for question q, with $C_{q,jk}$ the cost of deciding j when the truth is k (Section 7.6), and let $c_{r,q}$ be the cost of sending the question to review. Following Chow and Elkan [31, 64], the decision with least expected cost is

$$\delta_q^{\text{M2}}(x) = \begin{cases} \arg\min_j R_j(x) & \text{if } \min_j R_j(x) \leq c_{r,q}, \\ \perp & \text{otherwise,} \end{cases} \qquad R_j(x) = \sum_k p^{\text{cal}}(k \mid x, q)\, C_{q,jk}. \tag{31}$$

With unit error costs, this reduces to Chow's rule in Equation (8). M2 is optimal if the probabilities are calibrated and the costs are right. It is therefore only as good as the calibration, which is measured, and the costs, which are agreed. To add a guarantee, we pair M2 with selection with guaranteed risk [95]. On the calibration split, among thresholds that keep the bound on selective risk below a preregistered target, we pick the one with the highest coverage. If no threshold meets the target, the question is not automated.

M3: conformal prediction sets. On a separate conformal calibration split, we build a set $C_\alpha(x)$ for each question, with APS scores [187], RAPS for `event_type` [4], and Mondrian calibration for classes whose consequences are large [225]. The decision is

$$\delta_q^{\mathrm{M3}}(x) = \begin{cases} k & \text{if } C_\alpha(x) = \{k\} \text{ and } k \neq \texttt{UNKNOWN}, \\ \mathrm{LCA}\big(C_\alpha(x)\big) & \text{if } q = q_1,\ |C_\alpha(x)| > 1,\ \texttt{UNKNOWN} \notin C_\alpha(x),\ \mathrm{LCA} \neq \text{root}, \\ \bot & \text{otherwise.} \end{cases} \tag{32}$$

In words, M3 answers only when the set is a single definite class. For `asset_class`, a set whose members share a non-root ancestor returns that ancestor. Under exchangeability, the marginal guarantee Equation (9) implies that a singleton answer is wrong at most a fraction α of the time on average. M3 does not need calibrated probabilities to keep its coverage guarantee, which makes it robust to calibration errors. The guarantee is marginal, and it fails under shift. Class-conditional coverage on high-stakes classes and weighted conformal calibration under covariate shift address these limits in part [55, 211].

Conformal risk control. M3 controls miscoverage, but Zlef's real concern is cost. Conformal risk control [5] lets us bound expected cost instead. Let λ index a family of decision rules that become more cautious as λ grows, such as the M2 rule with review cost $c_r = \lambda$. Let $L_i(\lambda)$ be the realised cost on calibration record i, counting a review as cost zero for this purpose, so that L_i is non-increasing in λ. If the cost is bounded by B, choosing

$$\hat{\lambda} = \inf\Big\{\lambda : \frac{n}{n+1}\hat{R}_n(\lambda) + \frac{B}{n+1} \leq r^*\Big\}, \qquad \hat{R}_n(\lambda) = \frac{1}{n}\sum_{i=1}^{n} L_i(\lambda), \tag{33}$$

bounds the expected cost of automated decisions on a new exchangeable record by r^*. The target r^* is a cost-policy choice, to be agreed. The review volume this implies is then measured, not chosen.

Back-off to a coarser level. For `asset_class`, abstention need not be all or nothing. Define the expected cost of answering with node v as $R_v(x) = \sum_k p^{\mathrm{cal}}(k \mid x) C_{vk}$, where row $C_{v\cdot}$ is the cost of a partial answer at v, derived in Section 7.6. The back-off decision is

$$\delta^{\mathrm{BO}}(x) = \arg\min_{v \in \mathcal{Y}_{q_1} \cup \mathcal{V}_{q_1}^{\mathrm{int}} \cup \{\bot\}} R_v(x), \qquad R_\bot(x) = c_{r,q_1}. \tag{34}$$

A rule-of-thumb version picks the deepest node v with $p^{\mathrm{cal}}(v \mid x) \geq \tau_{\mathrm{depth}(v)}$, using thresholds per level. The cost version is preferred because it uses the same costs as the rest of the policy. When back-off returns `STB`, the rules layer receives `STB`. Under the convention of Section 7.3.2, every flag that would fire for some stablecoin leaf then fires at review severity or below, and a `DAT`-family flag asks for the reserve or mechanism fact that would decide the leaf. A partial answer thus sends the case to a narrower, faster review rather than a full one.

Table 46: Abstention mechanisms compared. All are to be evaluated on the same splits and costs (Section 8.6).

Mechanism	Needs	Guarantee	Main weakness
M1 `UNKNOWN` class	labels that use `UNKNOWN` consistently	none	ignores costs and model uncertainty between definite labels
M2 cost threshold	calibrated p; agreed costs; review cost	optimal if p and costs are right; SGR bound on selective risk	depends on calibration and on costs
M3 conformal set	exchangeable calibration data	marginal coverage $1-\alpha$; class-conditional if Mondrian	marginal only; fails under shift
Conformal risk control	bounded monotone cost	expected automated cost $\leq r^*$	needs bounded cost; marginal
Back-off (`asset_class`)	rules defined on internal nodes	inherits from M2 or M3	rule owners must agree the convention

Combining mechanisms. The mechanisms can be stacked, so that an item goes to review if any one of them abstains. We regard the combination as a candidate policy to evaluate, not a default, because stacking lowers coverage, and the extra review must be shown to buy a reduction in cost. Paper IV's research question RQ2 compared M1 with a reject threshold. This edition widens the comparison to M1, M2, M3, conformal risk control and back-off. [Research hypothesis] M2 or M3 dominates M1 on the risk–coverage curve weighted by cost.

7.6 Cost matrices from flag severity

A wrong answer matters only through what the rules layer does with it. We therefore derive each question's cost matrix from the flags that the rules would fire under the true answer and under the predicted one. We give the construction and no numbers. Severity weights, review costs and the jurisdiction mix are policy inputs, agreed with the first buyer and versioned with the rules.

Flag difference. Fix a question q, a jurisdiction context J and the other answers y_{-q}. For a candidate answer j and true answer k, let

$$\mathcal{M}_{jk} = \mathcal{R}(y_q{=}k) \setminus \mathcal{R}(y_q{=}j), \qquad \mathcal{S}_{jk} = \mathcal{R}(y_q{=}j) \setminus \mathcal{R}(y_q{=}k), \tag{35}$$

where $\mathcal{R}(y_q{=}k)$ abbreviates $\mathcal{R}_\rho((k, y_{-q}), J)$. $\mathcal{M}_{jk}$ holds the flags that are *missed* because the model answered j. $\mathcal{S}_{jk}$ holds the flags that are *spurious*. A flag that fires under both answers but with different severities counts as missed at the higher severity and spurious at the lower one.

Cost. With weights $m(\cdot)$ for missed flags and $s(\cdot)$ for spurious flags, indexed by severity, the cost of answering j when the truth is k in context J is

$$C_{jk}(J, y_{-q}) = \sum_{f \in \mathcal{M}_{jk}} m\big(\mathrm{sev}(f)\big) \; + \sum_{f \in \mathcal{S}_{jk}} s\big(\mathrm{sev}(f)\big). \tag{36}$$

The diagonal is zero. Two orderings are natural, and we propose them as constraints on the weights, not as values. Missing a flag costs at least as much as raising it spuriously at the same severity: $m(\sigma) \geq s(\sigma)$. Severity is ordered: $m(\text{block}) \geq m(\text{review}) \geq m(\text{inform})$, and similarly for s. A missed `block` flag, such as a sanctions hit, is the most expensive error. A spurious `inform` flag is the cheapest. The review cost c_r sits between them. The ratio of c_r to the missed-flag weights sets how much Zlef automates.

Example (illustrative). Suppose the true `asset_class` is `TKN.EQ`, a tokenized equity, and the model answers `GOV`. In a jurisdiction where the Paper V rules attach an `AST`-family flag about transfer restrictions on tokenized securities to `TKN.EQ`, and none to `GOV`, that flag is in $\mathcal{M}_{\texttt{GOV,TKN.EQ}}$. The error therefore costs $m(\text{sev})$ for that flag's severity, plus any other flags that differ. The reverse error, answering `TKN.EQ` when the truth is `GOV`, puts the same flag in $\mathcal{S}$ and costs $s(\text{sev})$. The matrix is asymmetric by construction. The example is illustrative. Whether such a flag exists, and at what severity, is for Paper V's rule set to state, not this section.

Aggregation over contexts. Costs depend on J and on the other answers y_{-q}. Zlef's labels are jurisdiction-free, so one model serves many contexts. We consider two aggregations. The *expected* cost averages Equation (36) over the distribution of jurisdiction contexts and other answers in the customer's traffic:

$$\bar{C}_{jk} = \mathbb{E}_{(J,y_{-q})}\big[C_{jk}(J, y_{-q})\big]. \tag{37}$$

The *context-specific* cost uses $C_{jk}(J, \hat{y}_{-q})$ at decision time, with the record's actual J and the current predictions for the other questions. The second is more accurate per record and keeps the model jurisdiction-free, because only the decision policy sees J. It is also more complex, and it couples questions through $\hat{y}_{-q}$. A conservative variant replaces the expectation in Equation (37) with a maximum over contexts. That variant suits a deployment that must meet the strictest corridor in its traffic.

Costs of partial answers and of review. For back-off to an internal node v of `asset_class`, the rules convention of Section 7.3.2 fixes which flags fire, at which severities. Equation (36) then gives C_{vk} directly. Because a partial answer turns some flags into review flags, C_{vk} usually lies between the cost of a correct leaf and the cost of full review. That is the reason to back off at all. The review cost $c_{r,q}$ is measured, not assumed. Its components are reviewer time, which the evaluation records (Section 8.6), plus the cost of reviewer error, which Dawid–Skene estimates of reviewer accuracy supply (Section 8.2). Review is not error-free, so c_r is not just a time cost.

Properties. The derived matrices are versioned with the rules. When a rule changes, the cost matrix changes and the decision policy changes. The model and its calibration do not, which is Elkan's point [64]. The matrices must satisfy Elkan's conditions for a reasonable cost matrix. No row may dominate another everywhere, or the dominated answer would never be chosen. Where a row does dominate, the rules make two answers indistinguishable, and the schema should merge them. That makes the cost matrix a useful check on the schema itself. If two labels never lead to different flags in any context, Zlef need not distinguish them for compliance purposes, though accounting uses may still require it.

7.7 Consistency across questions

The eight answers describe one payment, so some combinations are impossible. The set of consistency constraints $\mathcal{K}$ is the canonical list C1–C11 in Section 6.10; this section does not restate or extend it. Each constraint is a predicate over the answer vector y. For `travel_surface`, a constraint such as C8 (`touches_unhosted` = `yes`) is read on decisions as $\delta_{q_8,s}$ = `yes` and on the model's marginals as an event of probability p_s. Constraints are versioned with the schema. Each holds only in the stated direction: C1 does not say that every transfer without a change of beneficial owner is internal. Constraints involving `UNKNOWN`, `unknown` or an abstention are satisfied.

Joint decoding. If each question's distribution is treated as independent given x, the distribution restricted to consistent answers is

$$p_{\mathcal{K}}(y \mid x) = \frac{\prod_q p^{\text{cal}}(y_q \mid x, q)\, \mathbb{1}[y \in \mathcal{K}]}{Z_{\mathcal{K}}(x)}, \qquad Z_{\mathcal{K}}(x) = \sum_{y' \in \mathcal{K}} \prod_q p^{\text{cal}}(y'_q \mid x, q). \tag{38}$$

Constrained joint decoding finds the mode of $p_{\mathcal{K}}$, or applies the decision policy to its marginals. Its result is only a *proposal*: it is shown to the reviewer next to the raw outputs and never replaces them. The joint space is large, but each of C1–C11 links only two questions or two labels. Each links two questions or two labels. Exact inference is therefore cheap on the small graph they define, or the mode can be found by enumerating the few consistent combinations of the linked questions. Two cautions apply. First, the independence assumption is wrong for A4, whose heads share a representation, and is only approximate for others. Second, renormalising by $Z_{\mathcal{K}}$ changes the marginals, so marginals of $p_{\mathcal{K}}$ need their own calibration check. A model whose answers conflict is telling us something, and joint decoding can hide it.

Post-hoc check. The alternative leaves the per-question answers alone and measures inconsistency:

$$V_{\mathcal{K}}(x) = 1 - Z_{\mathcal{K}}(x), \tag{39}$$

which is the probability, under the independent product, that the answers violate some constraint. If the decided answers violate a constraint, or if $V_{\mathcal{K}}(x)$ exceeds a threshold to be agreed, the system emits a `DAT`-family flag at `review` severity. The flag names the violated constraint and the facts that would resolve it, for example evidence of common beneficial ownership. It does not silently correct either answer.

Policy. The two tools serve different ends, and the schema fixes how they combine (Section 6.10). Violations are never silently repaired. Whenever the raw decided outputs violate a constraint, the system emits a `DAT`-family flag at `review` severity, whatever flags the violation would affect. Joint decoding may be used to propose a consistent alternative for the reviewer, but the raw outputs and the violation stay on record. The violation rate per constraint is an evaluation metric in its own right (Section 8.6). A high rate suggests that the schema definitions or the labelling guidelines need repair.

7.8 Explanations kept apart from scores

Paper IV commits to keeping explanations apart from scores. We make this concrete with two kinds of explanation material and one rule.

Evidence pointers. An evidence pointer is a reference to an input fact, with its layer and source, that a measured attribution shows mattered to an answer. For a field i, let $x_{\setminus i}$ be the record with field i replaced by its layer-C absence marker. The attribution of field i to question q is the change in a proper score, or in the decision cost, when the field is removed:

$$a_{iq}(x) = \mathrm{KL}\big(p^{\mathrm{cal}}(\cdot \mid x, q) \;\big\|\; p^{\mathrm{cal}}(\cdot \mid x_{\setminus i}, q)\big), \tag{40}$$

or $R_{\delta_q(x)}(x_{\setminus i}) - R_{\delta_q(x)}(x)$ in the cost form. Removal by the layer-C marker has a clear meaning in Zlef: it is how a truly absent fact is presented. It is therefore not an out-of-distribution perturbation. The interface shows the top fields as pointers. An example is "memo text (layer B, counterparty-written) and invoice reference (layer B) moved this answer most". Each pointer links to the record itself. Attributions are computed from the scoring model, so they are faithful to it in the narrow sense of measuring the effect of removing a field. They do not claim to explain *why* the model reacts to that field. Other attribution methods may be compared, but none may be shown without a stated faithfulness test [125].

Generated rationales. A short natural-language rationale may be generated by a separate process, for display only. Its inputs are the record, the answers and the evidence pointers. It is labelled in the interface as a generated summary that may not reflect what drove the answer [215].

The rule. No explanation material, whether pointers or rationales, is ever an input to p_θ, the calibrator, the decision policy, the consistency layer or the rules layer. Concretely, the scoring path is a pure function of x and the versioned model, calibrator and policy. Its inputs are logged, and an audit can confirm that no explanation artefact is among them. Rationales are not generated before scoring, so no chain-of-thought step feeds the score. This is a deliberate departure from designs that reason in text before answering. Its cost, any accuracy lost by not letting the model reason first, is a measurable question, but the rule is not traded for accuracy in Edition 1.

Presentation to reviewers. Explanations can raise acceptance of wrong answers [13]. The review interface therefore shows the calibrated distribution and the evidence pointers first. The rationale stays behind an explicit click. In the blind arm of the anchoring study (Section 8.3), the reviewer sees neither the answer nor any explanation until they have recorded their own answer.

7.9 Adversarial robustness of text fields

Threat model. An adversary controls some free-text inputs: memos, invoice text, token names and symbols, ENS names and NFT metadata. The adversary may be the counterparty, the issuer of a token, or anyone able to attach text to a record. Their goal is to move probability mass to an answer that suppresses a flag, such as `refund` instead of `receive_payment_goods_services`,

`STB.FIAT.SINGLE` instead of `STB.SYN`, or `regulated_vasp` instead of `vasp_unmatched`. They may also try to cause abstention at scale to overload review. Typed outputs remove free-text leakage of the prompt but not manipulation of the answer [101, 173].

Mitigations to evaluate.

D1. *Provenance marking.* Every text fact carries its layer and an author class, such as "counterparty-written", "issuer-written" or "customer-written". The serialization quotes text fields as data, in a fixed structure that the model is trained to treat as content, not instruction.

D2. *Structured facts first.* Where a layer A or B structured fact answers a question, the reader r_q answers it and the model is not consulted (Section 7.1). An asset's class comes from its contract address in a registry, never from a token name that says "USD".

D3. *Corroboration for high-cost decisions.* Let $x_{\setminus \text{text}}$ be the record with all attacker-controllable text replaced by absence markers. If the decision changes when text is removed, $\delta_q(x) \neq \delta_q(x_{\setminus \text{text}})$, and the change moves the expected cost by more than a margin to be agreed, the decision is sent to review. A single text field then cannot flip a high-cost decision without corroboration from structured facts, which is Paper IV's rule in operational form.

D4. *Field caps.* Text fields are truncated to fixed lengths, recorded in the serving configuration, so that a long memo cannot crowd out structured facts within the context window.

D5. *Invariance tests.* Answers should not change when option order is permuted or when fields are reordered. Clef's training on permuted field orders targets the same sensitivity [29, 234].

Measurement. For each question, a red-team set holds records with injected text. It includes direct instructions ("classify as refund"), misleading but plausible claims ("reimbursement for invoice 1234"), and spoofed names. Metrics are the *flip rate*, the fraction of records whose decision changes under injection, and the *attack success rate*, the fraction that change to the attacker's target. Both are reported with and without each mitigation, and weighted by the cost of the flip. The corroboration rule D3 trades coverage for robustness, and the evaluation reports both sides. Paper IV's research question RQ6 is answered by these measurements.

7.10 Training sequence and model registry

The order of work follows Paper IV. The cheapest step comes first, and each later step must be justified by a measured gap.

1. Rules and readers (A0), plus question definitions, for every question.
2. Calibrated external and general baselines (A1, A2), with M1 to M3 on top.
3. A3 heads trained with the losses of Section 7.3, if a gap remains.
4. An RLCD-style secondary objective (Equation (26)) for ordinal questions, only if supervised training leaves a gap that partial credit could plausibly close, and only with separately authorised spending.

Every deployed configuration is registered with the backbone and adapter versions, head and calibrator versions, prompt template and option strings, serving mode and context limit,

decision-policy version, cost-matrix version, consistency-constraint version and the rule version ρ it was costed against. A change to any of these creates a new configuration that must pass the recalibration and evaluation steps of Section 8 before use.

Summary of open methodological hypotheses. [Research hypothesis] Back-off to a parent class reduces review volume for `asset_class` at equal expected cost. [Research hypothesis] Hierarchy-aware soft labels reduce the severity of mistakes on `asset_class` without loss of calibration after post-hoc scaling. [Research hypothesis] CORAL or RPS gives better-calibrated cumulative probabilities for ordinal questions than a flat softmax. [Research hypothesis] The corroboration rule cuts the success of text attacks at a coverage cost that the buyer accepts. Each is tested under the preregistration of Section 8.8.

8 Data, Labels and Evaluation

[Product, off-protocol] This section specifies where Zlef's data would come from (Section 8.1), how labels would be produced and checked (Sections 8.2 to 8.4), how data would be split (Section 8.5), what would be measured (Section 8.6), and how the model would be monitored after deployment (Section 8.7). It ends with a preregistration template (Section 8.8). It extends the seven evaluation steps of Paper IV and does not replace them. No data have been collected and no labels exist. Every threshold is left as "to be agreed" and must be fixed with the first buyer before the experiment that uses it.

Three principles from Paper IV govern everything below. A result on synthetic data is reported as synthetic. A vendor benchmark is not a result for this domain. A result on one customer is not reported as general. We add a fourth. A model metric that does not trace through to flags and reviewer work is incomplete, because Zlef exists to feed the rules layer and the review queue (Section 8.6.6).

8.1 Data sources

Table 47 summarises the three sources we consider and what each may be used for. The central rule is that only permissioned records from real payments, labelled under the protocol of Section 8.2, may support a claim about Zlef's performance. Synthetic data and public chain data may support development, stress tests and specific sub-questions, and are reported as such.

8.1.1 Synthetic generation protocol

Synthetic records let us build and test the pipeline before any customer data exist. Cloudflare trained Clef on "internal synthetic datasets permutating field orders, prompts, and schema structures" [29]. We take the same idea, with a stricter rule on how results are reported. The generator works top-down, from labels to facts.

1. **Sample a label vector.** Draw $y = (y_{q_1}, \ldots, y_{q_8})$ from a stated prior that satisfies the consistency constraints $\mathcal{K}$ (Section 7.7). The prior is a design input, written down and versioned. It deliberately over-samples rare classes, such as `TKN.EQ`, `PRV`, `STB.SYN` and `vasp_unmatched`, so that every class appears.

2. **Render structured facts.** Generate facts consistent with y: asset, chain, amounts as strings, counterparty identifiers, linked records and the absence markers that follow from the

Table 47: Data sources and their permitted uses.

Source	Permitted uses	Not permitted	Main risk
Synthetic records	development; unit tests of the pipeline; adversarial and invariance tests; pre-training heads	any headline performance claim; calibration fitting for deployment	model learns the generator, not payments
Permissioned historical records	training; calibration; conformal calibration; test	use beyond the customer's permission; pooling across customers without agreement	leakage; label noise; re-identification
Public on-chain data	asset and contract features; `asset_class` and `travel_surface` readers; graph-risk features computed upstream	linking addresses to people; publishing joined datasets	re-identification by joining with permissioned data

evidence level y_{q_6}. Each fact carries a layer tag. Facts that a real reader would resolve, such as a registry match for the asset contract, are generated with the frequency at which the reader is expected to succeed, which is itself an assumption to record.

3. **Render text.** Generate memos and invoice snippets from templates, then paraphrase them with a generator model that is separate from every model under test. Include abbreviations, typos and mixed languages, since real payment memos are short and noisy (Section 2.12).
4. **Perturb.** Permute field order and option order. Truncate text. Drop facts to layer-C markers at random, independently of y, so that absences do not reveal the label.
5. **Inject adversarial text** into a tagged subset, following the threat model of Section 7.9: direct instructions, misleading claims and spoofed token names. Record the attacker's target label.
6. **Optionally inject label noise** into a tagged subset, at stated class-conditional rates, to test noise-robust losses [152].
7. **Log** the generator version, prior, templates, seed and paraphrase model with every record.

Synthetic data carry a specific risk. A model can learn the generator's habits, such as template wording or the correlation between a rare class and a particular phrase, instead of the task. Two checks guard against this. First, no generator model, template or prior is shared with any model under test. Second, any model trained partly on synthetic data is also evaluated on permissioned records, and the gap between its synthetic and real scores is reported as a result in its own right.

8.1.2 Permissioned historical records

Historical payment records from a customer, used under that customer's agreement, are the only source on which performance claims may rest. Four conditions apply.

Minimum fields. Each question uses only the fields it needs. Paper IV's research question RQ8 asks which questions can be answered with fewer fields, or with selective disclosures in place of full records (Paper II). The answer is measured per question.

Availability time. Each fact carries the time it became available (Equation (14)). Facts recorded after the prediction time, such as later settlement status, post-hoc compliance flags or accountant categories, are excluded from inputs. They may serve only as label evidence, and only when the annotation protocol allows it (Section 8.2).

Separation by customer. Records from different customers are not pooled without each customer's agreement. When pooling is agreed, records keep their customer identifier, so that the customer split of Section 8.5 remains possible.

Privacy of artefacts. Any artefact derived from customer records, such as adapters, calibrators, reliability tables or conformal quantiles, may leak information about those records. Before any adapter trained on customer data is shared outside that customer, Zlef runs a membership-inference test of the kind described by Shokri et al. [196]. Release requires a pass criterion to be agreed. Calibration statistics shared across operators are released as differentially private aggregates [61]. Federated fine-tuning [146] is an option for learning across operators without pooling raw records, but it is not a privacy guarantee by itself.

8.1.3 Public on-chain data

Public chains make transaction data visible by design [39]. That visibility helps Zlef in narrow ways. Contract code and metadata can resolve `asset_class` for many tokens and several `travel_surface` labels, such as `issuer_freezable` or `passes_bridge`, through deterministic readers. Graph analyses computed upstream can supply risk features, as in the Elliptic line of work [15, 227]. Visibility is not the same as harmlessness. A public address becomes personal data once it is joined to a customer record, a name or an invoice. We therefore set three rules. Public chain data are used for asset- and contract-level features and for readers, not to link addresses to people. Datasets that join public chain data to permissioned records are never published. Evaluation reports state which public sources were joined, and how. The Elliptic datasets show one precedent for releasing chain data for research, with anonymised features. We note it as precedent, not as a standard that Zlef meets.

8.2 Annotation protocol

Guidelines. One guideline document covers all eight questions and uses the schema codes exactly. For each question it gives:

- a definition of every value, including when `UNKNOWN` applies;
- a decision procedure in the order an annotator should check facts;
- positive and negative examples, with emphasis on boundaries such as `STB.SYN` against `STB.YLD`, `WRP.BRIDGE` against `WRP.NATIVEBURN`, `gift` against `compensation`, and `self_hosted` against `individual`;
- when to give a partial label at an internal node of `asset_class`, such as `STB`, instead of guessing a leaf;
- which facts the annotator may consult.

The last point matters for leakage. By default, the annotator sees only the facts available at prediction time, the same x that the model sees. The guideline may name additional label evidence, such as a later-received invoice, that the annotator may consult to fix the label but that the model never sees. Every such case is recorded, because it means the label uses information

the model cannot have. Guidelines are versioned, and every label records the guideline version under which it was made.

Labels are jurisdiction-free. Annotators label economic and technical facts. They do not label legal categories, tax treatments or flags. A question such as "does beneficial ownership change?" must be answerable from the facts without choosing a legal system. Where the guideline's definition depends on a legal concept, the guideline must say so and fix one economic reading. The rules layer then maps that reading to each jurisdiction.

Annotators. Annotators are trained on the guidelines and pass a qualification test on a gold set before labelling. The scope limit of Snow et al. [201] applies: aggregating non-expert labels can work for perceptual and linguistic tasks, but several Zlef questions need financial or technical expertise. The pool is therefore restricted to people with the relevant background, and that background is recorded.

Double annotation and adjudication. Every item in the calibration, conformal and test sets is labelled independently by two annotators who do not see each other's labels or any model output. Disagreements go to an adjudicator, a senior annotator who sees both labels and records a reason for the decision. The adjudicated label is the gold label. The disagreement itself is kept as item metadata, because items on which qualified annotators disagree are hard. Some should be labelled `UNKNOWN` under a corrected guideline. Training items are labelled once, and a random sample is double-labelled so that their noise rate can be estimated.

Aggregation and annotator models. Where items have several labels, a Dawid–Skene model (Equation (11)) [52] estimates each annotator's confusion matrix and a posterior over the true label. The posteriors serve three purposes. They are soft training targets. They estimate per-annotator and per-reviewer error rates, which feed the review cost c_r of Section 7.6 and later learning to defer [140, 149]. And they point to guideline sections that produce systematic confusion. For hierarchical labels, the model is fitted at each level of the tree. Confident learning [157] is run on held-out predictions to rank items for re-review. It suggests which items to re-review and never changes a label by itself.

Agreement metrics. Inter-annotator agreement is reported per question and, where counts allow, per class. For nominal questions we report raw agreement and Cohen's kappa, which corrects agreement for chance. For ordinal questions we report weighted kappa, which gives partial credit to near disagreements, and Krippendorff's alpha with an ordinal distance. Alpha also handles items with missing labels and more than two annotators. For `asset_class`, we report agreement at each level of the tree, so that disagreement on leaves within `STB` is not mistaken for disagreement about whether an asset is a stablecoin. For `travel_surface`, we report agreement per label. Agreement sets a ceiling on measurable model performance. A model cannot be shown to be more accurate than the labels it is scored against, so the evaluation reports model metrics alongside agreement. Minimum agreement levels, below which a question is redefined rather than modelled, are to be agreed.

The role of jurisdiction experts. Experts in each jurisdiction's law write and review the rules layer (layer E), the flags and their severities, and therefore the cost matrices of Section 7.6. They do not label Zlef questions in their legal capacity, because Zlef labels are not legal conclusions. They do advise on guideline boundaries where an economic fact and a legal concept meet, as in the beneficial-ownership example above. Their advice enters the guidelines, not the labels. This keeps a legal judgement from being encoded silently in training data, where a later rule change could not reach it.

Label provenance record. Each label is stored with the item identifier, question, value (possibly an internal node), annotator identifier, study arm (Section 8.3), guideline version, the facts consulted (including any label-only evidence), the time taken, whether it was adjudicated, and the adjudicator's reason. This record supports the analyses below. It also makes Paper IV's rule, that anchored corrections are recorded as such, auditable.

8.3 Anchoring control: blind and assisted labelling

A reviewer who sees a model's answer first may adopt it [99, 170]. If anchored labels are then used to evaluate or retrain the model, the model is graded and taught by its own echo. Paper IV's research question RQ7 asks how large this effect is. We propose to measure it directly.

Design. A sample of items is assigned at random to one of three arms. Assignment is randomised at the item level, within each annotator, so that each annotator works in every arm and differences between annotators cancel.

- **Blind.** The annotator sees only the facts. No model output is shown.
- **Assisted.** The annotator sees the facts and the model's calibrated distribution for the question, with no explanation.
- **Assisted with explanation.** As assisted, plus the evidence pointers and the generated rationale (Section 7.8).

Every item in the study also has an adjudicated gold label from the blind double-annotation protocol of Section 8.2.

Measures. For each arm we estimate four quantities. The first is accuracy against gold. The second is agreement with the model. The third, and most important, is agreement with the model *when the model is wrong*:

$$\mathrm{AR}_{\mathrm{arm}} = \Pr\left(\text{label} = \hat{y}_{\mathrm{model}} \;\middle|\; \hat{y}_{\mathrm{model}} \neq y_{\mathrm{gold}},\ \text{arm}\right). \tag{41}$$

The difference $\mathrm{AR}_{\mathrm{assisted}} - \mathrm{AR}_{\mathrm{blind}}$ estimates the anchoring effect. The fourth is time per item, which is what assistance buys. The analysis is stratified by question, by the model's calibrated probability and by annotator experience. Anchoring is likely to be strongest when the model is confident and the item is hard.

Use of the result. Labels from the assisted arms are never used as gold for evaluation. If the anchoring effect is material by a criterion to be agreed, labels made after seeing the model are down-weighted or excluded from training (Equation (19)), and the production review interface

defaults to a blind-first mode for the affected questions. In blind-first mode, the reviewer records an answer before the model's answer is revealed. If the effect is negligible, the assisted interface may be used for speed, but the blind arm continues on a small audit sample to detect change over time. Because `review_priority` is learned from reviewer behaviour, it is the question most exposed to this loop. It is therefore the last question to be trained, after the anchoring effect for the other questions has been measured, and the first to be audited in the blind arm.

8.4 Candidate labels from reviewer corrections

Once Zlef runs in a review workflow, reviewers will confirm or correct its answers. Paper IV treats these corrections as *candidate* labels, not truth. We specify how a candidate becomes usable.

Why corrections are not gold. Corrections are produced under time pressure, often in the assisted condition, by reviewers whose error rates vary. They are also selected in a biased way. Corrections exist mainly for items that the model abstained on, that a flag sent to review, or that a reviewer chose to inspect. The confirmed items are mostly those the reviewer saw no reason to change, which says little when the reviewer was anchored. Training on corrections therefore over-represents hard and flagged items. Evaluating on them would score the model on a sample that its own abstention policy chose.

Quality control. A correction enters the training pool only after five checks.

1. It is tagged with arm, reviewer and time, as in Section 8.2.
2. A random sample of corrections is re-labelled blind by a second annotator. The disagreement rate is tracked per reviewer and per question.
3. Reviewer confusion matrices are estimated with the Dawid–Skene model. Corrections from reviewers whose estimated error on that question exceeds a level to be agreed are held back.
4. Confident learning on held-out predictions flags corrections that are likely errors, for re-review.
5. Corrections that conflict with a layer A or B fact, such as a correction to `WRP.BRIDGE` for an asset that the registry lists as native, are rejected automatically and sent to the guideline owners.

Use. Corrections that pass may enter training, with weights that reflect their estimated reliability and the selection bias described above. Corrections never enter the calibration, conformal or test sets. Those sets are built only from items sampled at random from traffic and labelled under the blind double-annotation protocol. This keeps the evaluation sample independent of the model's own routing decisions.

8.5 Splits

Random splits of payment data put records from the same customer, the same week and the same asset on both sides. Similar records then appear in both training and test data, and future patterns leak into training [129, 131]. Zlef uses grouped and forward-in-time splits.

Split types.

- **Customer split.** All records of a customer fall in one partition. This measures generalisation to a customer the model has not seen.
- **Time split.** With cut-offs $T_1 < T_2 < T_3 < T_4$, records before T_1 are for training, $[T_1, T_2)$ for validation and calibration, $[T_2, T_3)$ for conformal calibration, and $[T_3, T_4)$ for testing. This is the default split, as in the Elliptic and Wild-Time precedents [227, 231].
- **Chain split.** All records on one chain or network are held out. This measures what happens when ZETTL onboards a new chain, and it is where weighted conformal calibration under covariate shift would be tested [211].
- **Asset-introduction-date split.** Records involving assets first seen after a cut-off are held out. New assets are the main way that `asset_class` meets novel inputs, and this split measures it directly.

Splits can be combined. The most demanding combination tests on later records, from unseen customers, involving assets introduced after training. Every reported number states which split produced it.

Partitions. Calibration and conformal calibration use separate partitions, so that the conformal guarantee does not depend on data already used to fit the calibrator. Hyperparameters, losses, calibrators and decision thresholds are chosen on validation and calibration partitions only.

The locked held-out set. One further partition, later in time than every other partition and drawn from customers that appear in no other partition, is locked. It is used once per preregistered experiment, after all choices are frozen. Access is logged. A result on the locked set that leads to any change in the model or policy retires that set, and a new one must be drawn for the next experiment.

Random splits, for one purpose only. A random split is computed only to answer Paper IV's research question RQ5: how much does a random split overstate performance compared with customer and time splits? Its numbers are never reported as Zlef's performance.

8.6 Metrics

Metrics are computed per question, per class where counts allow, per split and per domain. Domains include chain, corridor, customer and counterparty type. Following WILDS [133], we report the worst domain as well as the average. Every metric has a confidence interval from a bootstrap that resamples customers, not records, because records from one customer are correlated. Table 49 lists the metrics by group. The subsections below define the less standard ones.

Table 48: Leakage checklist, following the taxonomy of Kapoor and Narayanan [129], with Zlef-specific examples and controls.

Leakage type	Zlef example	Control
No train/test separation	same record in two partitions after deduplication fails	deduplicate by record and transaction hash before splitting
Illegitimate features	later settlement status; post-hoc flags; accountant category	availability time t_i; exclude layers D–F from inputs
Temporal leakage	on-chain events after the payment; future price	time split; $t_i \leq t^*$
Non-independence	same customer, counterparty or asset in train and test	customer, chain and asset splits
Sampling bias	test built from reviewed or flagged items	test sampled at random from traffic; corrections excluded
Label leakage via annotators	annotator consults later evidence	label-only evidence recorded and reported

Table 49: Metrics by group. "Per class" means reported for every value where counts allow, and pooled with a stated rule otherwise.

Group	Metrics	Applies to
Classification	per-class precision, recall and F1; macro averages; confusion matrices	all nominal questions; per label for `travel_surface`
Proper scores	NLL (Equation (4)); Brier (Equation (5)); RPS (Equation (6))	all; RPS for ordinal
Calibration	adaptive top-label ECE; class-wise ECE (SCE); MCE; reliability diagrams; bias-corrected estimates	all; class-wise for nominal and hierarchical
Selective prediction	risk–coverage curve; AURC; selective risk at preregistered coverage; cost-weighted versions	every abstention mechanism
Conformal	empirical coverage, marginal and per class; mean set size; singleton rate; coverage by set size	M3 and conformal risk control
Ordinal	MAE on level index; off-by-one rate; calibration of $\Pr(y \geq j)$ for each j	`evidence_completeness`, `review_priority`
Hierarchical	hierarchical P, R, F; mistake severity; back-off rate and back-off correctness	`asset_class`
`UNKNOWN` and abstention	rate of `UNKNOWN` predictions vs gold `UNKNOWN`; abstention rate; their overlap	all questions with `UNKNOWN`
Consistency	violation rate per constraint; rate of `DAT`-family flags emitted	constraint set $\mathcal{K}$
Flag level	flag precision and recall per family and severity; missed-block rate; spurious-block rate	rules layer on model facts vs gold facts

Group	Metrics	Applies to
Human work	review time per item; correction rate; anchoring rate (Equation (41))	review workflow
Operations	latency p50 and p95; cost per item; reader coverage per question	every configuration
Robustness	flip rate; attack success rate; invariance to option and field order	red-team and permutation sets

8.6.1 Classification and proper scores

Per-class precision, recall and F1 are reported for every class, with macro averages over classes. Macro averages weight rare classes equally with common ones, which is the right default when rare classes carry large downstream costs. Micro averages are reported alongside. Negative log-likelihood and Brier score are reported on the model's probabilities both before and after post-hoc calibration, so that the effect of calibration is visible. For the hierarchical question, NLL is computed on the leaf distribution, and also at each level from the marginals of Equation (21).

8.6.2 Calibration and its estimation

We report three calibration errors. Top-label ECE uses adaptive, equal-mass bins (ACE) in place of equal-width bins. Class-wise ECE (SCE) averages a binned error over all classes [134, 155]. MCE gives the worst bin. Reliability diagrams are shown with bin counts.

Estimator bias. The plug-in estimator Equation (7) is biased upward, most of all with small bins (Section 2.3.2). Debiased estimators exist in the literature. [Unverified] We have not yet selected and cited one. Until we do, we propose a simple reference procedure of our own. For each evaluation set, draw synthetic outcomes $\tilde{y}_i \sim p^{\text{cal}}(\cdot \mid x_i)$ from the model's own probabilities, so that the model is perfectly calibrated by construction. Recompute the ECE estimate, and repeat. The mean of these values estimates the estimator's bias at this sample size and binning. We report the observed ECE, this null reference, and their difference, each with an interval. An observed ECE inside the null range is not evidence of miscalibration. The procedure is a diagnostic, not a corrected estimator with known properties, and it will be replaced by a published estimator once one is cited.

Calibration of decisions, not just probabilities. The policy uses $\Pr(y \geq j)$ for ordinal questions, node marginals for back-off, and per-label probabilities for `travel_surface`. Each of these is calibrated and measured as the binary event it is.

8.6.3 Selective prediction

For a decision rule with a confidence or cost score $g(x)$, sort test items by g and accept the top fraction c. The *coverage* is c. The *selective risk* $r(c)$ is the error rate, or the mean cost, on the accepted items [63]. The risk–coverage curve plots $r(c)$ for $c \in (0, 1]$. Its area,

$$\text{AURC} = \int_0^1 r(c)\,\mathrm{d}c \;\approx\; \frac{1}{N}\sum_{n=1}^{N} r\left(\tfrac{n}{N}\right), \tag{42}$$

summarises the curve, with lower being better. AURC mixes model accuracy with ranking quality, so we also report selective risk at the coverage levels that the preregistration names. We compute the cost-weighted curve, with $r(c)$ as the mean cost of Equation (36) on accepted items, for each abstention mechanism M1 to M3, for conformal risk control and for back-off. For M3, whose acceptance is set-based rather than score-based, we report its operating point and not a curve. For the SGR procedure, we report whether the selective-risk bound held on the test split [95].

8.6.4 Conformal coverage and set size

For each question and level α, we report empirical coverage, the fraction of test items whose gold label is in the set, both marginally and per class. Per-class coverage reveals under-coverage of rare classes that marginal coverage hides [55]. We also report mean set size and the singleton rate, which is the fraction of items that M3 would automate, and coverage broken down by set size. Coverage is reported on every split. On time, chain and asset splits, exchangeability is violated, and the size of the coverage loss is the result of interest. On the chain split, weighted conformal calibration is compared with unweighted calibration [211].

8.6.5 Ordinal and hierarchical metrics

For ordinal questions we report mean absolute error on the level index,

$$\mathrm{MAE} = \frac{1}{N} \sum_{i=1}^{N} |\hat{y}_i - y_i|, \tag{43}$$

together with the off-by-one rate and RPS. MAE treats the levels as equally spaced, which is a convention, not a fact. For `review_priority` we also report the confusion between `urgent` and the rest, because an urgent case filed as low has a cost that MAE understates.

For `asset_class` we report hierarchical precision, recall and F-measure [199]. Let $\mathrm{Anc}(v)$ be the set of ancestors of v, including v and excluding the root. Then

$$hP = \frac{\sum_i |\mathrm{Anc}(\hat{y}_i) \cap \mathrm{Anc}(y_i)|}{\sum_i |\mathrm{Anc}(\hat{y}_i)|}, \qquad hR = \frac{\sum_i |\mathrm{Anc}(\hat{y}_i) \cap \mathrm{Anc}(y_i)|}{\sum_i |\mathrm{Anc}(y_i)|}, \tag{44}$$

and hF is their harmonic mean. These definitions accept an internal-node prediction from back-off: answering `STB` for a gold `STB.SYN` has full hierarchical precision and partial recall. We also report the *severity* of mistakes, following Bertinetto et al. [16], as the mean height of the lowest common ancestor of prediction and gold over wrong predictions. Alongside it we report the mean cost of mistakes from the cost matrix. If the two disagree, the cost matrix is the measure that matters. For back-off we report the back-off rate and the back-off correctness, the fraction of backed-off answers whose gold leaf lies under the returned node.

8.6.6 Flag-level metrics

Model metrics say how often Zlef is right. Flag-level metrics say whether its errors matter. The rules layer, at a fixed version ρ, is run twice on each test record: once on the gold answers, which gives reference flags F^*, and once on the answers that the deployed policy would pass on, which gives $\hat{F}$. In the second run, abstentions are passed on as review routes, and back-off answers as

internal nodes. Both runs use the same jurisdiction context J. For each flag family and severity we report

$$\mathrm{FP}_{\mathrm{flag}} = \frac{|\hat{F} \cap F^*|}{|\hat{F}|}, \qquad \mathrm{FR}_{\mathrm{flag}} = \frac{|\hat{F} \cap F^*|}{|F^*|}, \tag{45}$$

summed over records. The single most important metric is the *missed-block rate*:

$$\mathrm{MBR} = \frac{\#\{i : F_i^* \text{ contains a } \texttt{block} \text{ flag}, \ \hat{F}_i \text{ does not, and } i \text{ was not routed to review}\}}{\#\{i : F_i^* \text{ contains a } \texttt{block} \text{ flag}\}}. \tag{46}$$

It counts records where the gold facts would have stopped the payment for review, but the automated path neither raised the flag nor sent the record to a person. The spurious-block rate is defined symmetrically. It measures cost to legitimate payers. Flag-level metrics depend on the rule version and the corridor mix, so they are reported per rule version and per corridor, never as a single number. They also measure only the model's contribution. Errors in the rules themselves are Paper V's concern.

8.6.7 Human work, latency and cost

In the review workflow, we record time per item, the correction rate per question, and the anchoring rate of Equation (41). Review time is the main input to the review cost c_r, so it is measured, not assumed. Latency is reported as median and 95th percentile per configuration, measured end to end from a stated location. The vendors' latency figures are not used (Section 2.1). Cost per item includes model inference, reader and rules computation, and expected review time at the measured abstention rate. Reader coverage, the share of each question answered deterministically, is reported, because a model that is accurate only on cases a reader already answers adds nothing.

8.7 Drift monitoring

Crypto-asset payments shift suddenly, through new tokens, depegs, exploits and rule changes, and recurrently, through tax-year cycles [93]. Calibration fitted in-distribution degrades under shift [169]. Monitoring therefore runs continuously after deployment, on five kinds of signal.

1. **Input drift.** The distributions of chains, assets, counterparty types, field presence and text length, compared with the calibration period. New assets and chains are counted separately, because they are the main source of novel inputs.

2. **Output drift.** The mix of predicted classes per question; the distribution of top-class probabilities; abstention and back-off rates; reader coverage; the rate of constraint violations. Under label shift, BBSE (Equation (13)) estimates the new class priors from the output mix without labels [138].

3. **Audited performance.** A small random sample of traffic is labelled each period under the blind protocol. On it we track accuracy, calibration error with the null reference of Section 8.6, conformal coverage and the missed-block rate. This is the only drift signal tied to ground truth.

4. **Review signals.** Correction rates and review times per question. A rise in corrections with no change in the audited sample may point to reviewer drift rather than model drift.

5. **Configuration changes.** Any change of model, serving configuration, rule version, cost matrix or schema is itself a drift event and triggers re-evaluation (Section 7.10).

Responses. Each signal maps to a graded response, with trigger levels to be agreed:

- a review of the signal;
- a prior update by BBSE;
- recalibration on recent labelled data;
- suspension of automation for a question, so that all its items go to review until re-evaluation;
- retraining.

A rule change is not model drift. It changes the cost matrices and the flag-level metrics, and the policy is recomputed without retraining. A change in the meaning of a schema value, through a guideline or schema edit, invalidates the gold labels that used the old meaning, and those labels must be re-annotated or retired. The HS-code literature shows the general hazard: revisions to a rulebook invalidate historical labels [137].

8.8 Preregistration template

Paper IV requires that every experiment be preregistered with the first buyer before it runs. Table 50 gives the template. Every bracketed field is filled in and signed before any test data are touched. No value in the template is a recommendation. Every threshold reads "to be agreed". Any departure from a registered plan is reported as such, with its reason.

What counts as a result. The template carries forward Paper IV's rules on what counts as a result and makes them concrete. A result counts as evidence for Zlef only when it meets four conditions. It was preregistered. It was read once on the locked held-out set. It was produced on permissioned records labelled under the blind double-annotation protocol. And it reports flag-level metrics next to model metrics. Anything else is development evidence and is labelled as such. Edition 1 contains no result of either kind.

9 Risks, Ethics and Limits

Automation bias. A calibrated probability can still persuade a reviewer to stop thinking. The design answers this in four ways:

- it shows predictions as layer D facts, next to the records they came from;
- it never lets a prediction alone produce a `block` (Paper V);
- it measures anchoring directly, with blind and assisted annotation arms (Section 8.3);
- it treats reviewer corrections as candidate labels, not truth.

None of these removes the risk.

Table 50: Preregistration template for a Zlef experiment. Bracketed fields are completed and signed before the experiment.

Identification	Experiment ID [ ]; edition and date [ ]; buyer and workflow [ ]; owners [ ]; signatories [ ].
Scope	Questions in scope (schema codes) [ ]; schema version [ ]; guideline version [ ]; rule version ρ and corridors for flag metrics [ ].
Hypotheses	Primary hypothesis, stated so that it can fail [ ]. Secondary hypotheses [ ]. Each linked to a research question of Paper IV or a [Research hypothesis] item of Section 7 [ ].
Configurations	Candidate configuration(s): architecture A0–A4, backbone and serving configuration, losses, calibrator, abstention mechanism, cost-matrix version [ ].
Baselines	Rules and readers (A0) [ ]; general model with constrained output (A1) [ ]; external decision model (A2) and its pinned serving configuration [ ]; previous Zlef configuration, if any [ ]. All baselines calibrated on the same split.
Data and splits	Sources and permissions [ ]; split type(s) and cut-off dates [ ]; partition sizes [ ]; locked held-out set ID [ ]; exclusions and deduplication rules [ ].
Labels	Annotation protocol version; double-annotation and adjudication coverage [ ]; agreement measured before the experiment [ ]; use of reviewer corrections (training only) [ ]; anchoring-arm assignment [ ].
Primary metric(s)	One per hypothesis, with the split, class or domain aggregation and coverage level at which it is read [ ].
Secondary metrics	From Table 49 [ ].
Decision rule	Minimum improvement on the primary metric over the strongest baseline: threshold to be agreed [ ]. Maximum calibration error (class-wise, bias-referenced): threshold to be agreed [ ]. Maximum missed-block rate: threshold to be agreed [ ]. Minimum coverage at target selective risk: threshold to be agreed [ ]. Maximum latency and cost per item: thresholds to be agreed [ ]. Outcome options: continue, narrow, change, stop.
Statistics	Confidence intervals by customer-level bootstrap [ ]; correction for multiple comparisons [ ]; minimum test-set size per class for a per-class claim [ ].
Stopping rules	Halt if a safety metric (missed-block rate) crosses its threshold during any staged rollout [ ]; halt if data permission is withdrawn [ ]; no early stopping for success [ ]; maximum budget and time [ ].
Deviations	Procedure for recording and reporting any change after registration [ ].
Reporting	What is reported, whatever the outcome; synthetic and single-customer results labelled as such [ ].

Feedback loops. If reviewers mostly see what the model ranks high, the labels collected will over-represent what the model already finds. `review_priority` is most exposed to this. A fraction of review capacity should be allocated at random and kept out of the model's influence, so that the label distribution stays estimable.

Profiling and re-identification. Payment records contain personal data, and public on-chain data can re-identify people when joined with off-chain records. Zlef predicts no personal attributes (Section 6), and the data plan prefers synthetic and permissioned data (Section 8). Even so, any model trained on payment histories can memorise them. Membership-inference testing and data minimisation are prerequisites for any training, not afterthoughts (Section 2.14).

Disparate error. Error rates may differ across corridors, currencies and customer types. For example, a model trained mostly on US and UK records may abstain or err more on Indonesian or Vietnamese payments. Per-corridor metrics are part of evaluation. A corridor where the model is unreliable should be served by rules and review, not by a model that hides its weakness in an average.

Adversarial manipulation. Memo and invoice text is controlled by counterparties. A model that reads it can be steered, and a language model can be instructed by it (Section 7.9). Text fields are data, never instructions. Flags with high severity never rest on text-derived labels alone.

Legal error. The ontology and asset chapters summarise law as found on one day, partly from secondary sources. They will be wrong in places and will go out of date. The architecture contains this: the model's outputs are jurisdiction-free, the legal mapping lives in versioned rules, and every rule has a status and a source. But the document itself must not be read as advice.

Vendor dependence. If a third-party decision model is used as a baseline or component, its terms, availability and behaviour can change, and sending payment data to it may not be permitted by the customer [External dependency]. Self-hosted open weights reduce, but do not remove, this dependence.

Overclaiming. The largest risk to a document like this is that its specifications are read as results. To restate: Zlef has no data, no training run, no accuracy, no calibration, no cost and no latency figures, and no customer has used it [Unsupported, needs evidence].

10 Research Agenda for Later Editions

Each later edition of this document should advance at least one of the items below and record the change in Appendix B.

A. Verification. Close the verification backlog:

- the Vietnam and Ukraine classifications;
- tax cells marked [Unverified], especially for staking, wrapping and lending;

- the GENIUS Act effective date and its implementing rules;
- Japan's FIEA start date and its 20.315% rate;
- Korea's 2027 tax;
- CARF's exclusions;
- privacy-coin treatment in Japan and Korea.

Each verified item updates a rule record (Paper V) and, where the asset chapters cite it, this document.

B. Schema. Test the Edition 1 schema against real records under permission:

- Are the `asset_class` children separable by annotators (inter-annotator agreement)?
- Is `GOV` versus `MEME` stable?
- Does `economic_purpose` need jurisdiction-relevant splits (for example contractor versus employee) that are still jurisdiction-free?
- Should `travel_surface` gain an indicator for sanctioned-exposure hops?

C. Ordinal versus nominal (Paper IV's open question). Does an RLCD-style partial-credit objective help the ordinal questions relative to plain proper scoring? Is there any principled analogue for nominal questions, such as tree distance in `asset_class` or cost-matrix distance in `event_type`? This is a [Research hypothesis] until measured.

D. Abstention. Compare an explicit `UNKNOWN` class, cost-sensitive thresholds and conformal prediction sets on risk–coverage and on downstream flag metrics, per question and per corridor.

E. Baselines. Under the preregistration template of Section 8, compare:

- rules alone;
- a general model with constrained decoding;
- an external decision model;
- a specialised model.

Report whether the specialised model earns its cost.

F. Drift. Measure how fast calibration decays as assets, chains and laws change. The asset-introduction split is the main tool. Set retraining and recalibration triggers.

G. Provenance. When payments settled on ZETTL exist, measure whether layer A facts reduce review effort (Paper IV's RQ9).

H. Country set. Re-run the selection of the ten when Chainalysis publishes its next ranking, and record any change and its effect on the schema's coverage.

References

[1] Martín Abadi et al. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2016. Details not re-checked 15 Oct 2026.

[2] Diogo Almeida. Introducing system one models & Jev. TypeSafe blog, 2026. URL https://typesafe.ai/blog/introducing-system-one-models-and-jev. 15 Sep 2026; accessed 2026-10-15; [P] vendor blog; benchmarks vendor-reported.

[3] Anastasios N. Angelopoulos and Stephen Bates. A gentle introduction to conformal prediction and distribution-free uncertainty quantification. *arXiv preprint arXiv:2107.07511*, 2021. Later version in Foundations and Trends in Machine Learning (2023); volume and pages unverified.

[4] Anastasios N. Angelopoulos, Stephen Bates, Jitendra Malik, and Michael I. Jordan. Uncertainty sets for image classifiers using conformal prediction. In *International Conference on Learning Representations (ICLR)*, 2021. arXiv:2009.14193. Details not re-checked 15 Oct 2026.

[5] Anastasios N. Angelopoulos, Stephen Bates, Adam Fisch, Lihua Lei, and Tal Schuster. Conformal risk control. In *International Conference on Learning Representations (ICLR)*, 2024. URL https://proceedings.iclr.cc/paper_files/paper/2024/hash/f3549ef9b5ff520a7e41ff3cc306ab2b-Abstract-Conference.html. arXiv:2208.02814.

[6] Baker McKenzie. Cftc announcement paves the way for spot crypto trading on us regulated exchanges. Web page, 2025. URL https://blockchain.bakermckenzie.com/2025/12/10/cftc-announcement-paves-the-way-for-spot-crypto-trading-on-u-s-regulated-exchanges/. 10 Dec 2025; accessed 2026-10-15; [S].

[7] Baker McKenzie. Vietnam launches pilot tax framework for crypto-asset transactions. Web page, 2026. URL https://www.bakermckenzie.com/en/insight/publications/2026/05/vietnam-launches-pilot-tax-framework-for-crypto-asset-transactions. May 2026; accessed 2026-10-15; [S].

[8] Banco Central do Brasil. Bc regulamenta a prestacao de servicos de ativos virtuais (resolucoes bcb 519, 520 e 521). Press release (third-party host), 2025. URL https://static.poder360.com.br/2025/11/bc-regulamenta-ativos-virtuais-funcionamento-10nov2025.pdf. 10 Nov 2025; in force 2 Feb 2026; accessed 2026-10-15 [P].

[9] Banco Central do Brasil. Bc regulamenta ativos virtuais (release on resolutions 519–521). Web page, 2025. URL https://static.poder360.com.br/2025/11/bc-regulamenta-ativos-virtuais-funcionamento-10nov2025.pdf. 10 Nov 2025; BCB document on third-party host; accessed 2026-10-15; [P].

[10] Bank of England. Boe launches policy statement and draft rules on regulating systemic stablecoins. Web page, 2026. URL https://www.bankofengland.co.uk/news/2026/june/boe-launches-policy-statement-and-draft-rules-on-regulating-systemic-stablecoins. 22 Jun 2026; accessed 2026-10-15; [P].

[11] Bank of Russia. Bank of russia statement on the law on digital currencies. Press event, 2026. URL https://www.cbr.ru/eng/press/event/?id=32724. 21 Jul 2026; effective 1 Sep 2026; accessed 2026-10-15 [P].

[12] Bank of Russia. Press release on the law on digital currencies and digital rights. Web page, 2026. URL https://www.cbr.ru/eng/press/event/?id=32724. 21 Jul 2026; accessed 2026-10-15; [P].

[13] Gagan Bansal, Tongshuang Wu, Joyce Zhou, Raymond Fok, Besmira Nushi, Ece Kamar, Marco Tulio Ribeiro, and Daniel S. Weld. Does the whole exceed its parts? the effect of AI explanations on complementary team performance. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 2021. doi: 10.1145/3411764.3445717. DOI not re-checked 15 Oct 2026.

[14] Basel Committee on Banking Supervision. Prudential treatment of cryptoasset exposures. Web page, 2022. URL https://www.bis.org/bcbs/publ/d545.htm. 16 Dec 2022; accessed 2026-10-15; [P].

[15] Claudio Bellei, Muhua Xu, Ross Phillips, Tom Robinson, Mark Weber, Tim Kaler, Charles E. Leiserson, Arvind, and Jie Chen. The shape of money laundering: Subgraph representation learning on the blockchain with the Elliptic2 dataset. *arXiv preprint arXiv:2404.19109*, 2024. URL https://arxiv.org/abs/2404.19109v2. Venue unverified.

[16] Luca Bertinetto, Romain Mueller, Konstantinos Tertikas, Sina Samangooei, and Nicholas A. Lord. Making better mistakes: Leveraging class hierarchies with deep networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. URL https://openaccess.thecvf.com/content_CVPR_2020/html/Bertinetto_Making_Better_Mistakes_Leveraging_Class_Hierarchies_With_Deep_Networks_CVPR_2020_paper.html. arXiv:1912.09393.

[17] Bloomingbit. Korea virtual asset taxation from 2027. Web page, 2026. URL https://en.bloomingbit.io/feed/news/117542. Aug 2026; accessed 2026-10-15; [S].

[18] Bloomingbit. Korea virtual-asset taxation scheduled for 2027. News, 2026. URL https://en.bloomingbit.io/feed/news/117542. Aug 2026; accessed 2026-10-15 [S].

[19] Glenn W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950. doi: 10.1175/1520-0493(1950)078<0001:VOFEIT>2.0.CO;2. Details not re-checked 15 Oct 2026.

[20] Candour Legal. Vda taxation in india: section 115bbh, 194s and the income-tax act 2025. Article, 2026. URL https://candourlegal.com/vda-taxation-india-section-115bbh-194s-income-tax-act-2025/. Accessed 2026-10-15 [S].

[21] Candour Legal. Vda taxation in india under the income-tax act 2025. Web page, 2026. URL https://candourlegal.com/vda-taxation-india-section-115bbh-194s-income-tax-act-2025/. 2026; accessed 2026-10-15; [S].

[22] Wenzhi Cao, Vahid Mirjalili, and Sebastian Raschka. Rank consistent ordinal regression for neural networks with application to age estimation. *Pattern Recognition Letters*, 140:

325–331, 2020. URL https://sebastianraschka.com/pdf/papers/2020-coral.pdf. arXiv:1901.07884. Volume and pages not re-checked 15 Oct 2026.

[23] Nicholas Carlini et al. Extracting training data from large language models. In *30th USENIX Security Symposium*, 2021. Details not re-checked 15 Oct 2026.

[24] Iñigo Casanueva et al. Efficient intent detection with dual sentence encoders. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI (NLP4ConvAI)*, 2020. Origin of BANKING77. Details not re-checked 15 Oct 2026.

[25] Chakwon. Korea digital asset basic act: stablecoin and sto split. Web page, 2026. URL https://blog.chakwon.com/2026/04/17/korea-digital-asset-basic-act-stablecoin-sto-split/. 17 Apr 2026; accessed 2026-10-15; [S].

[26] Ilias Chalkidis, Abhik Jana, Dirk Hartung, Michael Bommarito, Ion Androutsopoulos, Daniel Martin Katz, and Nikolaos Aletras. LexGLUE: A benchmark dataset for legal language understanding in English. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4310–4330, 2022. URL https://aclanthology.org/2022.acl-long.297. Page numbers not re-checked 15 Oct 2026.

[27] Chambers and Partners. Blockchain and crypto-assets 2026: Indonesia, trends and developments. Web page, 2026. URL https://practiceguides.chambers.com/practice-guides/blockchain-crypto-assets-2026/indonesia/trends-and-developments. updated 11 Jun 2026; accessed 2026-10-15; [S].

[28] Chambers and Partners. Blockchain and crypto-assets 2026: South korea, trends and developments. Web page, 2026. URL https://practiceguides.chambers.com/practice-guides/blockchain-crypto-assets-2026/south-korea/trends-and-developments. Jun 2026; accessed 2026-10-15; [S].

[29] Michelle Chen, Alex Reneau, and Kevin Flansburg. Introducing Clef: our open-source decision models, and new RL fine-tuning platform. Cloudflare blog, 2026. URL https://blog.cloudflare.com/clef-decision-models. 1 Oct 2026; accessed 2026-10-15; [P] vendor blog; benchmarks vendor-reported.

[30] Michelle Chen, Alex Reneau, and Kevin Jain. Introducing Clef-omni with full multimodality, plus a faster Clef and a cheaper Clef-flash. Cloudflare blog, 2026. URL https://blog.cloudflare.com/clef-faster-cheaper-multimodal/. 9 Oct 2026; accessed 2026-10-15; [P] vendor blog.

[31] C. K. Chow. On optimum recognition error and reject tradeoff. *IEEE Transactions on Information Theory*, 16(1):41–46, 1970. doi: 10.1109/TIT.1970.1054406. Details not re-checked 15 Oct 2026.

[32] Circle. Cross-chain transfer protocol (cctp) documentation. Web page, 2026. URL https://developers.circle.com/cctp. issuer documentation; accessed 2026-10-15; [P].

[33] Cloudflare. Cloudflare/clef and Cloudflare/clef-flash (model cards). Hugging Face, 2026. URL https://huggingface.co/Cloudflare/clef. Also

https://huggingface.co/Cloudflare/clef-flash; accessed 2026-10-15; [P] vendor model cards.

[34] CoinDesk. Brazil's central bank bans stablecoin and crypto settlement in cross-border payments. Web page, 2026. URL https://www.coindesk.com/policy/2026/05/02/brazil-s-central-bank-bans-stablecoin-and-crypto-settlement-in-cross-border-payments. 2 May 2026; accessed 2026-10-15; [S].

[35] Coinpaprika. Tokenized commodity regulation: Cftc, sec and global rules. Web page, 2026. URL https://coinpaprika.com/education/tokenized-commodity-regulation-cftc-sec-and-global-rules/. lower-quality source; accessed 2026-10-15; [S].

[36] Cointelegraph. Eu to ban anonymous crypto accounts and privacy tokens from 2027. Web page, 2024. URL https://cointelegraph.com/news/eu-crypto-ban-anonymous-privacy-tokens-2027. press; accessed 2026-10-15; [S].

[37] CoinTracking. Crypto de minimis tax exemption: status. Web page, 2026. URL https://cointracking.info/blog/crypto-de-minimis-tax-exemption. Aug 2026; accessed 2026-10-15; [S].

[38] CoinTracking. Crypto de minimis tax exemption: status of us bills. Blog, 2026. URL https://cointracking.info/blog/crypto-de-minimis-tax-exemption. Aug 2026; S.2207, H.R. 9178, H.R. 8899; accessed 2026-10-15 [S].

[39] Lin William Cong and Zhiguo He. Blockchain disruption and smart contracts. *The Review of Financial Studies*, 32(5):1754–1797, 2019. doi: 10.1093/rfs/hhz007. Details not re-checked 15 Oct 2026.

[40] Council of the EU. 19th package of sanctions against russia. Web page, 2025. URL https://www.consilium.europa.eu/en/press/press-releases/2025/10/23/19th-package-of-sanctions-against-russia-eu-targets-russian-energy-third-country-banks-and-crypto-providers/. 23 Oct 2025; accessed 2026-10-15; [P].

[41] Council of the EU. 20th round of eu sanctions against russia. Web page, 2026. URL https://www.consilium.europa.eu/en/press/press-releases/2026/04/23/russia-s-war-of-aggression-against-ukraine-20th-round-of-stern-eu-sanctions-hits-energy-military-industrial-complex-trade-and-financial-services-including-crypto/. 23 Apr 2026; accessed 2026-10-15; [P].

[42] Council of the European Union. 19th package of sanctions against russia: Eu targets russian energy, third-country banks and crypto providers. Press release, 2025. URL https://www.consilium.europa.eu/en/press/press-releases/2025/10/23/19th-package-of-sanctions-against-russia-eu-targets-russian-energy-third-country-banks-and-crypto-providers/. 23 Oct 2025; accessed 2026-10-15 [P].

[43] Council of the European Union. Russia's war of aggression against ukraine: 20th round of eu sanctions. Press release, 2026. URL https://www.consilium.europa.eu/en/press/press-releases/2026/04/23/russia-s-war-of-aggression-against-ukraine-20th-round-of-stern-eu-sanctions-hits-energy-military-industrial-complex-trade-and-financial-services-including-crypto/. 23 Apr 2026; accessed 2026-10-15 [P].

[44] Court of Justice of the EU. Case c-264/14 hedqvist, judgment of 22 oct 2015. Web page, 2015. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:62014CJ0264. 22 Oct 2015; accessed 2026-10-15; [P].

[45] Crypto Times. Sen. daines unveils crypto tax bill with stablecoin payment exemption. Web page, 2026. URL https://www.cryptotimes.io/2026/10/01/sen-daines-unveils-crypto-tax-bill-with-stablecoin-payment-exemption/. 1 Oct 2026; accessed 2026-10-15; [S].

[46] crypto.news. Vietnam crypto licences: 5 firms clear first review. Web page, 2026. URL https://crypto.news/vietnam-crypto-licenses-5-firms-clear-first-review/. as of 30 Aug 2026; accessed 2026-10-15; [S].

[47] CryptoRank. South korea sets november review for second-stage crypto law. Web page, 2026. URL https://cryptorank.io/news/feed/638c8-south-korea-sets-november-review-for-second-stage-crypto-law. 22 Sep 2026; low quality; accessed 2026-10-15; [S].

[48] CryptoTicker. Genius act stablecoin rules 2027. Web page, 2026. URL https://cryptoticker.io/en/genius-act-stablecoin-rules-2027/. effective-date summary; accessed 2026-10-15; [S].

[49] Cryptowire. Rbi doubles down on crypto prohibition as parliamentary panel report nears. Web page, 2026. URL https://www.cryptowire.in/blog/rbi-doubles-down-on-crypto-prohibition-as-parliamentary-panel-report-nears/. 2 Aug 2026; accessed 2026-10-15; [S].

[50] Cryptowire. Rbi doubles down on crypto prohibition as parliamentary panel report nears. Blog, 2026. URL https://www.cryptowire.in/blog/rbi-doubles-down-on-crypto-prohibition-as-parliamentary-panel-report-nears/. 2 Aug 2026; accessed 2026-10-15 [S].

[51] Mehul Damani, Isha Puri, Stewart Slocum, Idan Shenfeld, Leshem Choshen, Yoon Kim, and Jacob Andreas. Beyond binary rewards: Training LMs to reason about their uncertainty. In *International Conference on Learning Representations (ICLR)*, 2026. URL https://arxiv.org/abs/2507.16806v1. arXiv:2507.16806, v1 22 Jul 2025.

[52] A. P. Dawid and A. M. Skene. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 28(1):20–28, 1979. doi: 10.2307/2346806. Details not re-checked 15 Oct 2026.

[53] Dechert. Sec's crypto framework: the new token taxonomy. Web page, 2026. URL https://www.dechert.com/knowledge/onpoint/2026/3/sec-s-crypto-framework--the-new-token-taxonomy.html. 31 Mar 2026; accessed 2026-10-15; [S].

[54] Demarest. Bcb regulates the provision of virtual asset services. Web page, 2025. URL https://www.demarest.com.br/wp-content/uploads/2025/11/Demarest_BC_regulamenta_a_prestacao_de_servicos_de_ativos_virtuais_EN.pdf. 10 Nov 2025; accessed 2026-10-15; [S].

[55] Tiffany Ding, Anastasios N. Angelopoulos, Stephen Bates, Michael I. Jordan, and Ryan J. Tibshirani. Class-conditional conformal prediction with many classes. In *Advances in Neural Information Processing Systems 36*, 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/cb931eddd563f8d473c355518ce8601c-Abstract.html. arXiv:2306.09335.

[56] Directorate General of Taxes (Indonesia). Pmk 50/2025: dampak nyata penyesuaian aturan pajak bagi pemain kripto. Web page, 2025. URL https://www.pajak.go.id/id/artikel/pmk-502025-dampak-nyata-penyesuaian-aturan-pajak-bagi-pemain-kripto. article with personal-view disclaimer; accessed 2026-10-15; [P].

[57] Direktorat Jenderal Pajak. Pmk 50/2025: dampak nyata penyesuaian aturan pajak bagi pemain kripto. Article on pajak.go.id (author's personal-view disclaimer), 2025. URL https://www.pajak.go.id/id/artikel/pmk-502025-dampak-nyata-penyesuaian-aturan-pajak-bagi-pemain-kripto. PMK 50/2025 in force 1 Aug 2025; accessed 2026-10-15 [P].

[58] DLA Piper. Senate fails to advance the clarity act. Publication, 2026. URL https://www.dlapiper.com/en/insights/publications/2026/09/senate-fails-to-advance-the-clarity-act. Sep 2026; accessed 2026-10-15 [S].

[59] DLA Piper. Senate fails to advance the clarity act. Web page, 2026. URL https://www.dlapiper.com/en/insights/publications/2026/09/senate-fails-to-advance-the-clarity-act. Sep 2026; accessed 2026-10-15; [S].

[60] Pedro Domingos. MetaCost: A general method for making classifiers cost-sensitive. In *Proceedings of the 5th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 155–164, 1999. doi: 10.1145/312129.312220. Details not re-checked 15 Oct 2026.

[61] Cynthia Dwork and Aaron Roth. The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, 9(3–4):211–407, 2014. doi: 10.1561/0400000042. Details not re-checked 15 Oct 2026.

[62] eco.com. Mica and tokenized rwas: 2026 eu regulatory framework. Web page, 2026. URL https://eco.com/support/en/articles/15254022-mica-and-tokenized-rwas-2026-eu-regulatory-framework. lower-quality source; accessed 2026-10-15; [S].

[63] Ran El-Yaniv and Yair Wiener. On the foundations of noise-free selective classification. *Journal of Machine Learning Research*, 11:1605–1641, 2010. Details not re-checked 15 Oct 2026.

[64] Charles Elkan. The foundations of cost-sensitive learning. In *Proceedings of the 17th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 973–978, 2001. Details not re-checked 15 Oct 2026.

[65] ESMA. Markets in crypto-assets regulation (mica): overview and interim register. Web page, 2026. URL https://www.esma.europa.eu/esmas-activities/digital-finance-and-innovation/markets-crypto-assets-regulation-mica. register updated 7 Oct 2026; accessed 2026-10-15; [P].

[66] ESMA. Statement on the end of transitional periods under mica (esma75-113276571-1679). Web page, 2026. URL https://www.esma.europa.eu/sites/default/files/2026-04/ESMA75-113276571-1679_Statement_on_the_end_of_transitional_periods_under_MiCA.pdf. 17 Apr 2026; accessed 2026-10-15; [P].

[67] European Commission. Dac8: administrative cooperation in taxation for crypto-assets. Web page, 2026. URL https://taxation-customs.ec.europa.eu/taxation/tax-transparency-cooperation/administrative-co-operation-and-mutual-assistance/directive-administrative-cooperation-dac/dac8_en. accessed 15 Oct 2026; accessed 2026-10-15; [P].

[68] European Commission. Targeted consultation on the review of mica. Web page, 2026. URL https://finance.ec.europa.eu/document/download/62be7015-f066-4fac-b74e-71bacdbcc9f5_en?filename=2026-mica-review-targeted-consultation-document_en.pdf. 20 May 2026; closed 31 Aug 2026; accessed 2026-10-15; [P].

[69] European Commission, DG TAXUD. Dac8. Web page, 2026. URL https://taxation-customs.ec.europa.eu/taxation/tax-transparency-cooperation/administrative-co-operation-and-mutual-assistance/directive-administrative-cooperation-dac/dac8_en. Accessed 2026-10-15 [P].

[70] European Parliament and Council of the European Union. Regulation (eu) 2023/1113 on information accompanying transfers of funds and certain crypto-assets. EUR-Lex (metadata), 2023. URL https://eur-lex.europa.eu/eli/reg/2023/1113/oj/eng. Applicable 30 Dec 2024; accessed 2026-10-15 [P].

[71] European Union. Regulation (eu) 2023/1113 (transfer of funds regulation). Web page, 2023. URL https://eur-lex.europa.eu/eli/reg/2023/1113/oj/eng. metadata only; accessed 2026-10-15; [P].

[72] European Union. Regulation (eu) 2023/1114 on markets in crypto-assets (mica). Web page, 2023. URL https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32023R1114. Arts. 2 and 3 and Recital 11 read; accessed 2026-10-15; [P].

[73] European Union. Regulation (eu) 2024/1624 (amlr). Web page, 2024. URL https://eur-lex.europa.eu/eli/reg/2024/1624/oj/eng. metadata only; applies 10 Jul 2027; accessed 2026-10-15; [P].

[74] FATF. Update to recommendation 16 on payment transparency. Web page, 2025. URL https://www.fatf-gafi.org/en/publications/Fatfrecommendations/update-Recommendation-16-payment-transparency-june-2025.html. 18 Jun 2025; accessed 2026-10-15; [P].

[75] FATF. Targeted update on virtual assets and vasps. Web page, 2025. URL https://www.fatf-gafi.org/en/publications/Fatfrecommendations/targeted-update-virtual-assets-vasps-2025.html. 26 Jun 2025; accessed 2026-10-15; [P].

[76] FATF. Seventh targeted update on implementation of the fatf standards on virtual assets and vasps. Web page, 2026. URL https://www.fatf-gafi.org/content/dam/fatf-gafi/reports/7th-targeted-update-on-implementation-fatf-standards-vas-vasps-2026.pdf.coredownload.pdf. 16 Jul 2026; accessed 2026-10-15; [P].

[77] Federal Bureau of Investigation. North korea responsible for $1.5 billion bybit hack. Public Service Announcement PSA250226, 2025. URL https://www.ic3.gov/PSA/2025/PSA250226. 26 Feb 2025; accessed 2026-10-15 [P].

[78] Federal Reserve Board. Implementing the federal reserve board's responsibilities under the genius act. Web page, 2026. URL https://www.federalregister.gov/documents/2026/09/29/2026-19860/implementing-the-federal-reserve-boards-responsibilities-under-the-genius-act. 29 Sep 2026; not fetched; accessed 2026-10-15; [P].

[79] Fenwick. Sec statement on liquid staking. Web page, 2025. URL https://www.fenwick.com/insights/publications/sec-statement-on-liquid-staking. staff statement 5 Aug 2025; accessed 2026-10-15; [S].

[80] Financial Action Task Force. Targeted update on implementation of the fatf standards on virtual assets and virtual asset service providers. FATF report, 2025. URL https://www.fatf-gafi.org/en/publications/Fatfrecommendations/targeted-update-virtual-assets-vasps-2025.html. 26 Jun 2025; accessed 2026-10-15 [P].

[81] Financial Action Task Force. Outcomes fatf plenary, june 2026. FATF publication, 2026. URL https://www.fatf-gafi.org/en/publications/Fatfgeneral/outcomes-fatf-plenary-june-2026.html. Accessed 2026-10-15 [P].

[82] Financial Action Task Force. 7th targeted update on implementation of the fatf standards on virtual assets and virtual asset service providers. FATF report, 2026. URL https://www.fatf-gafi.org/content/dam/fatf-gafi/reports/7th-targeted-update-on-implementation-fatf-standards-vas-vasps-2026.pdf.coredownload.pdf. 16 Jul 2026; accessed 2026-10-15 [P].

[83] Financial Services Agency (Japan). Travel rule materials. Web page, 2025. URL https://www.fsa.go.jp/en/news/2025/20250425/02.pdf. 25 Apr 2025; accessed 2026-10-15; [P].

[84] Financial Services Agency of Japan. Travel rule in japan (fsa material). FSA publication, 2025. URL https://www.fsa.go.jp/en/news/2025/20250425/02.pdf. 25 Apr 2025; accessed 2026-10-15 [P].

[85] Financial Services Commission (Korea). Enforcement decree and regulation of the virtual asset user protection act. Web page, 2023. URL https://www.fsc.go.kr/eng/pr010101/81217. 11 Dec 2023; Act in force 19 Jul 2024; accessed 2026-10-15; [P].

[86] Financial Services Commission (Korea). Enforcement decree of the virtual asset user protection act. Press release, 2023. URL https://www.fsc.go.kr/eng/pr010101/81217. 11 Dec 2023; Act in force 19 Jul 2024; accessed 2026-10-15 [P].

[87] Financial Services Commission (Korea). Guideline on determining whether nfts are virtual assets. Web page, 2024. URL https://fsc.go.kr/eng/pr010101/82419?curPage=13. 10 Jun 2024; accessed 2026-10-15; [P].

[88] Financial Stability Board. High-level recommendations for the regulation, supervision and oversight of global stablecoin arrangements: final report. Web page, 2023. URL https://www.fsb.org/2023/07/high-level-recommendations-for-the-regulation-super

vision-and-oversight-of-global-stablecoin-arrangements-final-report/. 17 Jul 2023; accessed 2026-10-15; [P].

[89] FINMA. Ico guidelines (payment, utility and asset tokens). Web page, 2018. URL https://www.finma.ch/en/news/2018/02/20180216-mm-ico-wegleitung/. 16 Feb 2018; accessed 2026-10-15; [P].

[90] Forvis Mazars. Form 1099-da digital asset basis: what to know. Web page, 2026. URL https://www.forvismazars.us/forsights/2026/09/form-1099-da-digital-asset-basis-what-to-know. Sep 2026; accessed 2026-10-15; [S].

[91] Eibe Frank and Mark Hall. A simple approach to ordinal classification. In *Machine Learning: ECML 2001*, volume 2167 of *Lecture Notes in Computer Science*, pages 145–156. Springer, 2001. doi: 10.1007/3-540-44795-4_13. Details not re-checked 15 Oct 2026.

[92] Freshfields. Treasury proposes genius act gatekeeping rules with a path for foreign issuers. Web page, 2026. URL https://www.freshfields.com/en/our-thinking/blogs/a-fresh-take/treasury-proposes-genius-act-gatekeeping-rules-with-a-path-for-foreign-issuers-102nri4. Aug 2026; accessed 2026-10-15; [S].

[93] João Gama, Indrė Žliobaitė, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4): 44:1–44:37, 2014. doi: 10.1145/2523813. Details not re-checked 15 Oct 2026.

[94] Silvia García-Méndez, Milagros Fernández-Gavilanes, Jonathan Juncal-Martínez, Francisco J. González-Castaño, and Oscar Barba Seara. Identifying banking transaction descriptions via support vector machine short-text classification based on a specialized labelled corpus. *IEEE Access*, 2020. doi: 10.1109/ACCESS.2020.2983584. URL https://arxiv.org/abs/2404.08664v1. arXiv:2404.08664. Journal name inferred from DOI; unverified.

[95] Yonatan Geifman and Ran El-Yaniv. Selective classification for deep neural networks. In *Advances in Neural Information Processing Systems 30*, 2017. arXiv:1705.08500. Details not re-checked 15 Oct 2026.

[96] Yonatan Geifman and Ran El-Yaniv. SelectiveNet: A deep neural network with an integrated reject option. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *PMLR*, pages 2151–2159, 2019. arXiv:1901.09192. Details not re-checked 15 Oct 2026.

[97] Saibo Geng, Martin Josifoski, Maxime Peyrard, and Robert West. Grammar-constrained decoding for structured NLP tasks without finetuning. In *Proceedings of EMNLP 2023*, 2023. arXiv:2305.13971. Details not re-checked 15 Oct 2026.

[98] Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007. doi: 10.1198/016214506000001437. Details not re-checked 15 Oct 2026.

[99] Kate Goddard, Abdul Roudsari, and Jeremy C. Wyatt. Automation bias: A systematic review of frequency, effect mediators, and mitigators. *Journal of the American Medical*

Informatics Association, 19(1):121–127, 2012. doi: 10.1136/amiajnl-2011-000089. Volume, issue and pages not re-checked 15 Oct 2026.

[100] GOLAW. New stage of virtual asset regulation in ukraine. Web page, 2025. URL https://golaw.ua/insights/publication/novij-etap-regulyuvannya-virtualnih-aktiviv-v-ukrayini-shho-zminyuyetsya-dlya-biznesu-ta-investoriv/. law-firm note; accessed 2026-10-15; [S].

[101] Kai Greshake, Sahar Abdelnabi, Shailesh Mishra, Christoph Endres, Thorsten Holz, and Mario Fritz. Not what you've signed up for: Compromising real-world LLM-integrated applications with indirect prompt injection. In *Proceedings of the 16th ACM Workshop on Artificial Intelligence and Security (AISec '23)*, 2023. doi: 10.1145/3605764.3623985. URL https://arxiv.org/abs/2302.12173v2. arXiv:2302.12173; DOI not re-checked 15 Oct 2026.

[102] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *PMLR*, pages 1321–1330, 2017. arXiv:1706.04599. Details not re-checked 15 Oct 2026.

[103] HKMA. Press release: grant of stablecoin issuer licences. Web page, 2026. URL https://www.hkma.gov.hk/eng/news-and-media/press-releases/2026/04/20260410-4/. 10 Apr 2026; accessed 2026-10-15; [P].

[104] HKMA. Stablecoin issuers. Web page, 2026. URL https://www.hkma.gov.hk/eng/key-functions/international-financial-centre/stablecoin-issuers/. Cap. 656 effective 1 Aug 2025; accessed 2026-10-15; [P].

[105] HM Revenue and Customs. Tax treatment of cryptoasset loans and liquidity pools. Policy paper and draft legislation, 2026. URL https://www.gov.uk/government/publications/cryptoasset-loans-and-liquidity-pools/tax-treatment-of-cryptoasset-loans-and-liquidity-pools. Published 13 Jul 2026, updated 30 Sep 2026; effect from 6 Apr 2027 (draft); accessed 2026-10-15 [P].

[106] HM Treasury. Taxation of stablecoins: outcome. Call for evidence outcome, 2026. URL https://www.gov.uk/government/calls-for-evidence/cryptoasset-taxation-stablecoins/outcome/taxation-of-stablecoins. 13 Jul 2026; accessed 2026-10-15 [P].

[107] HM Treasury. Taxation of stablecoins: outcome of call for evidence. Web page, 2026. URL https://www.gov.uk/government/calls-for-evidence/cryptoasset-taxation-stablecoins/outcome/taxation-of-stablecoins. 13 Jul 2026; accessed 2026-10-15; [P].

[108] HM Treasury and HM Revenue and Customs. The taxation of decentralised finance involving the lending and staking of cryptoassets. Consultation page, 2025. URL https://www.gov.uk/government/consultations/the-taxation-of-decentralised-finance-involving-the-lending-and-staking-of-cryptoassets. Outcome 26 Nov 2025; accessed 2026-10-15 [P].

[109] HMRC. Cryptoassets manual crypto21200: staking. Web page, 2026. URL https://www.gov.uk/hmrc-internal-manuals/cryptoassets-manual/crypto21200. accessed 15 Oct 2026; accessed 2026-10-15; [P].

[110] HMRC. Cryptoassets manual crypto21250: airdrops. Web page, 2026. URL https://www.gov.uk/hmrc-internal-manuals/cryptoassets-manual/crypto21250. accessed 15 Oct 2026; accessed 2026-10-15; [P].

[111] HMRC. Cryptoassets manual crypto22100: disposals. Web page, 2026. URL https://gov.uk/hmrc-internal-manuals/cryptoassets-manual/crypto22100. accessed 15 Oct 2026; accessed 2026-10-15; [P].

[112] HMRC. Cryptoassets manual crypto22200: pooling. Web page, 2026. URL https://gov.uk/hmrc-internal-manuals/cryptoassets-manual/crypto22200. accessed 15 Oct 2026; accessed 2026-10-15; [P].

[113] HMRC. Tax treatment of cryptoasset loans and liquidity pools. Web page, 2026. URL https://www.gov.uk/government/publications/cryptoasset-loans-and-liquidity-pools/tax-treatment-of-cryptoasset-loans-and-liquidity-pools. 13 Jul 2026; draft legislation; accessed 2026-10-15; [P].

[114] House Committee on Ways and Means. President trump signs ways and means resolution overturning biden administration's burdensome irs defi broker rule. Press release, 2025. URL https://waysandmeans.house.gov/2025/04/10/president-trump-signs-ways-means-resolution-overturning-biden-administrations-burdensome-irs-defi-broker-rule. 10 Apr 2025; accessed 2026-10-15 [P].

[115] Hukumonline. Membedah pmk 50/2025 dan dampaknya terhadap perpajakan aset kripto di indonesia. Web page, 2025. URL https://www.hukumonline.com/berita/a/membedah-pmk-50-2025-dan-dampaknya-terhadap-perpajakan-aset-kripto-di-indonesia-lt68930d79e525e/. 2025; accessed 2026-10-15; [S].

[116] Income Tax Department, Government of India. Tds on payment for the transfer of virtual digital assets (vdas). Web page, 2026. URL https://www.incometaxindia.gov.in/w/tds-on-payment-for-the-transfer-of-virtual-digital-assets-vdas-. Accessed 2026-10-15 [P].

[117] Income Tax Department (India). Tds on payment for the transfer of virtual digital assets. Web page, 2026. URL https://www.incometaxindia.gov.in/w/tds-on-payment-for-the-transfer-of-virtual-digital-assets-vdas-. accessed 15 Oct 2026; accessed 2026-10-15; [P].

[118] InterWork Alliance. Token taxonomy framework. Web page, 2021. URL https://github.com/InterWorkAlliance/TokenTaxonomyFramework/blob/main/token-taxonomy.md. GitHub specification; accessed 2026-10-15; [P].

[119] IRS. Revenue ruling 2019-24 (hard forks and airdrops). Web page, 2019. URL https://www.irs.gov/pub/irs-drop/rr-19-24.pdf. 2019; accessed 2026-10-15; [P].

[120] IRS. Revenue ruling 2023-14 (staking rewards). Web page, 2023. URL https://www.irs.gov/pub/irs-drop/rr-23-14.pdf. 2023; accessed 2026-10-15; [P].

[121] IRS. Revenue procedure 2024-28 (per-wallet basis allocation). Web page, 2024. URL https://www.irs.gov/pub/irs-drop/rp-24-28.pdf. 2024; accessed 2026-10-15; [P].

[122] IRS. Corrections to the 2025 instructions for form 1099-da: de minimis rules. Web page, 2025. URL https://www.irs.gov/zh-hant/about-irs/corrections-to-the-2025-instructions-for-form-1099-da-de-minimis-rules-for-reporting-certain-sales-of-digital-assets-and-optional-reporting-methods. 2025 instructions; accessed 2026-10-15; [P].

[123] IRS. Frequently asked questions on virtual currency transactions. Web page, 2025. URL https://www.irs.gov/individuals/international-taxpayers/frequently-asked-questions-on-virtual-currency-transactions. accessed 15 Oct 2026; accessed 2026-10-15; [P].

[124] ITSA. International token classification. Web page, 2026. URL https://www.itsa.global/itc/. fetch failed; unverified; accessed 2026-10-15; [S].

[125] Alon Jacovi and Yoav Goldberg. Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4198–4205, 2020. arXiv:2004.03685. Details not re-checked 15 Oct 2026.

[126] Journal of Accountancy. Cryptoasset staking challenge dismissed as moot after government's concession. Journal of Accountancy, 2024. URL https://www.journalofaccountancy.com/issues/2024/sep/cryptoasset-staking-challenge-dismissed-as-moot-after-governments-concession/. Sep 2024; Jarrett v. United States; accessed 2026-10-15 [S].

[127] Journal of Accountancy. Cryptoasset staking challenge dismissed as moot after government's concession. Web page, 2024. URL https://www.journalofaccountancy.com/issues/2024/sep/cryptoasset-staking-challenge-dismissed-as-moot-after-governments-concession/. Sep 2024; accessed 2026-10-15; [S].

[128] Saurav Kadavath et al. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*, 2022.

[129] Sayash Kapoor and Arvind Narayanan. Leakage and the reproducibility crisis in machine-learning-based science. *Patterns*, 4(9):100804, 2023. doi: 10.1016/j.patter.2023.100804. Details not re-checked 15 Oct 2026.

[130] Daniel Martin Katz, Michael James Bommarito, Shang Gao, and Pablo Arredondo. GPT-4 passes the bar exam. *Philosophical Transactions of the Royal Society A*, 382(2270):20230254, 2024. doi: 10.1098/rsta.2023.0254. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC10894685.

[131] Shachar Kaufman, Saharon Rosset, Claudia Perlich, and Ori Stitelman. Leakage in data mining: Formulation, detection, and avoidance. *ACM Transactions on Knowledge Discovery from Data*, 6(4):15:1–15:21, 2012. doi: 10.1145/2382577.2382579. Details not re-checked 15 Oct 2026.

[132] K&L Gates. Crypto in 2026: the democratization of digital assets. Web page, 2026. URL https://www.klgates.com/thought-leadership/Crypto-in-2026-The-Democratization-of-Digital-Assets-1-29-2026. 29 Jan 2026; accessed 2026-10-15; [S].

[133] Pang Wei Koh et al. WILDS: A benchmark of in-the-wild distribution shifts. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *PMLR*, pages 5637–5664, 2021. arXiv:2012.07421. Details not re-checked 15 Oct 2026.

[134] Meelis Kull, Miquel Perelló-Nieto, Markus Kängsepp, Telmo Silva Filho, Hao Song, and Peter Flach. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with Dirichlet calibration. In *Advances in Neural Information Processing Systems 32*, 2019. arXiv:1910.12656. Details not re-checked 15 Oct 2026.

[135] Law Commission. The property (digital assets etc) act 2025 has received royal assent. Web page, 2025. URL https://lawcom.gov.uk/news/the-property-digital-assets-etc-act-2025-has-received-royal-assent/. Dec 2025; not fetched; accessed 2026-10-15; [P].

[136] Ledger Insights. Brazil imposes partial ban on stablecoins and crypto for cross-border payments and fx. Web page, 2026. URL https://www.ledgerinsights.com/brazil-imposes-partial-ban-on-stablecoins-crypto-for-cross-border-payments-and-fx/. 2026; accessed 2026-10-15; [S].

[137] Eunji Lee, Sundong Kim, Sihyun Kim, Sungwon Park, Meeyoung Cha, Soyeon Jung, Suyoung Yang, Yeonsoo Choi, Sungdae Ji, Minsoo Song, and Heeja Kim. Classification of goods using text descriptions with sentences retrieval. *arXiv preprint arXiv:2111.01663*, 2021. URL https://ar5iv.labs.arxiv.org/html/2111.01663.

[138] Zachary C. Lipton, Yu-Xiang Wang, and Alexander J. Smola. Detecting and correcting for label shift with black box predictors. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *PMLR*, pages 3122–3130, 2018. arXiv:1802.03916. Details not re-checked 15 Oct 2026.

[139] Wai Weng Lo, Gayan K. Kulatilleke, Mohanad Sarhan, Siamak Layeghy, and Marius Portmann. Inspection-L: Self-supervised GNN node embeddings for money laundering detection in Bitcoin. *arXiv preprint arXiv:2203.10465*, 2022. URL https://arxiv.org/abs/2203.10465. Journal version unverified.

[140] David Madras, Toniann Pitassi, and Richard Zemel. Predict responsibly: Improving fairness and accuracy by learning to defer. In *Advances in Neural Information Processing Systems 31*, 2018. arXiv:1711.06664. Details not re-checked 15 Oct 2026.

[141] Eric Martínez. Re-evaluating GPT-4's bar exam performance. *Artificial Intelligence and Law*, 2024. URL https://dspace.mit.edu/handle/1721.1/153986. Published online 30 Mar 2024; DOI unverified.

[142] Mattos Filho. Decripto: Receita federal and the oecd. Web page, 2025. URL https://www.mattosfilho.com.br/unico/decripto-receita-federal-ocde/. IN RFB 2.291/2025; accessed 2026-10-15; [S].

[143] Mattos Filho. Decripto: Receita federal and the oecd standard. Article, 2025. URL https://www.mattosfilho.com.br/unico/decripto-receita-federal-ocde/. On IN RFB 2.291/2025; accessed 2026-10-15 [S].

[144] Mayer Brown. Navigating the travel rule for uk cryptoasset businesses. Web page, 2023. URL https://www.mayerbrown.com/en/insights/publications/2023/09/navigating-the-travel-rule-for-uk-cryptoasset-businesses. Sep 2023; accessed 2026-10-15; [S].

[145] Mayer Brown. Navigating the travel rule for uk cryptoasset businesses. Insight, 2023. URL https://www.mayerbrown.com/en/insights/publications/2023/09/navigating-the-travel-rule-for-uk-cryptoasset-businesses. Sep 2023; accessed 2026-10-15 [S].

[146] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *PMLR*, pages 1273–1282, 2017. arXiv:1602.05629. Details not re-checked 15 Oct 2026.

[147] Matthias Minderer, Josip Djolonga, Rob Romijnders, Frances Hubis, Xiaohua Zhai, Neil Houlsby, Dustin Tran, and Mario Lucic. Revisiting the calibration of modern neural networks. In *Advances in Neural Information Processing Systems 34*, 2021. arXiv:2106.07998. Details not re-checked 15 Oct 2026.

[148] Monetary Authority of Singapore. Mas finalises stablecoin regulatory framework. Web page, 2023. URL https://www.mas.gov.sg/news/media-releases/2023/mas-finalises-stablecoin-regulatory-framework. 15 Aug 2023; accessed 2026-10-15; [P].

[149] Hussein Mozannar and David Sontag. Consistent estimators for learning to defer to an expert. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *PMLR*, pages 7076–7087, 2020. arXiv:2006.01862. Details not re-checked 15 Oct 2026.

[150] Rafael Müller, Simon Kornblith, and Geoffrey E. Hinton. When does label smoothing help? In *Advances in Neural Information Processing Systems 32*, 2019. arXiv:1906.02629. Details not re-checked 15 Oct 2026.

[151] Mahdi Pakdaman Naeini, Gregory F. Cooper, and Milos Hauskrecht. Obtaining well calibrated probabilities using Bayesian binning. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence*, pages 2901–2907, 2015. Details not re-checked 15 Oct 2026.

[152] Nagarajan Natarajan, Inderjit S. Dhillon, Pradeep Ravikumar, and Ambuj Tewari. Learning with noisy labels. In *Advances in Neural Information Processing Systems 26*, 2013. Details not re-checked 15 Oct 2026.

[153] National Tax Agency (Japan). Faq on tax treatment of crypto assets (revised december 2025). Web page, 2025. URL https://www.nta.go.jp/publication/pamph/pdf/virtual_currency_faq_03.pdf. 26 Dec 2025; accessed 2026-10-15; [P].

[154] National Tax Agency of Japan. Angoshisan-to ni kansuru zeimujo no toriatsukai ni tsuite (faq) [tax treatment of crypto assets: Faq]. NTA publication, 2025. URL https://www.nta.go.jp/publication/pamph/pdf/virtual_currency_faq_03.pdf. Revised Dec 2025; accessed 2026-10-15 [P].

[155] Jeremy Nixon, Michael W. Dusenberry, Linchuan Zhang, Ghassen Jerfel, and Dustin Tran. Measuring calibration in deep learning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2019. URL https://openaccess.thecvf.com/content_CVPRW_2019/html/Uncertainty_and_Robustness_in_Deep_Visual_Learning/Nixon_Measuring_Calibration_in_Deep_Learning_CVPRW_2019_paper.html. arXiv:1904.01685.

[156] Curtis G. Northcutt, Anish Athalye, and Jonas Mueller. Pervasive label errors in test sets destabilize machine learning benchmarks. In *NeurIPS Datasets and Benchmarks Track*, 2021. arXiv:2103.14749. Details not re-checked 15 Oct 2026.

[157] Curtis G. Northcutt, Lu Jiang, and Isaac L. Chuang. Confident learning: Estimating uncertainty in dataset labels. *Journal of Artificial Intelligence Research*, 70:1373–1411, 2021. arXiv:1911.00068. Details not re-checked 15 Oct 2026.

[158] Notabene. Travel rule: Brazil. Web page, 2026. URL https://notabene.id/world/brazil. 2026; accessed 2026-10-15; [S].

[159] Notabene. Travel rule in brazil. Jurisdiction page, 2026. URL https://notabene.id/world/brazil. Accessed 2026-10-15 [S].

[160] Notabene. Travel rule: Indonesia. Web page, 2026. URL https://notabene.id/world/indonesia. 2026; accessed 2026-10-15; [S].

[161] Notabene. Travel rule: India. Web page, 2026. URL https://notabene.id/world/india. 2026; accessed 2026-10-15; [S].

[162] Notabene. Travel rule in india. Jurisdiction page, 2026. URL https://notabene.id/world/india. Accessed 2026-10-15 [S].

[163] Notabene. Travel rule in indonesia. Jurisdiction page, 2026. URL https://notabene.id/world/indonesia. Accessed 2026-10-15 [S].

[164] Notabene. Travel rule in japan. Jurisdiction page, 2026. URL https://notabene.id/world/japan. Accessed 2026-10-15 [S].

[165] OCC. Implementing the genius act (notice of proposed rulemaking). Web page, 2026. URL https://www.federalregister.gov/documents/2026/03/02/2026-04089/implementing-the-guiding-and-establishing-national-innovation-for-us-stablecoins-act-for-the. 2 Mar 2026; accessed 2026-10-15; [P].

[166] OECD. International standards for automatic exchange of information in tax matters (carf and amended crs). Web page, 2023. URL https://www.oecd.org/en/publications/international-standards-for-automatic-exchange-of-information-in-tax-matters_896d79d1-en.html. abstract only; accessed 2026-10-15; [P].

[167] OECD. Crypto-asset reporting framework and amendments to the common reporting standard. Web page, 2023. URL https://www.oecd.org/tax/exchange-of-tax-information/crypto-asset-reporting-framework-and-amendments-to-the-common-reporting-standard.htm. 8 Jun 2023; accessed 2026-10-15; [P].

[168] Orrick. The clarity act stalls in the senate: what's next for digital asset regulation. Web page, 2026. URL https://www.orrick.com/en/Insights/2026/10/The-CLARITY-Act-Stalls-in-the-Senate-Whats-Next-for-Digital-Asset-Regulation. Oct 2026; accessed 2026-10-15; [S].

[169] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, D. Sculley, Sebastian Nowozin, Joshua V. Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model's uncertainty? evaluating predictive uncertainty under dataset shift. In *Advances in Neural Information Processing Systems 32*, 2019. arXiv:1906.02530. Details not re-checked 15 Oct 2026.

[170] Raja Parasuraman and Victor Riley. Humans and automation: Use, misuse, disuse, abuse. *Human Factors*, 39(2):230–253, 1997. doi: 10.1518/001872097778543886. Details not re-checked 15 Oct 2026.

[171] Patron Accounting. Crypto vda taxation under the income-tax act 2025 rules. Web page, 2026. URL https://www.patronaccounting.com/blog/crypto-vda-taxation-income-tax-act-2025-rules. 2026; accessed 2026-10-15; [S].

[172] Paul Weiss. Sec and cftc release interpretation on application of federal securities laws to crypto assets. Web page, 2026. URL https://www.paulweiss.com/insights/client-memos/sec-and-cftc-release-interpretation-on-application-of-federal-securities-laws-to-crypto-assets. 2026; accessed 2026-10-15; [S].

[173] Fábio Perez and Ian Ribeiro. Ignore previous prompt: Attack techniques for language models. In *NeurIPS ML Safety Workshop*, 2022. URL https://neurips.cc/virtual/2022/65627. arXiv:2211.09527.

[174] Perkins Coie. Stablecoin interest, yield and rewards: Occ proposes sweeping regulations. Web page, 2026. URL https://perkinscoie.com/insights/update/stablecoin-interest-yield-and-rewards-occ-proposes-sweeping-regulations-under. 2026; accessed 2026-10-15; [S].

[175] Plasma. Global stablecoin regulation map: United arab emirates. Web page, 2026. URL https://www.plasma.org/learn/tools/stablecoin-regulation-map/united-arab-emirates. interested party; accessed 2026-10-15; [S].

[176] Plasma. Global stablecoin regulation map: Brazil. Web page, 2026. URL https://www.plasma.org/learn/tools/stablecoin-regulation-map/brazil. interested party; accessed 2026-10-15; [S].

[177] Plasma. Global stablecoin regulation map: Switzerland. Web page, 2026. URL https://www.plasma.org/learn/tools/stablecoin-regulation-map/switzerland. interested party; accessed 2026-10-15; [S].

[178] Plasma. Global stablecoin regulation map: India. Web page, 2026. URL https://www.plasma.org/learn/tools/stablecoin-regulation-map/india. 20 May 2026; interested party; accessed 2026-10-15; [S].

[179] Plasma. Global stablecoin regulation map: Russia. Web page, 2026. URL https://www.plasma.org/learn/tools/stablecoin-regulation-map/russia. 20 May 2026; interested party; accessed 2026-10-15; [S].

[180] Plasma. Global stablecoin regulation map: Singapore. Web page, 2026. URL https://www.plasma.org/learn/tools/stablecoin-regulation-map/singapore. interested party; accessed 2026-10-15; [S].

[181] John C. Platt. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. In Alexander J. Smola, Peter Bartlett, Bernhard Schölkopf, and Dale Schuurmans, editors, *Advances in Large Margin Classifiers*, pages 61–74. MIT Press, 1999. Details not re-checked 15 Oct 2026.

[182] Portal Contabeis. Mp do iof perde validade e altera regras de tributacao financeira. Web page, 2025. URL https://www.contabeis.com.br/noticias/73398/mp-do-iof-perde-validade-e-altera-regras-de-tributacao-financeira/. 16 Oct 2025; accessed 2026-10-15; [S].

[183] PwC Legal. Micar: final guidelines on qualification of crypto-assets as financial instruments. Web page, 2025. URL https://legal.pwc.de/en/news/articles/micar-final-guidelines-on-qualification-of-crypto-assets-as-financial-instruments. final report 17 Dec 2024; accessed 2026-10-15; [S].

[184] Joaquin Quiñonero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D. Lawrence, editors. *Dataset Shift in Machine Learning*. MIT Press, 2009. Details not re-checked 15 Oct 2026.

[185] Maithra Raghu, Katy Blumer, Greg Corrado, Jon Kleinberg, Ziad Obermeyer, and Sendhil Mullainathan. The algorithmic automation problem: Prediction, triage, and human effort. *arXiv preprint arXiv:1903.12220*, 2019. URL https://alphaxiv.org/abs/1903.12220.

[186] Regulation Tomorrow. Official translations: guidelines on qualification of crypto-assets as financial instruments under mica. Web page, 2025. URL https://www.regulationtomorrow.com/2025/03/official-translations-guidelines-on-the-conditions-and-criteria-for-the-qualification-of-crypto-assets-as-financial-instruments-under-mica/. 19 Mar 2025; apply 18 May 2025; accessed 2026-10-15; [S].

[187] Yaniv Romano, Matteo Sesia, and Emmanuel J. Candès. Classification with valid and adaptive coverage. In *Advances in Neural Information Processing Systems 33*, 2020. arXiv:2006.02544. Details not re-checked 15 Oct 2026.

[188] Cynthia Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5):206–215, 2019. doi: 10.1038/s42256-019-0048-x. Details not re-checked 15 Oct 2026.

[189] Russin & Vecchi. Vietnam's crypto-asset regulation framework: from gray zone to a controlled pilot market. Web page, 2026. URL https://www.russinvecchi.com.vn/publication/vietnams-crypto-asset-regulation-framework-from-gray-zone-to-a-controlled-pilot-market/. 2026; accessed 2026-10-15; [S].

[190] Sattax. Hmrc defi tax rules: crypto cgt 2027. Web page, 2026. URL https://sattax.co.uk/2026/07/14/hmrc-defi-tax-rules-crypto-cgt-2027/. 14 Jul 2026; accessed 2026-10-15; [S].

[191] Sattax. Hmrc defi tax rules: crypto cgt from 2027. Blog, 2026. URL https://sattax.co.uk/2026/07/14/hmrc-defi-tax-rules-crypto-cgt-2027/. 14 Jul 2026; accessed 2026-10-15 [S].

[192] SEC. Press release 2026-30: Sec clarifies application of federal securities laws to crypto assets. Web page, 2026. URL https://www.sec.gov/newsroom/press-releases/2026-30-sec-clarifies-application-federal-securities-laws-crypto-assets. 17 Mar 2026; accessed 2026-10-15; [P].

[193] SEC. Press release 2026-76: Sec proposes new regulation crypto assets. Web page, 2026. URL https://www.sec.gov/newsroom/press-releases/2026-76-sec-proposes-new-regulation-crypto-assets. 18 Aug 2026; accessed 2026-10-15; [P].

[194] SEC and CFTC. Interpretation on application of federal securities laws to crypto assets, rel. nos. 33-11412; 34-105020. Web page, 2026. URL https://www.sec.gov/files/rules/interp/2026/33-11412.pdf. 17 Mar 2026; effective 23 Mar 2026; accessed 2026-10-15; [P].

[195] SEC Division of Corporation Finance. Statement on tokenized securities. Web page, 2026. URL https://www.sec.gov/newsroom/speeches-statements/corp-fin-statement-tokenized-securities-012826-statement-tokenized-securities. 28 Jan 2026; not fetched; accessed 2026-10-15; [P].

[196] Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference attacks against machine learning models. In *2017 IEEE Symposium on Security and Privacy (SP)*, pages 3–18, 2017. doi: 10.1109/SP.2017.41. Details not re-checked 15 Oct 2026.

[197] Sidley Austin. Sec staff unveils a playbook for tokenized securities. Web page, 2026. URL https://www.sidley.com/en/insights/newsupdates/2026/01/sec-staff-unveils-a-playbook-for-tokenized-securities. Jan 2026; accessed 2026-10-15; [S].

[198] Sidley Austin. Uk cryptoasset regulation: action points for 2026–2027. Web page, 2026. URL https://www.sidley.com/en/insights/newsupdates/2026/01/uk-cryptoasset-regulation-action-points-for-2026-2027. Jan 2026; accessed 2026-10-15; [S].

[199] Carlos N. Silla and Alex A. Freitas. A survey of hierarchical classification across different application domains. *Data Mining and Knowledge Discovery*, 22(1–2):31–72, 2011. doi: 10.1007/s10618-010-0175-9. Details not re-checked 15 Oct 2026.

[200] Skadden. Fit for purpose? the mica review consultation. Web page, 2026. URL https://www.skadden.com/insights/publications/2026/06/fit-for-purpose. Jun 2026; accessed 2026-10-15; [S].

[201] Rion Snow, Brendan O'Connor, Daniel Jurafsky, and Andrew Y. Ng. Cheap and fast – but is it good? evaluating non-expert annotations for natural language tasks. In *Proceedings of EMNLP 2008*, pages 254–263, 2008. Details not re-checked 15 Oct 2026.

[202] So & Sato. Japan crypto regulation guide. Web page, 2026. URL https://innovationlaw.jp/en/japan-crypto-regulation-guide/. law-firm guide; accessed 2026-10-15; [S].

[203] Sud.ua. Crypto exchanges to receive licences: Rada finalises bill. Web page, 2026. URL https://sud.ua/en/news/ukraine/364597-kriptobirzhi-poluchat-litsenzii-a-kriptovalyutu-legalizuyut-rada-finalizuet-zakonoproekt. undated, 2026; accessed 2026-10-15; [S].

[204] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016. arXiv:1512.00567. Details not re-checked 15 Oct 2026.

[205] Zhi Rui Tam et al. Let me speak freely? a study on the impact of format restrictions on performance of large language models. *arXiv preprint arXiv:2408.02442*, 2024. Details not re-checked 15 Oct 2026.

[206] TechTimes. Japan passes crypto law; etfs could arrive before tax rate drops to 20 percent. Web page, 2026. URL https://www.techtimes.com/articles/320705/20260716/japan-passes-crypto-law-etfs-could-arrive-before-tax-rate-drops-20-percent.htm. 16 Jul 2026; accessed 2026-10-15; [S].

[207] TFTC. South korea travel rule threshold eliminated. Web page, 2026. URL https://www.tftc.io/south-korea-travel-rule-threshold-eliminated-vasp-self-custody. Aug 2026; accessed 2026-10-15; [S].

[208] TFTC. South korea eliminates travel rule threshold; vasp and self-custody changes. News, 2026. URL https://www.tftc.io/south-korea-travel-rule-threshold-eliminated-vasp-self-custody. Aug 2026; accessed 2026-10-15 [S].

[209] The Block. Japan's sbi group launches jpysc. Web page, 2026. URL https://theblock.co/post/405952/japan-sbi-group-launches-jpysc. 24 Jun 2026; accessed 2026-10-15; [S].

[210] Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher D. Manning. Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback. In *Proceedings of EMNLP 2023*, 2023. arXiv:2305.14975.

[211] Ryan J. Tibshirani, Rina Foygel Barber, Emmanuel J. Candès, and Aaditya Ramdas. Conformal prediction under covariate shift. In *Advances in Neural Information Processing Systems 32*, 2019. URL https://papers.nips.cc/paper/8522-conformal-prediction-under-covariate-shift. arXiv:1904.06019.

[212] TokenTax. Crypto taxes in brazil. Web page, 2026. URL https://tokentax.co/blog/crypto-taxes-brazil. 2026; accessed 2026-10-15; [S].

[213] TokenTax. Crypto taxes in brazil. Blog, 2026. URL https://tokentax.co/blog/crypto-taxes-brazil. Accessed 2026-10-15 [S].

[214] Liam Toran, Cory Van Der Walt, Alan Sammarone, and Alex Keller. Scalable and weakly supervised bank transaction classification. *arXiv preprint arXiv:2305.18430*, 2023. URL https://arxiv.org/abs/2305.18430v1.

[215] Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. In *Advances in Neural Information Processing Systems 36*, 2023. URL https://neurips.cc/virtual/2023/poster/71118. arXiv:2305.04388.

[216] TypeSafe. TypeSafe documentation: API reference, confidence, noul, machine-learning primer. TypeSafe documentation, 2026. URL https://docs.typesafe.ai/confidence. Pages /api, /confidence, /primitives/noul, /introduction/machine-learning-primer; accessed 2026-10-15; [P] vendor docs.

[217] UK Government. The financial services and markets act 2000 (cryptoassets) regulations 2026, si 2026/102. Web page, 2026. URL https://www.legislation.gov.uk/uksi/2026/102/pdfs/uksi_20260102_en.pdf. made 4 Feb 2026; full commencement 25 Oct 2027; accessed 2026-10-15; [P].

[218] US Congress. Genius act, public law 119-27. Web page, 2025. URL https://www.govinfo.gov/content/pkg/PLAW-119publ27/html/PLAW-119publ27.htm. approved 18 Jul 2025; accessed 2026-10-15; [P].

[219] US Government. 31 cfr 1010.410 (recordkeeping and travel rule). Web page, 2026. URL https://www.ecfr.gov/current/title-31/subtitle-B/chapter-X/part-1010/subpart-D/section-1010.410. eCFR current; accessed 2026-10-15; [P].

[220] U.S. Securities and Exchange Commission and Commodity Futures Trading Commission. Application of the federal securities laws to certain types of crypto assets and certain transactions involving crypto assets (release nos. 33-11412; 34-105020). Joint interpretation, 2026. URL https://www.sec.gov/files/rules/interp/2026/33-11412.pdf. 17 Mar 2026; effective 23 Mar 2026; accessed 2026-10-15 [P].

[221] US Treasury. Press release sb0225: Garantex, grinex and a7 designations. Web page, 2025. URL https://home.treasury.gov/news/press-releases/sb0225. 14 Aug 2025; accessed 2026-10-15; [P].

[222] US Treasury. Press release sb0057: Tornado cash delisting. Web page, 2025. URL https://home.treasury.gov/news/press-releases/sb0057. 21 Mar 2025; accessed 2026-10-15; [P].

[223] Verkhovna Rada of Ukraine. News release on the virtual assets bill (first reading). Web page, 2025. URL https://www.rada.gov.ua/en/news/News/265708.html. 3 Sep 2025; accessed 2026-10-15; [P].

[224] Verkhovna Rada of Ukraine. Verkhovna rada adopts virtual assets bill no. 10225-d at first reading. News, 2025. URL https://www.rada.gov.ua/en/news/News/265708.html. 3 Sep 2025; accessed 2026-10-15 [P].

[225] Vladimir Vovk, David Lindsay, Ilia Nouretdinov, and Alex Gammerman. Mondrian confidence machine. Working Paper 4, On-line Compression Modelling Project, 2003. URL https://alrw.net/old/old.html. Year unverified (commonly cited 2003).

[226] Vladimir Vovk, Alex Gammerman, and Glenn Shafer. *Algorithmic Learning in a Random World*. Springer, 2005. doi: 10.1007/b106715. Details not re-checked 15 Oct 2026.

[227] Mark Weber, Giacomo Domeniconi, Jie Chen, Daniel Karl I. Weidele, Claudio Bellei, Tom Robinson, and Charles E. Leiserson. Anti-money laundering in Bitcoin: Experimenting with graph convolutional networks for financial forensics. *arXiv preprint arXiv:1908.02591*, 2019. URL https://arxiv.org/abs/1908.02591. KDD '19 Workshop on Anomaly Detection in Finance (venue not re-checked).

[228] WEEX. Japan's 20% crypto tax already enacted, start date not yet set. Web page, 2026. URL https://www.weex.com/learn/articles/japans-20-crypto-tax-already-enacted-start-date-not-yet-set-august-2026-qxihwz2kpzbmekews3cjepl4. Aug 2026; accessed 2026-10-15; [S].

[229] Brandon T. Willard and Rémi Louf. Efficient guided generation for large language models. *arXiv preprint arXiv:2307.09702*, 2023. URL https://arxiv.org/abs/2307.09702v2.

[230] Yamada & Partners. Fy2026 tax reform: separate taxation for crypto assets. Web page, 2026. URL https://www.yamada-partners.jp/reform/r8/k04-separate-taxation-for-crypto-assets. 2026; accessed 2026-10-15; [S].

[231] Huaxiu Yao et al. Wild-Time: A benchmark of in-the-wild distribution shift over time. In *NeurIPS Datasets and Benchmarks Track*, 2022. arXiv:2211.14238. Title wording and author list not re-checked 15 Oct 2026.

[232] Bianca Zadrozny and Charles Elkan. Obtaining calibrated probability estimates from decision trees and naive Bayesian classifiers. In *Proceedings of the 18th International Conference on Machine Learning (ICML)*, pages 609–616, 2001. Details not re-checked 15 Oct 2026.

[233] Bianca Zadrozny and Charles Elkan. Transforming classifier scores into accurate multiclass probability estimates. In *Proceedings of the 8th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 694–699, 2002. doi: 10.1145/775047.775151. Details not re-checked 15 Oct 2026.

[234] Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. Calibrate before use: Improving few-shot performance of language models. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *PMLR*, 2021. arXiv:2102.09690. Details not re-checked 15 Oct 2026.

[235] Lianmin Zheng et al. Judging LLM-as-a-judge with MT-Bench and chatbot arena. In *Advances in Neural Information Processing Systems 36 (Datasets and Benchmarks)*, 2023. arXiv:2306.05685. Details not re-checked 15 Oct 2026.

A Label Dictionary (Edition 1)

This appendix restates every Zlef value with a one-line definition, for annotators. Where the dictionary and the chapter text differ, Section 6 governs.

Q	Value	Definition
Q1	`NAT.L1` / `NAT.L2`	Protocol fee or staking unit of a base layer / of a layer that settles to another chain.
Q1	`STB.FIAT.SINGLE`	Issuer par-redemption claim in one fiat currency, reserve-backed.
Q1	`STB.FIAT.BASKET`	Par claim against a basket of fiat currencies.
Q1	`STB.COMM`	Claim on or reference to a commodity.
Q1	`STB.SYN`	Price target without an issuer reserve claim (collateral, hedge or algorithm).
Q1	`STB.YLD`	Stablecoin wrapper passing yield to holders.
Q1	`TKN.MMF` / `.EQ` / `.DEP` / `.DEBT` / `.RE`	Tokenized fund share / equity or tracker / bank deposit / debt / real estate or other RWA.
Q1	`GOV`	Governance or protocol-use rights; no redemption claim.
Q1	`MEME`	No function beyond trading and community.
Q1	`NFT.UNIQUE` / `NFT.FRAC`	Unique non-fungible token / series or fractionalised units.
Q1	`LST.STAKE` / `LST.RESTAKE`	Receipt for staked / restaked assets.
Q1	`WRP.CUST` / `.BRIDGE` / `.NATIVEBURN`	Custodial wrap / lock-and-mint bridge representation / burn-and-mint transfer.
Q1	`DFI.LP` / `.LEND` / `.DEBT`	Liquidity-pool share / lending receipt / debt position.
Q1	`PRV`	Protocol hides sender, receiver or amount by default or option.
Q1	`CBDC`	Central-bank money or e-money token outside MiCA crypto-asset scope.
Q1	`OTH`	Points, in-game items, unvested or claim rights.
Q1	`UNKNOWN`	Facts do not determine the class.
Q2	`acquire_fiat` / `dispose_fiat`	Buy with / sell for fiat currency.
Q2	`swap_crypto` / `swap_stable`	Exchange one crypto-asset for another / one stablecoin for another.
Q2	`pay_goods_services` / `receive_payment_goods_services`	Pay for / be paid for goods or services.
Q2	`compensation`	Receive pay for work.
Q2	`staking_reward` / `mining_reward`	Protocol reward for staking / mining.
Q2	`airdrop` / `fork`	Unsolicited distribution / new asset from a chain split.
Q2	`wrap_unwrap` / `bridge`	Convert to or from a wrapped representation / move across chains.
Q2	`lp_deposit` / `lp_withdraw`	Add to / remove from a liquidity pool.
Q2	`lend_out` / `lend_return`	Lend an asset / receive it back.
Q2	`borrow` / `repay`	Borrow an asset / repay it.
Q2	`gift`	Transfer without consideration (not family remittance purpose).
Q2	`internal_transfer`	Move between wallets of the same beneficial owner (evidence required).
Q2	`refund_reversal`	Return of an earlier payment.
Q2	`fee`	Fee paid in crypto.
Q2	`loss_theft`	Loss of control or theft.
Q2	`UNKNOWN`	Facts do not determine the event.
Q3	`service_purchase` / `goods_purchase`	Payment for services / goods (incl. contractors).
Q3	`refund`	Return of a prior payment.
Q3	`internal_transfer`	Same-owner movement.
Q3	`fee`	Charge for a service.

Q	Value	Definition
Q3	`payroll`	Pay to workers for work.
Q3	`loan`	Principal lent or repaid.
Q3	`distribution`	Payment to owners or members as owners.
Q3	`investment`	Purchase of an asset for investment.
Q3	`remittance_family`	Person-to-person support to family.
Q3	`UNKNOWN`	Facts do not determine purpose.
Q4	`yes` / `no` / `unknown`	Per leg: asset passes to another owner, a pooled protocol or a custodian for a receipt / same owner shown by records / undetermined.
Q5	`regulated_vasp` / `vasp_unmatched`	VASP found in a register / VASP-like service not matched in any register consulted (whether registration is required is a rules-layer question).
Q5	`self_hosted`	Wallet controlled by its owner's keys.
Q5	`smart_contract_protocol`	A contract, not a person or firm.
Q5	`bank_or_psp`	Bank or payment-service provider.
Q5	`individual` / `business`	Natural person / legal person, institutional type unknown.
Q5	`UNKNOWN`	Undetermined.
Q6	0 / 1 / 2 / 3 / 4	None / on-chain only / + counterparty identity / + underlying document / + reconciled to accounts.
Q7	`low` < `medium` < `high` < `urgent`	Review-queue order; never overrides flag severity.
Q8	seven indicators; annotated `yes`/`no`/`unknown`; model gives a marginal per indicator	`crosses_chain`, `passes_bridge`, `issuer_freezable`, `permissioned_transfer`, `privacy_enhanced`, `touches_mixer`, `touches_unhosted`.

B Edition Log

Edition	Date	Changes
1 (v0.5r5 draft)	15 Oct 2026	First edition. Legal ontology; asset types I–II with classification, tax and travel by jurisdiction; schema with eight questions; methodology; data and evaluation plan; risks; agenda. `travel_surface` annotated as yes/no/unknown (the draft spec had yes/no). Verification gaps: Vietnam, Ukraine, several tax cells (see Section 10, item A).
1 (v0.5r5 draft)	15 Oct 2026	Changes from Paper IV's illustrative schema (Section 6.13): "relationship" becomes a linked-record fact, not a Zlef question; "evidence sufficient" is subsumed by `evidence_completeness`; `economic_purpose` gains `goods_purchase`, `investment` and `remittance_family`; the rule that mechanics are read from records is generalised as P2. `counterparty_type` value `unregistered_vasp` renamed `vasp_unmatched`. Canonical consistency constraints C1–C11 fixed in Section 6.10.

Each later edition adds a row giving the edition number, the date, schema-version changes, legal sources re-verified or changed, and any results. Results are reported only with their preregistration.