Wrench Attack Coding Tool: Methodology

Nov 13, 2026 · El the Freedom Fairy (@el_freedomfairy)

Purpose and scope

The Wrench Attack Coding Tool is how I classify physical attacks on cryptocurrency holders. It takes each 2026 entry in Jameson Lopp's Known Physical Bitcoin Attacks database (147 cases at the time of writing) and codes it against my own definition and codebook (Dataset v2, 13 November 2026), so the cases can be counted, compared and used as evidence.

- **What it is for:** understanding who is targeted, how attacks are planned, how violent they get, who the offenders are, and how the crypto is taken.
- **Where the codes are used:** my research, and Map My Risk, which draws its example cases, figures and red flags from these codes.
- **What it is not:** a complete count of all wrench attacks (many are never reported), or a legal finding about any person or case.

Definition

> A wrench attack is any event in which a person is targeted with physical force, credible threat of physical force, real-world psychological coercion, or deception involving real-world presence, with intent to compel the transfer, disclosure, or surrender of cryptocurrency or cryptographic access credentials (private keys, seed phrases, PINs, or passwords).

A burglary where the victim is absent does not qualify. Forcing someone to open a safe or hand over a PIN is an old crime, and crypto can also be stolen without coercion. What is distinctive about wrench attacks is the combination: physical coercion meets direct control of cryptographic credentials, and the transfer is final and needs no institution. That removes the safeguards that limit similar crimes in traditional finance, such as withdrawal limits, reversals, fraud holds and a teller who might notice. Wrench attacks are a distinctive form of robbery and extortion, not an entirely new category of crime.

My definition is narrower than Lopp's inclusion rule. His database collects physical attacks connected to crypto, so it also holds cases where no crypto was sought, burglaries with nobody home, and killings where a crypto motive is only rumoured. I still code every case, but each one is marked against my definition:

- **Yes:** meets the definition, including attempts that failed.

- **No:** outside the definition (for example, no coercion of a person, or no crypto sought).
- **Unclear:** a crypto motive is suggested but not confirmed by the reporting.

Codebook

The full code legend, with every code explained, is in the companion PDF *How to read the codes*. Dataset v2 uses these codes:

- **Confidence:** A court or law-enforcement verified · B two or more independent credible media sources · C a single credible media report · D an unverified community report.
- **Type A, how the victim was chosen (multiple allowed):** A-OP opportunistic · A-TG targeted (deliberately selected before the attack) · A-RP repeat victimisation · A-PX proxy (a relative or associate used as leverage). A-TG/PX when the holder is targeted and household members are held as well.
- **Type K, how the attacker knew about the victim's crypto (one main code):** K-PUB public profile · K-LEAK leaked or bought data · K-INS insider or personal connection · K-TRADE the victim identified themselves in a trade · K-NONE no prior knowledge · UNKNOWN. When more than one applies, the main code follows the order INS, LEAK, PUB, TRADE.
- **Type B, planning:** B-1 none · B-2 semi-planned · B-3 fully planned · B-4 organised (B-3 plus a transnational network, division of labour, purchased or leaked data, or a remote orchestrator).
- **Type C, physical harm (highest level reached by any victim):** C-0 no physical harm · C-1 minor force · C-2 injury (including drugging) · C-3 grievous harm · C-4 fatal. The levels follow the injury thresholds used to separate common assault, actual bodily harm and grievous bodily harm.
- **Type L, loss of liberty (highest level reached):** L-0 none · L-1 brief restraint, or held for a length not reported · L-2 held for an hour or more · L-3 abducted (forced to move to another place).
- **Weapon:** Yes · No · UNKNOWN, with the type where reported: GUN · FAKE (replica or imitation) · BLADE · STUN · BLUNT · OTHER.
- **Type D, relationship of offender to victim:** D-S stranger · D-P pseudo-stranger (posing as police, a courier, a buyer) · D-A acquaintance · D-I genuine partner or family · D-R romance lure (offender built or faked a romantic or dating relationship to reach the victim) · D-B business associate.
- **Type E, access method (multiple allowed):** E-UW unlocked wallet · E-PW password or PIN · E-SP seed phrase · E-AT forced transaction · E-LN Lightning · E-MX multi-sig · FAILED (reason) when crypto was sought but none was obtained.

Coding rules: UNKNOWN means the reporting doesn't say. N/A means the field doesn't apply. PROVISIONAL means I have not yet reviewed the case; it is shown on the site as "Provisional · under review". Only Types A and E take more than one code.

Coding process

The tool is built through an iterative, human-in-the-loop process. AI does the volume work and the research; the judgement is mine.

1. **I set the rules.** I wrote the definition and codebook and instruct the AI with them.
2. **AI first pass.** The AI codes each case from Lopp's summary and the source he links.
3. **I review and recode.** I read the sources for each case, apply my own judgement, and change any code I disagree with.
4. **My reasoning goes back in.** My recodes are fed back to the AI, which states the rules it infers from my decisions (listed in the next section) so I can confirm or correct them.
5. **AI proposes codes for the rest.** Using those rules, the AI researches each remaining case beyond Lopp's single link and proposes codes. Each proposal is marked "pending your check", with a note on the evidence found and any rule applied.
6. **I decide.** I review every proposal, accept or change it, and mark it Reviewed. Then the cycle repeats as reporting or my rules change.

AI proposals never overwrite a case I have reviewed, and no code is final until I have reviewed it.

Status, 13 November 2026 (Dataset v2): all 147 cases have been recoded under the v2 codebook. I have fully reviewed 11 of them; the other 136 are provisional and use codes proposed under my rules, including the patterns from my reviews. 109 cases meet my definition, 31 fall outside it and 7 are unclear. The tool's figures use all cases, reviewed or not, and each unreviewed case is labelled provisional.

Decision rules learned from my reviews

These are the patterns in my own recoding. The AI applies them to the cases I have not yet reviewed, and I can change them at any time.

Confidence

- Upgrade with corroboration: B when two or more independent outlets report the case; A when there are court documents or an official police or prosecutor statement.
- A "police source" quoted by a newspaper is not an official statement, so it stays at B.

Type A, criminology

- A lure robbery (victim drawn to a meeting) is A-TG: the victim was chosen.
- When household or family members are held as well as the holder, write A-TG/PX.
- When only a relative is taken and the holder is absent, code A-PX.
- Don't use UNKNOWN when the victim was clearly chosen.

Type K, how they knew

- A lure robbery or peer-to-peer trade is K-TRADE: the victim identified themselves by offering to buy or sell.
- A friend, relative, partner, employee or business contact who knew about the crypto is K-INS.
- Wrong-target cases are coded by the source the attacker used, even though that information was wrong.
- If the reporting doesn't say how the attacker knew, code UNKNOWN; A-TG on its own does not show how the victim was found.

Type B, planning

- A masked crew with any sign of preparation is B-3.
- A wrong target or wrong address shows thin research, so B-2.
- A layered trap, a remote organiser, or leaked or purchased data is B-4.
- No operational detail at all is UNKNOWN.

Types C and L, harm and liberty, and weapons

- A weapon threat with no injury is C-0: the weapon is recorded in the Weapon field, not as harm.
- "Beaten" with no detail is C-2; badly beaten or hospitalised with serious injuries is C-3. Sexual violence is always C-3.
- An attempt foiled before any contact is C-0, L-0.
- For an in-person assault where the reporting doesn't mention restraint, L is UNKNOWN, not L-0.
- "Held" or "sequestered" with no length given is L-1; held for an hour or more is L-2; forced into a vehicle or another place is L-3.
- Weapon is No only when the reporting rules a weapon out; otherwise it is UNKNOWN.

Type D, relationship

- Arrested crews with no reported link to the victim, or crews plainly unconnected, are D-S; otherwise UNKNOWN.
- Posing as police, a courier or a buyer is D-P.
- A romance lure (a dating contact or a faked romantic relationship) is D-R. D-I is kept for genuine partners and family, such as the son in the Milan case.

Type E, access method

- N/A only when no crypto was sought.
- When no crypto was obtained, write FAILED with a short description of why, for example "FAILED (wrong address; nothing obtained)".

Sources and research

Lopp's entry and its single link are the starting point, not the end of the research. For each case I, or the AI working under my instructions, look for:

- official statements: police, gendarmerie and prosecutor press releases, indictments and court reports;
- local-language press where the attack happened (much of 2026 is French, but also Spanish, Russian, Turkish, Portuguese and others);
- wire services and national outlets;
- specialist crypto press, which often follows cases through to charges and trials.

Official sources are preferred. Several articles that only repeat one original report count as one source, not several, for the confidence rating.

Paywalls. Many of the outlets that report these cases first and in most detail are behind paywalls (for example Le Parisien, Ouest-France, La Provence, DNA and Kommersant). Where I could not read the original, I relied on other outlets' summaries of it. That can lose detail, such as how long a victim was held or how the target was found, and it limits how far I can verify each case.

Incomplete reporting. First reports are often thin and investigations are still open. Police withhold details, victims often don't come forward, and coverage outside English is harder to find. Many attacks are never reported at all.

For the figures in Map My Risk, my coded cases are cross-checked against published research (GART Research 2026; Ordekian et al. 2024; Chainalysis as a cross-check on family targeting), and where sources disagree the more cautious figure is used.

Limitations and status of codes

Reviews are limited by paywalls and by incomplete reporting. Because of these limitations, the codes given are subject to change. They have been assigned based on my judgement with the evidence available.

Other limitations to keep in mind:

- **One coder.** The judgement is mine. The AI applies my rules, so it is an assistant, not an independent second coder.
- **One database.** The cases come from Lopp's database, so they share its coverage and gaps.
- **UNKNOWN is not "none".** It means the reporting doesn't say, not that something didn't happen.
- **AI errors.** AI research and proposals can be wrong or out of date. Nothing is final until I have reviewed it.
- **Counts are a floor.** Unreported attacks don't appear, so the true numbers are higher.

Version history and review schedule

- **10 November 2026:** definition and codebook set.
- **11 November 2026:** one-page codebook published; first 14 cases reviewed by me; AI proposals for the remaining 133 (60 researched beyond Lopp's link); this methodology written. Later that day, D-R (romance lure) was added to Type D and the romance-lure cases were recoded.
- **13 November 2026 (Dataset v2):** definition wording revised; Type A-TG reworded; new Type K (how they knew), Type L (loss of liberty) and Weapon fields; Type C made a physical-harm scale only; all 147 cases recoded; 11 reviewed by me and my review patterns applied to the rest; coding tool and Map My Risk updated to Dataset v2.
- **Next:** my review of the remaining 136 provisional cases. Each review round becomes v2.1, v2.2 and so on.

Review schedule

- New Lopp entries are coded as they are added.
- Cases are rechecked when new reporting appears (charges, trials, official statements), especially those with UNKNOWN fields, an Unclear definition or C or D confidence.
- The decision rules are updated whenever my recoding shows a new pattern, and the codebook is reviewed at least once a year.