The Trust Problem: How Do Strangers Transact Without Authority?

Module 3: The Trust Problem

Prof. Dr. Joerg Osterrieder

Digital Finance — BSc Course

Module 3 — The Trust Problem: Why trust is the invisible infrastructure of every financial transaction.

Trust Issues

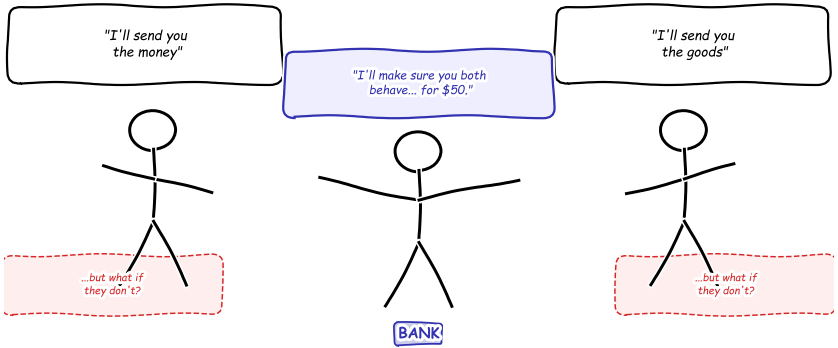

The fundamental trust problem: paying someone to trust for you.

How do you exchange value with someone you've never met and may never meet again?

Learning Objectives

After completing this lecture, you will be able to:

1. **Explain** why trust is the fundamental barrier to economic exchange between strangers [Understand]
2. **Describe** the double-spend problem and why digital trust requires special mechanisms [Apply]
3. **Compare** centralized, hybrid, and decentralized trust architectures with real examples [Analyze]
4. **Apply** game theory (prisoner's dilemma, mechanism design) to evaluate trust systems [Analyze]
5. **Evaluate** the business implications of trustless systems for financial intermediaries [Evaluate]

Bloom's levels covered: Understand, Apply, Analyze, Evaluate

Objectives follow Bloom's taxonomy: Understand → Apply → Analyze → Evaluate.

Bridge: From Cost and Access to Trust

Module 1 asked: *Why do financial services cost so much?*

Module 2 asked: *Who gets left out — and how do we include them?*

Module 3 asks: **How can strangers transact without trusting each other — or any authority?**

The common thread:

- Cost comes from **intermediaries** we trust
- Exclusion happens when **trust gatekeepers** say no
- What if we could **eliminate the need for trust** altogether?

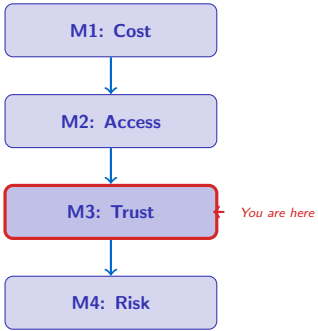

Cost, access, and trust are deeply linked — reducing trust requirements can lower costs and widen access.

The One-Sentence Problem

"How do two strangers exchange value

without trusting each other — or any authority?"

Three provocative sub-questions:

1. Why can't you just email someone €100 like you email a photo?
2. If banks are "trusted," why do they keep failing?
3. Is it possible to build a system where **nobody** needs to be trusted?

This single question drives all of Module 3 — and much of blockchain's reason for existing.

Trust in Everyday Life

Buying a Used Car

"Will it break down next week?"

Trust assumption: The seller is honest about the car's condition.

Mechanism: Inspection, warranty, reputation.

Hiring a Babysitter

"Will they take good care?"

Trust assumption: The sitter is competent and responsible.

Mechanism: References, background checks, cameras.

Lending €100

"Will they pay me back?"

Trust assumption: Your friend is reliable and solvent.

Mechanism: Social pressure, friendship, memory.

Pattern: Every transaction requires a **trust assumption** and a **mechanism** to enforce it.

Trust is so embedded in daily life that we rarely notice it — until it breaks.

How We Solve Trust Today

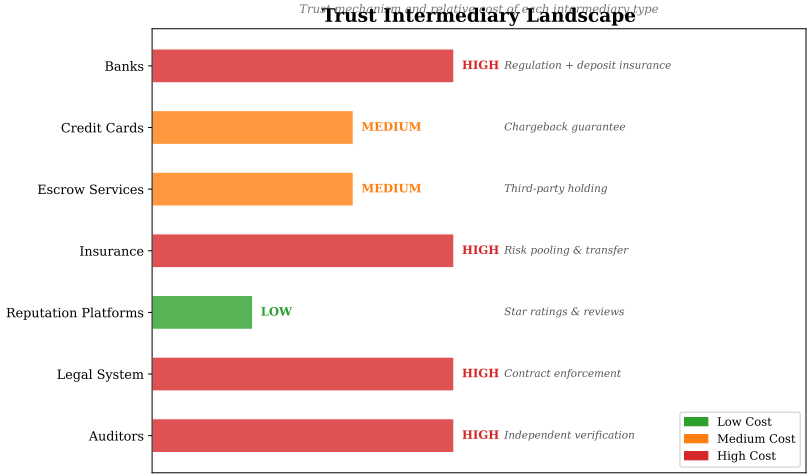

What you see: The ecosystem of trust institutions — banks, courts, regulators, auditors, rating agencies, and insurers, each addressing a different trust gap.

Modern economies use a patchwork of trust mechanisms — each adding cost but reducing uncertainty.

The Hidden Cost of Trust

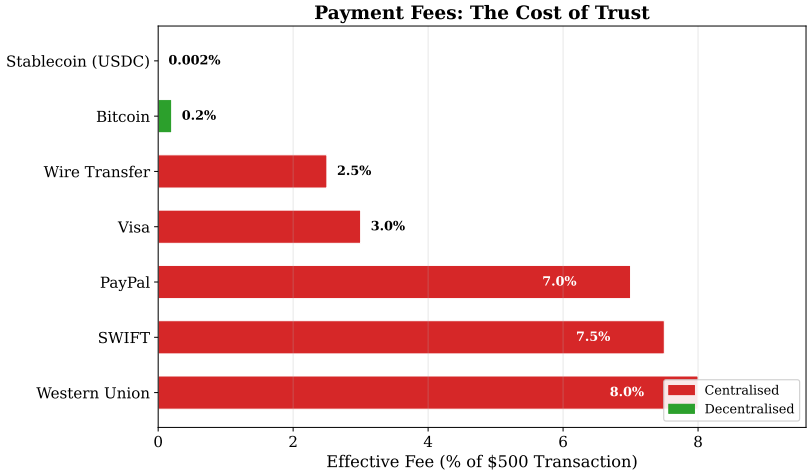

What you see: Fees, spreads, and overhead for providing trust — payment processing (1–3%), compliance (5–10% of revenue), auditing, insurance, and legal costs.

Trust has a price — and it is often hidden inside transaction fees, interest margins, and compliance budgets.

A World of Middlemen

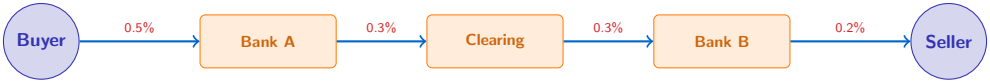

Total: 1.3% per transaction — for trust alone

Each intermediary exists because the buyer and seller do not trust each other directly. Each one takes a fee. The fees compound.

Every middleman charges for trust — cumulative costs can exceed 3% on cross-border payments.

When Trust Breaks Down

Three spectacular failures:

1. **2008 Financial Crisis**
   Banks rated AAA by trusted agencies collapsed. Lehman Brothers: $639B in assets vanished overnight.

2. **Wirecard 2020**
   German payment processor. Audited by EY. €1.9 billion simply did not exist.

3. **FTX 2022**
   "Most trusted" crypto exchange. $8 billion in customer funds misappropriated.

The Pattern

Trusted intermediaries can **fail**, **lie**, or **steal**. The more we concentrate trust, the bigger the damage when it breaks.

Key insight: Trust in institutions is *not* the same as safety. Trust means *exposure* to someone else's honesty.

History shows that "trusted" does not mean "safe" — concentrated trust creates concentrated risk.

The Trust Paradox

"The more you need trust,

the more expensive and fragile it becomes."

- High-value transactions require more verification → more intermediaries → higher cost
- Cross-border trades require trust across **multiple jurisdictions** → more friction
- Concentrating trust in a few institutions creates **systemic risk**
- The unbanked cannot *afford* the trust infrastructure the system demands

---

Trust is simultaneously essential to commerce and a source of cost, delay, and systemic fragility.

Physical vs. Digital: The Copy Problem

Physical Coin

- Hand it over → it's **gone**
- Cannot be in two places at once
- Scarcity is enforced by **physics**

Digital File

- Copy it → **both** have it
- Infinite copies at zero cost
- Scarcity must be enforced by **rules**

The fundamental asymmetry: Physics guarantees scarcity for physical objects. Digital objects need an **authority** to prevent copying.

You can email a photo to 1,000 people — but you cannot email €100 to 1,000 people. That's the problem.

The Double-Spend Problem

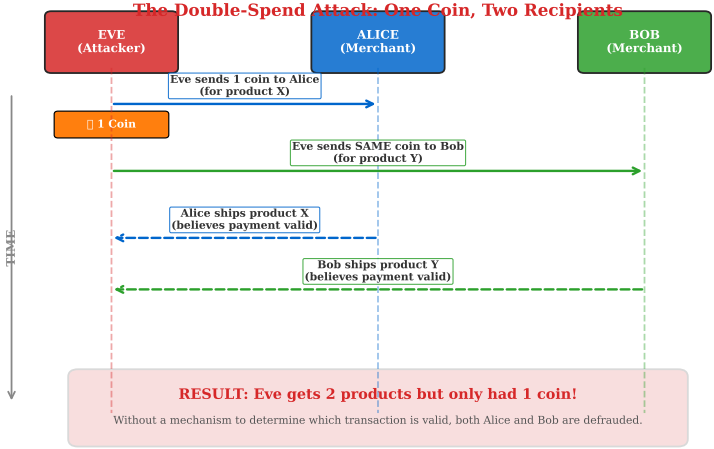

What you see: Eve sends the same coin to Alice and Bob. Without a central authority, both believe they received valid payment.

The double-spend problem is THE core challenge: without a ledger, digital money can be spent twice.

Why Digital Cash Is Hard

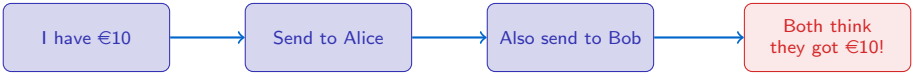

Who really got the money?

The paradox of digital cash:

- A digital coin is just a number — numbers can be copied
- Without a **single source of truth**, there is no way to know which payment is valid
- Traditional solution: a bank keeps the ledger and says "Alice got it, not Bob"

Digital cash requires someone (or something) to maintain a single, authoritative record of who owns what.

The Ledger Solution

Physical Notebook

Digital Database

Every bank is just a **ledger keeper**: "I'll track who owns what so nobody cheats."

The technology changes (stone tablets → paper → databases), but the role is the same.

A ledger is the oldest trust technology in human history — dating back 5,000 years to Mesopotamian clay tablets.

But Who Watches the Ledger-Keeper?

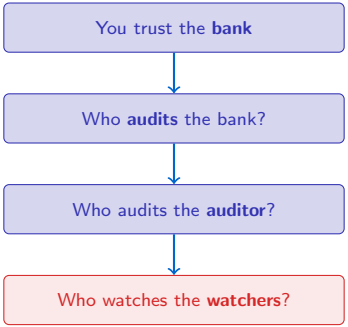

"Quis custodiet ipsos custodes?" — Juvenal, ~100 AD

Centralized trust creates an **infinite regress problem**: every authority needs a higher authority to verify it. At some point, you simply *have to trust someone*.

The recursive trust problem has troubled philosophers for 2,000 years — and it is still unsolved by centralized systems.

Three Centuries of Ledger Evolution

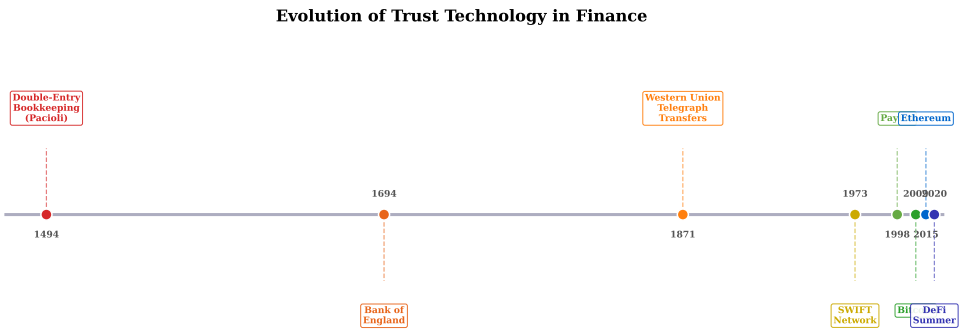

What you see: How ledger technology evolved from physical books (1700s) through centralized databases (1970s) to distributed ledgers (2009+), each wave expanding reach while shifting where trust resides.

Ledger technology evolves in waves — each reducing the trust required from individuals while expanding who can participate.

The Prisoner's Dilemma: A Story

The Setup: Alice and Bob are arrested for a crime. Police separate them and offer each the same deal:

- **Both stay silent** (cooperate): each gets **1 year**
- **One betrays, one is silent**: betrayer goes **free**, silent one gets **10 years**
- **Both betray** (defect): each gets **5 years**

They cannot communicate. Each must choose *without knowing* what the other will do.

The dilemma: Cooperation is best for *both*, but betrayal is best for *each*.

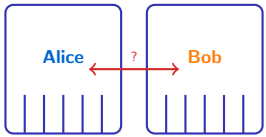

No communication

The prisoner's dilemma is the simplest model of why trust breaks down — and it applies directly to financial transactions.

The Payoff Matrix

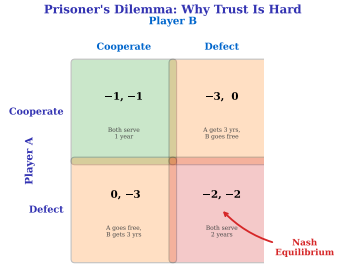

What you see: 2×2 payoff matrix for cooperate/defect. The Nash equilibrium (both defect) is highlighted — rational but suboptimal.

The dominant strategy is to defect — even though mutual cooperation yields a better outcome for everyone.

Nash Equilibrium: Why Rational People Cheat

Alice's reasoning (step by step):

1. If Bob **cooperates** → I should **defect** (0 years vs. 1 year)
2. If Bob **defects** → I should **defect** (5 years vs. 10 years)
3. **Therefore:** defect is always better, regardless of Bob's choice

Bob reasons identically. Both defect.

Key Insight

A **Nash equilibrium** (named after mathematician John Nash) is an outcome where no player can improve by unilaterally changing strategy.

Trust is irrational without enforcement.

Nash equilibrium explains why financial markets need enforcement mechanisms — goodwill alone is not enough.

Repeated Games: How Reputation Builds Trust

One-Shot Game

- You'll never see this person again
- Defection is rational
- No consequences for cheating

Example: Tourist trap restaurant — overcharge, you'll never return.

Repeated Game

- You interact many times
- Cooperation can emerge via **tit-for-tat**
- Reputation becomes valuable

Example: eBay ratings, Airbnb reviews — cheat once, lose future business.

Insight: Reputation systems turn one-shot games into repeated games — making cooperation rational. But they require a **platform** to keep score.

Reputation transforms the incentive structure — but creates dependency on whoever controls the ratings.

Mechanism Design: Engineering Trust

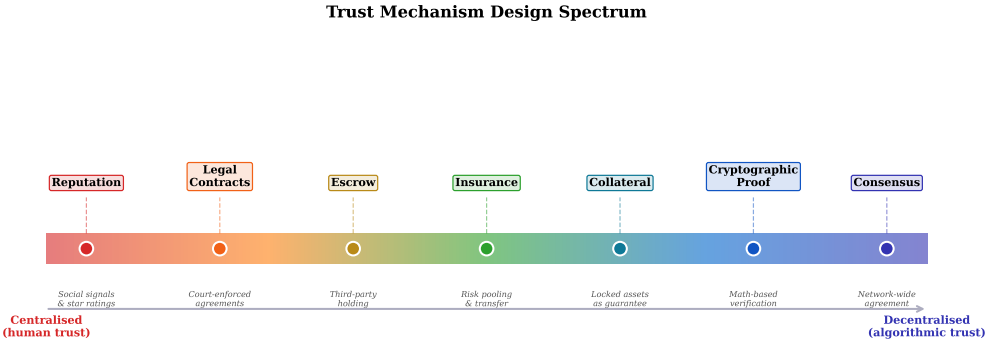

What you see: A spectrum of trust mechanisms — from pure reputation (low enforcement) through escrow and collateral to smart contracts (automated enforcement). Each trades off flexibility for reliability.

Mechanism design (a field of economics): design the rules so that **cheating costs more than cooperating**.

Mechanism design asks: can we engineer the rules so that self-interested actors produce a good outcome?

Think-Pair-Share: Design Your Own Trust Mechanism

Scenario

You run an **online marketplace**. Buyers and sellers don't trust each other. Buyers worry about receiving fake products. Sellers worry about chargebacks and fraud.

Three guiding questions:

1. What **rules** would you set to make the marketplace safe?
2. What **incentives** would you offer for honest behaviour?
3. What **penalties** would you impose for cheating?

Format: 2 min think individually → 2 min discuss with a partner → 2 min share with class.

Mechanism design is not just theory — every marketplace from eBay to Uber uses these principles daily.

From Theory to Practice

Game Theory Concept		Real-World Application
Prisoner's dilemma	→	Online marketplace fraud (buyer vs. seller)
Repeated games	→	Credit scores, eBay ratings, Uber driver ratings
Mechanism design	→	Escrow services, smart contracts, insurance
Byzantine fault tolerance	→	Distributed consensus in blockchain networks

Key pattern: Every trust solution is an answer to a game theory problem. The question is always: *how do we make cooperation the dominant strategy?*

Game theory is not abstract — it is the mathematical foundation of every payment system, marketplace, and blockchain.

The Byzantine Generals Problem

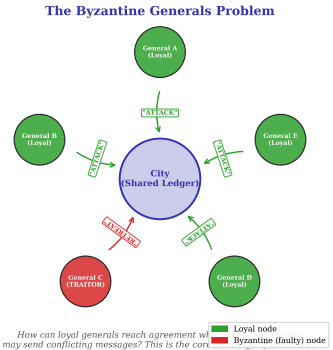

How can loyal generals reach agreement wh... may send conflicting messages? This is the cor...

What you see: Generals must agree on attack or retreat, but traitors may send false messages — reaching consensus despite unreliable actors.

The Byzantine Generals Problem (Lamport, 1982) formalizes consensus with untrusted participants — this is the problem Bitcoin solves.

Centralized Trust: Banks and Visa

How it works:

- **Single authority** maintains the ledger
- All participants trust that authority
- Regulated by government oversight
- Example: Visa processes ˜65,000 transactions/second

Pros:

- Fast and efficient
- Insured (deposit guarantees)
- Familiar to users

Cons:

- Single point of failure
- Fees and rent-seeking
- Can exclude users (the unbanked)

Centralized trust = fast and reliable, but requires **faith in the center**.

Centralized trust is the dominant model today — efficient when it works, catastrophic when it fails.

Hybrid Trust: Platforms and Marketplaces

The Platform Model:

- eBay, Airbnb, Uber act as **referees**
- They don't own the goods or services
- They provide **rules + reputation + dispute resolution**
- Trust is distributed between platform and community ratings

Reputation System Mechanics:

- Mutual ratings (buyer and seller rate each other)
- History builds a public track record
- Bad actors lose visibility and access
- Platform holds escrow during transaction

But: The platform itself must be trusted — it controls the rules, the data, and the escrow.

Platforms reduce person-to-person trust requirements but create a new trust dependency on the platform itself.

Decentralized Trust: The Blockchain Idea

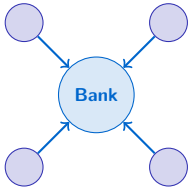

Central Ledger

One keeper, one copy, one point of failure.

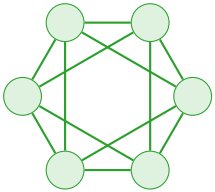

Distributed Ledger

Everyone keeps a copy. No single keeper.

The blockchain proposition: Replace the trusted authority with **mathematics and consensus**.

Decentralized trust eliminates the single point of failure — but introduces new challenges in coordination and speed.

The Trust Spectrum

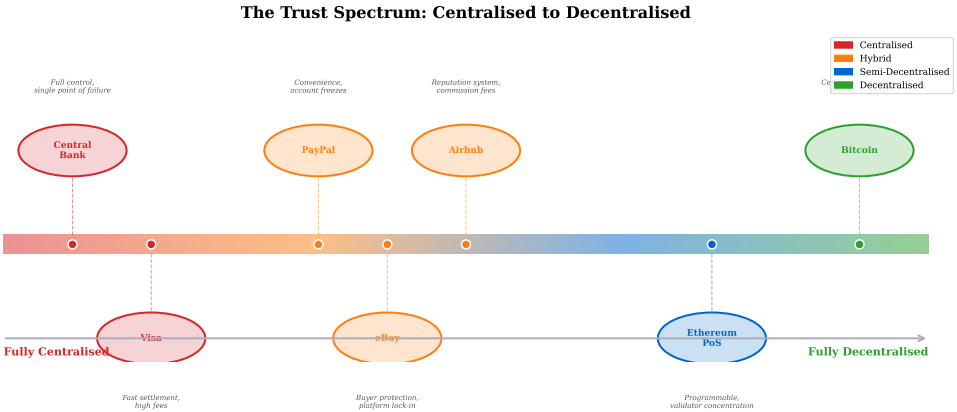

What you see: Trust is not binary (centralized vs. decentralized). Real systems exist on a spectrum — from fully centralized (banks) through hybrid (platforms) to fully decentralized (Bitcoin). Most innovations sit somewhere in between.

The future is not "centralized vs. decentralized" — it is finding the right point on the spectrum for each use case.

How Blockchain Works (30,000-Foot View)

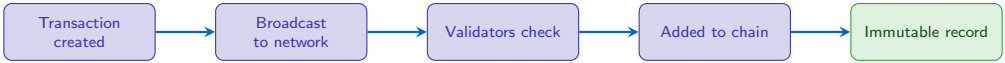

- No single validator can approve a transaction alone — the **network** must agree
- Once recorded, a transaction **cannot be altered** without redoing all subsequent work
- Trust shifts from "I trust the bank" to "I trust the math and the protocol"

No cryptography details here — those come in Lesson 3.1. This is the conceptual model only.

Blockchain replaces institutional trust with mathematical verification — the details come in the next four lessons.

Consensus Comparison

Consensus Mechanism Comparison

Property	Proof of Work (PoW)	Proof of Stake (PoS)	BFT Variants (PBFT/Tendermint)
Resource consumed	Electricity (computational)	Staked capital (economic)	Network messages (communication)
Sybil resistance	Cost of hardware + electricity	Cost of acquiring stake	Permissioned membership
Fault tolerance	Up to 50% hash power	Up to 33% of stake	Up to 33% of nodes
Finality	Probabilistic (~6 blocks)	Faster probabilistic (~2 epochs)	Deterministic (immediate)
Energy use	Very High	Very Low	Very Low
Throughput (TPS)	~7 TPS	~30 TPS	~1,000+ TPS
Decentralization	High (permissionless)	High (permissionless)	Low (permissioned)
Example	Bitcoin	Ethereum (post-Merge)	Hyperledger Fabric

What you see: Consensus mechanisms (PoW, PoS, delegated/Byzantine) compared on energy, speed, decentralization, and security. Each makes different trade-offs.

Consensus is the heart of any blockchain — it determines who gets to write the next page of the ledger.

The Blockchain Trilemma

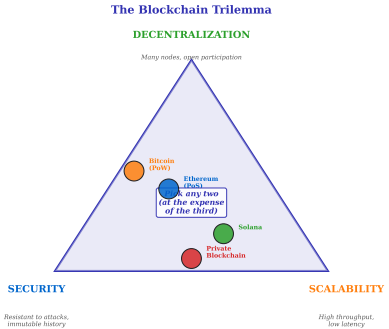

What you see: The "impossible triangle" — optimize at most two of three: decentralization, security, scalability. Bitcoin picks security + decentralization; Solana picks scalability + security.

The trilemma (coined by Vitalik Buterin) explains why no single blockchain dominates every use case.

The Cost of Trust: An Estimated ~$600 Billion Industry

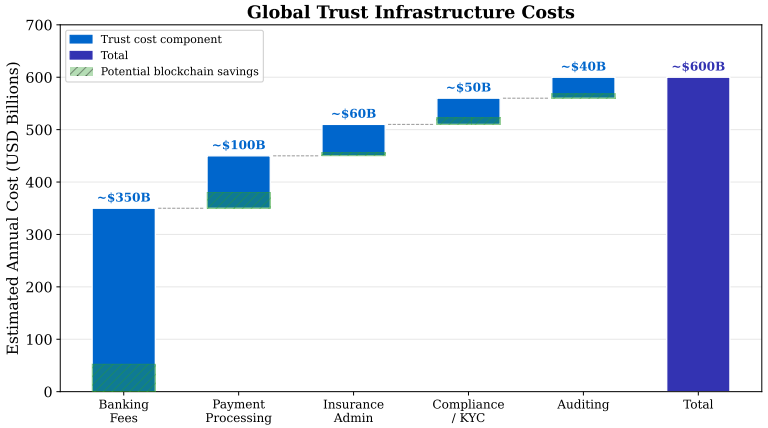

Illustrative estimates from industry reports. Green hatching = potential blockchain savings.

What you see: Global trust spending — banking fees (~$350B), payment processing (~$100B), insurance (~$60B), compliance (~$50B), auditing (~$40B).

Illustrative estimates from McKinsey, Thomson Reuters, IBIS World. Trust is one of the largest hidden costs in the global economy.

Disintermediation: Cutting Out the Middleman

Traditional Chain

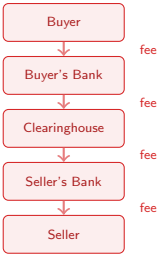

4 intermediaries, 3–5 days, \$30–50 in fees.

Blockchain Chain

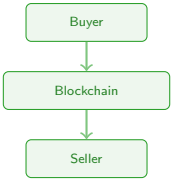

0 intermediaries, seconds to minutes, minimal fees.

Disintermediation (removing middlemen) is the core economic promise of blockchain technology.

Disintermediation does not eliminate trust — it replaces trust in people with trust in code and mathematics.

New Business Models Enabled by Trustless Systems

Application	Replaces	How
DeFi lending (Aave, Compound)	Bank	Automated lending pools, algorithmic rates
Tokenized real estate	Broker	Fractional ownership via tokens
DAOs (Decentralized Autonomous Organisations)	Board of directors	Token-holder voting, on-chain governance
Supply chain provenance	Auditor	Immutable record from farm to shelf
Self-sovereign identity	Government ID	User controls own credentials
Programmable money	Escrow agent	Smart contracts release funds on conditions

Trustless systems enable entirely new business models — not just cheaper versions of old ones.

Case Study: Cross-Border Payments

	SWIFT (Traditional)	Stablecoins (Blockchain)
Speed	3–5 business days	10 seconds – 2 minutes
Cost	$30–50 per transfer	<$0.01
Intermediaries	3–5 correspondent banks	Peer-to-peer
Transparency	Opaque (status unknown)	Fully transparent on-chain
Availability	Banking hours only	24/7/365
Minimum amount	Practically >$100	Any amount

Real impact: Migrant workers sending $200 home pay ~6.4% in fees via traditional channels (World Bank, 2024). Stablecoins reduce this to near zero.

Cross-border payments are the clearest use case where blockchain trust is already cheaper and faster than traditional trust.

Discussion: Who Benefits Most?

Class Activity

Pick **ONE** intermediary: bank, insurer, auditor, or payment processor.

Question: Could blockchain *fully* replace it within 10 years? Why or why not?

Consider:

- What trust function does this intermediary actually perform?
- Can that function be automated with code?
- What regulatory or social barriers exist?

Format: Show of hands for your chosen intermediary → 3-minute small-group discussion → 2-minute class share-out.

Not all intermediaries are equally vulnerable to disruption — the key question is whether their trust function can be coded.

Case Study: Supply Chain Transparency

Walmart + IBM Food Trust: Tracking food from farm to store shelf using blockchain.

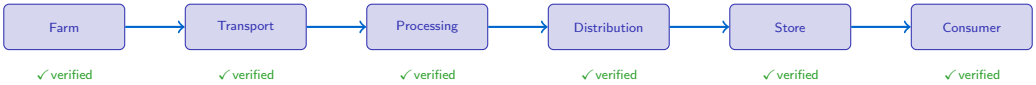

Result: Food origin tracing reduced from **7 days** to **2.2 seconds**.

Why it matters: In a contamination event (e.g., E. coli outbreak), speed of tracing saves lives and prevents waste from recalling unaffected products.

Supply chain blockchain replaces trust in each intermediary's paperwork with a shared, immutable record visible to all parties.

Who Wins, Who Loses?

Winners

- **Consumers:** lower fees, faster service
- **Developers:** new business opportunities
- **Unbanked:** access without gatekeepers
- **Regulators:** real-time transparency into transactions

Losers

- **Rent-seeking intermediaries:** business model threatened
- **Opaque businesses:** transparency exposes inefficiency
- **Fraudsters:** immutable records make fraud harder
- **Technophobic incumbents:** adapt or lose market share

The pattern: Trustless systems transfer value from *intermediaries who profit from opacity* to *users who benefit from transparency*.

Disruption is not about eliminating trust — it is about eliminating the premium charged by those who currently provide it.

What Blockchain Doesn't Solve

Five limitations to remember:

1. **Oracle problem** — Blockchain cannot verify facts from the real world (weather, prices, delivery confirmation) without trusting an external data source
2. **Governance** — Who decides to upgrade the protocol? Decentralized governance is slow and contentious
3. **Key management** — Lose your private key = lose everything. No "forgot password" button
4. **Regulation** — Legal uncertainty across jurisdictions. Which country's laws apply to a global ledger?
5. **Scalability** — Bitcoin: ~7 TPS. Ethereum: ~30 TPS. Visa: ~65,000 TPS. The gap is still large

Honest assessment: Blockchain shifts trust — it does not eliminate it.

Every technology has trade-offs. Blockchain reduces some trust costs while introducing new ones.

The Oracle Problem: Garbage In, Garbage Out

On-Chain (Trustless)

- Math guarantees correctness
- Immutable once recorded
- Transparent and auditable
- Example: token transfer, smart contract logic

Off-Chain (Trust Required)

- Weather data — who reports it?
- Asset prices — which source?
- Delivery confirmation — who verifies?
- Example: insurance payout trigger

A blockchain is only as trustworthy as its data inputs. An **oracle** (a service that feeds real-world data to a blockchain) reintroduces a trust dependency.

The oracle problem is the Achilles' heel of smart contracts — trustless code still needs trustworthy data.

The Trust Landscape in 2025

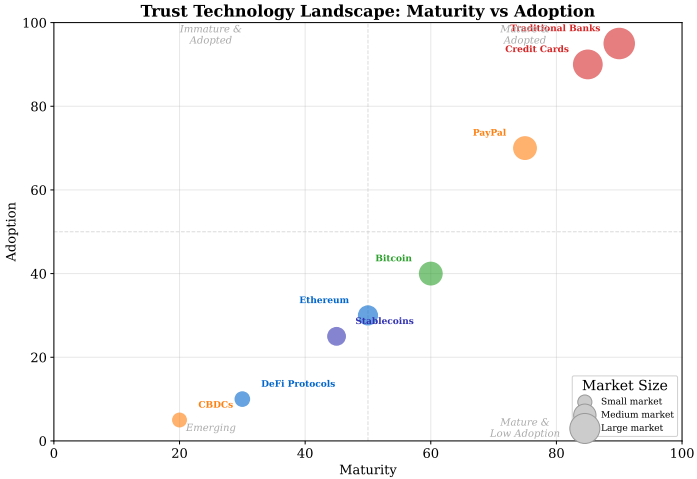

What you see: Maturity-vs-adoption map — from mature systems (banks, credit cards) to emerging solutions (DeFi, DAOs). Traditional systems still dominate.

Most trust innovation is still early-stage — the window for builders and entrepreneurs is wide open.

Key Takeaways

1. **Trust is the invisible tax** on every transaction — intermediaries charge for providing it
2. **The double-spend problem** makes digital trust uniquely hard — you can copy files but not coins
3. **Game theory explains** why we need enforcement mechanisms, not just goodwill
4. **Solutions exist on a spectrum** from fully centralized (banks) to fully decentralized (Bitcoin)
5. **Blockchain reduces trust costs** but introduces new trade-offs (oracles, key management, scalability)

Trust is being re-engineered — understanding the spectrum of solutions is essential for any career in finance.

What Comes Next: Your Module 3 Journey

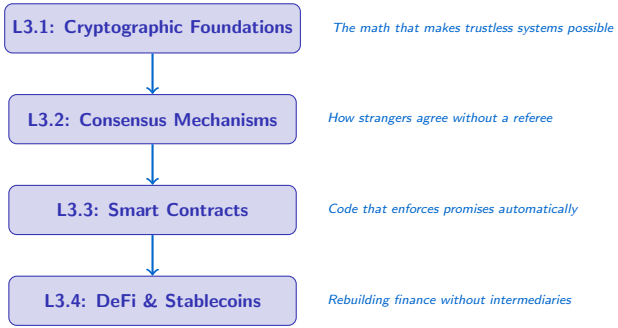

You now understand WHY. Next: the HOW.

Each lesson builds on the last — from cryptographic primitives to a fully functioning decentralized financial system.

Trust Is Being Rewritten

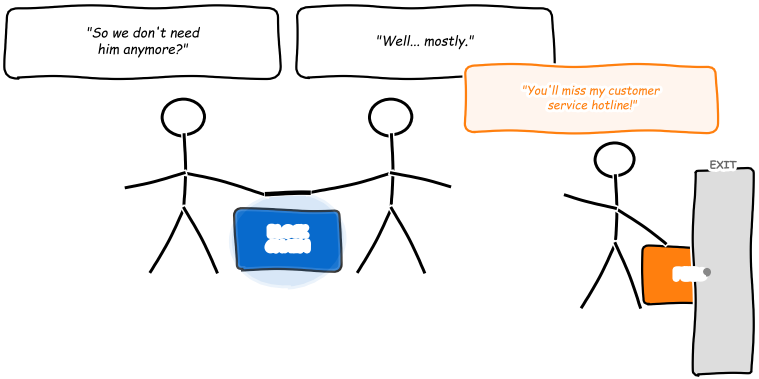

Blockchain: cutting out the middleman, but keeping the complexity.

Trust is being rewritten. Your generation will decide how.