

Dispersed Knowledge

A Series on the Economics of Information: Part 1

Steph Macurdy
Institute of Thermoeconomics

October 2026

Abstract

This paper reconstructs a central progression in modern economic theory through seven questions. The sequence begins with Hayek’s problem of coordinating dispersed knowledge; proceeds through asymmetric information, costly information acquisition, mechanism design, market microstructure, and information design; and ends with the question of how costly computation can transform private claims into publicly verifiable evidence. The argument is cumulative. Prices coordinate without reproducing all underlying knowledge, but asymmetric information can degrade exchange. Signals and mechanisms can elicit otherwise hidden information, but information must first be produced and its acquisition rewarded. Trading converts information into prices, while institutional rules determine what participants can observe and infer. Proof-of-work systems add a distinct capability of making a *narrowly specified costly computation* and the resulting protocol state *publicly auditable*. This does not make markets perfectly informed or establish the truth of arbitrary claims. It creates a distributed verifiable signal within a bounded policy regime.

1 Introduction: Seven Questions

The problem of rational economic order is often stated as if society possessed a complete inventory of resources, preferences, technologies, and constraints and merely needed to calculate an optimal allocation. Hayek’s objection was that this formulation assumes away the real problem. Relevant knowledge is dispersed among people, much of it is local or temporary, and no single participant possesses the whole (Hayek, 1945). Once this point is accepted, economic coordination becomes a problem not only of allocation but also of discovery, communication, incentives, and institutional design.

Seven questions organize the intellectual progression developed here.

- Hayek: How is dispersed knowledge coordinated?
- Akerlof–Spence–Stiglitz: What happens when information is asymmetric?
- Grossman–Stiglitz: Who pays to produce information reflected in price?
- Myerson–Milgrom–Weber: What rules induce strategic agents to reveal information?
- Kyle–Glosten–Milgrom: How does information become price through actual trading?
- Modern information design: Who chooses what agents can observe?

The seventh question extends the sequence from economic communication to computational verification.

How can costly computation transform private claims into publicly verifiable evidence?

The answer requires care. A computational proof does not automatically establish that an external proposition is true or that a market price contains all relevant knowledge. It can, however, establish that specified work was performed and that a state transition satisfied publicly known rules. The resulting evidence can be checked by participants who neither performed nor directly observed the original work. This capability extends the economics of information into the design of decentralized verification systems. Put briefly, the paper distinguishes making a claim from producing evidence that others can check. Throughout the discussion, Appendix A provides a running example of how private claims acquire credibility, consequences, and public verifiability.

1.1 How Is Dispersed Knowledge Coordinated?

Hayek's *The Use of Knowledge in Society* begins from the fragmentation of knowledge (Hayek, 1945). Individuals know particular facts of time and place—local scarcities, capabilities, preferences, substitutions, and opportunities. Because these facts are distributed and continually changing, rational economic order cannot depend upon transmitting a complete description of the economy to a central calculator.

Prices provide a partial solution. A price summarizes the consequences of many judgments without disclosing their entire content. A participant can respond to a rise in the price of an input without knowing whether it resulted from depleted supply, increased demand, transportation failure, or a newly discovered alternative use. The price reduces the amount of contextual knowledge required for coordinated adaptation.

This compression is useful precisely because it is incomplete. A price is not a lossless representation of the world. It is a scalar signal produced by heterogeneous beliefs, endowments, constraints, and actions. Hayek's account therefore yields the following transformation.

dispersed local knowledge $\longrightarrow$ decentralized action $\longrightarrow$ common price signal.

The unresolved question is whether this signal remains reliable when participants possess systematically different information. Hayek explains why decentralized coordination is necessary, but not all conditions under which market-signals are produced can be definitively audited or proved, independent of distortion and/or strategic manipulation. Appendix A makes this limitation concrete by distinguishing a publicly observable assertion from the private information or belief that produced it.

1.2 What Happens When Information Is Asymmetric?

The modern economics of asymmetric information shows that differences in knowledge can alter the composition of markets, the actions chosen by participants, and even the existence of equilibrium. Three canonical responses are adverse selection, signaling, and screening.

In Akerlof's market for used cars, sellers know more about quality than buyers (Akerlof, 1970). Buyers therefore offer a price based on expected rather than observed quality. That price may be too low for owners of high-quality goods, causing them to withdraw. Their withdrawal lowers average market quality, which further lowers the price. Information asymmetry consequently changes which goods are offered in a predictable sequence.

hidden quality $\longrightarrow$ average pricing $\longrightarrow$ high-quality exit $\longrightarrow$ market deterioration.

Spence reverses the direction of the problem by asking how an informed party can communicate a hidden type (Spence, 1973). A signal becomes informative when its cost or feasibility differs across types. Education, in the original model, can communicate productivity when it is less costly for more productive workers to acquire. Costliness alone is not sufficient. The cost must be correlated with the hidden characteristic. An equally wasteful action for every type would burn resources without separating anyone.

Rothschild and Stiglitz examine screening in insurance markets (Rothschild and Stiglitz, 1976). The uninformed party offers a menu of contracts designed so that differently informed types select different options. Instead of asking the informed agent to make a bare assertion, the institution makes revelation consequential. This yields a fundamental distinction between signaling and screening.

signaling: the informed party chooses a revealing action,
screening: the uninformed party designs choices that separate types.

These papers displace the idea that information is simply given. Economic outcomes depend on who knows what, whether knowledge can be credibly communicated, and how institutions connect messages to costs and consequences. The examples in Appendix A follow this distinction from cheap talk to signaling and screening, showing why a favorable claim becomes informative only when types face different costs or choices.

1.3 Who Pays to Produce the Information Reflected in Price?

Grossman and Stiglitz add the incentive to acquire information (Grossman and Stiglitz, 1980). Suppose a trader can pay a cost $c > 0$ to observe a signal about an asset's payoff. Trading on that signal moves the market price, allowing uninformed traders to infer part of the private information. If the price revealed the signal perfectly, an uninformed trader could receive its full benefit without paying c . The informed trader would lose the return that justified producing the information.

The contradiction can be stated compactly.

perfect price revelation $\implies$ no private informational advantage
 $\implies$ no incentive to acquire costly information.

Prices can therefore be highly informative without being perfectly informative. Grossman and Stiglitz describe an “equilibrium degree of disequilibrium”: enough information enters price to coordinate trading, while enough uncertainty remains to compensate those who acquire information. The fraction of informed traders is endogenous. If information is unusually profitable, more traders acquire it; their trading makes price more revealing, which reduces the informational advantage and discourages further acquisition.

This result extends Hayek in a decisive way. Hayek explains how a price can communicate dispersed knowledge. Grossman and Stiglitz ask what economic return sustains the discovery behind that price. Information embodied in a publicly observable price may be inexpensive to consume while remaining costly to produce.

The conclusion is not that inefficiency is always desirable or that every apparent mispricing finances useful discovery. It is narrower—perfect informational efficiency cannot generally be sustained when the underlying information is costly and privately produced. Appendix A carries

the same logic into its treatment of costly commitments, where credibility depends on what the claimant must risk, spend, or forgo.

1.4 What Rules Induce Strategic Agents to Reveal Information?

Once information is private, an institution cannot assume that participants will report it truthfully. Mechanism design therefore treats the rules of interaction as objects of deliberate construction.

Myerson's optimal-auction framework begins with privately known valuations and asks which allocation and payment rules make truthful reporting compatible with individual incentives (Myerson, 1981). The revelation principle allows the designer to study direct mechanisms in which truthful reporting is an equilibrium, but it does not make truthfulness costless or automatic. Truth emerges only because the rule connects a report to expected allocation and payment consequences.

Milgrom and Weber broaden auction theory to interdependent and affiliated values (Milgrom and Weber, 1982). One bidder's information may be relevant to the value faced by others. Auction formats therefore differ in how they aggregate and disclose information. Their linkage principle shows, under specified affiliation and regularity conditions, why linking payment more closely to information about value can improve expected revenue and reduce the informational disadvantage embodied in the winner's curse.

Myerson and Satterthwaite identify a boundary on institutional design (Myerson and Satterthwaite, 1983). With a buyer and seller who privately know their valuations, no mechanism can generally combine Bayesian incentive compatibility, voluntary participation, budget balance, and fully efficient trade. This result is important because it prevents mechanism design from being interpreted as omnipotent engineering. Private information creates real impossibility frontiers.

The lesson of this stage is that a market rule is not a neutral container. Its institutional choices can be summarized as follows.

what may be reported, what actions follow, and who pays;
who receives, and which reports become credible.

Mechanism design thus converts the knowledge problem into an institutional question; what mapping from private messages and observable actions to allocations and payments makes desired behavior individually rational? Appendix A develops this question as a six-stage process that moves from a private observation to a claim, a commitment, a verification rule, a consequence, and finally a public state.

6 How Does Information Become Price Through Actual Trading?

Mechanism design studies rules and incentives, but market microstructure examines the concrete process by which orders become prices. Kyle and Glosten–Milgrom provide two foundational accounts.

In Kyle's model, an informed trader submits orders alongside noise traders, while a market maker observes only aggregate order flow (Kyle, 1985). The informed trader wants to profit from private knowledge without revealing it immediately. The market maker sets price equal to expected value conditional on observed order flow. Price discovery is therefore a dynamic inference problem.

private signal $\longrightarrow$ strategic order $\longrightarrow$ aggregate order flow $\longrightarrow$ Bayesian price update.

Noise is not merely an accidental disturbance. It prevents the market maker from perfectly identifying informed trades and thereby permits the informed trader to earn a return. This reproduces the Grossman–Stiglitz incentive constraint inside a model of strategic execution.

Glosten and Milgrom model a specialist who quotes bid and ask prices while facing both informed and uninformed traders (Glosten and Milgrom, 1985). A buy order raises the probability that the trader possesses favorable information; a sell order raises the probability of unfavorable information. The market maker updates beliefs from the direction of trade. The bid–ask spread compensates for expected losses from trading against better-informed counterparties.

These models show that price is not information appearing spontaneously. It is generated through a sequence of commitments. Traders expose capital to gains and losses; intermediaries infer from those actions; and quotes adjust. Price is therefore the public consequence of private beliefs expressed through costly, risk-bearing action.

This also establishes a limitation. A trade demonstrates willingness to bear consequences at a particular price. It does not independently prove why the trader acted or whether the trader’s underlying belief is true. Market evidence is consequential but semantically incomplete. Appendix A uses this distinction to separate evidence of conviction from proof that the underlying claim is correct.

7 Who Chooses What Agents Can Observe?

Classical information economics often begins with a fixed information structure, where some agents observe one signal while others observe another. Modern information design makes the information structure itself a choice variable.

Blackwell’s comparison of experiments supplies the foundational ordering (Blackwell, 1953). One experiment is more informative than another when the less informative signal can be obtained by adding noise to the more informative one. This makes it possible to compare institutions by the decisions their signals enable.

Sims adds that acquiring and processing information consumes scarce attention (Sims, 2003). An agent does not merely update on a freely supplied signal; the agent allocates limited information-processing capacity. Rational inattention therefore broadens the information cost in Grossman–Stiglitz from research expenditure to the computational and cognitive cost of distinguishing among states.

Kamenica and Gentzkow’s Bayesian-persuasion framework asks how an informed sender who can commit to a disclosure policy will design the receiver’s information (Kamenica and Gentzkow, 2011). The sender selects a mapping from states to signals, thereby influencing posterior beliefs and actions. The sender cannot choose arbitrary posteriors. Their average must remain consistent with the prior. Yet within that constraint, the partition through which the state is disclosed becomes an instrument of persuasion.

Bergemann and Morris generalize this approach into information design (Bergemann and Morris, 2016, 2019). The analyst studies which outcomes can arise across information structures consistent with common fundamentals. Institutions no longer merely aggregate information; they determine who observes what, when observations arrive, how signals are correlated, and which distinctions remain hidden.

Modern platforms make this point concrete. Ranking systems, recommendation rules, disclosure policies, market interfaces, and data-access permissions affect the observations from which

participants form beliefs. The governing question can be stated directly.

Who designs the partition through which economic reality becomes observable?

This question prepares the transition to computation. Once an institution controls both the permissible messages and the verification procedure, it can define a new class of public evidence. Appendix A treats every verification rule as a partition between valid and invalid submissions, while emphasizing that anything the rule does not test remains outside the institution’s knowledge.

8 How Can Costly Computation Transform Private Claims into Publicly Verifiable Evidence?

The first six questions concern the movement from dispersed knowledge to economic signals. The seventh concerns a narrower but powerful transformation... How does one convert a proposed state transition into evidence that any participant can independently verify?

Nakamoto’s proof-of-work construction combines a public validation rule, a costly search process, a state-selection rule, and rewards for successful participation (Nakamoto, 2008). A computer assembles a candidate output and searches for an input that makes the output satisfy the protocol’s threshold condition. Finding such an input is probabilistic and computationally costly; validating it is relatively inexpensive. The asymmetry is straightforward:

costly to produce but cheap to verify.

The valid proof does not reveal the miner’s private motive, subjective valuation, electricity contract, hardware cost, or complete physical history. It establishes a narrower proposition, that a valid input was found for a publicly specified computational predicate, and the associated output satisfies the protocol’s validation rules. When embedded in a consensus protocol, this proof helps participants select among competing histories without relying on a single authoritative observer.

This can be interpreted using the earlier progression. Like a Spencian signal, proof of work is costly; unlike a loose social signal, its formal validity is mechanically checkable. Like a Myersonian mechanism, the protocol connects private action to rewards and allocative consequences. Like market microstructure, it aggregates strategic actions into a public state. Like information design, it determines exactly which facts participants observe and which proof conditions count as admissible evidence.

Yet the analogy has limits. Spence teaches that cost is informative only when it is appropriately connected to the hidden property. A proof of work proves work under a specified rule; it does not prove the truth of arbitrary off-chain data. An oracle report, asset-quality assertion, or claim about physical delivery requires an additional authenticated connection between the external event and the ledger input. Consensus can preserve and order a false claim just as effectively as a true one if the protocol lacks a valid test of the claim’s external content.

The economic security of proof of work also depends on incentives over time. Budish argues that permissionless consensus faces an economic constraint relating the recurring cost of honest operation to the value available from attack (Budish, 2018). Biais and coauthors model blockchain behavior as an equilibrium among strategic participants rather than as a purely cryptographic object (Biais et al., 2019). Cong and He similarly connect distributed consensus to competition, information distribution, and contracting (Cong and He, 2019). These works reinforce a central

lesson of mechanism design: cryptographic validity is necessary, but equilibrium behavior depends on payoffs.

Transaction fees and ordering introduce a second layer. Easley, O’Hara, and Basu connect mining incentives to congestion and fee formation (Easley et al., 2019). Daian and coauthors show that control over transaction ordering creates maximal extractable value and can generate destabilizing competition (Daian et al., 2020). The protocol therefore does more than verify isolated computations. It allocates scarce ordering rights and determines which strategic actions become final public history.

The strongest defensible conclusion is consequently not that proof of work makes a market perfectly informed. Grossman–Stiglitz has already shown why perfect informational efficiency is generally incompatible with costly private discovery. Nor does proof of work eliminate asymmetric information about external reality. Its distinctive contribution is the creation of a verifiable signal technology inside a strategically designed institution.

Proof of work transforms a privately performed, protocol-specified computation into public evidence that the computation’s validity conditions were satisfied and that a proposed ledger state may be evaluated under common rules.

The transformation proceeds through five steps.

private computational search $\longrightarrow$ costly commitment $\longrightarrow$ compact proof
 $\longrightarrow$ independent verification $\longrightarrow$ shared protocol state.

The word *evidence* (implied by **”proof”**) should be understood procedurally. The proof is evidence of compliance with a formal predicate and, through the consensus mechanism, evidence supporting a particular ordering of valid state transitions. Any stronger semantic claim requires additional assumptions about measurement, identity, data provenance, and the link between computation and the external world. Appendix A brings these limits together by separating identity, integrity, procedural, and semantic evidence.

9 Conclusion

The seven questions form a cumulative logistical framework of economic coordination. Hayek begins with knowledge dispersed beyond the reach of any central observer. Akerlof, Spence, and Stiglitz show that unequal knowledge changes exchange and creates demands for signaling and screening. Grossman and Stiglitz show that information reflected in price must remain worth producing. Myerson, Milgrom, and Weber show that rules determine whether strategic agents reveal what they know. Kyle and Glosten–Milgrom show how information enters prices through costly trades and Bayesian inference. Modern information design then makes the structure of observation itself an institutional choice.

Proof of work extends this lineage by separating the cost of producing a formal signal from the cost of checking it. It does not solve the entire economic knowledge problem, and it cannot transform every private assertion into truth. It solves a more precise problem: how mutually distrustful participants can verify that a specified computational condition was met and use that evidence to coordinate around a shared state.

Rational economic order is therefore neither an omniscient calculation nor a perfectly revealing price. It is an evolving institutional process that must coordinate partial knowledge, reward discovery, elicit credible action, govern observation, and correct error. Costly computation adds a

new instrument to that process by publicly checkable evidence of rule-bound work. Appendix A then consolidates the paper’s running example by tracing how a private claim can become credible public evidence without becoming an unquestionable statement of external truth.

A The Logistics + Implications of Private Claims

A private claim is an assertion made by a person or agent about something that other participants cannot directly observe or independently verify at the moment it is made.

The following examples illustrate the range of private claims.

- “I know this asset is worth more than its current price.”
- “This used car is high quality.”
- “I am a low-risk insurance customer.”
- “My production cost was \$100.”
- “This shipment arrived.”
- “I own these funds and authorize this payment.”
- “These transactions occurred in this order.”
- “I performed the required computation.”

The central problem is the uncertain relationship between a claim and reality.

A claim may be true, mistaken, exaggerated, or strategically false.

Consequently, a rational institution cannot treat every private claim as public knowledge.

A.1 A private claim has three components

A useful representation of a claim is the following.

$$C = (m, \omega, i),$$

The components have the following meanings.

- m is the message being asserted;
- ω is the underlying state of the world;
- i is the claimant.

For example, the message may assert that the car is high quality.

$$m = \text{“the car is high quality,”}$$

The corresponding state may take one of two values.

$$\omega \in \{\text{high quality, low quality}\}.$$

The claimant may observe ω , or possess a signal about it, while everyone else observes only m . The recipient must therefore estimate a conditional probability.

$$P(\omega \mid m),$$

the probability that the underlying state is ω , given the claim m .

But if saying m is free, every claimant may make the same favorable assertion regardless of the truth. The two types then generate the message with equal probability.

$$P(m \mid \omega = \text{good}) = P(m \mid \omega = \text{bad}),$$

The message consequently carries no information.

$$P(\omega \mid m) = P(\omega).$$

This is cheap talk.

A.2 Private does not necessarily mean secret

A private claim can be publicly spoken. “Private” describes its epistemic origin, not necessarily its audience.

This distinction can be expressed formally.

$$\text{publicly observable assertion} \neq \text{publicly verifiable fact}.$$

Someone can announce to the entire world that a shipment arrived. The message is public, but the underlying event may remain privately known.

Similarly, a trader’s order is publicly observable, but the information motivating that order remains private.

A.3 The logistics of turning a private claim into a credible signal

The transformation requires several stages.

Stage 1: A private state or observation

Let the relevant fact be represented by a state.

$$\omega \in \Omega.$$

The agent receives a signal generated conditional on that state.

$$s_i \sim P(s_i \mid \omega).$$

The signal may be imperfect. The agent does not necessarily know the true state; the agent may simply know more than others.

Private knowledge must therefore be distinguished from private belief.

$$\text{private knowledge} \quad \text{from} \quad \text{private belief}.$$

A trader may honestly believe an asset is undervalued and still be wrong.

Stage 2: The agent makes a claim

The agent translates the signal into a message.

$$m_i = g(s_i).$$

Claims can take many forms.

- verbal;
- contractual;
- a bid or offer;
- a trade;
- a signature;
- a submitted transaction;
- a sensor reading;
- a cryptographic commitment;
- a proposed block.

The message is not yet evidence merely because it has been transmitted.

Stage 3: The institution attaches consequences

Credibility generally requires the claim to carry a cost or consequence.

$$a_i = a(m_i),$$

The action a_i may involve several kinds of commitment.

- payment;
- collateral;
- risk exposure;
- reputation;
- legal liability;
- computational work;
- foregone alternatives;
- the possibility of losing a reward.

This shifts attention away from the agent's words alone. The more consequential question is what costs, risks, or other consequences the agent accepted by making the claim.

Stage 4: A verification rule is applied

The institution then applies a binary verification rule.

$$V(m_i, e_i) \in \{0, 1\},$$

where e_i is the submitted evidence.

The rule has two possible outputs.

$$V(m_i, e_i) = \begin{cases} 1, & \text{the formal verification conditions are satisfied,} \\ 0, & \text{the conditions are not satisfied.} \end{cases}$$

The verification rule must specify exactly what is being checked.
A digital-signature rule, for example, may apply the following check.

$$\text{Verify}(pk, m, \sigma) = 1.$$

This establishes that the signature corresponds to message m and public key pk . It does not independently establish that the signer's statement is factually correct.

Stage 5: The result becomes public evidence

If others can independently repeat the test, private evidence becomes publicly checkable.

$$(m_i, e_i, V) \longrightarrow \text{common verification result.}$$

Participants no longer need to trust the claimant's bare assertion. They can evaluate the claim under a shared rule.

Stage 6: The institution assigns consequences

The verification result then changes the institutional state.

$$S_{t+1} = T(S_t, m_i, e_i),$$

where T is the state-transition rule.

That state change can produce several consequences.

- accepting or rejecting a transaction;
- transferring funds;
- releasing collateral;
- updating a price;
- adding a block;
- imposing a penalty;
- paying a reward;
- recording ownership.

The full sequence can be summarized as follows.

$$\begin{aligned} &\text{private observation} \rightarrow \text{claim} \rightarrow \text{commitment} \\ &\quad \rightarrow \text{verification} \rightarrow \text{consequence} \rightarrow \text{public state.} \end{aligned}$$

A.4 Four different levels of evidence

These should not be collapsed.

Identity evidence

Who made the claim?

A digital signature can establish control of a private key.

$$\text{Verify}(pk, m, \sigma) = 1.$$

It proves authorization relative to the key. It does not necessarily prove the legal or physical identity of the key holder.

Integrity evidence

Has the message been altered?

A cryptographic hash can commit to exact data.

$$h = H(m).$$

Later, anyone can check whether a presented message matches the commitment. The hash does not prove the message was true when committed.

Procedural evidence

Were specified rules followed?

A proof of work, zero-knowledge proof, or protocol-valid transaction may establish that a formal predicate was satisfied.

$$V(m, e) = 1.$$

This is strong evidence about rule compliance.

Semantic evidence

Is the claim about external reality true?

The relevant external claims include the following examples.

- the shipment actually arrived;
- the electricity was generated renewably;
- the asset really exists;
- the temperature was actually 20°C;
- the reported production cost is accurate.

Cryptography alone cannot generally establish this. The system therefore needs a trusted or incentive-compatible connection to the external world.

- sensors;
- auditors;
- legal attestations;
- redundant observers;
- prediction markets;
- authenticated measurement devices;

- dispute and slashing procedures.

These distinctions form a hierarchy of increasingly demanding claims.

authenticated $\nRightarrow$ unaltered
 $\nRightarrow$ procedurally valid
 $\nRightarrow$ externally true.

A better rendering is that each level answers a different question; none should silently substitute for another.

A.5 Where the major economic theories fit

Akerlof: the claim is not credible

A used-car seller may claim that a car is good.

$m = \text{“this is a good car.”}$

But low-quality sellers can make the same claim. If claims are costless and unverifiable, buyers cannot separate types. They price the average quality, producing adverse selection.

Spence: the claim is supported by differential cost

The agent takes an action that is easier or more valuable for one type than another.

Credibility requires a difference in signaling costs across types.

$c(a \mid \text{good type}) < c(a \mid \text{bad type})$.

The signal separates types because imitation is more costly for the wrong type.

Rothschild–Stiglitz: the institution screens the claim

Instead of trusting “I am low risk,” the insurer offers contracts that different risk types prefer differently.

The agent reveals information through contract selection.

private type $\rightarrow$ contract choice $\rightarrow$ inferred type.

Grossman–Stiglitz: the claim is expressed through costly trading

The trader does not merely announce that an asset is undervalued. The trader instead commits capital.

private signal $\rightarrow$ position $\rightarrow$ profit or loss.

The trade is evidence of conviction, although not proof of correctness.

Myerson: the rules make truth-telling advantageous

A mechanism can make truth-telling optimal by satisfying an incentive constraint.

$$U_i(\text{truthful report}) \geq U_i(\text{misreport}).$$

The claim becomes credible because truthfulness is incentive-compatible.

Kyle and Glosten–Milgrom: claims appear indirectly in order flow

A trade is an indirect claim about value. Market makers infer from that order.

$$\text{order} \rightarrow P(\text{informed trader} \mid \text{order}) \rightarrow \text{updated price}.$$

But noise prevents complete inference.

Information design: the institution chooses the observable evidence

The designer chooses the signal structure.

$$\pi(s \mid \omega).$$

That determines which aspects of the state become observable and which remain hidden.

Proof of work: the claim is constrained by a computational predicate

A miner effectively makes a specific, testable claim.

“This candidate block is valid, and I found a nonce satisfying the current difficulty target.”

Anyone can test the claim against the difficulty target.

$$H(\text{block header}) < T.$$

The claimant cannot cheaply fabricate a passing output without performing the probabilistic search, while verification is inexpensive.

A.6 What proof of work turns public

Proof of work begins with several privately performed activities.

- the miner chooses hardware;
- buys electricity;
- constructs candidate blocks;
- repeatedly hashes private candidate inputs;
- discovers a qualifying nonce.

The miner needs to publish only the successful result.

(block, nonce).

Other nodes verify the target condition.

$$H(\text{block header}) < T.$$

They also verify the block's transactions and state-transition rules.

What becomes public is not the miner's complete physical history. It is evidence sufficient to validate the protocol claim.

The result is a compact public certificate of private work.

private production process $\rightarrow$ compact public certificate.

More precisely, proof of work supports the following three claims.

1. The submitted header satisfies the target.
2. Producing a satisfying header probably required repeated computational trials.
3. The proof is attached to a particular candidate ledger history.

It does not establish any of the following claims.

- which machine performed the work;
- where the electricity came from;
- the miner's exact energy expenditure;
- the miner's accounting cost;
- that external data in a transaction are true;
- that the resulting token has a particular market value.

A.7 The major implication: verification replaces some trust

Without public verification, participants may need to trust several intermediaries.

- the claimant;
- an auditor;
- a central database;
- a settlement intermediary;
- a government registry;
- a market maker.

A publicly checkable proof can replace some of this institutional trust with a common verification procedure.

trust the speaker $\longrightarrow$ check the proof.

But it never eliminates trust altogether. Trust instead moves into other parts of the system.

- the protocol specification;
- the software implementation;
- the cryptographic assumptions;
- the measurement interface;
- the incentive design;

- the governance process.

The more accurate claim is narrower.

Verification relocates and narrows trust; it does not abolish it.

A.8 The second implication: claims become composable

Once evidence is public and machine-verifiable, other contracts can condition on a successful verification.

$$V(m, e) = 1 \implies \text{execute subsequent action.}$$

A simple sequence illustrates this composability.

$$\begin{aligned} &\text{valid payment} \rightarrow \text{release collateral} \\ &\quad \rightarrow \text{transfer title} \rightarrow \text{update accounting.} \end{aligned}$$

This makes evidence composable. One verified state transition becomes an admissible input to another.

That is a major difference between an ordinary costly signal and a cryptographically verified claim. A reputation, education credential, or trade may require interpretation. A protocol-valid proof can trigger deterministic downstream rules.

A.9 The third implication: verification itself is a partition

Every verification rule partitions submissions into two sets.

$$\mathcal{M} = \mathcal{M}_{\text{valid}} \cup \mathcal{M}_{\text{invalid}}.$$

The verifier recognizes only the distinctions represented by the predicate V .

Anything the predicate does not test remains outside its knowledge.

For example, consider a transaction for which the verifier returns one.

$$V(\text{transaction}) = 1$$

That result may establish several facts.

- the signature is valid;
- the funds exist;
- the transaction obeys protocol rules.

But if V does not test whether a physical product was delivered, protocol acceptance says nothing about delivery.

This limitation connects directly to the idea of partition choice.

A verification system knows only through the distinctions
its rules are capable of making.

Designing the predicate is therefore designing the boundary between what the institution can and cannot verify.

A.10 The deepest formulation

A private claim becomes credible public evidence only when three structures align.

epistemic connection + incentive compatibility
+public verifiability.

The epistemic connection requires the evidence to be related to the claimed fact. Incentive compatibility requires honest behavior to be more rewarding than manipulation. Public verifiability requires independent participants to be able to apply the same test and obtain the same result.

Proof of work strongly supplies the third property and partially supplies the second through rewards and attack costs. It supplies the first only for a narrow endogenous claim: satisfaction of the protocol’s computational condition.

The following statement summarizes the argument.

A private claim is an assertion whose underlying state is not directly observable to all participants. Institutions make such claims economically credible by connecting them to costly commitments, incentive-compatible consequences, and shared verification rules. Proof of work is a special case in which a privately performed computational search produces compact public evidence that a protocol-defined condition has been satisfied. It proves compliance with that condition, not the truth of every external proposition associated with it.

References

- George A. Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3):488–500, 1970.
- Dirk Bergemann and Stephen Morris. Information design, Bayesian persuasion, and Bayes correlated equilibrium. *American Economic Review*, 106(5):586–591, 2016.
- Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.
- Bruno Biais, Christophe Bisière, Matthieu Bouvard, and Catherine Casamatta. The blockchain folk theorem. *Review of Economic Studies*, 86(5):1849–1875, 2019.
- David Blackwell. Equivalent comparisons of experiments. *Annals of Mathematical Statistics*, 24(2):265–272, 1953.
- Eric Budish. The economic limits of Bitcoin and the blockchain. NBER Working Paper No. 24717, 2018.
- Lin William Cong and Zhiguo He. Blockchain disruption and smart contracts. *Review of Financial Studies*, 32(5):1754–1797, 2019.
- Philip Daian, Steven Goldfeder, Tyler Kell, Yunqi Li, Xueyuan Zhao, Iddo Bentov, Lorenz Breidenbach, and Ari Juels. Flash Boys 2.0: Frontrunning in decentralized exchanges, miner extractable value, and consensus instability. In *2020 IEEE Symposium on Security and Privacy*, pages 910–927, 2020.

- David Easley, Maureen O'Hara, and Soumya Basu. From mining to markets: The evolution of Bitcoin transaction fees. *Journal of Financial Economics*, 134(1):91–109, 2019.
- Lawrence R. Glosten and Paul R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.
- Sanford J. Grossman and Joseph E. Stiglitz. On the impossibility of informationally efficient markets. *American Economic Review*, 70(3):393–408, 1980.
- Friedrich A. Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, 1945.
- Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.
- Paul R. Milgrom and Robert J. Weber. A theory of auctions and competitive bidding. *Econometrica*, 50(5):1089–1122, 1982.
- Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- Roger B. Myerson and Mark A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 29(2):265–281, 1983.
- Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. 2008.
- Michael Rothschild and Joseph E. Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *Quarterly Journal of Economics*, 90(4):629–649, 1976.
- Christopher A. Sims. Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690, 2003.
- Michael Spence. Job market signaling. *Quarterly Journal of Economics*, 87(3):355–374, 1973.