Asymmetric Information Markets:
A Coordination Mechanism for Latent Network Value

BOUNTI10 SAS
Paris, France · contact@bounti.fr · bounti.fr

Working paper · v1.0 · May 2026

Abstract. Every individual holds a substantial inventory of relational and contextual information — who they know, what is happening in their network, what is materially relevant to which counterparty — that has measurable economic value to other parties. Yet this inventory remains, in aggregate, almost entirely unmonetized. The reasons are structural: the holder has no immediate use for most of the information themselves, lacks a discovery mechanism to learn who values it, faces a disclosure paradox that frustrates direct negotiation, bears reputation hazard, and in many cases is selected against in equilibrium for sharing because it erodes a scarcity-based competitive position. Coordination costs (negotiation, contract, payment, dispute mechanism) further dominate the expected payoff for any single transaction. We characterize this as an asymmetric-information market with chronic under-clearing. We propose a lightweight coordination protocol with three primitives — explicit demand-side pricing, opt-in atomic supply, and escrow-conditional release on declared success criteria — that lowers each of the five identified frictions enough to clear a meaningful subset of the latent market. We discuss the protocol's boundary conditions, including its strict exclusion of regulated intermediation activities, and the case categories where the asymmetry is highest.

1. Introduction

The classical analysis of information goods (Arrow [1]; Shapiro and Varian [4]) treats information as a public good with high fixed-production cost and zero marginal-replication cost. This framing captures the economics of patents, software, and academic publication, but it does not adequately describe a different category: distributed, atomic, person-bound information whose value is precisely its non-replicability and its contextual specificity to a moment in time.

A direct contact at a target customer; the early signal that a particular family is preparing the sale of a property before listing; the knowledge that a former colleague is open to a senior role move; the awareness that a mid-market company is quietly soliciting acquisition interest — these information units share three properties. They have no fixed production cost; they accrue passively as a byproduct of professional life. They are bound to a specific holder and resist replication without a chain of custody. And their value to a counterparty often exceeds, by orders of magnitude, the value to the holder, because the holder has no use for the information themselves.

We argue that this distributed information stock — what we shall call the information bank carried by each individual — is among the largest under-monetized economic resources in modern networks, and that its under-monetization is not accidental but the product of a small set of identifiable frictions. The remainder of the paper formalizes the thesis (§2), enumerates the frictions (§3), reviews prior coordination attempts and their limits (§4), proposes a coordination protocol (§5), is explicit about what the protocol is not designed to do (§6), characterizes case categories where the value asymmetry is highest (§7), and discusses limitations and open questions (§8).

2. The information bank thesis

Define the information bank of an individual i as the set I b of facts known by i that are (a) not in public circulation, (b) currently relevant to at least one other party, and (c) verifiable by i either by direct

knowledge or by a one-step social proximity to a verifier.

We posit, without formal proof, the following empirical claim. For any working professional, the cardinality of I b is dominated not by facts of high private value to i, but by facts whose value to other parties exceeds their value to i. The intuition is straightforward: each professional accumulates on the order of 10^2 active relationships and 10^3 weak ties (Granovetter [2]), each of which is a continuous source of small information signals. Most of those signals are useful to someone — a vendor, a recruiter, a buyer, an investor, an acquirer — but only a vanishing fraction are useful to the holder.

If this claim holds, the marginal value of an information bank to its holder approaches zero quickly, while the marginal value to a well-matched counterparty does not. This is an asymmetry in the strict sense: the holder is the only party in the market who can supply unit u, but is also the party who values u least. By Akerlof's [3] reasoning the asymmetric-information market does not clear at the social optimum without an external coordination mechanism. Unlike the lemons market, however, the failure mode here is not adverse selection (the supply side is not strategically withholding quality) but rather chronic non-supply: the holder simply never offers u for transaction, because no efficient mechanism exists to do so.

2.1 Three intrinsic properties of distributed information.

The under-clearing identified above is rooted in three properties of information as an economic good. We state them plainly because the rest of the paper depends on them.

(i) Once revealed, information loses scarcity-derived value. A unit u sold to counterparty C is now held by both seller and buyer. The seller cannot unsell it, and the buyer can in principle re-supply it at zero cost. The value of u to the next prospective buyer collapses immediately upon the first disclosure, in proportion to the leakage probability. The seller therefore cannot use repeat-sale pricing to recover their reservation cost across multiple buyers, as is standard for replicable goods.

(ii) Information is hard to price ex ante. The buyer cannot determine whether u is worth paying for until after they observe u. This is the classical inversion: in most markets the buyer evaluates the good before purchase; in an information market the act of evaluation is itself the consumption of the good. Stated price discovery — bidding, posted offers, auctions — therefore presupposes a description of u rich enough to value it but not so rich that it substitutes for u itself, a condition that is rarely satisfiable for atomic, person-bound units.

(iii) Information has near-zero marginal replication cost. Once disclosed in any verifiable form, the unit propagates at no incremental cost to any party with downstream interest. Property (iii) compounds property (i): the seller's loss of pricing power on subsequent disclosures is not bounded by the seller's own behaviour but by the behaviour of every party in the disclosure chain.

Taken together these three properties produce the apparent paradox at the heart of the latent-information market: why would any rational holder pay the cost of revealing information that loses its value the moment it is revealed? Without a coordination mechanism that decouples payment from disclosure, the holder's dominant strategy is to retain the information indefinitely. The protocol of §5 is designed precisely to break this deadlock by guaranteeing payment as a contractual function of the buyer's validation rather than as a discretionary act after disclosure has already occurred.

2.2 Why information is rarely sold directly: the indirect-monetization equilibrium.

The properties of §2.1 explain why most economically valuable information is not traded as a good in its own right. Direct sale is structurally fragile: the seller cannot unsell, cannot price ex ante, and cannot prevent replication. The empirical equilibrium is that information is monetized indirectly, embedded in a wrapper whose payoff depends on the information without making the information itself the unit of exchange. We distinguish three dominant wrappers.

(A) Trading on information rather than selling it. The dominant model in financial markets, formalized by Kyle [9]: an informed agent does not disclose a private signal; they take a position in an asset whose price will move when the signal is impounded by the market. The payoff is the difference between the position's entry and exit price, not a transfer from a counterparty paying for the signal. Under this wrapper, properties (i)–(iii) do not bite, because the agent never reveals u to anyone. The market itself extracts the informational rent through price impact. This wrapper is highly efficient where it applies but is bounded by what regulators classify as insider trading or market abuse, and is therefore inaccessible to most categories of relational information.

(B) Access-based monetization. The agent does not sell u directly; they sell access to a situation in which u exists or can be obtained. Examples include introducers, deal-flow brokers, expert networks, scout programmes for venture funds, and curated communities. The unit transacted is not "the information" but "membership in the room where the information lives" — typically operationalized as an introduction, an invitation, a meeting, or a referral. Property (i) is sidestepped because the room is replenished continuously; property (ii) is mitigated because the buyer prices the access rather than the underlying signal; property (iii) is partially neutralized because the access itself is harder to copy than the signal it produces. The protocol of §5 is, by construction, a Type B mechanism.

(C) Bundling with a complementary good or service. The information is embedded in software (the Bloomberg terminal; market-data feeds), advisory engagements (consulting; equity research), or products that incorporate it (credit ratings, analytics platforms, expert-network calls). The buyer does not pay for the information per se; they pay for the system that produces, contextualizes, or operationalizes it. Properties (i)–(iii) are partially absorbed by the wrapper: the bundle is rivalrous (subscription seats, billable hours), the price is set on the bundle rather than the signal, and the value lives partly in the production process rather than the raw output.

Direct sale of information does occur, but only under restrictive conditions. The information must be (a) time-sensitive — alpha decays before propagation can erode it; (b) exclusive — distribution is contractually constrained; (c) verifiable — either intrinsically or via the seller's reputation. The clearest examples are hedge-fund research subscriptions, expert-network calls (GLG, Third Bridge), and alternative-data feeds. Outside these conditions direct sale is rarely a stable equilibrium.

A useful summary distinction is between selling information and using information. Selling information is a fragile business: each transaction depreciates the asset and prices are difficult to defend. Using information — through a position, an access channel, or a bundled service — is a durable advantage, because the wrapper survives the disclosure of any individual signal. The protocol presented in §5 inherits the durability of the access wrapper while making its supply side permissionless: any holder can opt in atomically without committing to ongoing intermediation, and the buyer pays for an outcome rather than for a continuing service relationship. We argue this is the principal mechanical contribution of the design.

3. Why information markets under-clear

We identify five distinct frictions, each of which independently suppresses voluntary supply.

3.1 Discovery cost. The holder of u does not know which counterparty values u; the counterparty who values u does not know which holders possess it. Random search across a network is intractable at scale, and broadcast search via existing professional networks (LinkedIn, alumni groups) is rate-limited by social cost: the holder cannot continuously signal "I have valuable information, who wants it" without depreciating their standing.

3.2 Disclosure paradox (Arrow [1]). The holder cannot establish u's value to the counterparty without disclosing enough of u to extinguish its scarcity. A direct negotiation — "I know someone at X; what is it worth to you?" — leaks the existence and identity of the contact, after which the counterparty can

pursue independent discovery and refuse to pay.

3.3 Reputation hazard. Sharing network-derived information carries an asymmetric downside for the holder. If the introduction goes well, the value accrues to the introduced parties; the holder receives at most a token reciprocation. If the introduction goes poorly, the holder absorbs the social cost from the named network member. The expected utility of an unstructured introduction is, for many holders, weakly negative.

3.4 Coordination cost. Even when the parties locate each other, an actual transaction requires negotiation of terms, contractual definition of "successful delivery," payment mechanics, dispute escalation, and identity verification. For a single information unit valued at €1 000–5 000, these fixed costs frequently dominate. Repeated transactions amortize the costs but presuppose a stable bilateral relationship — itself the result of prior introductions.

3.5 Negative scarcity incentive. In zero-sum framings — a recruiter, a deal broker, a fund-allocator — the holder's competitive position depends on the rivalrous control of their information bank. Sharing destroys the hedge. Even when the holder is willing in absolute terms, the agent is selected against in equilibrium for sharing, because their peers who do not share retain the competitive advantage.

The combined effect is that the great majority of economically-valuable units of distributed information are never offered for transaction. They expire silently as the counterparty fills their need elsewhere or abandons it.

4. Prior coordination attempts and their limits

Several mechanisms have addressed subsets of the frictions in §3.

4.1 Search and brokerage intermediaries. Executive search firms, real estate agents, M&A advisory boutiques, and licensed deal brokers solve the discovery problem via consolidation. They aggregate demand on one side and curate a roster of supply on the other. They operate at high unit cost — typically 20–35% of the underlying transaction value — and require professional licensing in most jurisdictions, which constrains the supply side to a small population of credentialed intermediaries. They do not democratize supply.

4.2 Professional networks. LinkedIn and analogous platforms solve part of the discovery problem at zero unit cost. They contain no native pricing mechanism and no enforcement layer. The "ask my network for an intro" workflow generates effort with no expected reward and is subject to the reputation hazard of §3.3. Empirically the resulting flow is dominated by reciprocal-altruism dynamics among close ties; weak-tie information remains substantially unmobilized.

4.3 Informal referral fees. Bilateral handshake agreements where party A pays party B for a successful introduction. These contracts have no neutral arbiter, no price discovery (the fee is set by negotiation between parties who already know each other), and no escrow. Disputes resolve via reputational sanction inside an existing relationship, which presupposes the relationship and reproduces the discovery problem the mechanism aims to solve.

4.4 Crypto-economic coordination. Prediction markets, decentralized oracle networks, and on-chain reputation primitives address truth verification and trustless settlement. They do not address the latent-supply problem of information that no one knows is being demanded, because the demand side has no efficient mechanism to express a request directed at a specific holder population without first identifying that population.

No prior mechanism prices atomic, holder-specific information units at the moment a counterparty is willing to pay for them, with sufficiently low coordination cost to make the median transaction worthwhile.

5. A coordination protocol

We propose a protocol with three primitives. We describe it abstractly; an open implementation is given by Bounti [7].

5.1 Demand-side explicit pricing. A buyer (the principal) publishes a request for an information unit with a fixed monetary bounty B and a deadline T. The request specifies the success criteria — what would constitute valid delivery — without disclosing the buyer's identity in cases where that disclosure would itself leak commercial information. The bounty is denominated in fiat and held in escrow at the moment of request.

5.2 Opt-in atomic supply. Any individual in the network can declare that they hold a candidate information unit and apply to deliver it. The application is opt-in, asynchronous, requires no prior negotiation, and can be withdrawn until the buyer accepts it. Multiple suppliers may apply concurrently. The buyer selects one and the others are released without obligation. This eliminates ex-ante negotiation between unknown parties.

5.3 Escrow with success-conditional release. The bounty B is held by the protocol from request time until validation. It releases to the supplier only on the buyer's explicit confirmation that the success criteria of §5.1 were met. Disputes default to non-payment, biasing the protocol against bad-faith suppliers. The buyer's symmetric incentive to behave in good faith is enforced by reputation feedback: refused validations on demonstrated deliveries are visible across future requests.

The combined effect of the three primitives is that:

• Discovery is solved by the visibility of the request (§5.1). The buyer states their need; any holder who recognizes themselves in the description self-identifies.

• The disclosure paradox (§3.2) is solved by escrow and success criteria. The supplier can deliver the unit with the guarantee that payment is contractually conditional on validation, not on the buyer's discretion after disclosure.

• Reputation hazard (§3.3) is reduced by formalizing the introducer's role: a documented term sheet replaces an informal favour, and the introducer's economic incentive is visible to all parties.

• Coordination cost (§3.4) is amortized across all transactions by the protocol's standard contract template, jurisdictional law-of-origin clause, and standardized dispute escalation.

• The negative scarcity incentive (§3.5) is partially addressed but not eliminated. The holder receives a measurable cash payment B × (1 " <b’À where <` is the protocol fee. For a non-trivial subset of cases this exceeds the expected loss from the eroded competitive hedge. For the remainder it does not, and these holders rationally do not participate. The protocol does not eliminate this constraint; it makes it explicit and individually evaluable.

The protocol does not abolish any of the five frictions globally. It lowers each one enough that the threshold for a transaction to clear is met for a measurable fraction of latent information units. Below this threshold the unit remains untraded; this is the protocol's revealed inefficiency boundary.

5.4 Inverted price discovery and the resolution of the information paradox.

The information paradox of §2.1 — that a unit u cannot be priced before it is observed but loses value the moment it is observed — is the central obstacle to direct trading of distributed information. Conventional price-discovery mechanisms (auctions, posted offers, bilateral negotiation) presuppose that the buyer can either evaluate u prior to commitment or rely on a credible signal of u's quality. Neither holds for atomic, person-bound, contextual information. The protocol resolves the paradox by inverting the chronology of price formation. The buyer commits to a denominated price B before any information is exchanged, the bounty is escrowed irrevocably for the duration T, and the supplier

observes the commitment before electing to deliver. At no point in the sequence does the buyer set or revise a price after observation; pricing is structurally pre-disclosure.

This inversion is non-trivial because it shifts the pricing problem to the only party in the market who can solve it. The seller cannot price u: they do not know the buyer's use-case, the marginal value of u in that use-case, or the buyer's outside options. The buyer, by contrast, knows their own use-case exactly. They are not asked to value u in the abstract; they are asked to value the outcome that observing u would enable — the closed sale, the hire that fills a vacancy, the term sheet, the LOI. The bounty is therefore not a price on the signal but a price on the marginal economic gain the buyer attributes to the realization of the signal in their specific context. We claim — and this is the central proposition of the design — that this contextual, use-case-grounded number is the most defensible expression of the fair price of u that a market can produce. The same unit u, supplied to two different buyers, will attract two different bounties, and this is the correct outcome: u has different value-in-use across buyers, and an efficient mechanism must reflect that.

Mapping the inversion onto the three properties of §2.1 makes the resolution explicit. Property (i), that revealed information loses scarcity-derived value, is neutralized by decoupling payment from disclosure: the contractual obligation to pay is fixed at request time and is conditional only on the buyer's validation against published success criteria, not on anything the buyer learns after observing u. The supplier therefore bears no revelation-without-payment risk that the conventional disclosure paradox imposes. Property (ii), the inability to price ex ante, is reassigned from the seller (who cannot solve it) to the buyer (who can), and the buyer is forced to express the answer publicly and irrevocably before any further step in the protocol. Property (iii), near-zero marginal replication cost, is sidestepped by scoping the protocol's deliverable to the introduction itself rather than to the underlying signal — a Type B (access-based) construction, in the taxonomy of §2.2. An introduction is a relational event between specific parties at a specific moment; it is not a transferable artefact and does not propagate at zero cost. The protocol thus monetizes the wrapper that survives properties (i)–(iii), not the signal that does not.

A second-order property of the inversion is that the published bounty acts as a public market signal even when no transaction occurs. If no qualified supplier opts in at price B, the protocol has revealed — without further action — that B is below the population's clearing price for u in that context. The buyer is informed and may revise upward; the supply side has communicated reservation pricing without engaging in any bilateral disclosure. This is materially different from a conventional posted-offer mechanism, where non-supply is indistinguishable from non-discovery. Under the protocol, non-supply against an observed published bounty is a positive market datum: the addressable supply has decided not to deliver at this price.

A further consequence, less formal but practically significant, is that the act of escrowing the bounty at request time is itself a credibility signal. The buyer is not stating an aspirational willingness to pay; they have already moved the funds. This converts the request from an expression of interest into a binding offer and, when communicated by the supplier to the introduced third party, legitimizes the supplier's role as a paid intermediary in a way that ad-hoc referral arrangements cannot. The introduced party, observing that the requesting buyer has placed €B into escrow, receives a sharper signal of seriousness than any unpriced introduction can convey.

6. Boundary conditions and out-of-scope domains

We are explicit about where the protocol does not apply. Information transactions in regulated industries are excluded by construction. In the European Union and most G20 jurisdictions, the following intermediation activities require professional licensing or registration: real estate transactions (in France, the loi Hoguet of 1970); regulated investment advice and capital-markets

intermediation (AMF in France; FCA in the UK; SEC and FINRA in the United States); insurance brokerage (ORIAS in France); credit intermediation (ACPR); and licensed legal, tax, and accounting services. The protocol of §5 is engineered for the introduction itself as the deliverable. Where the deliverable is advice, fiduciary intermediation, or transaction execution on behalf of a party, professional licensing applies and the protocol is not a substitute.

This is a meaningful constraint, not a marketing detail. Approximately 60–70% of high-value introduction flow in mature economies passes through regulated channels. The remaining 30–40% — commercial introductions in software, talent referral, partnership and distribution, M&A target sourcing, and non-regulated advisory — defines the protocol's addressable surface. Within that surface the protocol's contract instrument explicitly excludes regulated activity, requires the principal to certify non-regulated subject matter at request time, and renders the supplier's representation as introducer-only with no advisory or representation authority.

7. Case categories and value densities

We sketch four representative case classes, with order-of-magnitude estimates of the value asymmetry observed in the addressable surface.

7.1 B2B enterprise sales introductions. The buyer is a vendor seeking a warm introduction to a procurement decision-maker at a target customer. Bounty range: €500–5 000 per introduction. Holder's private value of the contact is typically zero — no direct sale opportunity for the holder — and the buyer's value if the deal closes is 10–100× the bounty. The asymmetry is structurally clean.

7.2 Senior talent referrals. A company hiring a senior or executive role. Bounty range: €2 000–10 000 per successful hire. Holder's private value is zero or weakly negative (loss of access to the candidate as a future personal hire). Buyer's implicit value is 5–10× the bounty when compared to standard executive search retainers. The market currently uses search firms; the protocol competes by being self-serve and by tapping a wider supply pool than any individual firm's curated roster.

7.3 Fundraising introductions. A founder seeks a warm introduction to a venture capital fund. Bounty range: €2 000–10 000 per introduction that converts to a term sheet. Holder's private value is zero unless the holder is fund-internal, in which case sharing is misaligned and excluded by the protocol's reputation mechanics. The introducer must explicitly disclaim investment advice; the protocol's contract instrument enforces this.

7.4 M&A target sourcing. A buy-side principal seeks acquisition targets matching a thesis (sector, revenue band, geography). Bounty range: €5 000–20 000 per qualified target leading to a non-binding letter of intent. Holder's private value is zero unless the holder is an active competitor for the same target, in which case the holder is rationally excluded.

The common pattern across categories is that holders' private value sits at the low end of the asymmetry distribution and buyers' value at the high end. The bounty captures a small fraction of the buyer's expected gain but is large enough to clear the holder's reservation price plus expected social cost.

8. Discussion and limitations

8.1 Gaming. Suppliers may attempt to deliver fake or low-quality introductions in pursuit of payment. The protocol mitigates this via buyer-side validation and rated supplier history. It does not eliminate gaming; it reduces the steady-state rate to a level absorbed by the price of the bounty.

8.2 Mispricing. In thin markets buyers may post bounties below the clearing level. The supply side's response — non-participation — is a passive but informative price signal. We expect price discovery to converge as the protocol scales.

8.3 Thin-market problems. For highly niche requests no qualified supplier may exist within the protocol's active user base. This dominates the failure distribution at small scale and attenuates with adoption. It is an intrinsic limit of any matching protocol with a finite participant set.

8.4 Cannibalization vs. expansion. A second-order question is whether the protocol cannibalizes existing intermediation channels (search firms, deal brokers) or expands the total volume of introduction transactions. We expect both, with a structural bias toward expansion. The lower per-transaction coordination cost pulls into circulation a stock of latent information units that previously did not justify the friction of any existing channel — the introductions that would otherwise have expired silently.

8.5 Ethics of monetizing personal-network information. We note that the alternative — the present world — does not preserve such information in anyone's interest. It expires unused. Pricing it does not extract it from a previously protected reserve; the price creates the incentive that brings the unit into existence as a tradeable object. The privacy-preserving primitive — opt-in supply, holder-controlled disclosure, refusal cost-free — remains intact.

9. Conclusion

The dispersed information held by individual professionals is among the largest under-monetized economic stocks in modern networks. We have argued that under-monetization is structural, not contingent: it follows from the joint operation of discovery cost, the disclosure paradox, reputation hazard, coordination overhead, and the negative scarcity incentive. We have proposed a coordination protocol — explicit demand-side bounties, opt-in atomic supply, and escrow with success-conditional release — that reduces each of the five frictions enough to clear a meaningful subset of the latent market. We have stated the regulatory boundary plainly: the protocol applies to the introduction itself, not to regulated intermediation, and approximately 30–40% of the high-value introduction surface is therefore addressable. Within that surface, the case categories of B2B sales, senior talent, fundraising, partnerships, and M&A sourcing exhibit value asymmetries of one to two orders of magnitude between buyer and holder, supporting bounties that clear the holder's reservation price.

We expect the addressable surface to grow as professional networks digitize further, as the ratio of latent-to-realized information value continues to widen, and as the standardization of the contract instrument reduces the marginal cost of each subsequent transaction. The protocol is permissionless on the supply side, which we believe is the precondition for clearing latent markets at scale.

References

[1] K. J. Arrow, "Economic Welfare and the Allocation of Resources for Invention," in The Rate and Direction of Inventive Activity: Economic and Social Factors, NBER, 1962, pp. 609–626.

[2] M. S. Granovetter, "The Strength of Weak Ties," American Journal of Sociology, vol. 78, no. 6, pp. 1360–1380, 1973.

[3] G. A. Akerlof, "The Market for Lemons: Quality Uncertainty and the Market Mechanism," Quarterly Journal of Economics, vol. 84, no. 3, pp. 488–500, 1970.

[4] C. Shapiro and H. R. Varian, Information Rules: A Strategic Guide to the Network Economy. Harvard Business School Press, 1999.

[5] S. Nakamoto, "Bitcoin: A Peer-to-Peer Electronic Cash System," 2008. Available: https://bitcoin.org/bitcoin.pdf

[6] R. H. Coase, "The Nature of the Firm," Economica, vol. 4, no. 16, pp. 386–405, 1937.

[7] BOUNTI10 SAS, "Mission term sheet, version 2026-05," bounti.fr.

[8] O. E. Williamson, "Transaction-Cost Economics: The Governance of Contractual Relations," Journal of Law and Economics, vol. 22, no. 2, pp. 233–261, 1979.

[9] A. S. Kyle, "Continuous Auctions and Insider Trading," Econometrica, vol. 53, no. 6, pp. 1315–1335,

1985.

BOUNTI10 SAS is the operator of the bounti.fr coordination platform. This paper is descriptive and does not constitute investment, legal, or tax advice. Comments to research@bounti.fr.