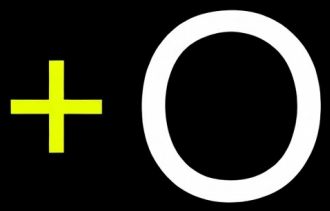
Abstract Economics

The Theory of Abstract Pricing

Value spaces, offer spaces, exchange records and the capabilities of AI agents

A foundational research monograph

A mathematical and economic study of how opportunities for value exchange become discoverable, composable and executable offers.



Conditional proofs and synthetic computations. Foundational research draft; not externally peer reviewed.

FIRST EXPANDED EDITION IN ENGLISH

Concepts proposed within the Offerato project

4 OCTOBER 2026

100 PAGES

FRONT MATTER

Abstract

This study develops a framework for abstract economics, with abstract pricing as its principal application. It starts from a value space defined by a structured environment and the elements interacting within it, then distinguishes the possibility of exchange, the evaluation of its outcomes and access to information about it. A product is defined through its causal role in enabling value exchange among parties that may be human, digital or physical. This generalization provides a common descriptive language while requiring the measures and mechanisms of each application to remain explicit and open to criticism.

The study defines Promoi as the space of offers associated with a product assortment. An offer specifies eligibility, authority, resources, effects and commitments, and may be composed with other offers when its semantics and constraints permit. A kup is the abstract record of a value exchange, with a distinction between the event, the recorded content and the carrier that stores it. This definition neither assumes that a complete real-world record of every exchange can be recovered nor equates evidence of authorization with evidence of fulfillment.

The monograph presents preliminary conditional results on the invariance of the optimum under a change of representation alone, bounds on the loss from offer-menu compression, the effect of truthful evidence on decisions guaranteed to be feasible, and the insufficiency of marginal accessibility to characterize discovery quality. In studying agents, it stipulates intelligence as the ability to exchange value, then measures this capability relative to specified tasks, resources and interfaces. It also derives exchange-capacity bounds that depend on resources, temporal structure and concurrency.

The economic analysis separates prices and transfers from real surplus. It shows when changing delivery terms or reducing search costs can create mutual gains. It examines seller participation, the funding of commitments, private incentives and restrictions on simultaneously achieving efficiency, truthfulness and budget balance. The results are situated within the literature on multi-attribute offers, compositional contracts, bounded rationality and the economics of information.

Small synthetic computations check capacity examples and offer-compression bounds. A separate design sets out agent and market experiments that have not yet been conducted. The candidate research contribution is to characterize the cost of evidence sufficient to select or compose a valid offer, and to design representations and access policies that reduce that cost subject to a declared loss bound. Definitions alone establish no claim of priority; the framework must earn its usefulness through reproducible results against strong alternatives.

Keywords: abstract economics; abstract pricing; value spaces; offer spaces; exchange records; mechanism design; information economics; AI agents.

FRONT MATTER

Scope and contribution

This monograph develops a proposed framework for abstract economics, with abstract pricing as its principal economic application. A value space consists of a structured environment and its elements. Exchange possibilities, local valuations, information access, and execution constraints are represented separately. The broad definition of a product as a causal reason for exchange permits applications beyond conventional goods while requiring an explicit causal model and a declared evaluation perspective.

Promoi denotes an offer space associated with a product assortment. Offers specify eligibility, authority, resource use, outcomes, and commitments; composition is partial and depends on these semantics. A kup is an abstract record of a value exchange, independent of its carrier. Events, record contents, and storage media are distinguished. This definition does not imply that every physical exchange leaves a complete, permanent, or recoverable trace.

The manuscript establishes elementary conditional propositions concerning representation neutrality, resource-preserving offer compression, decision loss under partial evidence, and exchange capacity. It develops operational profiles for the stipulated concept of intelligence as the ability to exchange value. These profiles are relative to tasks, available policies, observation interfaces, resources, and time. They are not presented as validated universal measures of cognitive intelligence.

The pricing analysis separates monetary transfers from real surplus and studies multidimensional terms, seller participation, funding constraints, bargaining, and incentive compatibility. Existing research on configurable offers, compositional contracts, information economics, and bounded agents provides the surrounding literature. Current commerce protocols and agent benchmarks are treated as substantive prior work rather than dismissed as payment-only infrastructure.

Two small computational exercises examine finite synthetic worlds and offer-compression instances. They do not constitute market evidence or evaluations of language-model agents. A separate experimental program proposes controlled comparisons of evidence acquisition, offer representation, information accessibility, and execution. Its central candidate contribution is a characterization of the cost of sufficient evidence for a feasible decision, supported by algorithms, limits, and reproducible evaluation.

Keywords: abstract economics; abstract pricing; value spaces; offer composition; exchange records; mechanism design; information acquisition; AI agents.

FRONT MATTER

Contents

The page references below apply to this numbered 100-page edition. The exposition moves from the general structure to the economic specialization, then to information, agents and evaluation. Core symbols are introduced at the outset; local symbols are explained on their first use within each chapter.

Content	Pages
Cover abstracts and reading guide	1-5
Chapter 1 Foundations of abstract economics	6-19
Chapter 2 Offer spaces composition and compression	20-33
Chapter 3 The microeconomics of abstract pricing	34-47
Chapter 4 Exchange records and access to evidence	48-61
Chapter 5 Agent capability and value-space capacity	62-75
Chapter 6 Literature evaluation and experiments	76-89
Synthesis, gains, generalization and reproducibility	90-97
Scientific and technical references	98-100

For economists, pages 34-47 give the direct analysis of utility, profit, surplus and incentives. Pages 20-33 explain what changes when offer terms are represented independently of product identity. For technical readers, pages 48-75 develop the semantics of evidence, capability and time, while Chapter 6 sets out an implementation and evaluation programme.

Readers assessing the scientific contribution should examine the propositions on representation neutrality, compression bounds, non-recoverability and capability measures together with their assumptions. Some results are elementary or close to established results; their role is to make the framework precise. Candidate novelty lies in the questions about costs, algorithms and guarantees that still require development and comparison.

Front matter and references are included in the total page count. This English PDF retains the 100-page structure of the Arabic manuscript. The first abstract is translated; the following scope statement retains the original English overview. Page numbers refer to the PDF edition. Word-processing software may repaginate an editable version after changes to fonts or print settings.

FRONT MATTER

Terminology and reading conventions

Term	Meaning in this monograph
Value space ⚹	A pair consisting of a structured environment E and elements A
Value	Local evaluation or structural effect under a declared perspective
Product	A causal role in enabling value exchange
Promoi	An offer space associated with a product assortment Q
kup	An abstract exchange record independent of its carrier
Accessibility ρ	Information reaching elements under specified conditions
Exchange capability	Tasks an element can accomplish within its constraints

Definition introduces a stipulated term or mathematical structure. Proposition denotes a result derived under specified assumptions, and proof presents its argument. Hypothesis identifies a claim requiring a test. Synthetic examples are constructed for analysis and are not attributed to real markets or users. Proposed experiments are protocols for future work; computations already performed are identified explicitly in Chapter 6.

Equations use chapter-specific labels to facilitate cross-reference. Decision safety means compliance with the constraints declared in the model; it does not imply an unconditional guarantee outside that model or the absence of every real-world risk. Expected value refers to an expectation under an explicitly specified distribution. Probabilistic assumptions do not substitute for missing data.

The text uses author-year citations, with bibliographic details and source links at the end. References distinguish peer-reviewed articles, preprints and technical specifications. The review cutoff is 4 October 2026. Software protocols change over time, so claims about them are tied to the cited version or page.

Intelligence has a broad stipulated meaning within this project: the ability to exchange value without requiring consciousness or intention. Measuring that ability nevertheless requires specifying what an element can actually do. A researcher's available choices must not be attributed to a rock or plant, and transaction volume alone is not evidence of decision quality.

The symbol ⚹ denotes a value space, with ⚹ = (E, A). The same symbol is used in definitions, equations and subscripts throughout this monograph.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

The subject of abstract economics and the limits of its claims

This study proposes abstract economics as the name of a research programme concerned with the conditions under which value exchange is possible, and with ways of representing, measuring and organising exchange across environments and elements whose structures may differ profoundly. Its starting point is to ask which relations allow one element to affect another's possibilities, which constraints govern that influence, and which information makes particular exchanges discoverable and achievable. The programme does not initially presuppose money, markets or human participants. Each application must nevertheless specify what counts as value, exchange and an outcome.

Abstract pricing is the principal economic specialisation within this programme. It examines the analytical separation between a product's identity and the terms under which it becomes the subject of a particular exchange: consideration, eligibility, deadlines, rights, risks and the parties' commitments. This separation allows the same product to be associated with multiple offers and different products to share an offer structure. Expressing a condition, however, does not make it fundable or executable, and introducing a new representation does not itself demonstrate the creation of an additional economic benefit.

The study distinguishes three levels of claim. The first is definitional, specifying the meanings of value space, product, abstract purchase window and exchange record. The second is deductive, deriving results from explicit assumptions about resources, information and time. The third is empirical, testing whether a particular offer representation or discovery policy improves measurable outcomes. Keeping these levels separate prevents a linguistic choice from becoming a purported natural law, or a conditional theorem from being treated as an established description of every market.

The usefulness of abstraction lies in transferring the structure of a problem between apparently different applications. Reserving a store's inventory and allocating processor time may share a constraint against concurrent use of the same resource. Ownership rights, the effects of delay and the possibility of copying the output may nevertheless differ. A scientifically valid transfer therefore requires more than substituting names: it requires showing that the transformation preserves the relevant constraints, information and outcomes. Any distinction omitted from the representation must either be immaterial to the question or be treated as a limitation of the model.

The methodological ambition is to formulate small problems with precise answers within a broad framework. The programme's success depends on explaining previously unclear differences, proving useful bounds, and improving algorithms or experiments under fair comparisons. A useful result might establish that a particular improvement is impossible under fixed conditions, or show that an omitted resource accounts for the entire observed improvement. In this sense, abstract economics bears the ordinary burden of proof in economic and computational research; the breadth of its title does not confer scientific standing.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

Conversion, scarcity and their relation to existing theories

Economic analysis often begins with constrained choice, but a constraint can extend beyond a monetary budget. Time, energy, memory, communication and legal assurance are limited, and some actions may remain unavailable regardless of the consideration offered. The framework therefore proposes describing what elements can do with available resources before assigning prices to them. This order does not make prices secondary in economic applications. It clarifies that the meaning of a price depends on the exchanges an environment can support and the rights under which they are permitted.

The mathematical theory of resources offers a related language by studying resource conversion and composition relative to operations that are admissible or free within a specified model (Coecke, Fritz and Spekkens, 2016). This study draws on that structural perspective without identifying every conversion relation with utility or price. We represent the possibility of converting a resource description x into a resource description y by a primitive comparison relation relative to an explicitly stated set of operations. A resource may denote a system state or an available capability, rather than necessarily a transferable good.

$$x \succcurlyeq_{\mathrm{conv}} y \quad \Longleftrightarrow \quad x \longrightarrow y \text{ is admissible} \tag{F1}$$

If the model permits inaction and its admissible operations are closed under sequential composition, the relation is reflexive and transitive. The proof is immediate: inaction converts a resource into itself, and composing two admissible operations produces a third admissible conversion. This structure is a preorder. Two distinct resources may be mutually convertible, so they need not be distinguished in terms of convertibility. The manuscript makes no novelty claim for this result; it uses it to identify what should be preserved when abstracting different kinds of exchange.

Transitivity may fail, however, if the description omits the first operation's effect on the environment. The second conversion might require a machine consumed by the first, or an authorisation that has expired. The remedy is to include the relevant resources and contextual state in the descriptions of the initial and final conditions, or explicitly acknowledge that the relation is not transitive. Failure of the mathematical structure then exposes an incomplete model, rather than justifying the imposition of a property the application does not possess.

Scarcity alone does not generate preferences, and convertibility alone does not provide a criterion of justice. Two people with identical resources may prefer different outcomes; an operation may benefit one party and harm another. The theory therefore separates physical or technical possibility, institutional permission and preference-based choice. These levels can intersect in a single decision without being reduced to one another. Abstract pricing becomes relevant when constructing exchange terms within that intersection and analysing the information, incentives and resources that determine the actual outcome.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

Value space as a structured environment and elements

We define a value space as the pair ⚹ consisting of an environment E and a set of elements A. An element may be a person, institution, computational agent, physical object or living system. Membership does not require intention or an internal representation of price. The environment is the structure specifying the elements' positions and relations, shared resources and admissible interactions. Elements are described by properties relevant to the question under study. The symbol ⚹ denotes this environment-element pair throughout the monograph.

This two-component definition does not imply that the environment is an empty name or that the elements are merely a list of names. An environmental description may contain a state space, a transition relation, observation channels and a temporal structure. Each element's description may specify its local states, available actions and observations of the environment. These remain components of the descriptions of E and A, rather than a third foundational entity imposed alongside them. Abbreviated notation for these components facilitates proof and computation without altering the conceptual definition of the space.

$$⚹ = (E, A), \qquad A = \{a_1, a_2, \ldots\} \tag{F2}$$

We distinguish the described space from an observer's knowledge of it. An environment may be fully specified in a simulator while an agent sees only partial messages. Alternatively, the researcher may not know every transition of a real market. In the first case, the agent's ignorance is a controlled part of the experiment. In the second, the model approximates an external environment and requires assessment of fit and uncertainty. Confusing these cases makes missing data appear to be an environmental law, or gives the agent the designer's knowledge at no cost.

The division between environment and element is an analytical choice that must be stated. Depending on the problem, a platform may be represented as a decision-making element or as part of an environment implementing fixed rules. Changing this choice requires preserving the relevant actions, information and constraints. Moving a capability from an element into the environment does not justify attributing all its results to the element's intelligence. A fair comparison traces the source of computation, resources and authorisation in both representations.

Elements can be composite and nested: a company contains agents, an agent uses another service, and a plant system interacts through multiple parts. The level of analysis and the boundaries of the parties must then be specified to avoid counting one interaction as several independent occurrences. The model must also state whether relations are fixed or can change through exchange. Creating a channel or changing a party's eligibility can alter future opportunities. Recording a single product transfer may therefore be insufficient; one must trace the effect on the environment's structure and on the actions that subsequently become possible.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

States, histories and admissible transitions

A state represents the information a model needs to determine its possible subsequent evolution. In a simple store it includes inventory and outstanding commitments; in a computing network it includes load, memory and reserved tasks. The state need not be fully observable, nor need its representation be unique. Sufficiency is relative to purpose: if two histories assigned the same represented state permit different executions, the model has omitted relevant information, such as a previously consumed entitlement or an unsettled obligation.

We use a transition relation T connecting an initial state, a described event and a subsequent state. A history is a sequence of states and events consistent with this relation. The symbols s, e and h denote a state, an event and a history, respectively, while n is the length of a finite history. Not every transition event is a value exchange: it may represent the expiry of a deadline, a failed message or an internal update. A separate criterion determines which events belong to the exchanges studied in the problem.

$$h = (s_0, e_1, s_1, \ldots, e_n, s_n), \qquad (s_{i-1}, e_i, s_i) \in T \tag{F3}$$

The relation is more general than a deterministic transition function, since the same state and action may admit multiple outcomes. Probabilistic analysis requires a probability kernel over outcomes, together with its informational conditions. A worst-case guarantee requires specifying who selects among admissible outcomes and what information that selector possesses. Existential attainability means that a successful path exists. Guaranteed capability means possessing a policy that succeeds against every adverse evolution allowed by the model. These are distinct requirements.

An interaction may involve many parties concurrently. It can be represented by one joint event if its internal detail is immaterial, or by a partially ordered event network if the order of messages and commitments matters. Converting that network into a computational sequence may assist implementation, but it does not give the artificial ordering additional causal meaning. The model should specify whether independent operations can be reordered without changing the outcome, rights or information.

This language allows complete execution to be defined as an admissible history satisfying explicitly stated completion conditions. A seller's promise, a buyer's authorisation, payment acceptance and delivery of the benefit may occur in different events; achieving one does not establish completion of the others. Nor does an error at the end of a history erase obligations created earlier. Where relevant, the state description should therefore include the effects of failure, compensation and cancellation, so that the study measures interpretable economic outcomes rather than merely counting success messages produced by the agent.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

Deriving minimal rules from a description of the space

Minimal rules can be defined from the environment and elements without presupposing an independent list of rules, provided their description contains the structure needed for inference. We first fix a formal language L capable of expressing the relevant properties, and let D denote the stated description of E and A. Depending on the problem, this description includes types of states, possible transitions, and properties of resources and observations. A derived rule is a sentence true in every model compatible with that description, rather than one the researcher expects from the environment's name. Here, sentences of the language mean closed formulae with a truth interpretation within a model, not merely items of vocabulary.

The following formulation defines minimal R as the set of semantic consequences of the description; the entailment operator means truth in all compatible models. The word minimal indicates that no particular commitment extending beyond the given description is introduced. It does not mean the smallest axiom list by number of sentences, nor guarantee that a finite list generates the entire set. The set may contain infinitely many consequences, and several logically equivalent foundational formulations may differ in form or length.

$$R_{\min}^{\mathcal{L}}(E,A) = \{r \in \mathcal{L} : D(E,A) \models r\} \tag{F4}$$

Foundational proposition: if the description is consistent, every rule derived in this way is true in any realisation of the space that satisfies the description. The proof is simply the definition of semantic entailment: a compatible realisation belongs to the collection of models in which the rule must hold. The proposition's importance is methodological. A rule's validity depends on the description's correspondence to the environment and on the language's meaning. Inference does not thereby acquire the power to correct false inputs or determine information absent from the description.

The description must have at least one model. If its statements are contradictory, the conventional definition of entailment makes every sentence follow vacuously, and the rules cease to describe a possible environment. In applications, contradictions may be detected with constraint solvers in suitable domains, addressed by identifying a reliable consistent subset, or treated through an explicitly declared alternative logic. These approaches must not be interchanged silently, because changing the logic changes the meaning of the proof.

Rules also depend on the expressive power of the language. A language without a representation of time cannot formulate an offer-expiry rule; one representing inventory only as a total quantity may fail to distinguish reserved units from available units. The language and description should therefore be published with the results. Removing R from the list of foundational components does not remove the need to specify semantics. It makes a rule's origin traceable and reveals whether it follows from the structure or is an optional policy added to a particular version of the environment.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

Limits of inference and learning rules from observation

Names for the elements and environment do not suffice to derive laws of transformation. To establish this, consider a constructed logical example rather than a physical experiment. There are two elements, each with state zero or one, and the environment begins with the first in state one and the second in state zero. The shared description states only these facts. In one realisation we add a transition transferring the one from the first element to the second. In another, we add a transition copying the one to the second while leaving the first unchanged.

Both models satisfy the same description of names and initial states, but the number of units is conserved in the first and increases in the second. Conservation of that number is therefore not a logical consequence of the description: if it were, it would hold in both realisations. Likewise, neither permission nor prohibition of copying follows without additional information. This example does not contradict deriving rules from structured E and A. It identifies precisely the kind of structure that must be included for the inference to be justified and checkable.

Inferring laws from observations is a different problem. A researcher may observe thousands of number-preserving transformations and propose a conservation law, yet finite data do not automatically exclude a rare, unobserved transition. Learning produces a hypothesis with a confidence level or domain of validity, depending on assumptions about sampling, temporal change and the model class. It becomes a certain logical consequence only when sufficient assumptions enter the description. One must then distinguish what the observations contributed from what those assumptions supplied.

This distinction can be managed by representing a set of realisations consistent with the evidence rather than prematurely declaring a single law. New information may narrow that set or reveal that the earlier description was inadequate. The formulation has immediate relevance to economic agents: the absence of a discount-use record does not establish that the discount was unused, and a seller's silence does not establish inventory availability. Permissible commitments depend on the conclusions supported by the evidence and the type of guarantee required, while accounting for the cost of obtaining further information.

Optional policies have a different origin as well. A purchasing limit imposed by a platform is not a natural law deduced from its being a platform; it is a documented institutional property of the current environment. It can be included in the description, its consequences derived, and a modified environment studied in which it is removed or replaced. This treatment distinguishes what is necessary in every realisation of an abstract space, what is fixed in a particular implementation, and what is statistically common in specified data. Each conclusion should clearly identify the level to which it belongs.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

Value as effect, ordering and numerical measurement

Value in this framework is neither synonymous with money nor a numerical property inherent in everything. We begin with a describable effect: a change in a resource, right, possibility or state. We then specify an evaluative perspective Γ and an element a, representing state evaluations by a function into a suitable domain L. The domain may be an ordered set, a vector of attributes or the real numbers. The perspective is a parameter of measurement and interpretation within the space; it need not become a third foundational component alongside the environment and elements.

If the domain provides only an ordering, one can say that an outcome is at least as preferred as another without subtracting values or calculating a percentage increase. If evaluations are vectors, such as cost, time and risk, two outcomes may be incomparable because one is better in one dimension and worse in another. Numerical measurement requires justification for its unit and for the interpretation of differences. Assigning a number to every state does not, by itself, make adding those numbers or comparing them across people economically meaningful.

$$v_{a,\Gamma} : S \longrightarrow L_{a,\Gamma} \tag{F5}$$

Evaluation may depend on an entire history rather than its final state alone. Medicine arriving after it is needed is not equivalent to arrival beforehand, even when final ownership is identical; a repeatedly interrupted service differs from a continuous service with the same average. The state domain must then be replaced by a suitable history domain, or the state expanded to retain relevant temporal information. The model must specify whether evaluation occurs before or after uncertainty is resolved, and how it handles probabilistic outcomes or circumstances lacking reliable probabilities.

The formulation accommodates elements without intentions without attributing psychological preferences to them. For a plant, the researcher may choose survival or growth as an explicit evaluative perspective; for a stone, one may measure energy, stress or displacement. These are measures of interaction effects. They do not establish that the stone desires an outcome or that the plant computes numerical utility. Claims about goal-directed behaviour require further evidence or application-specific assumptions; they do not follow from using the word value.

The study does not assume conservation of aggregate value. An environment may conserve a particular physical quantity while economic utility changes with knowledge, timing or distribution. Exchange may benefit two parties or harm a third; calling a process an exchange does not establish its social benefit. Each problem therefore declares its outcome criterion and the parties included, separating resource transfer from changes in evaluation. When a numerical welfare sum is used later, it is an analytical choice conditional on explicit units and weights, rather than a general law of value spaces.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

The product as a causal role in exchange

We adopt the definition of a product as any causal reason for value exchange between two or more parties, whatever their nature. This includes an object, information, a service, an arrangement or a channel enabling interaction; the cause need not be a good offered for sale. Being a product is thus a role within a specified relation, rather than a fixed material category. Information may be a product in a knowledge exchange, an input into producing a service in another exchange, and an incidental effect with no causal role in a third.

Using the word cause requires specifying the context and mechanism of influence. Causal analysis distinguishes observing an association between variables from the effect of changing one through a controlled intervention. This distinction provides an appropriate methodological basis for the proposed definition (Pearl, 1995). In a probabilistic model admitting interventions, a change in the distribution of outcome Y when variable X is set to two different values can provide evidence of that variable's causal effect in the specified context. Pr denotes a declared probability, and do denotes intervention within the model rather than merely observing a realised value.

$$\Pr(Y \mid \mathrm{do}(X = x)) \neq \Pr(Y \mid \mathrm{do}(X = x')) \tag{F6}$$

The equation illustrates an operational criterion for an effect, not an exhaustive definition of every possible cause. Averages may conceal opposing effects, and two causes may each suffice to produce an outcome, so that removing one fails to reveal its effect in the particular occurrence. Some environments also do not permit actual intervention. The researcher then needs a mechanistic model or assumptions identifying effects from observations. Such assumptions must be disclosed: a purchase occurring alongside an advertising message is insufficient to attribute the purchase to that message.

The definition's breadth does not require every background condition in the universe to be called a product in every study. We specify the system boundary, the level of detail and the causes to be analysed, acknowledging that background causes remain fixed in the comparison. In a plant-light interaction, radiation flow can be studied as causing a valuable transformation under an explicit biological criterion, without treating the sun as a seller or attributing commercial intent to the plant. In a collision between bodies, momentum transfer may be studied as an exchange effect under a specified physical model, without assuming that every collision is beneficial.

Economically, causation alone establishes neither ownership nor a right to set a price. Public information or shared infrastructure may contribute to exchange without any party having the right to sell it or claim a return from it. The market specialisation therefore supplements the product's causal role with descriptions of rights, transferability and contracting. This distinction permits analysis of demand sources and access channels without turning every causal contribution into a financial entitlement, or confusing evidence linking an event to a channel with evidence that the channel changed its occurrence.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

Exchange events, parties and effects

An exchange event is a unit of analysis involving at least two distinct parties and a relation of influence relevant to value under an explicit criterion. The general meaning does not assume that both parties provide equal or simultaneous consideration, consent to the event, or benefit from it. Flow may be one-way in the chosen description, and the effect may arise through an intermediary. Voluntary market exchange is a more specialised category requiring definitions of rights, acceptance and consideration. Its conditions are not imposed on all natural interactions.

We represent an event by an occurrence identifier, a set of parties, a state before and after the event, and a context specifying the relation and the product's causal role. The identifier distinguishes occurrences with similar descriptions, such as two purchases of the same value on different days. U denotes the set of parties and c the context, which may specify the interaction channel, relevant constraints and measurement boundaries. This representation need not imply that any single observer knows every component. It is an analytical description of the event; the information available to each element remains a separate question.

$$e = (\iota_e, U_e, s^-, s^+, c_e), \qquad |U_e| \geq 2 \tag{F7}$$

The economic or causal question determines which parties are included. In a shipment, the seller and buyer may be the contractual parties while the carrier participates in the execution network. The event may also affect local residents through emissions not priced in the contract. We therefore distinguish participants in execution, holders of obligations and affected elements. Restricting transaction analysis to a small group does not remove externalities. They must be acknowledged when the evaluation criterion becomes social or encompasses the wider environment.

A kup record is content representing an event or some aspect of it, carried by a mind, a server or a physical trace. Differences in carrier do not alter the logical distinction between an exchange occurring and its description. Multiple records may concern one occurrence; partial records may conflict; a trace may be insufficient to reconstruct the entire event. Consequently, the number of records is not the number of exchanges, and a recorded promise is not evidence of delivery. Each application requires an explicit relation between recorded content and the facts it establishes.

This separation assists the analysis of composite exchanges. Reserving, paying for and using a service may be counted as one transaction or several component exchanges, depending on the counting criterion. Either choice is permissible if declared beforehand and kept fixed across comparisons. Changing the granularity to increase the count changes the measure rather than improving the space. Reversals, cancellations and compensation should be recorded according to their actual effects on rights and resources, so that retransmitting a failed request or backing up a receipt is not counted as new economic value.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

The economic specialisation: rights, acceptance and contracts

An economic structure arises when particular outcomes become subjects of choice or negotiation under resource and rights constraints. A party may possess a benefit it can provide but lack authority to commit it, or have contractual authority while being unable to execute without a third party's resource. The model therefore distinguishes technical capability, authorisation, ownership and consent. A contract adds conditions specifying what the parties must do and what happens upon failure, whereas an admissible transition describes what the environment actually allows.

In an explicitly specified numerical utility model, individual rationality can be defined by comparing the utility of participation with that of a non-participation alternative, after accounting for relevant costs and risks. We denote the contractual outcome by x, the alternative by x_0, and party i's utility by the function u. The alternative should be realistic given the available information and opportunities, rather than chosen conveniently to make the contract appear beneficial. When a justified numerical utility is unavailable, the condition can instead be expressed through a preference ordering.

$$u_i(x) \geq u_i(x_0) \tag{F8}$$

Individual rationality guarantees neither truthful disclosure of preferences, the intermediary's budget balance, nor maximum welfare for the parties. These are distinct properties requiring separate assumptions and mechanisms. An offer may also appear acceptable before uncertainty is resolved while its realised outcome is undesirable. The guarantee must specify whether it applies before or after information is revealed, and whether it concerns expected utility or every possible outcome. Concealing these distinctions within the expression fair offer leaves the claim unspecified.

Abstract pricing treats consideration as part of a describable bundle of commitments. A price may vary with the delivery deadline, order volume, an accompanying service or an eligibility condition. A single structure can represent a discount, instalments or a priority right without claiming to have invented those practices. The potential research contribution lies in systematically analysing which can be composed, which consume shared resources, and how an agent recognises ordering effects and conflicts before making a commitment.

Publishing an offer must remain distinct from a seller adopting it, a buyer discovering it, and its successful execution. An agent may propose a linguistically intelligible condition that no provider accepts, or a provider may accept a condition incompatible with another outstanding commitment. The price paid alone also measures neither the buyer's gain nor the seller's cost. Evaluating improvement requires comparing complete outcomes: product quality, time, risks, recourse rights, and the costs of search, verification and execution. This specialisation supports an economics of offers within value space while keeping institutional assumptions visible and open to modification.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

Temporal structure and the unit of exchange counting

Time in an abstract space is part of the environmental description; it need not be a global clock using familiar units. It may consist of discrete rounds, continuous duration or a precedence relation between events. In distributed systems, the happened-before relation helps distinguish precedence supported by communication from an artificial ordering imposed by a logical clock (Lamport, 1978). This distinction is useful when information is insufficient to determine whether two commitments are concurrent or sequential from the standpoint of execution.

If a strict precedence relation is chosen, it is irreflexive and transitive by the model's assumptions. The notation $e_1 \prec e_2$ means that the first event precedes the second in the specified structure. Two events may remain incomparable. This does not mean they have the same physical timestamp; it means the description has not ordered one before the other. A window J can be defined as a specified portion of this structure, provided its boundaries and the events counted within it are made clear.

$$e_1 \prec e_2,\ e_2 \prec e_3 \implies e_1 \prec e_3, \qquad e \not\prec e \tag{F9}$$

The number of exchanges cannot be divided by time without an explicit duration measure and unit. Where only steps exist, one can discuss completions within a specified number of rounds; where precedence alone is defined, one can study dependency-chain lengths and parallelism. Transferring a result about exchanges per second to a world without seconds requires a conversion of units and semantics, rather than relabelling a graph's axis. Some elements may also have local time distinct from the environment's overall measure.

Capacity additionally requires a fixed counting policy. We specify what makes an exchange complete, how repeated occurrences are distinguished, and what weights apply to different types, if any. Counting the exchange of a large file as one event and then counting each fragment separately changes the quantity without necessarily changing capability. One can therefore declare a standard work unit or report parallel measures, such as completed tasks and the volume of service provided. Heterogeneous units should not be added without a justified conversion.

Some abstract environments require constraints preventing unbounded completions within a bounded window. A positive minimum cost per completion and a non-renewing budget provide a simple numerical bound. By contrast, permitting a repeatable cycle that produces a new completion without consuming time or a limited resource may make capacity infinite in the model. This is a consequence of chosen assumptions, not evidence of a real physical capability. If an environment permits cyclic temporal relations, partial-order results no longer apply automatically; alternative semantics must be developed before those results are used.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

Abstract worlds with different resource laws

Value spaces can be designed for environments that do not follow familiar physical laws, provided those laws are presented as explicit construction assumptions. The initial scientific purpose is to test how results depend on the structure of resources, time and information. No claim that such worlds exist outside the model is required. We might ask whether an allocation algorithm succeeds because of scarcity, because a resource is transferable, or because copying information leaves the original resource unconsumed at the chosen level of description.

Consider two constructed examples sharing two elements, each with a binary state. In the transfer world, a transition converts the condition in which the first possesses a capability and the second lacks it into the opposite condition. In the copying world, the second acquires the state while the first retains it. The arrows denote different relations, not two options automatically permitted within one environment. The digits are not monetary or energy units; they indicate possession of an abstract capability. Any cost or usage right must be added explicitly within the environmental description.

$$(1,0) \longrightarrow_{\text{transfer}} (0,1), \qquad (1,0) \longrightarrow_{\text{copy}} (1,1) \tag{F10}$$

Under this transition alone, the number of possession states is conserved in the transfer world and increases in the copying world. This observation determines neither utility nor price. A copy may be useless to an element already possessing the capability, or require a key, expertise or a complementary resource. A third world could permit copies while consuming a scarce authorisation, making the content reproducible but the right constrained. The example illustrates how some market complexity arises from differences between the properties of an object and those of permission to use it.

Communication structure can also change while conversion rules remain fixed. A world with direct connections between all elements differs from one mediated by a single gateway, even with identical aggregate resources. The intermediary's position, capacity and policy then help determine the feasible exchanges. An agent's superiority in one world does not establish general superiority: success may depend on a channel unavailable to its counterpart. Families of environments and the scope of comparison must therefore be specified before measuring the ability to exchange value.

Synthetic experiments permit one factor to change while others remain fixed, which is difficult in a real market. Such control does not remove the problem of model selection. A result holds within the constructed environment; transferring it to an economic application requires matching the essential constraints. This study proposes using such worlds to derive counterexamples and theoretical bounds, then testing candidate mechanisms in environments closer to application. If a result fails to transfer, identifying the missing property is itself an important scientific outcome, rather than renaming the world to make it appear consistent with the theory.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

Equivalent representations and invariant economic outcomes

A transaction can be described in natural language, a table of terms, a program or a state network. Changing the representation may improve understanding and computation without automatically changing the set of possible economic outcomes. To clarify this, we define equivalence relative to a specific problem: a bijection between admissible histories in two representations preserves the resources, rights and information required by that problem, together with each outcome's score. Every linguistic detail need not be retained; every distinction that can affect choice, execution or measurement must be.

Representation-invariance proposition: if a bijection matches every feasible history in the first model to a feasible history in the second, preserving objective values and constraints, the best achievable objective values are equal. H_1 and H_2 denote the sets of eligible histories, f the transformation, and W the objective criterion in each representation. The formulation uses a supremum to include cases in which no best solution is attained. This is a general structural result; the manuscript does not present it as a new discovery.

$$\sup_{h \in H_1} W_1(h) = \sup_{h \in H_1} W_2(f(h)) = \sup_{g \in H_2} W_2(g) \tag{F11}$$

For the proof, every value achieved by a history in the first set is achieved by its image in the second, so the second supremum is at least the first. Applying the inverse bijection yields the opposite inequality, establishing equality. If the values differ in an application, either a preservation condition fails or the computation or measurement is incorrect. This observation helps identify the source of improvement rather than attributing it to abstraction as an unspecified force.

The proof concerns ideal possibilities and does not suffice to compare bounded agents. Extracting a condition from prose may be time-consuming even when it is immediately visible in a table. To establish equal agent capability, the transformation must preserve available policies, processing costs and information sequences, rather than final outcomes alone. Probabilistic environments also require preservation of the relevant outcome distributions; adversarial guarantees require preservation of the adversary's available choices. A bijection between terminal states does not imply these stronger conditions.

Abstract pricing can accordingly produce justified improvements by reducing search or verification costs, enabling terms that previously could not be expressed or adopted, or improving coordination among dispersed resources. These effects must be separated empirically. Comparing two representations of the same offers measures interface and computational effects; adding offers measures expanded opportunities; adding information that establishes eligibility measures knowledge effects. An application may legitimately combine them, but attributing improvement to one scientific mechanism requires a design that separates their contributions and accounts for their costs.

CHAPTER 1 FOUNDATIONS OF ABSTRACT ECONOMICS

Testability and the limits of the foundational theory

General definitions organise language and are not falsified in the same way as predictive hypotheses. Intelligence may be defined as the ability to exchange value within a space, but the definition's scientific usefulness depends on tasks, resources and measures that distinguish capabilities and explain outcomes. A product may likewise be defined through its causal role, but establishing that a particular element played that role in a particular occurrence requires a model, experimental design or appropriate evidence. The conceptual construction therefore remains incomplete until linked to procedures specifying what counts as success or failure.

Every applied claim should take a comparative form: under a declared family of environments, a fixed task set and a specified budget, a proposed method achieves a stated result relative to a reference method. That result might be lower decision loss, fewer failed commitments, faster search or more economical resource use. More exchanges are insufficient if their quality declines, and a lower price is insufficient if the study omits time, risk and verification costs. The criterion must also identify whether it concerns an individual, a group of parties or the wider environment.

We propose separating mechanisms experimentally: hold products and offers fixed when testing representation, representation fixed when testing accessibility, and information fixed when testing selection policy. Where complete separation is impossible, report the result as a combined effect accompanied by sensitivity analysis. Small instances that can be solved exactly should check implementation correctness, while larger instances assess scalability; their different purposes should be explicit. These are experimental-design proposals. These pages make no claim that an experiment has been executed or market data obtained.

Mathematical proof requires review independently of successful simulation. An algorithm may succeed on every tested sample and fail on an admissible instance that was never generated. Conversely, a proof may be correct while its assumptions fail to describe the intended application. Researchers should therefore publish assumptions, counterexamples and limits to generalisation alongside positive results. Negative results are particularly valuable when they show that missing information makes a guaranteed decision impossible, or that compression fails to preserve offer-resource constraints under a particular composition.

This manuscript remains a foundation for a research programme, rather than an announcement of a proven universal law of economics, physics or intelligence. Progress is measured by accumulating useful definitions, checkable conditional results and reproducible comparisons. The move to abstract pricing is justified when the foundations support specifying an offer space, describing the parties' commitments, and measuring how information and resources affect execution. Extension to other environments depends on preserving that structure or developing a new one while making explicit which conditions of inference have changed.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

Promoi: the offer space associated with a product assortment

The economic specialization of value-space theory distinguishes a product's identity from the terms under which it can be exchanged. The physical good may remain unchanged while the delivery date, payment method, right of return, quality guarantee, and quantity conditional on a group purchase vary. Each combination of these terms has different consequences for participants and resources. A catalogue of products and individual prices therefore cannot represent all exchange opportunities, even when the number of products is small and known in advance.

We use Promoi to denote the offer space associated with a specified product assortment, written O_Q. Here Q is a declared assortment containing product identities and the attributes needed to distinguish the subject of exchange. Space means an organized collection of offers and the compatibility, substitution, and composition relations among them; it does not automatically imply a metric or topological structure. A distance may subsequently be defined if it serves a clear analytical purpose, such as measuring differences in requirements or the cost of moving between particular offers.

$$O_Q = \{o \in \mathcal{O} : \ \mathrm{scope}(o) \subseteq Q\} \tag{P1}$$

In the equation above, the ambient set O represents the offer language chosen for the model, while the scope function identifies the products to which an offer relates. At the outset, the researcher specifies whether ancillary products, such as delivery and insurance, belong to the assortment or are external services invoked by an offer. This choice affects relations between spaces. Composing two offers associated with different assortments may produce an offer in the space of their union, provided that the intermediate services and resources actually needed for execution are included.

Membership in this space does not by itself make an offer a binding contract. An offer may be a buyer's proposal, a published template, an offer adopted by a seller within specified limits, or a commitment accepted by both parties. Multiple entities may adopt a single template with different subsidiary conditions, producing economically distinct instances. Both template identity and adopted-instance identity should therefore be retained: similarity supports search, whereas executability depends on the entity, the instance, and the actual state.

The literature on multi-attribute auctions already studies configurable offers containing alternative attributes, discount rules, and allocation across multiple suppliers; Bichler and Kalagnanam (2005) provide one example. This research builds on that precedent and asks how an offer space can be organized, compressed, and accessed under limited resources. The required scientific distinction concerns results, guarantees, and tests, rather than merely introducing a new term for programmable offers.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

An offer as a typed specification of commitments

Studying offer composition requires a structure that exposes conflicts before commitment. A brief description, such as a discount with free delivery, may conceal differences in the beneficiary, the funding party, and the date of entitlement. We therefore define a typed offer through its product scope, authorized entity, eligibility, activation conditions, resource consumption, outcomes, commitments, and validity window. This is an extensible modelling choice; it does not assert that every exchange in the general value space takes the form of a commercial offer.

$$o = (Q_o, \alpha_o, g_o, h_o, a_o, \psi_o, c_o, J_o) \tag{P2}$$

The symbol α specifies the authority required to issue and adopt the offer; g denotes the participant's eligibility test; and h denotes an activation condition depending on the environment's state and history. The component a describes resources consumed or reserved, ψ specifies permitted outcomes, and c specifies participants' commitments. Finally, J is the validity window under the environment's definition of time. In the general model, these components may be functions or relations; in finite computational models, they can be represented by finite data and tests.

Monetary and resource types must be distinguished. Currency units, computing time, inventory units, and single-use rights are not quantities that can be added without specifying their types. A resource vector aggregates quantities of the same type within each coordinate, and comparisons are componentwise. Similarly, the outcome of promising a service differs from delivering it, just as refundable credit differs from conditional or non-transferable credit. Sound composition requires the representation to retain these distinctions.

In a synthetic example, a computing provider offers ten hours for eighty monetary units, conditional on overnight operation and limited to one beneficiary. Execution requirements are not exhausted by the price: there is a ten-hour operating quota, a timing constraint, and a right of acceptance that cannot be copied for independent beneficiaries. If a sponsor covers ten units of the price, an additional financial resource tied to the sponsor's budget is involved. Neither the appearance of a logo nor the provider's signature alone establishes that this budget is available.

A typed representation benefits the agent by enabling it to request missing information instead of guessing. If eligibility is known but the validity of authorization is not, the next question concerns authority; if authority is established but quota availability is unknown, the question concerns the resource. Each rejection can also carry an explanation identifying the responsible condition. Such explanations describe and document the decision; they do not establish that an exchange has been completed unless the represented events occur and are verified through the environment's procedures.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

Publication, adoption, discovery, and execution are distinct states

An account of the offer space must distinguish actions that create a description from actions that create or discharge a commitment. A participant publishes a template, making its information available for circulation; a seller may adopt it subject to product, country, and budget limits. A buyer discovers it when its information arrives, after which compliance with its conditions can be assessed. This sequence of statements does not prescribe a single mandatory sequence of events: a private proposal may arrive before any public release, and an expired offer may be discovered after execution for an earlier beneficiary.

Analytical distinctions between offer states; states may overlap and change over time.

State	What it establishes	What it does not establish
Published	An available description	Any seller's commitment
Adopted	Adoption by a specific entity	Availability of all resources
Discovered	Information receipt	Understanding or eligibility
Reserved	Conditional resource allocation	Completed delivery
Executed	Required events occurred	Satisfaction of all parties

Representing these states as independent properties is more useful than placing each offer in a single category. An offer may be published, adopted, and discovered yet currently unexecutable because inventory is exhausted. It may be physically executable but unauthorized for the agent attempting to use it. A cached copy may also remain available after withdrawal. Research comparisons must therefore identify the instance and reference state, rather than relying solely on an offer's name or product-page title.

This distinction is directly related to the meaning of kup. A publication record documents publication; a record of the parties' acceptance documents acceptance; and a delivery record documents the available evidence of delivery. All may be interconnected parts of an exchange history, but their evidentiary functions are not automatically interchangeable. A single purchase interface may display them together, while the exchange record remains a broader concept than either the interface or the medium that stores it.

When several sellers adopt one template, usage limits must be identified as local to each seller or shared across sellers. The phrase single use may mean once per store or once per beneficiary throughout the network. These meanings generate different sets of possible exchanges. A local counter is insufficient for the second meaning when concurrent acceptance is possible, even if every local record is internally consistent.

The UCP discounts specification recognizes discount application order and cases requiring additional eligibility verification (Universal Commerce Protocol, 2026). This study uses those functions as a practical illustration of the importance of explicit semantics, then extends the analysis to shared resources and multiple parties. Extending a description remains distinct from proving correct execution: a protocol specifies messages and states, whereas a guarantee requires assumptions about entities, data, and reservation and verification mechanisms.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

Sequential composition and the non-commutativity of offers

An offer may alter the state in which a subsequent offer tests its conditions. Changes may concern the cart total, balance, beneficiary's eligibility, or remaining time. We therefore begin with partial sequential composition: the first offer executes, and the second is then checked in the resulting state. If the first cannot execute, or its outcome conflicts with the second's conditions, composition is undefined in that state. Undefined composition is not treated as a zero discount unless the system's semantics explicitly make that choice.

$$T_{o_2 \circ o_1}(s) = T_{o_2}(T_{o_1}(s)) \tag{P3}$$

Here s is a sufficiently informative state, and T is a partial transition function in the deterministic case. With multiple possible outcomes, the function is replaced by a transition relation or a probability kernel, with an explicit interpretation of expectations and risk. Composition is read from right to left: the first offer is applied, followed by the second. The state must contain all resources and commitments relevant to the next step; checking compatibility on price alone is invalid if a right or hidden constraint changes during execution.

A synthetic numerical example illustrates non-commutativity. Let the original price be 120 units, with one offer providing a 20 percent discount and another a fixed reduction of 15 units. Assume that the fixed reduction does not make the amount negative, that the percentage discount applies to the current amount, and that combining the two offers is authorized. Applying the percentage discount first yields 81, whereas applying the fixed reduction first yields 84.

$$0.8 \times 120 - 15 = 81 \neq 0.8 \times (120 - 15) = 84 \tag{P4}$$

Changing the order may have effects beyond the final amount. If free delivery requires a cart total of at least 100 units after discounts, reducing the price first may remove delivery eligibility. If the threshold is calculated before discounts, this consequence disappears. Describing offers as combinable does not resolve either case: the calculation base and testing order must be specified. Similarly, purchasing a subscription may confer a right to a subsequent offer, whereas reversing the order does not create the subscription retroactively.

Any algebraic law of composition requires proof under specified semantics. Transition-function composition may be associative when the complete state is preserved, yet regrouping offers operationally may change reservation timing, monetary rounding, or failure handling. In that case, the associativity of function composition does not by itself establish equivalence between two operational procedures. Compositional contract languages study contract meanings through event histories and ways of completing them (Andersen et al., 2006), providing a basis for formalizing this distinction between a compact expression and an execution history.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

Partial composition: parallel offers, alternatives, and bundles

Composing two offers in parallel differs from choosing one of them, and both differ from requiring a complete bundle. Parallel composition incurs commitments in both branches, accounting for overlapping resources and possible event orderings. An exclusive alternative permits one branch under a declared choice rule. A bundle links value or eligibility to the completion of several components and may require a joint completion condition. Replacing these operations with an undifferentiated notion of addition loses essential information about commitments.

In a simple static model with independent resources, checking aggregate consumption suffices when the only constraints are additive constraints with known capacities. An offer's resource consumption is represented by a nonnegative vector, compared coordinate by coordinate. This check is nevertheless a limited numerical condition, not a complete definition of compatibility. Resources may be available while composition remains impossible because authorized beneficiaries differ, delivery dates conflict, or joint use is expressly prohibited.

$$a(o_1) + a(o_2) \leq B \tag{P5}$$

For a synthetic example, two offers require three and five units from inventory with a capacity of seven. Each offer is individually executable, but simultaneous execution is impossible. If they use two independent inventories, these must not be collapsed into a single coordinate merely because the product names match. Two offers may also use the same unique voucher; the voucher right must be represented as a resource of capacity one so that the conflict appears in the model rather than being concealed by an acceptable monetary total.

A different problem arises when a bundle spans two sellers. Reservation of the first component may succeed and reservation of the second fail. The mechanism may prevent commitment until reservation is complete, allow partial execution, provide compensation, or grant the buyer a specified cancellation right. Each choice has costs and consequences for resources and incentives. A unified interface or shared offer identifier does not imply atomic bundle completion; the guarantee depends on the parties' procedures and what the environment can enforce.

Exclusive alternatives require specification of who holds the choice and when it can be exercised. A buyer's right to choose at delivery differs from a seller's substitution right, even if the alternatives have the same expected value under one model. A choice right may reserve inventory for several alternatives before the final branch is known, making flexibility more costly than the delivered good itself. The offer space therefore records the difference between resources ultimately consumed and resources that must remain available throughout the commitment period.

This analysis supports a classification of conflicts into resource, eligibility, authority, timing, and contractual-semantic conflicts. Each may require a different remedy: a resource conflict calls for reservation or reallocation, whereas an authority conflict requires valid authorization. An agent's ability to identify the cause of failure is a measurable outcome alongside successful exchange, because repairing an offer depends on identifying the constraint that actually prevented it.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

Transfer of rights and the limits of offer-product independence

Separating an offer from product identity enables template reuse and the transfer of certain rights, but does not make every offer a tradable asset. Informational copyability is not equivalent to transferability of entitlement. Someone may send another person an image of a personalized voucher while the authorized beneficiary remains unchanged. A seller may permit one transfer for a fee, prohibit transfer after reservation, or tie the right to membership. These differences concern the right's content and enter the offer's structure before any analysis of its price.

We define valid exercise of a right through four conjunctive tests: the entitlement's origin is established; the exercising participant is currently authorized; the right remains unspent; and execution conditions still hold. These tests are not synonymous with a digital signature. A signature may help attribute a message to an entity, but alone establishes neither continuing inventory availability nor the absence of a conflicting commitment elsewhere. In other environments, evidence may come from an institutional record, participants' memories, or inspectable natural traces.

$$\mathrm{valid}(o, i, s) = \mathrm{origin}(o, s) \wedge \mathrm{auth}(i, o, s) \wedge \mathrm{unspent}(o, s) \wedge \mathrm{feas}(o, s) \quad \text{(P6)}$$

Here i identifies a participant in the model and s is the reference state in which the tests are evaluated. The agent may not know some function values, so the system distinguishes true, false, and unresolved. A kup record becomes a possible source of information about a transfer or exercise of a right, not a logical substitute for the event. Conflicting records require a specified resolution mechanism and its epistemic limits; an absent record does not automatically imply that the right has not been used.

In a synthetic example, a seller issues one hundred rights to purchase at a discount and permits each right to be transferred to another beneficiary once. Circulating five copies of the same description creates five instances of information receipt, not five purchase rights. Representing each right by a resource identity with capacity one can prevent its allocation to different buyers within the computational solution. Enforcing this in a distributed system requires verification and reservation that connect decisions to actual resources at commitment time.

A discount right must also be distinguished from a right to obtain the good. One contract may guarantee a discount conditional on continuing product availability, while another guarantees a reserved unit at a specified price. Their marketing interfaces may resemble each other, but their allocation of risk differs substantially. The theory therefore compares complete bundles of rights and commitments, including failure, refund, and transfer conditions. Their prices may differ while the underlying good remains unchanged, illustrating a precise economic meaning of separating the offer from the product.

The meanings of authorization and enforcement depend on the institutional environment. This formulation states conditions within an operational model; it does not establish general legal validity across environments. That limitation is analytically useful: transfer validity can be tested against declared rules, after which changes in those rules can be studied for their effects on liquidity, scarcity, and verification costs, without conflating a conditional mathematical result with a universal institutional judgment.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

Representation neutrality and claims of value creation

Proposition: a structure-preserving recoding leaves the optimum unchanged

Suppose two representations of one offer space admit a bijection between their feasible execution histories. The bijection preserves participants, outcomes, resources, timing, preferences, and the access and computation costs relevant to the decision. It also preserves what each party observes, when it observes it, and its available policies. If recoding preserves all these elements, the optimum under a declared criterion is unchanged. This foundational proposition prevents economic improvement from being attributed merely to renaming or data formatting. Worst-case guarantees additionally require preservation of the adversary's choices and information.

$$\sup_{\pi \in \Pi_1} W_1(\pi) = \sup_{\pi' \in \Pi_2} W_2(\pi') \tag{P7}$$

On each side, Π is the policy set available in the corresponding representation, and W is the same criterion with its meaning transported through the bijection. The proof is immediate: every policy in the first representation has a counterpart producing the corresponding history with the same cost, information, and outcome, and hence the same value. The second side is therefore at least as large as the first. Applying the inverse correspondence gives the reverse inequality and establishes equality. A probabilistic model must also preserve outcome distributions, not merely the set of possible outcomes.

The proposition explains how representation can nevertheless be useful. If structured representation reduces the required number of queries, facilitates discovery of a compatibility relation, or enables a policy previously infeasible within the budget, a neutrality assumption has changed. Introducing an acceptable contract that was previously unavailable changes the opportunity set. Turning an ambiguous description into a definite commitment changes information or enforceability. The relevant channel should be measured rather than simply describing the outcome as value created by abstraction.

In a synthetic example, one interface presents every option in a long document, while another presents the same options in queryable fields. An ideal, unconstrained solver may reach the same decision in both cases, while an agent with limited context fails in the first and succeeds in the second. The experimental outcome concerns information-processing costs and the policies feasible under the constraint, rather than improvement in an unconstrained economic optimum. This distinction supports fair comparisons of AI agents.

A neutrality experiment holds products, participants, offer instances, resources, and objectives fixed, changes only the representation, and includes the cost of converting to it. A separate experiment may expand the contract space while holding representation fixed. This design distinguishes interface effects from the effects of new opportunities and the algorithm. Differences in outcomes can then be attributed to an identified variable, with confidence bounds and the costs of data integration and validation reported.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

A finite allocation model for offer-space compression

We now introduce a finite model in which the effect of compression can be proved. There are finitely many buyers, each with a finite set of complete available bundles. A solution selects one bundle or an opt-out choice for each buyer. A bundle may contain several products, services, and conditions, but its evaluation is given at this stage. Complete bundles avoid assuming additive values for internal components that may be complementary or competing. Aggregation across buyers, however, assumes a declared numerical criterion with additive contributions.

$$\mathrm{OPT}_{\mathcal{E}} = \max_{e_i \in \mathcal{E}_i} \left\{ \sum_{i=1}^{n} w_i(e_i) : \sum_{i=1}^{n} a_i(e_i) \le B \right\} \tag{P8}$$

For each buyer, the set E contains acceptable bundles and includes a null option that consumes no resources and adds no value. The vector a describes nonnegative resource consumption, while B gives shared capacities. The quantity w scores a bundle under the experiment's criterion, which might be utility measured in monetary terms, profit, or a weighted service score. The theory does not assume that interpersonal aggregation of utilities is always justified; it makes aggregation an explicit assumption of the following result.

Eligibility, authority, and individual conditions enter the construction of available sets. Shared constraints, such as inventory, budgets, and unique vouchers, enter the resource vector whenever additive representation is appropriate. Any other shared constraints must be added and preserved by each substitution used in the proof. Omitting a constraint that is difficult to encode does not make it unimportant: it changes the problem and may make the result inapplicable to the intended system.

A small synthetic example has two buyers and a subsidy budget of six units. The first buyer has a bundle scoring nine and consuming four units, and another scoring seven and consuming two. The second has a bundle scoring eight and consuming four units, and another scoring six and consuming two. Selecting both higher-scoring bundles is infeasible because it consumes eight units. Selecting the first buyer's higher-scoring bundle and the second buyer's lower-scoring bundle yields fifteen points using six units, improving on selecting both lower-scoring bundles.

A higher score does not imply that participants will reveal preferences truthfully or accept the required payment. The model treats its data as specified and addresses constrained selection. Data elicitation, incentives, and voluntary participation remain separate questions. Nor do multiple instances of a representative bundle eliminate repeated consumption: each allocation to a separate buyer is counted again. The algorithm does not treat a stored representative as a free inventory unit that can be used without limit.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

A resource-preserving offer-cover proposition

Definition and proposition: representative substitution and allocation loss

For each buyer, choose a representative subset of the original bundles and a map assigning every bundle to a representative for that same buyer. Substitution must preserve eligibility and all constraints not represented by resources, must not increase any resource consumption, and may reduce the bundle's score by no more than a known bound. This bound may vary across buyers. We do not yet assume that the representative set is small or easy to construct; these are additional questions that the definition of a cover does not settle.

$$a_i(\phi_i(e)) \le a_i(e), \qquad w_i(\phi_i(e)) \ge w_i(e) - \varepsilon_i \tag{P9}$$

The map φ performs substitution, and ε is an allowable nonnegative loss measured in the same units as the score. We explicitly require that simultaneous substitution of all selected bundles preserve every shared constraint outside the resource vector; the basic model satisfies this condition when no such constraints exist. Under these assumptions, the optimum using representatives is no greater than the original optimum, but its loss is bounded by the sum of individual substitution bounds.

$$\mathrm{OPT}_{\mathcal{E}} - \sum_{i=1}^{n} \varepsilon_i \le \mathrm{OPT}_{\mathcal{R}} \le \mathrm{OPT}_{\mathcal{E}} \tag{P10}$$

Proof. Take an optimal allocation for the original problem and replace each buyer's bundle with its representative. Aggregate resource consumption does not increase in any coordinate, so it remains within capacity. The assumptions preserve all other constraints, making the resulting allocation feasible for the compressed problem. Summing the score inequalities shows that its score is at least the original optimum minus the sum of errors. The compressed optimum is no worse than this feasible allocation, establishing the lower bound. The upper bound follows because representatives are original bundles.

In a separate synthetic example, the first buyer has a bundle scoring nine and consuming four units, with an available representative scoring eight and a half and consuming four. The second has a bundle scoring six and consuming two units, with an available representative scoring five and three quarters and consuming two. Resource allocation remains feasible after substitution, and the score decreases by three quarters. This example does not establish an algorithm capable of finding suitable representatives in every market. The conclusion also changes if even a small consumption increase is permitted at a saturated budget.

The proposition is an elementary, familiar form of substitution-based approximation argument, not a claim of an independent discovery in algorithm theory. Its role here is to make the conditions for compressing Promoi inspectable: what the representative preserves, for whom it is suitable, and where error is counted. Potential contributions lie in constructing small covers for realistic offer classes or proving that a particular economic condition precludes a cover of the desired size.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

Combining compression, estimation, and optimization errors

Offer scores are not always known exactly. An agent may estimate delivery time, the value of a feature, or the cost of an incompletely specified service. We therefore distinguish the true score defined in the experimental model from the estimated score used by the solver. Assume a uniform error bound for each buyer across all its representatives, while eligibility and resource consumption remain correctly known. This distinction is crucial: score error affects choice quality, whereas resource error may make an apparently feasible decision impossible to execute.

$$|\widehat{w}_i(r) - w_i(r)| \le \eta_i \qquad (r \in \mathcal{R}_i) \tag{P11}$$

The quantity η denotes a nonnegative error in score units. Suppose the solver returns a feasible allocation whose estimated score is at most δ below the best estimated score. The following bound combines three distinct losses: reducing the offer space, using estimates rather than the specified values, and solving the compressed problem only approximately. It is a deterministic bound conditional on all its assumptions, not a promise about the performance of any language model involved in estimation.

$$W(\widehat{x}) \ge \mathrm{OPT}_{\mathcal{E}} - \sum_i \varepsilon_i - 2\sum_i \eta_i - \delta \tag{P12}$$

To prove the bound, compare the output with the true optimal allocation within the representative set. Estimation may overstate the output's score by at most the sum of η, so this amount is subtracted when moving from estimated to true score. Apply the solver's guarantee of losing at most δ relative to the estimated optimum. Estimation may also understate the true optimum's score by another sum of η, which must be subtracted a second time. Finally, apply the preceding offer-cover bound to obtain the stated inequality.

In a synthetic example with ten buyers, a compression bound of half a point per buyer, an estimation bound of one tenth of a point per buyer, and a solver error of one point, the guaranteed loss is at most eight points: five from compression, two from estimation, and one from optimization. Actual loss may be substantially smaller, but the example does not establish that. Separating these bounds identifies where further expenditure may help: better estimation does not remove large compression loss, and a better solver does not repair missing suitable representatives.

When estimation bounds are derived from random data, the probability that they hold jointly over the offers the algorithm may select must be specified. A guarantee for one randomly selected offer is insufficient after adaptive selection from a large list. If estimated scores also determine the cover, its construction error must be analysed rather than assuming that its true-score conditions hold. Subject to these qualifications, the bound supports experimental design and scrutiny of the decision pipeline without turning prediction-error bounds into unproved execution guarantees.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

A small cover with one resource and bounded scores

A simple case admits a direct small-cover construction. Fix one buyer and a class of offers agreeing in all non-resource constraints. Assume one scalar resource and scores between zero and an upper bound V. The offer set is finite and known, and the error tolerance ε is positive. Partition the score axis into buckets of width ε, retaining the least resource-consuming offer in each nonempty bucket. Also retain the opt-out choice, whether it already appears among the bucket representatives or is added separately.

$$k(o) = \left\lfloor \frac{w(o)}{\varepsilon} \right\rfloor \tag{P13}$$

An offer and its representative in the same bucket differ in score by less than the bucket width, or by at most ε depending on endpoint conventions. Because the representative minimizes consumption, it requires no more of the resource than the original offer. The cover conditions therefore hold, provided the class genuinely agrees in eligibility and all other constraints. If several classes cannot safely be merged, apply the construction separately within each and add the resulting cover sizes.

$$|\mathcal{R}| \le \left\lfloor \frac{V}{\varepsilon} \right\rfloor + 2 \tag{P14}$$

The bound counts the possible bucket indices and the additional opt-out choice. It may be loose if an existing representative already includes opting out. The count does not depend directly on the number of original offers, but it depends on the score range and compression accuracy. Construction time also depends on accessing data and computing scores and resource consumption: a small output is not costless to produce. If preferences are private and have not been elicited, knowing offer prices alone does not supply the scores needed for this construction.

In a synthetic example, three bundles lie in the score bucket from ten to less than twelve. Their scores are ten, eleven and a half, and eleven and nine tenths; their resource consumptions are four, six, and five units, respectively. The representative is the bundle scoring ten and consuming four units. Replacing the highest-scoring bundle loses less than two points and saves one resource unit. This representative should not be described as best for every buyer; it is appropriate only for the buyer, criterion, and class for which the cover was constructed.

This case connects Promoi compression to the literature on approximating trade-off frontiers, including Papadimitriou and Yannakakis (2000). Our formulation, however, permits no increase at all in resource use and allows additive error only in score. Multi-objective approximation may allow controlled deterioration in several coordinates under a different definition. This distinction explains why a small-size guarantee for a multi-objective problem cannot be transferred to strict budget constraints without examining the direction of approximation and available capacity slack.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

A compression limit under exact two-resource dominance

Counterexample: all offers may have to be retained

The one-resource construction does not automatically extend to two resources. Construct a synthetic family of N offers for one buyer, each scoring one, together with an opt-out choice scoring zero. All share eligibility, authority, and other constraints. Offer j consumes an increasing fraction of the first resource and a decreasing fraction of the second, with the two fractions summing to one. We require a cover using representatives from existing offers and preserving each resource without any increase.

$$a(o_j) = \left(\frac{j}{N+1}, 1 - \frac{j}{N+1}\right), \qquad w(o_j) = 1 \tag{P15}$$

Indices range from one to N. If offer k consumes no more of either resource than offer j, comparison in the first coordinate requires k to be no greater than j. Comparison in the second requires k to be no smaller than j. The indices must therefore be equal. No offer is dominated by another in consumption of the two resources: the family is an antichain under the partial order of resource dominance.

$$a(o_k) \le a(o_j) \implies k \le j \text{ and } k \ge j \implies k = j \tag{P16}$$

If the allowed score error is less than one, the opt-out choice cannot represent any offer because it loses the entire score. Since no other offer can represent a given offer while maintaining exact dominance in both resources, the cover must retain all N offers. Thus, under these conditions, the number of representatives cannot be bounded by score range and error alone. The number N can grow arbitrarily even though every offer has the same score.

A synthetic three-offer instance has resource vectors consisting of one quarter and three quarters, one half and one half, and three quarters and one quarter. With a budget equal to one offer's resource vector, that offer is feasible and each other offer violates one constraint. Deleting a particular offer can therefore reduce the best available score from one to zero under a suitable budget. The example illustrates the importance of the shape of the trade-off frontier, rather than the number of textual attributes in offer descriptions.

The result's scope is as important as its conclusion. It concerns covers drawn from original offers, exact dominance, and a guarantee covering every offer in the family. It does not establish impossibility for all algorithms, implicit representations, randomized methods, or compression tailored to a single known budget. New contracts or mixtures of divisible resources may allow a different structure, but change the choice space. The example identifies the failure of one compression route and motivates finding additional structures under which positive results become possible.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

Resource slack, complementarity, and private preferences

A designer may try to circumvent the two-resource example by permitting a small consumption increase for each representative. Local tolerance, however, does not ensure aggregate allocation feasibility. If each buyer's increase is bounded componentwise by a nonnegative vector γ, then n substitutions may increase total consumption by nγ. Compression respects capacities if the initial allocation leaves this slack in every coordinate and all remaining constraints continue to be preserved.

$$a_i(\phi_i(e_i)) \le a_i(e_i) + \gamma \tag{P17}$$

$$\sum_i a_i(e_i) \le B - n\gamma \implies \sum_i a_i(\phi_i(e_i)) \le B \tag{P18}$$

The comparison is then against the best allocation under the tightened budget, unless a separate bound is established on the loss caused by tightening. A very high-scoring bundle may consume the entire capacity; any positive reduction in capacity excludes it and leaves only the opt-out choice. A small resource increase therefore does not imply a small value loss relative to the original problem. Stronger conclusions require assumptions such as divisibility, nearby alternatives, or reserve resources, with the contribution of each assumption explicitly stated.

Another limitation arises when compressing bundle components instead of complete bundles. A battery and a particular charger may each be nearly useless alone yet highly valuable together. Preserving each component's standalone score does not guarantee preservation of the pair's value. The pair can be included as a complete bundle in the model, but the number of bundles may grow rapidly. Logical bidding languages already examine concise expressions of complementarity and substitutability, as in Boutilier and Hoos (2001). A contribution here therefore requires an identified compressible class or a new guarantee with clear assumptions.

Nor can one representative be guaranteed to preserve quality for all buyers with unrestricted preferences. One buyer may prefer speed and another repairability at the same price. Covers may be buyer-specific, defined for declared classes of valuation functions, or constructed from attribute vectors with suitable guarantees. Each choice requires different information and may reveal private data. Elicitation costs, error risks, and the consequences of data sharing should be included in the practical evaluation of a compression algorithm.

Finally, successful compression does not establish truthful reporting. A seller may choose a description that makes its offer the representative of an attractive bucket, split an offer into copies to increase visibility, or conceal a constraint affecting execution. This study therefore separates correct optimization on assumed inputs from the incentives generating those inputs. Verification rules and inspectable evidence may reduce manipulation, while strategic truthfulness requires a separate mechanism analysis.

CHAPTER 2 OFFER SPACES AND ABSTRACT PRICING

From Promoi structure to a testable research contribution

The preceding results define a research programme for offer spaces beyond multiplying discount formats. The scientific problem begins with an offer family whose execution semantics are declared, specifies the available information and its acquisition cost, and identifies constraints that compression or composition must preserve. The questions are how far search can be reduced, what decision loss results, and which properties enable or obstruct an effective solution. This connects the elementary algebra of offers to actual systems with limited agents and resources.

A first proposed test fixes the assortment, eligibility, resources, and bundle scores, then compares the complete space, a resource-preserving cover, and a summary based solely on linguistic similarity. Measure representation size, query count, construction cost, allocation quality, and constraint-violation rates. Use an exact solver on small instances to measure loss, separating the time required to compute the benchmark from the tested agent's budget. This is not a report of an executed experiment: it is a research design that can be replicated once its data and software are published.

Instances should include families favourable to compression and families that challenge it. The former include one resource and few eligibility classes; the latter include multi-resource trade-off frontiers, complementary bundles, and unique rights. Selecting only cases favourable to the proposed algorithm makes the result a description of sample selection. Adverse cases help explain performance limits and reveal whether the algorithm stops safely when it cannot preserve executability.

A second test separates expansion of the contract space from improvement in its representation. In one condition, opportunities remain unchanged and only information organization varies; in the other, new contracts are added while search tools remain fixed. A third factor may vary the time or query budget. This design estimates interactions: expanding the space may not benefit a bounded agent if search costs rise by more than the value of the opportunities, while effective compression may restore the benefit of expansion within the tested domain.

Promoi has specific functional connections to the wider theory. Environment E supplies resources, transitions, and timing; elements A supply exchanging and authorized entities; and accessibility ρ describes which offer information reaches each element. The kup records provide evidence of offer adoption, reservation and consumption of rights, and exchange outcomes. Agent capability appears in its choices, query sequence, and executable commitments under a declared budget. These concepts cannot replace one another, because an offer's quality, clarity, validity, and accessibility are different properties.

An assessable contribution may be an algorithm with error and cost bounds, a characterization of a class admitting small covers, an impossibility result that limits expectations, or reproducible empirical evidence of a specified improvement. The proposed terminology remains a way of organizing these questions. Scientific acceptance of the field depends on explicit assumptions, rigorous comparisons, and an independent researcher's ability to reconstruct the results and test them in another offer space.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

The economic specialization of value space

Abstract pricing enters a value space when its elements can establish describable commitments, those commitments are constrained by resources, and a mechanism determines the consideration exchanged between parties. The general theory does not require monetary consideration. Where explicitly stated, however, the economic analysis in this chapter assumes a common unit of account and the possibility of expressing specified benefits as local monetary equivalents. This modelling assumption enables derivation; it is neither a universal definition of value nor a reduction of every relationship to money.

Within this specialization, we define abstract pricing as the design of consideration and commitment terms as composable variables, holding product identity or the description of the intended transformation fixed whenever the comparison requires it. The monetary price becomes one component of an offer that may include delivery, warranties, service priority, refund conditions, commitment duration, and resources or information supplied by the buyer. Promoi associates these offers with a product assortment. Separate descriptions do not make an offer economically independent of product cost or the seller's ability to fulfil it.

There is a direct precedent in multi-attribute auction research: Bichler and Kalagnanam studied configurable offers, configuration and discount rules, and allocation problems involving multiple suppliers. We therefore treat programmable offer representation as an existing foundation. The contribution sought here concerns information, resource, and execution constraints within the space (Bichler and Kalagnanam, 2005).

Our first analytical distinction separates three operations. Representation changes how opportunities are described. Design changes the set of opportunities actually available. A market mechanism changes who chooses, who pays, and who bears risk. An accessible interface can reduce search cost without creating a new contract; a new contract can introduce a cheaper delivery option; and a different mechanism can redistribute a fixed surplus. These effects should not be conflated and the entire improvement attributed to price abstraction.

A scientific comparison must hold constant the parties, resources, environmental state, and definition of completion. If feasible histories, information, costs, and preferences remain identical, relabelling offers does not change the optimum for an agent with unlimited computation. A bounded agent may nevertheless benefit from a shorter representation or faster verification. This yields a measurable question: how much improvement comes from expanding exchange possibilities, and how much from making existing possibilities usable?

This chapter uses explicit utilities and profits to trace these distinctions. It treats acceptance, publication, and execution as distinct events: a buyer's proposal does not bind the seller; the seller's acceptance does not establish completed delivery; and a kup record of payment does not by itself establish service quality. The economic model thus remains connected to the structure of events and evidence defining the value space.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Consideration vectors and the buyer's total cost

A posted price provides only partial information when payment timing, delivery quality, or ancillary obligations vary. We represent consideration as a vector of dated monetary flows and non-monetary terms. Its coordinates cannot simply be added: units of time differ from units of money, and information disclosure is not inherently a monetary quantity. Reduction to a scalar requires a particular party's preferences and explicit valuation rules. Some offers may remain only partially ordered when that party declines the relevant trade-offs.

$$\mathbf{p}(o) = (p_0, \ldots, p_T, \tau, q, r, d) \tag{E1}$$

The offer denotes the complete contractual arrangement. Payments are the amounts due in specified periods; the remaining coordinates refer to delivery time, quality, rights, and disclosure data. A payment may be negative if it is a guaranteed flow to the buyer. A reward conditional on uncertain eligibility or subsequent performance is instead recorded as a state-contingent variable, rather than automatically treated as a present cash discount.

In a simplified model without uncertainty, we define the buyer's equivalent total cost as follows. Discount factors express the buyer's own valuation of payments over time. The other terms are equivalents of search, verification, delay, and privacy costs. Addition is meaningful here only because the model assumes that these quantities can be expressed in the same unit of account.

$$C_i(o) = \sum_{t=0}^{T} \delta_i^t p_t + c_i^{\text{search}} + c_i^{\text{verify}} + c_i^{\text{delay}} + c_i^{\text{privacy}} \tag{E2}$$

This definition does not specify a universal risk cost. With probabilistic outcomes, analysis begins with utility over outcomes and an explicit probabilistic belief, then computes a certainty equivalent if one exists. Adding an expected loss and an arbitrary risk premium may count the same effect twice. Liquidity constraints also differ from time preference: an individual may prefer immediate payment in the calculation yet be unable to make it, excluding the offer from the feasible set before comparison.

In a synthetic example, an immediate payment of 90 competes with two payments of 50, one now and one in the next period. With a discount factor of 0.9, the second arrangement has a present value of 95 before fees or other effects. Nevertheless, it may be the only feasible arrangement for a buyer currently holding no more than 60. Ability to pay is distinct from having a lower calculated cost.

An interface should retain the details underlying its ranking rather than display a single opaque score. A purchasing agent can explain that it accepted a specified delay for a particular saving, or rejected a cheaper offer because its disclosure condition was unauthorized. A subsequent kup records evidence of the actual commitment and outcome. The user's preferences and their conversion into costs remain part of the valuation context, which may differ between tasks.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Demand, outside options, and willingness-to-pay bounds

Demand within Promoi is a choice among complete arrangements, including abstention, waiting, or obtaining an alternative outside the platform. Let the buyer's outside-option utility be the reference utility. Positive product value does not imply a desire to purchase: another opportunity may be better, or commitment costs may exceed the gain. Constructing the choice set must account for what the buyer knows and what the buyer authorizes the agent to use.

The local model assumes quasi-linear utility, certain knowledge of attributes, and a common scale for monetary equivalents. Gross value is the value of the selected configuration before its price; ancillary cost comprises delay, search, and verification. The buyer then chooses a utility-maximizing feasible offer, provided its utility is no lower than that of the outside option.

$$u_i(o) = v_i(z(o)) - p(o) - c_i(o), \qquad u_i(o) \geq u_i^0 \tag{E3}$$

The feasible offer set incorporates budget, eligibility, and authorization constraints, not desire alone. If the outside option is a known purchase from another seller, comparison includes the cost of accessing it. The value of waiting requires a model of future opportunities and cannot simply be assumed zero in a study of dynamic pricing. Correct abstention may demonstrate agent competence more clearly than a high purchase rate.

Holding other attributes fixed, the participation condition yields an upper bound on an acceptable price. In this model, an increase in ancillary cost lowers that bound by the same amount, while an improved outside option reduces the buyer's willingness to accept the current offer. This is an algebraic consequence of quasi-linearity, not a general law of demand elasticity.

$$\overline{p}_i(z) = v_i(z) - c_i(z) - u_i^0 \tag{E4}$$

For a synthetic illustration, assume a value of 100, a delay cost of 4, and outside-option utility of 12. The maximum acceptable price is 84. An offer priced at 82 leaves utility of 14, an improvement of only 2 over the outside option. Describing the saving as 18 on the basis of product value alone confuses gross benefit with the gain from switching away from the option actually available.

Moving to market demand requires a distribution of types, preferences, and constraints. One user's response cannot identify a demand curve, and publishing a new offer may change demand for others through substitution or complementarity. Nor does one observed choice reveal an entire utility function: multiple preferences can produce the same decision. Experiments must therefore distinguish synthetic preferences known to the simulator, user preferences estimated under uncertainty, and authorization specifying what an agent may infer or request.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Seller profit and the cost of offer adoption

A seller faces two distinct decisions: whether to adopt an offer template and whether to accept a particular exchange under it. Each additional order may be profitable after technical integration, while expected orders remain insufficient to recover integration's fixed cost. A template may also be attractive overall yet unsuitable when inventory or labour capacity is congested. We therefore distinguish adoption cost, marginal execution cost, and opportunity cost. Price and demand are assumed positive wherever elasticity and margin ratios are calculated.

In a single-order model, collected consideration is actual revenue after identifying the recipient of each flow. Costs include production, delivery, verification, expected servicing, and scarce resources used. A fee paid to a party within the market boundary is a transfer to that party, but remains a deduction from seller profit. The alternative value of resources depends on their opportunity for use in another order, rather than an arbitrary amount appended to every transaction.

$$\Pi_s(o) = p_s(o) - c_s^{\mathrm{prod}}(o) - c_s^{\mathrm{fulfil}}(o) - c_s^{\mathrm{verify}}(o) - f_s(o) \tag{E5}$$

This expression gives local accounting profit under the selected cost boundary. If rejecting the transaction permits another profit, utility is compared with that alternative, or opportunity cost is included once in economic cost. Including it in both terms double-counts the deduction. The optimal response may be to change quantity or timing rather than lower price, because some resource constraints can be relaxed over time and others cannot.

For a synthetic adoption analysis, suppose setup has a fixed cost of 200 and each new order provides a known net margin increase of 4. More than 50 additional orders are needed to exceed setup cost within the measurement horizon. Additional is essential: if those orders would already have occurred on more profitable terms, the offer may divert existing sales rather than expand them. A controlled comparison should estimate this diversion and its effects on other orders.

A common template may let several sellers share the cost of understanding it and developing its tools, while adding data-standardization and exception-handling costs. Subscribing to a template also differs from authorizing its publisher to change future terms. The adopted version, quantity and budget limits, and acceptance period must be specified. These are not administrative details outside the theory: they determine which exchanges a seller can commit to.

For the special case of a seller facing differentiable demand that decreases with price and constant marginal cost c, differentiating profit, defined as unit margin times demand, gives the following condition for an interior optimum. Elasticity measures the relative responsiveness of demand to price. This condition is necessary here; by itself, it neither guarantees maximization nor applies automatically to multi-attribute offers.

$$\frac{p - c}{p} = \frac{1}{\varepsilon(p)}, \qquad \varepsilon(p) = -\frac{pD'(p)}{D(p)} \tag{E6}$$

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Synthetic example: creating and redistributing surplus

We construct a numerical example to separate offer design from price reduction. Assume one buyer and one seller, quasi-linear utility, no fees or external parties, and certain delivery. Product value before the timing effect is 100 units, and production costs 60. Fast delivery costs 15; slow delivery costs 5 but imposes a utility loss of 4 on the buyer. All numbers are synthetic and represent neither customer nor market data.

Example E: differences generated by terms and prices, in a hypothetical monetary unit.

Offer	Buyer utility	Seller profit	Joint surplus
Fast; price 90	10	15	25
Slow; price 82	14	17	31
Fast; price 82	18	7	25

Under the first offer, the buyer receives 10 and the seller earns 15 after production and shipping. Reducing the price to 82 while retaining fast shipping increases buyer utility to 18 and reduces seller profit to 7; their sum remains 25. The cash saving transfers existing surplus from seller to buyer. It may be desirable to the buyer, but it does not increase this transaction's joint surplus.

Switching to slow delivery instead saves 10 in real cost and costs the buyer 4 in utility, producing a net increase of 6. At a price of 82, buyer utility improves by 4 and seller profit by 2 relative to the original fast-delivery offer. Both parties gain because the commitment design permits distribution of a real gain, not because the price has been relabelled.

We can identify the price interval that makes both parties weakly better off than under the first offer. The buyer requires a price no higher than 86; the seller requires one no lower than 80. Every price between these bounds distributes the additional surplus between them, with strict improvements in the interval's interior. Efficiency alone does not select a point: bargaining, outside options, and the offer-proposal procedure influence distribution.

If implementing the slow arrangement introduces a further real cost of 7, such as additional verification or specialized handling, its surplus falls to 24, one below the benchmark. It is then impossible to improve both parties' positions using internal monetary transfers alone. If those 7 units were instead a fee to a third party included within the welfare boundary, its costs and benefits would need analysis rather than treating the entire fee as a resource loss. The example explains why accounting boundaries and all operating costs of abstract pricing must be disclosed.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Welfare, transfers, and the limits of utility aggregation

The preceding example demonstrates both the capabilities and limits of economic accounting. With quasi-linear utility, transferable money, and all payment recipients included, internal transfers cancel on aggregation. What remains is the valuation of outcomes minus resource consumption and real costs. This accounting cancellation neither establishes that monetary distribution is unimportant nor licenses aggregation of heterogeneous utility scales.

$$W(x) = \sum_i v_i(x_i) - C_{\text{real}}(x) - D_{\text{external}}(x) \tag{E7}$$

The allocation denotes the set of executed exchanges. Value represents gross outcome valuations before transfers and costs included in the other terms, expressed in the assumed common unit. Real cost denotes consumed resources. The final term represents external effects measured in a compatible unit, where these can be estimated and fall within the study's scope. If privacy or safety effects cannot be measured in this way, they should appear as separate constraints or indicators rather than be forced into a monetary sum.

A price reduction may improve distribution without increasing total surplus. Additional liquidity may also enable a previously infeasible transaction, changing the feasible set rather than merely redistributing surplus. Cancellation of price in the sum for an individual transaction therefore does not imply that price is irrelevant to market efficiency: it affects participation, allocation, search, and investment. The local result holds execution, parties, and resources fixed.

The Pareto criterion provides an alternative that does not require utility aggregation: an allocation improves if no relevant party is worse off and at least one is better off. However, more parties make this condition harder to satisfy, and it may leave most alternatives unordered. The theory must not confuse incomparability with indifference. Reporting a frontier of alternatives, with an account of who gains and who bears costs, is preferable.

In a platform study, a service fee may appear as a loss to buyer and seller when the platform is excluded, but as platform revenue when it is included. Operating resources nevertheless remain a real cost, and net platform profit forms part of the distribution. The same distinction applies to third-party-funded discounts: the funder, what it receives in return, and alternative uses of its resources must be identified. Moving an obligation to a less visible account does not create a general increase in value.

We propose reporting three parallel levels in evaluations: separate outcomes for each party, an economic aggregate under explicit assumptions, and constraints whose violation cannot be offset by a higher score. This helps prevent an agent from obtaining a high score by exhausting a scarce resource or accepting an unauthorized commitment. It also makes claims about abstract pricing open to scrutiny: did the mechanism improve execution, distribution, choice quality, or several dimensions at once?

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Offer menus, self-selection, and privacy

A seller can publish a menu of offers from which buyers select what suits them. Self-selection distinguishes preferences without requiring the seller to know each preference beforehand. A larger menu does not, however, guarantee buyer benefit: the seller may design terms to extract a higher willingness to pay or make the basic option less attractive. We therefore distinguish expanded effective choice from merely adding buttons and labels.

In a synthetic model, two types value the same product at 100, but waiting for slow delivery costs the patient type 2 and the urgent type 20. The seller offers fast delivery for 90 and slow delivery for 75. The patient type obtains utility of 23 from slow delivery and 10 from fast delivery, choosing slow. The urgent type obtains 5 from slow delivery and 10 from fast, choosing fast. This result assumes equal product value and no superior outside option.

For the menu to induce the intended self-selection, no type should prefer the offer assigned to another type. Let a type describe preferences and constraints, and let its intended offer be a function of that type. Two separate conditions then express no gain from choosing another type label and willingness to participate relative to the outside option.

$$u_i(o(\theta),\theta) \geq u_i(o(\widehat{\theta}),\theta), \qquad u_i(o(\theta),\theta) \geq u_i^{0}(\theta) \tag{E8}$$

These are choice conditions for a particular menu. They do not establish truthful information provision in every subsequent negotiation or by every agent. Eligibility may depend on a separate right requiring proof, and duplicate accounts may grant access to offers not intended by the designer. The study's set of deviations must encompass the choices actually available; restricting it to changing one answer conceals strategic behaviour enabled by the interface.

An offer may be provided in return for data sharing, but privacy is a resource or constraint specified by the user. An agent cannot assume that every user will sell any information at a sufficiently high price. When disclosure is prohibited, incompatible offers are infeasible. When conditionally permitted, disclosure cost, purpose, and duration enter the evaluation. An estimate that the discount is profitable cannot substitute for missing authorization.

A menu's effect is measured against a usable available alternative, accounting for search and comprehension and separating preference differences from income or bargaining-power differences. A structured menu may assist buyers' agents in comparison, while excessive disclosure may weaken their future bargaining position. Design should therefore examine the information necessary to execute an offer separately from information commercially useful to the seller but unnecessary for fulfilment.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Bundles and economic complementarity between offers

An offer space distinguishes syntactically valid composition, economically beneficial composition, and executable composition. A data format may permit two offers to be combined even though one prohibits using the other, or both consume the same unique right. A combination may also be feasible but undesirable to the buyer. The composition function must therefore check eligibility, resources, and the joint outcome; the decision then requires a separate economic valuation.

Bundle value need not equal the sum of component values. In a synthetic analytics example, dataset access has limited value without compatible computing capacity, whereas their combination completes the task. For a given base set, we define the complementarity effect as the value of adding both components minus their separate additions, with the base value corrected for. Positive and negative signs distinguish local complementarity from local substitution.

$$\Delta_{ab} v(S) = v(S \cup \{a, b\}) - v(S \cup \{a\}) - v(S \cup \{b\}) + v(S) \tag{E9}$$

This is a local measure of a particular party's preferences; it does not determine demand in every context. Two subscriptions may be substitutes for one use but complements when backup service is needed. Bundle value also depends on timing: computing available today and data arriving after the quota expires do not complete the same task. Promoi must preserve temporal relationships, not merely service names.

Suppose, synthetically, that data alone is worth 2, computing alone 3, and their compatible combination 12. If their prices are 4 and 5, the buyer rejects either service individually but accepts the bundle with utility of 3. An algorithm filtering offers by standalone attractiveness would discard a useful solution. If the computing service is technically incompatible, however, the value of 12 cannot be used. This illustrates why valuation must be linked to an actual compatibility certificate.

Partial commitment across sellers creates another problem: the buyer may pay for data and subsequently fail to reserve computing. A design may require completion of both reservations, provide an agreed cancellation option, or use an intermediary that bears the risk. Every remedy has a cost and execution limits; labelling a bundle atomic does not make it so. A kup record distinguishes completion of each exchange from completion of the composite task, because failure of the latter does not necessarily erase partial exchanges.

Algorithmic studies involving complementarity must state the assumptions justifying additive evaluation. If each buyer is entitled to one bundle and its complete value has been calculated, buyer weights may be added within the stated welfare framework. Adding scores for individual offers and inferring bundle quality may instead fail regardless of each score's accuracy. This distinction is particularly important when compressing Promoi or selecting a limited number of offers for a language-model agent.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

The platform and participation on both sides of the market

An offer platform needs participation by buyers and sellers, and the value of participation on either side may depend on the other side's presence. Rochet and Tirole showed that the structure of price allocation between the two sides matters in platform markets, and that attracting both sides is part of business-model design (Rochet and Tirole, 2003). We take this as a starting point, without assuming that every additional participant creates positive benefits or that expansion alone yields a stable equilibrium.

In a proposed preliminary model, two participation rates describe buyer and seller adoption. Two functions determine each side's response to the value of the other side's offers, platform fees, and integration costs. A mutually consistent state is one in which each side's participation matches its response to the other side. This defines a fixed point; it does not automatically predict user behaviour.

$$n_b = D_b(n_s; f_b, \rho), \qquad n_s = D_s(n_b; f_s, c^{\text{adopt}}) \tag{E10}$$

With continuity and response functions mapping the closed participation domain into itself, fixed-point existence can be established in an appropriate formulation. Uniqueness requires a stronger restriction. If the two responses have bounded Lipschitz constants whose product is below one, the composite response is a contraction, implying uniqueness and convergence for ideal sequential updating. This is a conditional mathematical application, not a guarantee for a market with deadlines, strategic effects, and delayed information.

Multiple participation states may arise when each side requires a minimum level of activity on the other, making platform launch a coordination problem. This differs from an absence of economic value: joint participation may be profitable for everyone, yet individual switching costs prevent it. Simultaneous commitments or limited compatible groups could be tested, with their operating costs accounted for rather than assuming that temporary subsidies will finance themselves in the future.

Offer quality matters more than the raw count. A thousand offers inapplicable to the user's country or task are not equivalent to two executable offers. More sellers may increase the chance of finding a suitable alternative while also increasing verification cost and noise. Effective participation can therefore be defined by the number of compatible opportunities or their distribution across tasks, alongside separate measures of eligibility, accessibility, and reliability.

For a proposed application such as Offerato, this chapter assumes neither an existing customer network nor proven network effects. An appropriate test begins with a group of sellers and a defined assortment, then measures whether a shared template reduces joining costs and creates additional buyer opportunities. Fees, integration costs, new demand, and diversion from existing channels should be compared: platform profit can rise even as one side's utility falls, and the reverse is also possible.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Funding commitments and liquidity over time

An offer may be profitable at the end of the horizon yet impossible to finance midway through it. This occurs when cashback falls due before revenue arrives, several orders draw on a shared budget, or expected payments are not final. We therefore require a constraint on every prefix of the execution history, not merely an aggregate month-end budget. The monetary constraint remains separate from inventory, energy, and time constraints.

Assume an ordered record of relevant financial events. The initial balance is funding already available; collected revenue consists of final inflows authorized for use; disbursements are payments made; and reserves are amounts allocated to obligations that cannot be spent again. We choose an explicit reserve rule appropriate to the obligation. In the deterministic model, reserves cover all outstanding amounts payable.

$$L_t = B_0 + \sum_{\tau \le t} R_\tau - \sum_{\tau \le t} D_\tau - H_t \ge 0 \tag{E11}$$

This is a simple accounting condition when reserves represent every deterministic obligation. Paying an obligation reduces reserves and increases disbursements by the same amount, leaving available headroom unchanged by reclassification alone. Adding a new obligation immediately reduces headroom. Reversible or refundable revenues cannot be included in full as final funding without an appropriate reserve or specified guarantee.

In a synthetic example, a platform holds 100 units, has accepted two cashback obligations of 60 each due that day, and expects revenue of 50 in two days. Its final budget appears to have a surplus of 30, but its obligations exceed currently available funding by 20. Counting expected revenue does not resolve the timing problem. Design can limit acceptances, defer maturity by prior agreement, or provide explicit financing; each choice affects the parties' utilities differently.

A low average redemption rate does not justify guaranteed promises above the funding limit when the model claims fulfilment in every case. A probabilistic constraint is possible if demand probabilities, dependence, and the acceptable default level are specified, but it provides a different guarantee requiring evaluation. Accepting a small failure probability is not equivalent to deterministic protection. Order independence is also an assumption that may fail when an event increases redemption for everyone.

Here kup has an economic function: it establishes, according to available evidence, what has been paid, what has been performed, and what remains due. Preventing double spending nevertheless requires reservation and coordination at decision time, because two accurate records after the event cannot prevent simultaneous commitments against the same balance. With multiple entities, reservation authority and responsibility during delays or outages must be identified, and these authorities and costs incorporated into the offer definition.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Mechanism design and the limits of joint guarantees

Selecting a high-surplus allocation is insufficient to design a market. Participants with private information may alter their reports to obtain better terms. A mechanism is defined here as a rule mapping permitted messages to an allocation, payments, and commitments. The information environment specifies what each party knows and when; the deviation set specifies what it can change, including withdrawal or delayed messages where permitted. The result discussed below assumes utility quasi-linear in money and risk neutrality.

We distinguish dominant-strategy from Bayesian incentive compatibility. The former compares truthful and alternative reports for every profile of others' reports; the latter compares expected utilities under specified beliefs. Individual rationality likewise differs before a type is known, after it is known, and after outcomes are realized. Neither compatibility nor participation guarantees should be stated without identifying the temporal and informational level established by the proof.

$$u_i(M(\theta_i, \theta_{-i}); \theta_i) \geq u_i(M(\widehat{\theta}_i, \theta_{-i}); \theta_i) \tag{E12}$$

The displayed condition is the dominant-strategy version for a direct mechanism. Utility includes payments and costs according to the participant's model, and comparison must cover the stated domain of types and reports. Strong budget balance equates aggregate incoming and outgoing payments, whereas no deficit allows the mechanism to retain a surplus. Neither condition establishes service quality or the accuracy of inputs.

The Myerson-Satterthwaite result imposes an important limit in the classical bilateral-trade model: one seller, one buyer, one indivisible good, and independent private values with positive densities on intervals whose interiors overlap, allowing both trade and no trade. Under their assumptions, ex-post efficiency cannot be combined with Bayesian incentive compatibility, interim individual rationality, and no expected deficit. This restriction concerns a specified class of environments; it does not establish impossibility for every exchange design (Myerson and Satterthwaite, 1983).

Expanding Promoi thus requires declaring the intended trade-off: approximate efficiency, a weaker incentive requirement, explicit funding, a restricted type domain, or additional verifiable information. Each choice must be tested under the new assumptions. Agent intelligence does not eliminate private-information problems; it may enable more sophisticated manipulation, while assisting computation and verification only when suitable evidence is available.

An initial experiment can hold simulator-known valuations fixed to study execution and search. This does not establish truthfulness once valuations become strategic reports. We therefore propose separating allocation-solving experiments from incentive experiments: the former compare quality and cost; the latter generate deviations and seek profitable ones. A profitable deviation refutes the guarantee within the model. Failure to find one in a limited test is not a general proof of incentive compatibility.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Bargaining and the feasible price interval

Design can determine how to divide the surplus of a known feasible exchange. Nash's bargaining analysis provides a normative formulation for selecting an outcome from a utility set with a specified disagreement point. In its symmetric form, the solution is associated with maximizing the product of the parties' gains above that point (Nash, 1950). Here we use a simplified monetary model and a weighted variant as an explicit selection rule, without automatically attributing bargaining weight to agent intelligence.

Assume a fixed product configuration, net buyer value before price of V, seller economic cost C, and outside-option utilities b and s. Buyer utility is V minus price, and seller profit is price minus C. Trade satisfying both participation conditions exists when the surplus remaining after the outside options is non-negative. Acceptable prices then lie in the following interval.

$$C + s \le p \le V - b, \qquad S = V - C - b - s \ge 0 \tag{E13}$$

Assume positive surplus and a buyer weight strictly between zero and one. The weighted bargaining rule maximizes the product of the buyer's gain raised to that weight and the seller's gain raised to the complementary weight. Substituting surplus minus buyer gain for seller gain, then differentiating the logarithm of the objective, yields a buyer gain equal to its weight times total surplus.

$$p^{*} = V - b - \alpha S, \qquad 0 < \alpha < 1 \tag{E14}$$

This derives the chosen rule on a linear monetary feasible set; it does not prove that actual negotiations converge to it. Delay costs, private information, or a risk of bargaining breakdown may change the outcome, and transfers may be indivisible. With different utility units or non-linear utilities, analysis should use the utility set itself rather than directly apply the price formula.

In a synthetic example, net value is 100, cost is 70, the buyer's outside option is 5, and the seller's is 3. Available surplus is 22, and the price interval runs from 73 to 95. With equal weights, the price is 84, giving each party 11 above its outside option. Improving the buyer's outside option changes the interval and its division even if the product and physical cost remain fixed.

The formula also clarifies the role of new offer construction: C can be reduced through less costly timing, or V increased through a better match, expanding bargaining surplus. Improving an agent's ability to extract a larger share may instead benefit its principal without enlarging surplus. A bargaining agent should therefore be assessed by its principal's gain, the completion rate of feasible exchanges, and bargaining cost, while keeping the distribution of power and information explicit rather than concealing it in one success score.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Allocation, shadow prices, and the integrality gap

Consider a set of complete bundles from which each buyer chooses at most one. A bundle's weight is its stated value under the study objective, and its resource vector specifies consumption. We add one capacity constraint per resource and one constraint per buyer. For an indivisible, single execution, choice variables are binary integers. Linear relaxations admit fractional values for computation and analysis, but these do not automatically represent executable sales. This analysis assumes finitely many buyers, bundles, and resources, with each bundle assigned to one buyer.

$$\max_{x \ge 0} \sum_{o} w_o x_o \quad ; \quad \sum_{o} a_{ro} x_o \le B_r, \qquad \sum_{o \in O_i} x_o \le 1 \tag{E15}$$

The expression presents the linear relaxation. The first index identifies a resource, the second a bundle, and each bundle subset contains the offers for one buyer. If the feasible set is non-empty and the optimum bounded, linear duality gives an upper bound equal to the relaxation value. Non-negative multipliers correspond to resource constraints, with additional multipliers for buyer constraints.

$$\min_{\lambda, \mu \ge 0} \sum_{r} B_r \lambda_r + \sum_{i} \mu_i \quad ; \quad \sum_{r} a_{ro} \lambda_r + \mu_i \ge w_o \tag{E16}$$

A resource multiplier can be interpreted as a local shadow value within this model: it describes a bound or marginal change in the programme's value under appropriate sensitivity conditions. It does not automatically yield a market price supporting an integral allocation or guaranteeing participation by all parties. The weight may represent a social objective distinct from participants' preferences, and complementarities may prevent support through separate resource prices.

A simple synthetic example has three resources, each with unit capacity, and three bundles using resource pairs: first and second, second and third, and first and third. Each bundle has weight one and belongs to a different buyer. At most one bundle can be executed because every pair of bundles shares a resource. The relaxation nevertheless chooses one half of each bundle, using each resource exactly once and attaining 1.5 compared with the integral optimum of 1.

These calculations establish a gap between relaxation and execution, not every claim concerning equilibrium existence or impossibility. Nor can the fractional solution be interpreted as a feasible lottery over integral solutions: every integral outcome has weight at most 1, so no lottery over them can have mean 1.5. Capacity feasibility in expectation must be distinguished from feasibility in every realization.

In practice, shadow prices can rank scarcity or guide proposals towards less resource-intensive offers, after which the result undergoes a full execution check. Researchers must state whether a guarantee concerns the integer problem, an upper relaxation, or an approximation to either. Incentive claims require separate analysis because an improved allocation algorithm does not by itself specify payment rules or truthful reporting.

CHAPTER 3 THE MICROECONOMICS OF ABSTRACT PRICING

Competition, dynamics, and an economic evaluation programme

Promoi changes as new offers are published, existing offers expire, and agents learn one another's responses. Solving a static allocation problem therefore does not establish market stability over time. A seller may lower its price in response to a competitor, who then changes delivery timing or eligibility, altering the comparison itself. Price becomes part of a strategy over multiple attributes, while revision costs and attention limits affect response speed.

Market state can be defined by a vector of active offers, reserved resources, and available information, with an update rule determining the next state. Before discussing convergence, one must specify whether updates are simultaneous or sequential, responses optimal or approximate, and preferences and resources fixed. Equilibrium existence does not guarantee its attainment; simultaneous responses or delayed information may produce cycles. These are model possibilities to test, not a claim of inevitable cycling in every abstract market.

Agency adds an internal incentive problem: a platform may choose a recommendation that maximizes commission rather than the buyer's offer quality. We therefore distinguish the user's objective, the agent's objective, and the service operator's objective. If the system promises improved buyer utility, commissions and ranking conditions must either be included in the experiment or held fixed. Otherwise, changed choices may reflect commercial bias rather than improved discovery or negotiation.

We propose three separate interventions for evaluating an application such as Offerato. The first changes offer representation while holding opportunities and information fixed. The second changes the feasible contract set by adding a configuration with a known cost. The third changes surplus allocation while holding feasible execution fixed. For each intervention, the experiment measures search time, real cost, each party's utility, fulfilment rate, and new demand relative to the alternative. Separating interventions permits attribution to a specified cause.

Adoption studies must track setup costs, objections, and revisions before the first transaction, because restricting analysis to users who completed a purchase excludes failures. Seller comparisons must account for inventory, quality, and service differences rather than attribute every difference to the pricing mechanism. Work can begin with synthetic data and known preferences to validate calculations, followed by limited, explicitly scoped experiments when real parties are available. This manuscript reports no unperformed field results.

The proposed criterion of economic success is not the number of published offers or the sum of nominal discounts. It is the availability of suitable, executable exchanges after search, verification, and coordination costs, with explicit outcome distribution and observance of resource and authorization constraints. This formulation permits comparison with fixed prices, simple menus, or conventional bargaining on the same tasks. Simpler alternatives may outperform in settings with few attributes and little uncertainty; identifying the limits of benefit is part of the theory's scientific value.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

The event, the content of kup, and its carrier

This study defines kup as the abstract record associated with a value exchange, independently of the medium carrying that record. Its carrier may be a participant's memory, data on a server, or a physical trace that can be interpreted within a model of the environment. This definition is deliberately broad, but requires three levels to be distinguished: the occurrence of the event, the content attributed to it, and the medium through which some of that content can be retained or read. The existence of one level does not imply that another is complete or accessible to every element.

$$\kappa_e = r(e), \qquad y_{a,c} = d_{a,c}(c) \tag{K1}$$

Here e denotes an exchange event distinguished at the chosen descriptive granularity, and κ denotes its record content. The mapping r is an analytical encoding that specifies which properties of the event the model selects; it does not posit a natural device recording everything. The symbol c denotes a carrier, while d describes how element a reads that carrier. The output y is what is actually read: it may be part of the content, a noisy description of it, or unavailable if reading fails. We therefore assume neither that the encoding is invertible nor that every abstract content has a permanent carrier.

In a synthetic example, one person hands a tool to another. Both remember the handover but disagree about its date. There is one event, shared information about some aspects, and disagreement about others. In a physical example, friction between two bodies may leave a surface change supporting limited inferences about their interaction. The trace alone does not establish every detail of the past, and different histories may fit the same observation. Here kup organises what the theory attributes to the exchange and what an observer can recover, without treating every physical change as a complete archive.

In economic applications, record content may include the participants, subject matter, commitments, and outcome, according to the task the record serves. A purchase window remains an interface for creating and displaying some of this information and linking it to actions; it is not the general definition of the record. Multiple purchasing interfaces may refer to the same event. A proposal that never becomes a purchase may also be recorded. Its fields must state that status explicitly, because recording an intention is not equivalent to recording completion.

This distinction yields a practical design requirement: every technical architecture should state what it records, what it infers, and what it leaves unknown. Record quality is relative to decision questions. A limited, accurate record may be more useful than a broad description that fails to distinguish observations from conjectures. This is a measurable operational property, not an absolute judgement about the carrier's value or nature.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

Exchange identity, duplicate copies, and partial records

Counting value exchanges requires distinguishing an event's identity from the identity of a document describing it. A system can send ten receipts for one transaction, while one document can summarise ten independent transactions. If message or row counts replace exchange counts, the measured capacity of the space can be increased merely by copying data. We therefore choose a counting policy that specifies the unit of exchange, the point at which a transaction is considered complete, and the rules for aggregating or decomposing compound events.

$$N_{\eta}(J) = \left|\{\iota_e : e \in \mathcal{E}_{\eta}(J)\}\right| \tag{K2}$$

The symbol η denotes the event individuation policy, and the set in braces contains events completed within window J under that policy. The symbol ι denotes event identity in the model. This notation does not posit a universal identifier measurable in every environment. An engineered system may have an authority that assigns identifiers, whereas deciding whether two observations concern the same event may remain an inferential problem elsewhere. Evaluation should test both false merges, where distinct transactions are combined, and false splits, where a single transaction is counted twice.

Retransmitting a request after a connection failure exposes this problem. The sender wants to recover the previous transaction's result, but the recipient may interpret the message as a new request. A clearly scoped idempotency key can link attempts to one logical operation, while retaining the result needed for reuse. Its effectiveness depends on the key's lifetime, the execution boundary that honours it, and storage and recovery limits. Merely adding a field named identifier does not automatically make every effect in a distributed system occur once.

Parties may hold different partial records: the buyer knows that approval was given, the provider knows that a resource was allocated, and the payment party knows that an amount was settled. These fragments do not conflict merely because they differ, but sufficient links are needed to establish that they concern the intended transaction. A genuine conflict arises if reliable records attribute outcomes that cannot coexist within the model. Scientific analysis begins by specifying compatibility rules and possible causes of disagreement, including delay, retries, and incorrect linkage, rather than equating more records with more events.

A partial ordering of events helps distinguish causal precedence from local timestamps, an established distinction in distributed systems (Lamport, 1978). We use it when combining kup fragments without assuming a shared absolute clock. If a system introduces a total order for execution, it should identify this as an organisational choice consistent with the required constraints. An ordering number alone does not establish a causal influence between every pair of adjacent events.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

Exchange stages and what their evidence establishes

An economic record cannot be assessed independently of the stage it describes. In everyday usage, purchase completed can mean several things, whereas an agent needs to know whether terms were proposed, accepted, resources allocated, or a service delivered. To prevent unjustified transitions between these meanings, we propose a vocabulary of transaction stages. It admits application-specific refinements and does not claim that every natural or economic exchange passes through every stage, or follows a single sequence.

A proposed semantic distinction; actual verification depends on each application's rules.

Stage	What its record may establish	What is not established automatically
Proposal	Proposed terms and their source	Agreement of the parties
Authorisation	The scope of an agent's permitted action	Availability of seller resources
Reservation	Resource allocation for a specified period	Final fulfilment
Commitment	Creation of a specified obligation	Receipt of the benefit
Fulfilment	An outcome meeting a criterion	Satisfaction with every effect
Refund	Specified compensation or repayment	Erasure of the earlier event

A simple application may combine two stages in one atomic transition, while another separates them. Contracts requiring joint execution must be specified so that the system does not record bundle success after completing only one component. A refund creates a subsequent event that changes the state; it does not rewrite history as though the original commitment never occurred. The analysis therefore preserves explicit links among original events, corrections, and compensating events, and reports net or gross outcomes according to the research question.

AP2 describes chains of authorisation linking user instructions to cart contents and payment (Parikh and Surapaneni, 2025). This existing example illustrates the importance of establishing authority and scope. In our model, proving that an external service was delivered remains a further question of observation and fulfilment criteria. A valid signature on instructions is not logically sufficient to conclude that the intended effect occurred, unless assumptions connecting those instructions to reliable execution are added.

A preprint by Jiang and colleagues analyses agent payment protocols across authorisation, settlement, and fulfilment, stressing consistency among these stages (Jiang et al., 2026; preprint). This overlap helps delimit our proposed question: what stage-specific information does an agent need to make a sound economic decision under shared resource constraints? We do not attribute the invention of stage separation or protocol auditing to kup.

In experiments, each transition should generate evidence identifying its type, subject, authorised party, and the state it describes. An agent can plan incorrectly even when every signature is valid, for example by treating a short reservation as a permanent guarantee or a plan as a completed service. Separating stages makes such failures classifiable and measurable, and helps identify the decision rule requiring correction.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

Recovery limits: attribution and causation

We define an environmental history as a sequence or structure of events permitted by its model. An element's observation is what it can extract from that history using its available means. Observations may include readable records, direct traces, and messages. This definition does not assume that the element sees the entire history; it permits identical observations in different histories. Such a situation establishes a simple logical limit on what any algorithm for recovering the past can achieve.

$$\mathrm{Obs}_a(h_1) = \mathrm{Obs}_a(h_2), \qquad g(h_1) \neq g(h_2) \tag{K3}$$

Proposition: guaranteed recovery is impossible from an indistinguishable observation

If the preceding relation holds, no function of the observation alone can return the property g correctly in both histories. The proof is immediate: identical inputs produce the same function output, while the required values differ. Allowing internal randomness does not permit probability-one correctness in both cases, since the output distribution is identical for the same input. The proposition allows a best guess under an explicitly stated prior distribution, but prevents that guess from being recast as unconditional, certain recovery.

Suppose, in a synthetic example, that a carrier records only the total amount paid. The same total could arise from one expensive unit or two cheaper units. If the property sought is the number of units, reading the total is insufficient. Another invoice or fixed knowledge of the product assortment might help, but that help comes from additional information or an explicit restriction of the model. Recovery limits depend on the question: the same record may establish the total but not the number of products or the source of demand.

Attributing an event to a channel also differs from establishing that channel's effect on its occurrence. Recording that a buyer arrived through a particular link establishes a tracking relationship if the system is reliable, but does not by itself show that the purchase would not have occurred without the link. Linking a product or information item to an exchange in kup identifies the association asserted by the record. Treating that element as a cause additionally requires a causal definition and an appropriate counterfactual comparison or intervention model. This distinction matters when allocating advertising revenue, agent rewards, or referral rights.

Data design should therefore separate observation, inference, and causal-claim fields. The first may be direct, the second conditional on a model, and the third disputed. Combining them into one narrative makes a record appear more certain than the information warrants. The recovery proposition identifies where measurement must improve: it distinguishes cases in which the limitation is missing information, so greater computational reasoning power alone cannot resolve it.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

Evidence sets and compatible worlds

We now adopt a restricted economic setting in which an agent's decision can be analysed mathematically. Assume a finite set of possible transaction situations, a finite set of decisions, and a state that remains unchanged during deliberation. Each situation specifies decision-relevant details, such as the remaining discount allowance, the validity of authorisation, inventory availability, or a service's value to the buyer. We need not include every fact about the universe, only those that could change feasibility or the ranking of outcomes for the task under study.

$$\Theta(K) = \{\theta \in \Theta_0 : \theta \models K\}, \qquad \Theta(K) \neq \varnothing \tag{K4}$$

The set Θ_0 contains the initially admissible situations, and K is the accepted evidence held by the agent. It may include readable kup fragments, facts from an independent source, or known environmental constraints. The satisfaction relation means that a situation agrees with the evidence under the model's language and acceptance criteria. Each situation is called a possible world in this limited modelling sense; it need not be a separate physical universe. The validity of the result remains conditional on the initial set including the actual situation.

We assume that accepted evidence is correct within its scope and that the actual world has not been excluded. If messages are merely untrusted assertions, their content should not be inserted as certain fact. Instead, the model may represent probabilities of their truth or scenarios in which they are false. An empty set of worlds signals a contradiction or an error in the model or evidence. We do not interpret an empty intersection as complete knowledge permitting every decision; the guarantee is suspended until the inconsistency is resolved.

In a synthetic example, an earlier record states that a discount voucher was unused when issued but says nothing about subsequent use. The set contains both a situation where the voucher remains valid and one where it has been consumed. A recent, properly scoped attestation can narrow that set, provided its issuer can know the relevant state. Reading the old record repeatedly adds none of the missing information. This yields a testable distinction between data volume and the number of decision-relevant situations excluded.

The baseline model deliberately assigns no probability distribution to worlds; its results are conservative across every remaining situation. Probabilistic beliefs and their updating can be added later, but the probabilities require justification or calibration and imply different decision criteria. We also separate evidence acquisition cost from the value of information once received. Evidence may improve inference while being uneconomic to obtain because of time, fees, or consumption of an allowed privacy budget.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

Decisions guaranteed feasible by the evidence

Each possible world has a set of feasible decisions. A decision may purchase one offer, select a conditional bundle, allocate part of a resource, or abstain. Feasibility here means satisfying the modelled constraints, including eligibility, authorisation, resources, and deadlines. It does not mean that the decision is economically optimal or fair to every party; those properties require separate criteria. In this chapter, abstaining from creating a new commitment is assumed to be permitted in every world.

$$F_{\mathrm{safe}}(K) = \bigcap_{\theta \in \Theta(K)} F_{\theta}, \qquad \varnothing_{\mathrm{act}} \in F_{\mathrm{safe}}(K) \tag{K5}$$

The symbol F denotes the feasible decision set in a specified world; the final symbol denotes the action of abstaining, not the empty set. We call the intersection the set of decisions guaranteed feasible by the evidence. Guaranteed is conditional on the model, the truth of the evidence, and an unchanged state. The set can be small even when an excellent exchange exists in the actual world, because the agent cannot yet distinguish that world from one that prevents the exchange.

Proposition: execution feasibility within a correct model

If the actual world belongs to the compatible-world set and a decision is selected from the preceding intersection, that decision is feasible in the actual world. Membership in the intersection implies membership in every remaining world's feasible set, including that of the actual world. This is an elementary logical result, not a novelty claim. Its practical value lies in exposing the assumption at which a guarantee fails: the actual world may be absent, may change after the decision, or the interface may execute a different action from the one verified.

In a synthetic example, a bundle requires two units of a resource and another option requires one. The evidence establishes that either one or two units are available, and eligibility conditions are known. The guaranteed set includes the one-unit option and abstention but excludes the bundle. Correct evidence confirming two units in the unchanged state admits the bundle. If, instead, obtaining new information also secures an additional resource through a new reservation, the possibilities have changed. That effect must be separated from improvement due to knowledge alone.

This criterion is stringent when many rare possibilities remain. An application may instead choose a probabilistic criterion accepting a specified failure risk, but that is a different design choice whose level and probability source must be disclosed. A frontier can compare outcome quality with non-execution risk rather than compressing both into an opaque number. Empirical evaluation should count false acceptances and conservative rejections together: rejecting every offer may provide formal safety while failing to deliver a useful economic service.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

The decision gap caused by incomplete evidence

We now select an explicit numerical value criterion for a particular agent or decision-maker, such as buyer utility net of cost or an agreed service score. The criterion must retain the same meaning and unit across the worlds being compared. This definition neither automatically aggregates different people's utilities nor establishes a universal value measure for the general value space. Assume also that values are bounded and decision sets finite and nonempty, so an optimal decision exists in each world.

$$\mathrm{OPT}_\theta = \max_{x \in F_\theta} W_\theta(x) \tag{K6}$$

$$\mathcal{G}(K) = \min_{x \in F_{\mathrm{safe}}(K)} \max_{\theta \in \Theta(K)} [\mathrm{OPT}_\theta - W_\theta(x)] \tag{K7}$$

Here W is the score of decision x in world θ, and OPT is the best score available with knowledge of that world. The difference measures the opportunity loss from choosing the decision rather than making a fully informed choice. Among guaranteed-feasible decisions, G is the smallest attainable worst-case loss across worlds not excluded by the evidence. This explicitly applies a worst-case regret criterion. The formulation does not claim to invent regret; it applies that criterion to a specified structure of offers, evidence, and execution constraints.

In a synthetic example, options X and Y are feasible in two worlds. Their scores are eight and two in the first, and two and eight in the second; abstention scores zero. Without distinguishing information, each option has worst-case regret of six units, whereas abstention has regret of eight. The gap is therefore six when decisions are restricted to deterministic choices. Knowing the actual world permits the better choice and reduces the gap to zero. Randomised decisions use a different criterion and may improve expected outcomes; they are not implicitly included in the deterministic decision set.

This definition separates execution feasibility from choice quality. A guaranteed-feasible decision may be costly or slow; an apparently valuable decision may be impermissible in one world. Multicriteria frontiers can be reported when quality, time, and privacy cannot acceptably be reduced to one score. The numerical score used here is a local analytical choice, not a redefinition of value in general as money or a single measure of satisfaction.

The fully informed comparator remains an informational bound, not an equally situated experimental competitor to an agent paying for queries. We therefore report the gap alongside evidence costs and resource use, and compare operational algorithms under the same initial knowledge and budget. Loss from incomplete information must also be distinguished from loss due to approximate solution of the computational problem: knowing every constraint does not ensure that the agent has found the best feasible decision.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

Information monotonicity and its conditions

A useful property of the preceding measure can be established without assigning probabilities to worlds. Hold fixed the underlying decision set, value functions, resources, deadline, and execution rules in every world. Suppose correct evidence eliminates some candidate worlds without changing the actual world or its constraints. Costs once the evidence has arrived are treated as fixed and outside this comparison. Under these conditions, we can give a precise meaning to the statement that additional information does not harm robust choice.

$$\Theta(K') \subseteq \Theta(K) \implies \mathcal{G}(K') \leq \mathcal{G}(K) \tag{K8}$$

Proposition: correctly narrowing worlds cannot increase the decision gap

Let x* attain the minimum defining the old gap. Since it is feasible in every old world, it remains feasible in the smaller new set and thus remains among the guaranteed-feasible choices. Its maximum regret also cannot increase when evaluated over fewer worlds with value functions and optimality benchmarks fixed. Finally, the new problem can select that same decision or one that performs better under its criterion. These observations imply that the new minimum cannot exceed the old one, as required.

The proof identifies two distinct conditions: the set of epistemically admissible decisions expands, while the set of situations over which the worst outcome is measured contracts. Improvement need not be strict. Evidence may eliminate a world that affected neither the decision nor its worst regret. The gap can remain constant after many queries and then fall when one decisive fact arrives. Information should therefore be assessed by its effect on the task and its cost, not solely by message counts or record lengths.

The result no longer applies if a query takes long enough for an offer to expire, if its fee consumes the execution budget, or if it discloses information to a competitor who changes behaviour. These are not counterexamples to the proof: they change quantities held fixed by its assumptions. Nor does increasing each message's individual delivery probability establish an actual narrowing of worlds in every trial. Joint dependence, content, or reliability may change. The probabilistic case requires a comparison of specified channels, information, and policies, not replacement of precise conditions with a general slogan.

The proposition can serve as a simulator consistency check. If a correct fact is added without cost to a frozen state and the exactly computed optimal gap rises, there is an error in constructing the situations or solving the problem. By contrast, greater loss for an approximate agent may reveal a weak policy or distraction without contradicting the property of the optimal decision problem. This distinction locates the source of failure rather than automatically attributing it to information itself.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

Decision certificates and evidence acquisition costs

A decision certificate is evidence justifying both the feasibility of a particular choice and its quality within a declared tolerance. Certificate is a semantic concept here, not necessarily a single document, a cryptographic proof, or a short expression. It may combine kup fragments, published conditions, and independent query results. Its scope of worlds, value criterion, and loss tolerance must be specified; otherwise, verifiable becomes a description without a determinate question to verify.

$$\mathrm{Cert}_{\varepsilon}(x, K): \quad x \in F_{\mathrm{safe}}(K), \quad \max_{\theta \in \Theta(K)} [\mathrm{OPT}_{\theta} - W_{\theta}(x)] \leq \varepsilon \tag{K9}$$

The symbol ε denotes the largest permitted loss under the declared numerical criterion. A certificate can be sufficient without being the smallest possible certificate, and checking its quality may require costly computation. This separates the existence of sufficient evidence from the ease of finding or checking it. An application may use a valid upper bound on the optimal value instead of solving every world exactly, but loose bounds may prevent certification of a decision that is in fact good.

$$C_{\varepsilon} = \inf_{\pi \in \Pi_{\varepsilon}} \max_{\theta \in \Theta_0} \sum_{j=1}^{T(\pi, \theta)} c(q_j) \tag{K10}$$

Here π is a deterministic query policy, q is a query selected using previous answers, and c is its nonnegative cost in one specified unit. The set Π contains terminating policies that stop in every world with evidence permitting a certificate at the required tolerance. The symbol T counts the policy's queries in the relevant world. Cost is defined as infinite if no policy satisfies the requirement. The expression defines an optimisation problem; it neither supplies a general solution algorithm nor assumes that a finite domain is necessarily small.

In a synthetic example, asking whether a voucher is valid may eliminate most options, whereas requesting another product description leaves the decision unchanged. The optimal plan changes when the first question is costly, or when an inexpensive answer reveals which source should be contacted next. Adaptive policies therefore play a different role from selecting a fixed document list in advance. It may also be economically preferable to stop early and accept a small loss rather than acquire all obtainable information.

Monetary cost, elapsed time, and personal data disclosure should not be added without conversion factors meaningful to the decision-maker. Some can be objectives and the others constraints, or a trade-off frontier can be presented. Measurement should separate query fees from the cost of processing and verifying answers. Query efficiency is already an established research topic. The potential contribution here concerns particular classes of offers and execution evidence, and the bounds that can be proved for them.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

When is additional information necessary?

The need to query is not merely a preference for a better explanation; the decision structure itself may require it. Consider two worlds the agent cannot distinguish using its present information. In each, specify the decisions meeting both feasibility and the required loss tolerance. If those acceptable-decision sets have an empty intersection, no deterministic policy can make one final decision from the current information and succeed in both worlds. It receives the same input and therefore makes the same choice.

In a synthetic example, options X and Y are available. X scores one and Y zero in the first world, with the scores reversed in the second. Requiring loss below one unit in each world forces a deterministic policy to determine which world holds. A distinguishing fact might come from a record of a right's consumption, an eligibility attestation, or information about the user's request. Rephrasing existing information, however strong the linguistic capability, cannot create a distinction absent from the observation.

Randomised decisions require separate treatment. In this example, selecting X and Y with equal probability gives expected loss of one half in each world. The deterministic example therefore does not establish the impossibility of every probabilistic guarantee. It also shows, however, that distinguishing information permits zero loss, and that accepting randomisation changes the criterion rather than recovering the missing fact. An experiment must state whether it measures guaranteed success, average performance, or the probability of loss exceeding a threshold.

$$b^{q} \geq L \quad \Longrightarrow \quad q \geq \lceil \log_{b} L \rceil \tag{K11}$$

The relation gives a numerical bound in a restricted model. There are L worlds whose acceptable-decision sets are pairwise disjoint. Each query has at most b possible answers, with b greater than one, and a deterministic policy succeeds in every world after at most q queries. Its answer tree has at most b raised to the power q leaves. The worlds must reach different leaves; otherwise they share a final decision unacceptable in at least one of them. Counting leaves gives the bound.

This is an elementary information bound. It does not cover a query with unboundedly many answers or a message encoding the entire world without its size being charged. Nor does the number of worlds always equal the number of leaves required: many worlds can accept the same decision. This suggests a sharper research question: what is the smallest partition of worlds admitting an acceptable decision in each cell, and which available queries realise it at bounded cost? This determines the value of organising kup records relative to the decisions they are intended to support.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

Accessibility ρ as a property of information and channels

Accessibility describes whether a specified item of information can reach a specified element within a window of the environment's time. It must be distinguished from state reachability in a dynamical system; the concepts differ even when related in an application. Receiving a message also does not imply understanding it, believing it, or holding authority to act on it. An agent may clearly see an offer yet remain ineligible, or receive a kup fragment insufficient to identify the exchange stage reached.

$$\rho_{\mu,\alpha}(m, a; J) = \Pr_{\mu,\alpha}\{m \in \mathrm{Rec}_a(J)\} \tag{K12}$$

Here m denotes a semantically specified information item or message, a the receiving element, and J the time window. Rec contains what the element receives during that window under an explicit reception criterion. The distribution μ specifies uncertain circumstances, while α specifies the policy for publication, routing, and requesting information. Without these components, ρ is not a meaningful probability. In a deterministic environment, a zero-one delivery indicator or a non-probabilistic accessibility relation may be used instead.

The policy may depend on history: an agent requests evidence after seeing an offer, and a seller publishes different terms according to eligibility. Accessibility must therefore be identified as measured before any query, under a fixed policy, or under the best policy within a budget. Low delivery probability need not indicate a network fault. It may result from privacy policy, access permissions, or the fact that the information was never requested. Each cause calls for a different intervention.

Information must also be defined so that duplicate copies do not inflate accessibility. Several messages may carry the same fact, while one message may contain several facts affecting different decisions. We measure receipt of the content needed by the task, or explicitly state that the unit is the message. Where messages are encrypted or uninterpretable by an element, delivery of the carrier can be distinguished from access to usable content. Decoding conditions should be documented rather than assumed.

In practice, three successive levels can be measured: the carrier delivery rate, the rate of correct meaning extraction, and the rate of correct decisions given that meaning. A system may improve the first while still failing at the second because of an ambiguous format, or at the third because of a weak decision policy. This decomposition supports comparisons among agents in the same environment and among environments serving the same agent. It also prevents higher message-opening rates from being converted into unsupported claims of improved value-exchange ability.

An element's ability to request information remains part of its actual action repertoire in the model. We do not attribute to it queries the researcher can perform outside that repertoire, or confuse what the simulator knows with what the policy can access during execution.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

Joint information dependence and the value of discovery

Individual accessibilities alone do not determine decision quality, because information matters through its combinations. Some offers are substitutes, so discovering either suffices. Others are complements with practical value only when discovered together and meeting their composition conditions. The distribution of received information sets, or at least their task-relevant dependence structure, must therefore be represented. Economic improvement cannot be inferred from a higher delivery probability for each information item considered in isolation.

Consider a synthetic example calculated directly. Two substitute offers each yield one unit of utility, with no additional benefit from taking both. Under the first channel, only the first offer arrives with probability one half, and only the second arrives with probability one half. Each offer's marginal delivery probability is one half, but a useful offer always arrives, so expected utility is one. Under the second channel, both offers arrive together with probability six tenths, and neither arrives otherwise. Each marginal probability has risen, but expected utility has fallen.

$$\mathbb{E}[U_1] = 1, \qquad \mathbb{E}[U_2] = 0.6 \tag{K13}$$

The example assumes neither fees nor errors in understanding messages. The decline comes solely from changing their joint dependence. If instead the offers were complements and utility arose only when both were discovered, the ranking would reverse: the first channel would never produce that benefit, while the second would produce it with probability six tenths. Marginal accessibilities alone therefore provide no universal total ordering of the channels. The task structure and economic relationships among messages or their offers must be specified.

The same reasoning applies to evidence. Authorisation evidence without eligibility evidence may leave a decision unresolved, whereas either of two alternative attestations from acceptable authorities may suffice. Accessibility design should account for this structure when allocating communication budgets. Publishing more copies of redundant information may consume the channel without excluding another world. A short message about a shared resource can have a greater effect on the guaranteed-feasible decision set.

A controlled experiment could hold message counts, costs, and marginal delivery probabilities fixed while varying joint dependence and timing. It would compare choice quality, the evidence-related decision gap, and execution costs across tasks involving substitutes, complements, and mixtures. This is a proposed experiment, not a field result. If a language-model agent is used, wording should remain as constant as possible so that the test does not become a comparison of presentation styles. The design isolates the effect of communication structure on value-exchange ability rather than associating that ability with total information volume received.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

Evidence reliability, privacy, and verification costs

Readable evidence need not have true content. A message may be signed by a known party yet concern an outdated state, a different scope, or a claim that party has not checked. We therefore distinguish message provenance, transmission integrity, the source's authority regarding the relevant fact, validity duration, and linkage to the decision. These remain different properties even if an interface combines them into one visible trust marker. A general numerical score should not suppress these distinctions where commitment validity depends on them.

In the preceding model, accepted evidence was treated as correct. It can be extended to explicitly model errors, defining failure events such as incorrect eligibility verification or linking an attestation to the wrong transaction. If bounds on their probabilities are known under the test conditions, a conservative bound follows for the probability that at least one occurs. This analyses risk within a specified model; writing unmeasured probabilities as numbers does not make them established facts.

$$\Pr\left(\bigcup_{j=1}^{n} B_j\right) \le \sum_{j=1}^{n} \Pr(B_j) \le \sum_{j=1}^{n} \delta_j \tag{K14}$$

The symbols B denote the defined failure events, and δ their valid probability bounds in the model. Independence is not required, although overlapping events can make the bound loose. Success rates obtained under tests unrepresentative of operating conditions cannot automatically be transferred. For an adaptive policy selecting evidence in response to answers, the guarantees must remain valid under that selection. Otherwise, selection bias can make actual risks greater than isolated test figures suggest.

Privacy constrains which evidence may be requested. A seller may need proof that an age or geographical condition holds without learning other personal information. The model can specify the requirement as establishing a logical predicate while leaving its implementation to a suitable technical architecture. Concealing data may, however, limit auditability or increase verification cost, so the trade-off must be measured. We do not assume that a proof technique automatically resolves the correctness of the original data on which its proof depends.

For applications, we propose separating evidence acquisition cost from the cost of verifying its relevance and validity, and from retention cost. A better-structured record may shorten checking time without reducing data received. A scoped attestation may reduce disclosure while increasing issuance time. Successful design satisfies the decision criterion together with privacy and resource constraints. More evidence is not an end in itself: the objective is sufficient knowledge that can be justified and used at an acceptable cost.

CHAPTER 4 EXCHANGE RECORDS AND EVIDENCE

Stale evidence, reservations, and the static model's limits

Evidence-based guarantees become more difficult when the state changes during querying and execution. An agent may read that a unit is available, only for another agent to buy it before the commitment is submitted. The information was correct when read, yet the decision fails. This time-of-check/time-of-use gap must be represented explicitly when applying the results to agent markets. A more accurate historical description is insufficient when the resource can be consumed by others after that description.

A reservation can grant an agent a specified resource right for a declared period. Alternatively, verification and commitment can be linked in one transition executed as a single operation by the resource authority. Reservations, however, change the state, occupy capacity temporarily, and impose an economic cost on others. We therefore distinguish learning that a resource is available from acquiring a right that prevents conflicting consumption. Both may improve execution, but their mechanisms differ and require separate measurement.

$$\sum_{i \in C_t} a_i + \sum_{i \in H_t} a_i \le B_t \tag{K15}$$

The set C contains outstanding commitments, and H contains active reservations not yet converted into commitments; these sets are disjoint. The quantity a is nonnegative resource consumption, or its vector, and B is the total available capacity at that time under the state accounting. With multiple resources, the inequality is interpreted componentwise. A reservation is removed from its set when converted to a commitment, preventing double counting. This is a proposed local invariant: each execution transition must be proved to preserve it, and external resource updates must enter the model correctly.

Reservations require expiry, release, and recovery policies after disruption. Parties may also reserve strategically to obstruct competitors. Maintaining the resource invariant therefore does not establish market efficiency or sound incentives. Consistency of one seller's resource does not ensure consistency of a bundle spanning another seller. The application must specify whether partial completion is acceptable or further coordination required. kup records the stage-specific events needed for inspection, but collecting them in a record does not itself create coordination.

This yields a research question connecting the chapter's sections: what is the least evidence, and the least temporary resource entitlement, an agent needs to implement a good decision before its window expires? Work can begin with a fixed state to test evidence sufficiency, then introduce controlled inventory changes, message delays, and competition for reservations. The share of unused reservations and their cost to other participants should be reported. This delineates the validity of static results and makes their limits a starting point for a testable temporal model, rather than extending a guarantee without proof.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Intelligence as the ability to exchange value

A stipulated definition and its scope

This study defines intelligence as the ability to exchange value within a value space, and its degree as the extent of that ability. This foundational choice establishes the subject of inquiry. It is broader than customary psychological usage and permits natural and artificial elements to enter a common description. It does not, by itself, attribute consciousness, intention, or linguistic understanding to every element participating in an exchange relation. Subsequent scientific claims concern what the definition explains and makes measurable, rather than the breadth of the term alone.

We must therefore distinguish the possibility of mutual influence from skill in choosing among its pathways. A stone has particular responses to impact or a temperature difference; a plant has mechanisms of transformation and adaptation; and a software agent has policies for searching, comparing, and committing. All can be included in a value space, but their action sets, observations, and memory differ. If a researcher grants an element interventions it cannot perform, the measurement captures the researcher's control over the experiment and incorrectly attributes that control to the element under study.

There is a methodological connection with definitions relating intelligence to goal achievement across environments, as in Legg and Hutter, and with Russell and Subramanian's analysis of optimality under architectural and computational constraints. These works do not establish the validity of the exchange-based definition. They support the need to specify environments, tasks, and constraints before making comparisons. Klyubin and colleagues' concept of empowerment provides a related reference for studying influence, although the breadth of influence alone determines neither exchange quality nor the alignment of participants' interests (Legg and Hutter, 2007; Russell and Subramanian, 1995; Klyubin et al., 2005).

The economic specialisation requires an explicit success criterion. A resource transfer may count as exchange in a structural sense while constituting an unsuitable transaction for a particular buyer. Abstention may be the best outcome for a task requiring minimum quality or prohibiting actions beyond delegated authority. Accordingly, the proposed intelligence is not measured by purchase counts, payment volume, or the number of messages between elements. We need tasks that specify acceptable exchanges and distinguish their completion from activity that does not serve them.

The unit of analysis comprises an element, an initial state, a time window, a budget, and a task set. Capability is relative to this unit, rather than an absolute number residing in the element independently of its environment. An agent may consequently outperform another in a particular offer space without that superiority extending to spaces that withhold information, alter resources, or require different policies. A claim that a general factor operates across environments requires an independent transfer test after local measures have been constructed.

The definition gains empirical content when its measures predict outcomes independent of the tests used to construct them. Merely renaming every interaction intelligence, without a measurable distinction, adds no explanatory power to the theory.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Capability sets and available actions

Let the task set be specified, with each task defining a condition that can be checked against an interaction history. The condition might require obtaining a service within a budget, transferring a resource to another element, or creating a commitment compatible with existing rights. Denote the initial state by s, the time window by J, and the budget by b. The state must include information affecting actual feasibility, whereas the element sees only the observations its environment makes available.

$$\mathcal{C}_{a,\divideontimes}(s;J,b) = \{\theta \in \mathcal{T} : \exists \pi \in \Pi_a(s;J,b),\ \exists h \in H_{\divideontimes}(s,\pi),\ h \models \theta\} \tag{A1}$$

The equation defines the existential capability set of element a in space ⋇. Here Π denotes the policies actually available, H the admissible histories that can result from following a policy, and T the task set; the final relation expresses task satisfaction on a history. This definition establishes the possibility of success in at least one history. It establishes neither guaranteed success nor a high probability of success. Choosing this formulation makes the weaker meaning of capability explicit before stronger formulations are developed.

For a passive element that does not choose between actions, the policy set may contain only one response mechanism. That mechanism is not thereby a conscious decision, but it allows us to describe how the element can participate in exchange under specified conditions. A purchasing agent, by contrast, may possess several tools and policies, while its authority to invoke them remains constrained by delegation and available funds. A policy does not belong to its set merely because a researcher can program it outside the adopted constraints or execute it with additional resources.

Capability sets support a natural partial ordering. If, under identical conditions, one element's set contains all tasks in another's, the former has broader capability in the specified sense. The two sets may instead overlap without either containing the other, with each element capable of tasks unavailable to its counterpart. This situation cannot be converted into overall superiority before task importance is specified or success on one task is permitted to compensate for failure on another.

The task's level of abstraction must also be controlled. If obtaining information is described as one task, then divided into a thousand characters with each counted as a separate success, the result changes without a substantive change in capability. The task list must therefore preserve stable semantic units and distinguish the original task from the steps used to complete it. This requirement makes comparisons reproducible and prevents capability from being inflated by changing the recording method or dividing an effect into arbitrary parts.

The cost of discovering or executing a policy is counted whenever it falls within the task constraints. Excluding planning costs can place a mathematically feasible policy outside the budget available to the element in practice.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Possibility, probability, and guarantee

The meaning of capability depends on the type of uncertainty. An environment may permit several states without a probability distribution, provide a known random mechanism, or include a strategic party choosing its actions. The first case supports statements about possibility or guarantee. The second permits success probabilities under a specified distribution. In the third, a guarantee requires specifying what the other party can do, what it knows, and when it chooses. Combining these cases into one number conceals decisive assumptions.

$$q_a(\theta \mid s, J, b, \mu) = \sup_{\pi \in \Pi_a} \mathsf{P}_{\mu,\pi}(h \vDash \theta) \tag{A2}$$

The function q represents the best success probability permitted by the adopted policy set when μ specifies the law governing relevant states and transitions. The supremum need not be attained by any particular policy; using it therefore does not assume an optimal policy exists. The probability does not become interpretable simply because a symbol has been written. The sources of randomness, their dependence on actions, and the stability of their estimates across tests must be stated.

$$\mathcal{C}^{\mathrm{rob}}_{a,\divideontimes} = \{\theta : \exists \pi \in \Pi_a,\ \forall h \in H_{\divideontimes}(s,\pi),\ h \vDash \theta\} \tag{A3}$$

The second expression describes guaranteed capability relative to the specified set of admissible histories. The guarantee remains relative to that set: if the model excludes server failure or resource cancellation, it does not cover those events. Every admissible policy must also have at least one admissible history, preventing a vacuous logical success from arising from a policy that cannot run. When probabilistic knowledge is incomplete, a lower probabilistic guarantee can instead be constructed over a specified class of distributions.

A synthetic example illustrates the difference. Suppose a resource is available in one of ten equally likely states and the element can obtain it whenever available. The element may possess existential capability for the task even though its success probability is no greater than one tenth. If it can reserve an acceptable substitute in the remaining states, its guaranteed capability changes. Reading an old record indicating that the resource existed does not alter its actual availability, although it may change the agent's belief or search method.

Empirical reports must specify which quantity is being estimated. A sample success rate estimates performance for a particular policy and test distribution; it does not prove guaranteed capability or establish the best achievable probability. Statistical intervals and repeated random seeds help describe empirical uncertainty, but cannot replace missing model assumptions. This distinction is essential when agent results are used to justify economic commitments or spending limits.

Overall success probability must also be separated from success conditional on selecting easy cases. Reporting success conditional on resource availability does not measure the ability to handle its absence or find an acceptable substitute.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Intelligence scores and the order of information and choice

When an application requires a single score, a finite task set and nonnegative weights summing to one can be selected, and task success probabilities aggregated. These weights are not natural properties imposed on every space: they represent the purpose of evaluation. They may reflect request frequencies, outcome importance, or an announced service policy. They should be fixed before agents are compared, while individual task scores remain available so that the aggregate cannot conceal important differences.

$$I_w^{\mathrm{ann}}(a) = \sum_{\theta \in \mathcal{T}} w_\theta \sup_{\pi \in \Pi_a} \mathrm{P}_{\mu,\pi}(h \vDash \theta), \quad \sum_\theta w_\theta = 1 \tag{A4}$$

In this formulation, the task is announced to the element before its policy is chosen, so optimisation occurs inside the sum. The comparison permits a different policy for each task while keeping resource constraints and the announcement procedure fixed. The score is suitable, for example, for testing an agent that receives a clear purchase request in each episode and then chooses a search and execution plan. It cannot be interpreted as the ability to satisfy demands the agent did not know about or that left no trace in its observations.

$$I_w^{\mathrm{hid}}(a) = \sup_{\pi \in \Pi_a} \sum_{\theta \in \mathcal{T}} w_\theta \mathrm{P}_{\mu,\pi}(h \vDash \theta) \;\le\; I_w^{\mathrm{ann}}(a) \tag{A5}$$

When the task is selected without being revealed and one common policy must be used, optimisation moves outside the sum. The inequality follows because the best common policy cannot exceed the sum of the task-specific optima. The proof is immediate: for every task, any policy's success is bounded above by the corresponding supremum. Multiply by the nonnegative weights and sum, then take the supremum over common policies. If observations reveal the task later, that revelation must be represented within the common policy and the observation law.

The inequality can be strict. In a synthetic example, an element has two mutually exclusive actions; the first task succeeds only under the first action and the second only under the second. With equal weights, capability is one when the task is announced and one half when it is hidden, provided no signal distinguishes it. This shows that better access to the request specification may improve performance without changing the algorithm or increasing physical resources.

A measured score for an actually deployed policy replaces the supremum with that policy's outcome. The difference between the two scores is an implementation, learning, or search gap; it should not automatically be attributed to a shortage of environmental opportunities. Computing optimal capability is difficult in practice. One can report a lower bound from a successful policy and an upper bound from an optimistic solver, then narrow the gap rather than present current model performance as the ultimate capability limit.

If changing the policy requires training or installing new tools, that change must either count towards the budget or be explicitly prohibited. One element cannot be granted free redesign and then compared with another running a fixed program.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Why capabilities have no universally neutral numerical ranking

Particular care is needed when comparing elements that succeed in different kinds of exchange. One agent may excel at offer discovery but perform poorly at checking delegated authority, while another exhibits the opposite pattern. Their performance profiles contain useful information, but combining them into a total ordering requires a decision about the trade-off between tasks. The general structure of a value space does not supply that decision automatically, since value itself is relative to a specified perspective and conditions of use.

$$u = (1, 0), \quad v = (0, 1), \quad I_w(u) = w_1, \quad I_w(v) = w_2 \tag{A6}$$

Proposition: the principle of preferring higher success on each task does not, by itself, imply a strict weight-independent ordering of the two displayed profiles. To prove this, assign the larger weight to the first task: the first profile is preferred. Assign the larger weight to the second task: the preference reverses, while the weights remain nonnegative and sum to one. Equal weights produce equal scores. This is a specific argument against claiming neutrality for a ranking derived from conflicting profiles; it does not mathematically prohibit an explicitly stipulated score.

Componentwise comparison remains possible. A profile that is no worse on any task and strictly better on a task with positive weight receives a higher score under every linear aggregation satisfying those conditions. When profiles conflict, it is preferable to publish the frontier of nondominated profiles and let the user determine weights or acceptance constraints. For example, verification success may have to exceed a noncompensable minimum before search cost or price quality can enter the comparison.

The same issue arises in generalisation across environments. Choosing one hundred easy environments and one difficult instance is not neutral; duplicating an environment under several names also changes its implicit weight. The environment family, its unit of diversity, and its sampling procedure must therefore be specified. Rankings can be compared as weights vary over an announced range, showing whether a conclusion is stable or depends on a particular weighting of one task.

In the economic specialisation, ranking agents by raw financial returns is insufficient. An agent may receive a larger budget, wider authority, or an easier customer base and generate higher returns without possessing a better exchange policy. Higher profit for one party may also accompany a worse outcome for another. Comparisons should therefore use task-specific performance scores while reporting resources and the distribution of outcomes among parties. Additional social judgements remain normative choices distinct from the structural definition of intelligence.

A sensitivity analysis can also identify the range of weights preserving a particular preference. The result is then a conditional, testable recommendation rather than a number suggesting that differences in intended use have disappeared.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

When do additional channels expand capability?

Adding an information channel or an exchange mechanism may appear always to increase capability. This conclusion holds in a particular model, but fails if the addition makes existing actions more expensive or slower, changes other parties' behaviour, or forces additional messages to be processed. An optional extension must therefore be defined precisely: the element can ignore the new channel, its previous policies remain executable, and their resource requirements, outcome laws, and success conditions are preserved.

$$\Pi_a \hookrightarrow \Pi_a^{+} \quad \Longrightarrow \quad \mathcal{C}_{a,\divideontimes} \subseteq \mathcal{C}_{a,\divideontimes^{+}}, \qquad q_a^{+}(\theta) \geq q_a(\theta) \tag{A7}$$

Conditional monotonicity proposition: if existing policies and histories are preserved under a corresponding embedding, the existential capability set does not shrink. If each old policy's outcome distribution is also preserved, the best success probability does not decrease. The proof reuses the old policy in the extended environment. Its successful history remains available in the existential case, and its probability remains unchanged in the probabilistic case, while optimisation takes place over a set containing the old policies' images and additional choices. With fixed weights, monotonicity also holds for the aggregate score.

Increasing each message's marginal accessibility is insufficient for this proof. Consider a synthetic example with two substitute offers, either yielding value one. In the first environment, one offer always arrives, with each arriving with probability one half. In the second, both arrive together with probability 0.6, and neither arrives otherwise. Each offer's arrival probability has increased, yet the probability of obtaining at least one opportunity has fallen from one to 0.6. The cause is altered dependence between arrivals, not agent weakness.

The availability of a best possible policy also does not imply that a deployed policy will improve when tools are added. The agent may choose an unreliable channel or spend its time comparing further alternatives. We therefore distinguish monotonicity of the opportunity set from monotonicity of learned performance. The former concerns the model; the latter requires an algorithm or experiment controlling task distributions, tool costs, and the decision budget.

This result provides a practical criterion for designing Promoi. If a system expands its offer space while seeking to preserve existing possibilities, it must retain an operational route that supports them, including valid rights, reservation conditions, and interface costs. If the expansion changes competitive equilibrium or server loads, its effects constitute a new problem outside the proposition's scope. The assumptions thus delineate design boundaries rather than turn monotonicity into a general slogan about connectivity.

A channel may impose reception costs even when its content is ignored, for example by filling memory or congesting a communication port. In that case, optionality does not hold and the comparison requires a new resource model.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Separating agent effects from environment and resource effects

Empirical capability should be tested through a design that separates performance components. We propose a matrix crossing several agents, environments, and specified budgets, with repeated trials in each cell on tasks comparable in difficulty and distribution. This tests whether an agent's advantage is stable or depends on favourable offer wording or an access channel providing extra information. Explanatory tests should vary one factor at a time before examining interactions with sufficient cases. Estimation requires reference coding or zero-sum constraints on main effects and on interaction rows and columns to identify the coefficients.

$$Y_{ijr} = \alpha + \beta_i + \gamma_j + \delta_{ij} + \varepsilon_{ijr} \quad \text{(A8)}$$

The equation is a proposed statistical description of outcome Y. The index i identifies the agent, j the environment, and r the repetition; β represents agent differences, γ environment differences, δ their interaction, and ε the residual. This is neither a law of value spaces nor proof of causality. For binary outcomes, a suitable probabilistic model may replace linear regression while preserving the crossed design. A causal interpretation of differences requires control of assignment and accompanying factors.

Resources include money, computing time, tool calls, context size, execution deadlines, and access rights. These dimensions should not all be replaced by a single price unless an exchange rate is specified and remains valid for the comparison. An agent may be unable to buy additional time regardless of its funds, and data rights may prevent it from substituting one service for another. It is therefore useful to report success curves against each budget, or a multidimensional performance-cost frontier, rather than a compact ratio that conceals constraints.

A fair comparison begins with equivalent tasks, authority, and information. If the representation of Promoi is the factor being studied, the feasible contracts and actual resources remain unchanged while their organisation varies. A separate experiment can expand the feasible contracts to measure offer innovation. Separating these experiments is essential: improving an interface and adding a new economic opportunity are different causes of better performance, with different costs and scopes of generalisation.

Test information must not leak into policy training or selection. Families of conditions and structures should be divided into development and test sets, with evaluation on combinations not used for tuning. Correct abstentions should be reported alongside successful purchases, since a higher transaction rate may conceal acceptance of terms that cannot be executed. These procedures constitute a proposed experimental programme. The manuscript does not assume that particular agents have passed the tests or that any statistical effect has already been established.

The policy for retaining memory between episodes must also be stated. A persistent information cache may improve performance at a cost paid earlier. This advantage must either be standardised across agents or measured as part of their ability to learn over repeated interactions.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Exchange capacity and the unit of event counting

An element's ability to complete tasks differs from the capacity of an entire space to accommodate exchanges. Each agent might be able to purchase a resource individually even though all cannot purchase it simultaneously. Relative to an initial state, a window, resource constraints, and a semantic counting rule, capacity is defined as the supremum of distinct completed exchange events across admissible histories. This requires separating an event from its recorded copies and specifying what makes two events economically or structurally distinct.

Capacity does not increase because a kup record is copied to ten servers or retrieved ten times. Nor may one transaction be divided into an arbitrary number of messages and counted as additional transactions. Conversely, two repeated service operations may constitute distinct events if each has a separate request and completion under the study's criterion. Occurrence identity and the completion criterion therefore form part of the measurement protocol, while the record's physical carrier remains a means of obtaining evidence about the event.

$$C_{\curlywedge}(s;J,B) = \sup_{h \in H_{\curlywedge}(s;J,B)} N_{\mathrm{done}}(h) \;\le\; \lfloor B/c \rfloor \tag{A9}$$

Budget proposition: suppose every counted operation irreversibly consumes at least c units of one resource, the total available amount is B, and c is positive. Capacity then cannot exceed the integer part of B divided by c. The proof is that completing n operations consumes at least n times c; total consumption cannot exceed B, giving the bound. The quantities are measured in resource units per event and total resource units, so the ratio is a dimensionless count.

The bound requires genuine incremental consumption for each event, with no reuse within the window. In this form, it does not apply to a resource recovered after service or replenished during it, nor to another recording of the same event. Multiple-resource bounds can be obtained by taking the tightest valid bound when each resource has a separate minimum-consumption condition. Dependence between resources, however, may yield a stronger bound than separate comparisons.

High capacity does not imply high benefit. A resource may circulate repeatedly between two elements, increasing event counts without increasing inventories or participants' welfare. A single composite transaction may produce a greater effect than one hundred small transfers. Capacity should therefore be reported alongside measures of semantic completion, specified value, and cost. Event counts help analyse system loads and bottlenecks, but do not replace outcome evaluation or measures of an agent's ability to choose appropriate exchanges.

The initial state also affects capacity: previously reserved inventory differs from an equal quantity of unrestricted inventory. Pending rights and existing commitments must therefore be recorded rather than reducing the state to the physical number of units alone.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Capacity when time is a partial order

Not every environment requires a global clock measuring duration in seconds. We may know that one event preceded another because it supplied a prerequisite or transmitted information, while no order is known between independent events. Capacity can then be studied through the precedence structure. Lamport provides a classical framework distinguishing event order from the relations generating it in distributed systems. We use that distinction to establish counting bounds without assuming global synchronisation (Lamport, 1978).

Assume a finite set of distinct completion events and a partial order on them. A chain is a set whose events are pairwise comparable; its length is its cardinality. An antichain is a set with no two comparable events. Let H denote the maximum chain length and W the maximum antichain cardinality. The former measures dependency depth; the latter measures the width of independence under the chosen relation. Width does not necessarily imply sufficient resources to execute all those events physically at the same time.

$$N_{\text{done}} \leq H\,W \tag{A10}$$

Depth-width proposition: under these assumptions, the number of events is at most H times W. Assign each event a rank equal to the length of the longest chain ending at it. If one event precedes another, the second has a greater rank because the chain ending at the first can be extended. Events of the same rank therefore form an antichain containing at most W events. There are at most H ranks; summing their bounds proves the result.

This elementary proposition about partial orders clarifies capacity when no shared numerical duration is available. If the admissibility constraints impose uniform depth and width bounds on all histories, the result becomes a bound on the space's capacity within the study's scope. Measuring H and W from one completed history, however, does not establish that every other history satisfies the same bounds; it describes only that history.

The bound may be loose when shared resources prevent the maximum width from being attained at every layer. Adding an artificial order between independent events also changes H and W even though their number remains fixed. Precedence must therefore be derived from actual rules or explicitly introduced as a scheduling device. If a clock becomes available, layers can be associated with durations under further assumptions; dividing event counts by undefined time does not produce a scientifically meaningful rate.

Computing the layers is an implementable procedure when the event list and precedence links are known. Inferring those links from partial records remains a separate information problem and may yield several compatible orders.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

A finite capacity model with an abstract clock

To obtain a computable result, specialise the value space to a finite transition graph with N states. Assume that the state is sufficient to determine admissible transitions: any budget, rights, or memory affecting whether a subsequent action is permitted must be included. Each transition represents an operational step and carries two binary labels: k for the abstract clock increment and q for new semantic completion. A step may occur without advancing this clock, and the clock may advance without completing an exchange.

$$C(s,T) = \sup_{h:\ \sum_{e\in h} k(e)\le T} \sum_{e\in h} q(e), \qquad k(e), q(e)\in\{0,1\} \tag{A11}$$

The supremum is taken over finite paths starting at s, with T a nonnegative integer. The label q assumes that a step completes at most one new event under the counting rule. If a path repeats a cycle, its completion labels must represent genuinely distinct occurrences, not copies of an earlier occurrence. Occurrences may be distinguished by their positions even when the system returns to the same operational state, provided the history of their identities does not affect the transition law.

$$C(s,T) = \infty \iff \exists \gamma :\ \sum_{e\in\gamma} k(e) = 0,\ \sum_{e\in\gamma} q(e) > 0 \tag{A12}$$

Productive-cycle proposition: capacity is unbounded if and only if a cycle satisfying the two displayed conditions can be reached using at most T units of clock advancement. The existential statement in the equation includes this reachability restriction. Sufficiency follows by repeating the cycle arbitrarily many times after the access path: the clock does not advance, while the number of completions increases. The assertion concerns the absence of an upper bound across finite paths, not the physical execution of infinitely many operations in a single run.

The converse depends on the finite number of states and the nonnegative labels. Without a productive zero-clock cycle, every zero-clock cycle can be removed without losing completions, after which the remaining segments have bounded length. The proof on the following page supplies a numerical bound and establishes the converse. If the state hides an exhaustible resource or a prohibition on repeating a cycle, the graph is not a sufficient description of the environment and the equivalence does not apply.

This proposition claims no general law of time or nature. The clock is an explicit model variable and may count negotiation rounds or simulation ticks. If a history-dependent condition prohibits certain sequences, that condition must be represented in the state, or the path class must be restricted and the proof reconsidered. The specialisation is useful because it identifies precisely when unboundedness results from a testable modelling assumption.

Multiple local clocks cannot be reduced to this binary-increment representation without a specified rule preserving the required constraints. A different aggregation may change computed capacity even when the operational events remain the same.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Proof of the finite bound and diagnosis of model omissions

Assume now that no productive cycle with zero clock advancement is reachable within the time budget. Take any admissible path consuming m clock units, where m is at most T. Since clock labels are binary, the path contains m clock-advancing transitions, separated by segments that do not advance the clock. There are m plus one such segments, some of which may be empty.

Within any zero-clock segment, a repeated state identifies a cycle. By assumption, that cycle cannot have a positive completion total; since its labels are nonnegative, the total must be zero. Remove the cycle while preserving the segment's endpoints and completion count. Repeating this deletion produces a simple path that visits no state twice. Its length is therefore at most N minus one, and its completion count cannot exceed that length because each transition completes at most one event.

$$C(s,T) \le (T+1)(N-1) + T = N(T+1) - 1 \tag{A13}$$

To complete the proof, add the completions from at most T plus one zero-clock segments, then include at most T clock-advancing transitions. This yields the displayed bound for every path and hence for capacity. The proof permits arbitrarily long paths of steps that neither advance the clock nor complete anything: those steps do not increase capacity, and their cycles can be deleted for counting purposes. Completion may also occur on a clock-advancing transition, accounting for the additional T term.

The bound is tight within this class. In a synthetic example, arrange N states in a chain and let each transition between consecutive states have zero clock advancement and one completion. Add a transition from the last state to the first that consumes one clock unit and completes one event. Starting at the first state, traverse the chain and return T times, then traverse the chain once more. This attains the bound exactly, with no complete zero-clock cycle. When N equals one, a self-loop that advances the clock and completes an event suffices.

The result can test a simulator. If its count exceeds the bound while all assumptions hold, the counter duplicates events, the declared graph omits relevant state, or the labels do not match their definitions. If a productive zero-clock cycle appears, one should examine whether it is intentional abstraction or the consequence of omitting a cost or delay. Correcting the model may require reservation state, resource consumption, or clock advancement; a correction should not be selected merely to force agreement with a prior expectation.

A short expression for the bound does not ensure easy computation: representing rights and resources may require an enormous state space. Proof validity must therefore be distinguished from the practicality of constructing or exploring the graph within a realistic budget.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Contention, reservation, and distributed authority

A new difficulty arises when the analysis moves from one agent to elements acting concurrently. Two agents may read kup records consistent with resource availability and then both commit to using the same unit. The problem need not be the record's truth when created, but a state change between reading and commitment. The Promoi model must therefore distinguish a discoverable offer, an offer eligible in a particular state, a reserved right, and an actually completed outcome.

$$\sum_{i=1}^{n} a_r(o_i)x_i \le B_r, \qquad x_i \in \{0,1\} \tag{A14}$$

The equation expresses the constraint on resource r when several offers are selected: a denotes offer consumption, B available capacity, and x acceptance. Each agent may satisfy the constraint individually while their aggregate violates it. Local checking is therefore insufficient when offers draw on a shared balance or inventory, unless rights have been allocated in advance to separate quotas that keep the aggregate within the bound. Such a guarantee requires specifying quota ownership, transfer, and prevention of repeated use of the same right.

One possible design grants the agent a reservation of limited duration before final commitment, specifying the resource identifier, quantity, and expiry conditions. Other designs allocate consumable units of rights or use a coordinator to arbitrate conflicting requests. These mechanisms are known in principle. The research question here concerns their effects on the offer space, access costs, and completion capacity. No single mechanism is assumed suitable for every environment or resource type.

Combinations spanning providers require analysis of partial failure. If one component is purchased successfully and another fails, the first may become useless or retain independent value. The offer must specify whether this outcome is acceptable, whether linked commitment is required, or whether defined cancellation and compensation are available. An intent record does not prove component completion, a delegation signature does not prove service delivery, and monetary compensation does not necessarily restore lost time or opportunity.

Operational safety is defined here as respecting the specified resource, authority, and commitment limits. The possibility of progress means that acceptable requests can be completed under stated access conditions. A system may preserve all limits by rejecting every request, yet exhibit little exchange capability. The two properties must therefore be measured separately, alongside waiting time, reservation cost, and the fraction of resources held idle. Comparisons must show how these costs are distributed among buyers, providers, and the platform, since improving one party's outcome may shift a burden to another.

Repeated denial of service must also be examined: a system may complete many requests while continually postponing one participant's. Scheduling fairness and waiting-time bounds are additional properties that do not follow from reservation safety alone.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Three synthetic environments with different laws

Small models can test whether definitions depend on the rules of a particular market. We present three synthetic worlds, each containing four elements with binary states and discrete rounds; no element participates in more than one event per round. State one denotes possession of a resource or information, depending on the world, and zero denotes its absence. The following results are proofs for these artificial rules, not empirical measurements or physical descriptions of stones and plants.

$$10 \longrightarrow 01: \quad C_{\text{complete}}(T) = 2T, \qquad C_{\text{star}}(T) = T \tag{A15}$$

In the first world, every pair of elements permits resource transfer according to the displayed transformation. Initially, two elements possess the resource and two do not. At most two events occur per round because each uses two distinct elements. The bound is attained by matching each holder with a non-holder and reversing the transfers in the following round. It is therefore attainable in every round, yielding twenty events in ten rounds while the number of resource units remains two.

The second world retains the same transfer rule and initial state but has a star topology: every exchange passes through the central element. Since the centre participates only once per round, the bound becomes one event. It is always attainable. If the centre holds the resource, a peripheral element lacks it; if the centre lacks it, a peripheral element holds it, because the total remains two units. Changing connectivity thus yields ten events in ten rounds without changing the conservation law or element count.

$$10 \longrightarrow 11: \quad C_{\text{copy}}(1) = 1, \qquad C_{\text{copy}}(T) = 3 \quad (T \geq 2) \tag{A16}$$

In the third world, connectivity is complete again, but exchange copies information rather than transferring a resource, and initially only one element knows it. With only one source in the first round, there is one newly informed recipient. In the second, the two informed elements can inform the remaining two in parallel. All elements then know the information, so no permitted transition remains. The total is three newly informed recipients regardless of any further extension of the window. Repeating a message to an already informed element is not a new completion.

The comparison shows that greater replicability may accelerate goal attainment while reducing subsequent event counts. Activity rates alone therefore cannot rank intelligence or value. It also shows that changing environmental laws or connectivity produces different capacities without assuming money or human preferences. The laboratory can be extended with forgetting, resource replenishment, or partially ordered time, but each addition changes the problem and requires the bounds to be recalculated under the same completion definition.

Binary states carry no implicit preferences, and transfer is not treated as mutual benefit merely because it occurs. If an economic perspective is assigned to the elements, their acceptance and rejection criteria must be defined before event quality is evaluated.

CHAPTER 5 INTELLIGENCE AND EXCHANGE CAPACITY

Service agents and the limits of economic generalisation

Specialising the framework to AI agents permits applications to computing resources, service interfaces, data, and delegated tasks. A product is a cause of exchange and may consist of a right to invoke a model, an analytical result, a processing deadline, or a commitment to complete a task. Promoi describes different offers for accessing these capabilities, while kup records facts concerning commitment, consumption, or completion according to its available content. This structure does not imply that every digital outcome can be sold or that every piece of information may lawfully be copied.

Consider a synthetic example: an agent needs data analysed by a deadline. One provider offers a faster, more expensive service; another offers a lower price with a call limit; a third conditions service on data-use rights. Combining several services may make the outcome possible but adds transfer, verification, and coordination costs. The agent compares paths by task outcome, resources, and authority, so the lowest advertised price is not automatically best. Capability changes if new eligibility evidence arrives or another provider's capacity is reserved.

The proposed research hypothesis is that organising offers and evidence according to their effects on decisions may improve completion under a fixed search budget. It can be tested by comparing policies with equivalent information and authority while changing only how access is organised. A separate experiment examines expansion of feasible contracts or the introduction of reservations. If performance does not improve after integration and verification costs are counted, successfully representing the mechanism in a simulator does not establish an advantage for abstract pricing.

The analysis can then extend to a market of negotiating agents representing different interests. This introduces strategic behaviour, information concealment, reservations used to obstruct competitors, and trade-offs between disclosure and privacy. Neither a capability definition nor a truthful record alone resolves these issues. Every claim about incentives, efficiency, or fairness requires a separate model of the parties, their preferences and information, and their decision mechanism. Execution feasibility must remain distinct from the desirability of the resulting distribution.

The framework extends to families of abstract environments when their elements, transitions, observations, time structure, and resources can be specified. An element's intelligence range becomes a collection of capability profiles across that family, with the environment-selection rule published rather than a neutral universal ranking assumed. Theorems transfer only when mappings between environments preserve their assumptions: a proof for conserved resources cannot be transferred to a world that copies them freely, and a probabilistic guarantee requires preservation of the outcome law. Generality thus rests on specific, testable relations, with abstract pricing providing a precise economic domain in which to investigate and develop them.

In an environment that does not admit this kind of causal order, one can begin with compatibility relations between states instead of paths. Capacity results based on the graph and clock do not transfer to such an environment without reformulation and proof.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Abstract Pricing and Existing Commerce Protocols

The theory must locate itself in relation to existing structures before claiming any contribution. The technical materials discussed in this chapter are considered as they stood on 4 October 2026. An operational specification for exchanging messages is not, by itself, a theorem about efficiency, just as a successful simulation does not establish merchants' willingness to participate. We therefore distinguish conceptual precedents, implementation precedents, and comparable mathematical or empirical results.

In March 2026, OpenAI announced the expansion of ACP to product discovery, with merchants sharing product feeds and promotions with the discovery environment. Abstract pricing therefore cannot base its contribution on the claim that current protocols address payments alone. The proposed object of evaluation is the cost of discovering useful, executable terms within a specified offer space, rather than merely transferring a catalogue into a conversational interface (OpenAI, 2026).

The UCP discount specification defines the order in which discounts are applied, the allocation of their amounts, and certain eligibility and verification conditions. This directly overlaps with the representation of one aspect of offer composition. Discount ordering is therefore treated as a reference case for semantic compatibility, while the proposed model investigates how decision quality changes when offers share scarce resources or their validity depends on distributed evidence (UCP, 2026).

AP2 links an agent's authorization to user instructions, the cart, and payment through verifiable evidence. This helps distinguish permission to act from the outcome produced by the action. This manuscript examines exchange records in terms of their content and readability, without assuming that a signature on payment instructions alone proves that an external service was completed. This is a methodological inference within our model, rather than a criticism of a particular empirical claim (Parikh and Surapaneni, 2025).

These protocols are proposed as execution connectors in a later experiment. The theoretical analysis operates on semantics independent of field names, so that the scientific claim does not depend on a particular interface version. Integration must test whether translating an offer preserves eligibility, budgets, deadlines, and outcomes: syntactic compatibility is weaker than preservation of feasible decisions. A translation may fail in one of these respects while still successfully exchanging messages.

The potential contribution therefore lies in a specific result about decisions under information and execution constraints. Examples include a bounded-loss guarantee after compressing offers, or identifying the evidence needed to rule out an infeasible composition. Renaming a discount, an authorization, or a record may provide shared terminology within the project, but does not establish scientific priority. Every comparative report should document the protocol version, the date examined, and the functionality actually enabled.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

What Do Existing Commerce Benchmarks Measure?

Comparison with agent benchmarks establishes whether the proposed programme addresses a difficulty that can be isolated. ShoppingBench provides shopping tasks whose intentions include budget management, voucher use, and searches involving multiple products; the repository version reports acceptance at AAAI 2026. It can inform the selection of tasks that reflect user intent rather than product-name matching alone. Its reported model results must not be transferred to other, untested versions (Wang et al., 2025/2026).

The ACWorld preprint proposes an environment for buyer and merchant agents, validating actions before updating transaction state and recording interactions for auditable evaluation. A record and a verification engine in our platform are therefore insufficient as a standalone contribution. The relevant comparison concerns the constraints and hidden facts involved, and whether the experiment isolates the cost of obtaining evidence needed for decisions under those constraints (S. Fan et al., 2026).

The E-Commerce Bench preprint evaluates business operation over a long horizon, including sourcing, negotiation, order fulfillment, and cash flow in a simulated environment. Its scope differs from a bounded decision experiment, but provides a reference when extending the analysis to commitments that accumulate over time. Store simulation and multi-round negotiation cannot be presented as unexplored areas (W. Fan et al., 2026).

Comparison of research scopes. The final column describes a proposed programme, not a claim that the subject is entirely absent from the reference.

Reference	Declared scope	Effect to isolate
ShoppingBench	Shopping intentions	Effect of offer terms
ACWorld	Auditable interaction	Evidence sufficiency for decisions
E-Commerce Bench	Long-horizon business operation	Accumulating resource commitments

This comparison does not justify inventing a new task set merely because it uses different terminology. Tasks should reveal a behavioral difference that evaluation of the final purchase alone would miss. An agent may achieve the same outcome after two queries or two hundred, or commit to an offer that happens to be available even though its evidence was insufficient to establish availability. The former measures the cost of reaching a decision; the latter measures decision discipline relative to available knowledge. A large catalogue alone does not demonstrate structural diversity in constraints: most products may share one decision pattern, whereas a small problem may contain intricate dependencies among conditions.

A fair comparison also requires a published mapping between our tasks and the benchmark tasks we can represent. If a shared resource or hidden history prevents such a mapping, the difference should be described precisely and a variant without that element supplied. Increased difficulty after adding the element becomes evidence of its effect within the simulator; a claim of realism still requires independent evidence from an application.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

A Candidate Research Gap: Evidence That Changes Decisions

The candidate research gap concerns the cost of obtaining information sufficient for a high-quality, executable exchange decision. This differs from collecting as much data as possible: a single fact about an exhausted resource allocation may rule out thousands of compositions, while further descriptions of similar products may leave the decision unchanged. A query's value is determined here by its ability to narrow the possible states of the world or improve choice within available resources.

Preference-elicitation research has already investigated reducing the communication required for resource allocation. Huang and colleagues use language-model-based proxies to represent participants' preferences and communicate in natural language within an auction setting; the chapter appeared in the WINE 2025 proceedings in 2026. We therefore do not claim novelty for using a language model to select fewer questions. Our focus is the validity of commitments and its relationship to eligibility, resources, and partial records (Huang et al., 2026).

A preprint by Jiang and colleagues formally analyzes several agent payment protocols, linking authorization to transaction stages and their effects. The existence of this work reinforces the need for specified properties and an explicit threat model when proving consistency. The proposed contribution of kup in this chapter concerns whether the available content is sufficient for a decision and for recovering the information required from it; it is not reduced to asserting that a record is signed or immutable (Jiang et al., 2026).

The first project isolates three sources of uncertainty: the buyer's preferences, the truth of an eligibility condition, and resource availability. Two can be held fixed while the third is studied. The experiment begins with declared objectives and resources that remain fixed during the decision, making the condition-querying problem explicit. Changing resources and time-limited reservations are then introduced. A guarantee from the static model must not be transferred to the dynamic model without an additional assumption about evidence remaining valid or a mechanism preserving the resource.

The contribution criterion can be expressed as a counterfactual question: could an established preference-elicitation or constrained-search method achieve the same result if given the same tools? If so, the contribution becomes an application or benchmark, rather than a new algorithm. If a lower bound on queries is established and attained by a method within a specified class, a theoretical contribution can be defended even when the method uses familiar computational structures.

Impossibility examples should also be examined. If two states agree on every accessible piece of evidence but differ in exchange validity, no policy can guarantee that it distinguishes them using that evidence alone. This restriction identifies what improving a language model cannot repair. The useful intervention is then to modify the space's interface or supply additional evidence, with the cost of that modification included in the economic comparison.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Specific, Falsifiable Hypotheses

A hypothesis register should be established before the main comparisons, specifying the unit of analysis, budget, and primary metric. The first hypothesis is that querying according to evidence's effect on the decision reduces query cost at a specified level of decision quality and execution validity, compared with declared search methods. It does not predict superiority in every environment: equally useful evidence, or a high cost of computing the next query, may make the proposed method slower.

The second hypothesis is that a compressed representation of the offer space preserves choice quality within a known bound in classes satisfying the resource- and value-preservation conditions. The bound is investigated theoretically within its class, while algorithmic performance is evaluated using different generators. The third hypothesis predicts that dependence between evidence arrivals changes performance even when each item's marginal accessibility is fixed. The fraction of messages delivered is therefore insufficient to explain agent success.

Evidence sufficiency is defined through the set of states compatible with knowledge K. Guaranteed decisions are valid in all these states under the static model's assumptions. We therefore use the following expression, where θ is a possible state and F is its set of executable decisions:

$$F_{\mathrm{safe}}(K) = \bigcap_{\theta \in \Theta(K)} F_{\theta} \tag{X1}$$

The decision not to commit is included to avoid assuming that a guaranteed transaction always exists. An experimental objective can be to minimize the cost of evidence needed to reach a declared threshold of decision loss. A policy that rejects every transaction is not successful merely because it violates no constraint: the quality of opportunities forgone must be measured. Similarly, a policy is not rewarded for obtaining high gains through commitments outside its authorization.

The fourth hypothesis concerns congestion: accessibility opportunities may increase while completion rates fall as more agents compete for the same resource. This is examined by holding total supply fixed, increasing demand, and then introducing reservations. The fifth concerns capability measurement: an agent's capability profile should explain performance differences on new tasks after controlling for budget and interface quality; otherwise the index merely summarizes training outcomes.

Confirmatory tests are separated from exploratory analysis. A researcher must not change the primary threshold after observing results and then present that choice as having preceded the experiment. When a hypothesis fails, its failure domain and the method's cost within it are reported, since a lack of superiority may be the strongest evidence that an environment does not justify the additional complexity. The improvement required to justify adoption is declared before testing. A small, measurable difference may be worth less than the cost of integrating the system, and thus fail to support adoption even if it can be reported as a statistical result.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Architecture of an Offer-Space Simulator

The simulator begins with a machine-readable environment descriptor specifying states, transitions, resources, time, and observation. This descriptor is separated from the agent's policy so that different policies can run in the same space. Except in a special reference mode, the agent does not receive the complete true state; it receives observations determined by the access interface. The environment retains the ground truth needed to verify eligibility, resources, and outcomes.

An offer description specifies the product or assortment identity, publisher, adopting party, constraints, resource consumption, commitments, and completion outcome. Conditions can be expressed in a small typed language instead of unrestricted general-purpose code. Types include quantities, durations, exercisable rights, and verification states. Incompatible units are rejected before price optimization: the system must not add a service balance to a count of products as if they were the same resource.

The action cycle consists of proposing a decision, independent verification, reservation where needed, execution, and issuance of an outcome record. The natural-language parser translates user intent into a structured proposal, but grants no authority and changes no authorization limit. The guard rules should remain inspectable even if the language model changes or misinterprets an offer description. Ambiguity in a material condition is returned as a request for evidence or a reason to abstain.

The engine assigns an exchange event an identity distinct from that of a record copy describing it. Several records may concern the same event, and some may be delayed or conflict. Precedence relations and the events on which a decision depended are preserved, together with local time where the environment provides it. A complete causal order is not inferred merely from the order in which messages reach a central server; some events may remain unordered relative to one another.

Execution must handle retried requests without counting a new exchange when a request retransmits the same operation. This does not resolve every form of semantic duplication: a client may create two identities for one commitment. The guards therefore also check consumed rights and allocations. Cancellation or expiration of a reservation produces a new recorded state, so an old reservation record cannot establish that the resource remains reserved.

Two separate simulation interfaces are exposed: one used by the agent and another by the evaluator. The latter contains ground truth and metric verification without leaking them to the former. This separation follows from the experimental objective: measuring decisions under limited information while retaining the ability to explain errors afterward. The observation distributor or delay model can be changed without altering the offer definition or final metric.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Task-Set Design and Controlled Comparisons

The task set begins with a small finite world in which every decision problem can be solved by exact search. Tasks are grouped into families: selecting an independent offer, composing two complementary offers, verifying a single-use right, trading off a deadline against price, and competing for a shared allocation. This organization explains failures by structure rather than relying only on an average across problems requiring different skills. Tasks with no suitable transaction are included to test correct abstention.

A factorial design varies representation quality, evidence completeness, offer-space size, and the tightness of the resource constraint as independently as possible. When comparing a structured interface with a textual interface, semantics, offers, and budgets remain identical. An expansion of Promoi, if tested, is introduced as a separate factor. This distinguishes improvements attributable to lower comprehension costs from those attributable to an actual increase in opportunities.

Comparators include an agent reading offer descriptions, the same agent with structured tools, a similarity-based search policy, an established query policy suitable for the problem class, and the proposed method. All receive the same access opportunities and declared query and computation limits. A solver that knows the complete state is presented as an informational reference bound; its performance gap relative to the agent is not interpreted as a purely algorithmic shortcoming, because the gap also reflects its additional information.

The textual policies and examples available to each model are fixed before testing. If the new method uses a verification tool, the structured comparator also receives that tool; otherwise the algorithm's effect is confounded with access to a constraint guard. A separate comparison of complete system packages from the user's perspective is permissible, provided it is identified as such and not attributed to a single component. Data preparation and rule-authoring costs are included in the later application analysis.

Different seeds from the same generator do not establish generalization. Compositional families absent from development are held out, such as transitions from independent to complementary conditions, or from one resource to two. Performance within and outside the development distribution is reported separately. Some linguistic formulations are also held out, with their semantics known to the evaluator, to test whether the method learned the logic of the conditions or patterns of phrasing.

Every task file is linked to a reference solution or a bound certificate when a complete solution is unavailable. In larger problems, the best solution a researcher finds is not called optimal without a proof or suitable computational certificate. The reference solver's time limit must also be recorded. Problem difficulty may reflect the cost of ideal search, which is part of abstract pricing when the size of the offer space itself impedes decision making.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Defining Metrics and Separating Sources of Loss

The proposed primary metric is the quality of completed exchange within the constraints, with commitment validity and the cost of reaching the decision reported separately. The method's outcome on task t is denoted by value W on a declared scale for that task. This may measure buyer utility or a specified surplus in a synthetic model, but does not justify aggregating different parties' values beyond the assumptions permitting such aggregation. Party-specific metrics are therefore also reported when interests conflict.

Decision loss is defined relative to a reference with a declared information level. With a full-information reference, we calculate the difference between the problem's optimum and the realized value. Comparisons between policies with the same available knowledge instead use outcomes from policies subject to the same interface and budget. The first reference combines information loss and search loss, so these differences should not receive a name implying that they measure the same thing.

$$L_t = W_t^{\mathrm{ref}} - W_t^{\mathrm{method}} \tag{X2}$$

If scales differ across tasks, normalization is chosen using a value range specified before the experiment. Loss must not be divided by a reference value that may equal zero; one may use the difference between declared bounds or report outcomes in their original units within each family. Monetary costs, computation time, and query counts are recorded separately, because fewer text tokens do not necessarily imply lower integration costs or shorter decision delays.

Validity indicators include the rate of invalid commitments, completion after commitment, repeated exercise of single-use rights, and correct abstention. The denominator matters: completion among accepted transactions can rise when a policy rejects most feasible opportunities. The report therefore compares coverage and precision and presents the cost of opportunities forgone. It also distinguishes rejection because no transaction exists from rejection because validity cannot be established. If a numerical violation penalty is incorporated into the score, its value is declared in advance and sensitivity is examined. Violation counts remain visible so that high benefits on other tasks cannot conceal them.

When comparing two methods on the same tasks, we use paired differences, where d is the outcome difference for a task and n is the number of independent units of analysis. The mean estimates the effect size for the declared task set:

$$\overline{d} = \frac{1}{n} \sum_{t=1}^{n} d_t \tag{X3}$$

The mean is accompanied by the distribution of differences and an appropriate uncertainty interval, not a p-value alone. Tail metrics, such as the ninety-fifth percentile of elapsed time, are reported when delay affects validity. Exact enumeration results for a specified finite world are not a sample from that instance and require no confidence interval to establish the computed count. Generalization to other worlds remains an entirely separate question.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Executed Computational Experiment: Capacity in Three Finite Worlds

For this manuscript, exact computational enumeration was performed for three synthetic worlds using a Python program that neither calls language models nor collects market data. Each world contains four elements, each with a one-bit state. At each step, operations between disjoint pairs are permitted: no element participates in two operations during the same step. One exchange event is counted per completed operation, not per record copy describing it. This verifies a numerical model; it does not measure human or agent intelligence.

In the first world, every pair is connected. A unit moves from an element in state one to an element in state zero, changing the former to zero and the latter to one. The initial state contains two units. The second world retains the transfer rule and initial state, but uses a star network centered on one element. The third has a complete network and starts with one unit of information. An operation copies that information to an uninformed element while the sender remains informed; retransmission to an already informed element is not counted as a new success.

Maximum completed exchanges: actual results of exact enumeration in the specified synthetic worlds.

World	One step	Two steps	Ten steps
Transfer; complete network	2	4	20
Transfer; star network	1	2	10
Copying new information	1	3	3

The program enumerated the sixteen states, constructed all permitted transitions including inaction, and applied a dynamic-programming recurrence for horizons from zero to ten steps. It evaluated 528 world-state-horizon combinations. It checked conservation of units under transfer and verified that, under copying, the number of informed elements increased by the number of successful copies. The code and results, including a code fingerprint, were retained with the experiment files for recomputation.

Structure explains the differences without changing the elements' local capabilities. The complete network permits two parallel transfers per step, whereas the star's center limits the count to one. With copying, information reaches every element after two steps, leaving no new informing events to count even though communication remains possible. The number three reflects the chosen definition of completion, not an inability of the network to send further messages.

A count of twenty does not mean that net value was created twenty times: units can move back and forth. The experiment demonstrates why exchange throughput must be separated from growth in a stock or utility. These results establish no general law of spaces; they check the capacity engine's consistency with the rules of three fully enumerated examples. Capacity in any other environment requires definitions of a new event, what each operation consumes, and what its temporal relations permit.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Executed Computational Experiment: Single-Resource Offer Compression

A second experiment used two hundred finite synthetic instances generated with the fixed seed 20261004. Each instance has three buyers, each with eight offers plus the option not to purchase. An offer consumes an integer quantity of one resource from 1 to 20 and yields an integer score from zero to forty. The shared budget ranges from 5 to 30. All these numbers are synthetic generation parameters, not estimates of market behavior.

A solution selects one offer or no purchase for each buyer and sums scores subject to the shared resource cap. Compression partitions each buyer's offers by the integer part of the score divided by three, retaining the offer with the lowest consumption in each class and choosing the higher score when consumption is tied. The no-purchase option remains available. This construction ensures that an offer's representative consumes no more resources and loses no more than three score units.

$$0 \le W^{*} - W^{*}_{\text{compressed}} \le 3 \times 3 = 9 \tag{X4}$$

Exhaustive search over combinations before and after compression. Scores have no monetary unit.

Metric	Computed result
Mean full optimum	89.29
Mean compressed optimum	89.18
Mean loss; maximum loss	0.11; 3
Instances with no loss	185 of 200
Mean options retained	20.44 of 27

The program checked the consumption and score condition for every offer, then enumerated the full and compressed optima. No instance exceeded the theoretical bound, and compression retained approximately 75.7% of the options on average. This fraction is neither a guaranteed target nor a general size bound: offers may occupy different classes, leaving no useful compression. Nor does this enumeration measure language-agent speedups or establish lower operating costs. Integers were drawn from uniform distributions within the declared ranges. All 729 original combinations per instance were examined before excluding budget violations. The results file contains each instance's offers and representatives, allowing inspection without regeneration.

The evidence is limited to correct implementation of the construction in these instances and its computed effect under known, additive scores. There are no differing eligibility conditions within our classes and no complementary benefits across buyers. Introducing those structures requires partitioning classes by compatibility or revising the proof. When values are private and unknown, the system cannot apply the same partition free of cost; value elicitation becomes part of the research question. Minimizing consumption of one resource also does not guarantee simultaneous reductions in two resources.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Statistical Inference and Limits to Generalization

The methodology distinguishes three types of results. The first is exact enumeration of a specified finite space, as in the two experiments in this chapter. The second is estimation of mean performance over a declared distribution of synthetic tasks. The third is an effect on users or institutions in a field setting. More synthetic tasks do not automatically convert a result into evidence of the third type. Nor can field uncertainty be reduced to computational uncertainty removable by rerunning the simulator.

Future language-model evaluations should include repetitions across tasks and model randomness. Methods are paired on the same task and, as far as possible, the same environmental shocks, while the randomness of each call is recorded. Merely specifying a seed does not guarantee exact pairing if the model service does not provide complete control. In that case, the limits of control are described and sufficient repetitions are used to examine observed variation without claiming deterministic reproducibility.

A task family sharing a single template must be treated cautiously: one hundred paraphrases of the same constraint are not one hundred independent economic structures. Uncertainty intervals can be constructed by resampling families when that is the level of dependence. Reports state the number of families, tasks per family, and model runs, allowing readers to distinguish an abundance of textual data from genuine structural diversity in the problems.

The confirmatory experiment's size is chosen after a pilot study estimates variation and identifies the smallest practically important effect. Pilot and test data remain separate. No universal task count guaranteeing statistical power is prescribed here; requirements depend on effect size, outcome variation, and dependence. If cost limits cause early termination, the stopping reason is explained, and researchers avoid selecting a stopping point because it produces the desired difference.

Primary results are reported by task family as well as in aggregate, alongside failure rates, time, and cost. With many comparisons, primary analyses are declared in advance or multiplicity is accounted for, rather than selecting the best metric after inspection. Small, unstable effects call for further validation or narrower claims. A null result with a narrow interval around a practically negligible difference can help avoid unproductive complexity.

Statistical significance alone is insufficient for an economic decision. An algorithm may modestly reduce queries while raising maintenance costs or increasing delays in rare cases. The gain is compared with those costs and the application's acceptable risk level. Publication should remain defensible even if superiority is not established: a well-characterized failure class may reveal a fundamental limitation of the abstraction.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Reproducibility, Data, and Privacy

The reproducibility package comprises the environment descriptor, task generator, policy versions, verification engine, decision logs, and metric-extraction code. Every run is accompanied by a version identifier, resource and time parameters, and seed identifiers. A final-outcome record is insufficient to explain an error concerning an earlier condition; the evidence the agent could actually observe is therefore retained separately from the ground truth available to the evaluator.

Methodological tests cover model properties such as remaining within budget, preventing a single-use right from being exercised twice, and preserving event identity when a request is retransmitted. They also check that removing information from the agent's interface does not remove it from the simulator's ground truth, and that an unadopted offer does not become a commitment merely by appearing in search results. These tests concern the experiment's semantics, rather than file integrity alone.

The first phase uses fully disclosed synthetic data, enabling regeneration and reducing the risk of confusing gaps in research knowledge with restrictions on access to commercial data. These data do not represent the preference distribution of real users. When an experimental collaboration becomes available, the information needed for each question is identified rather than copying every record. A limited eligibility proof may suffice for a decision without revealing the basis of eligibility or a participant's complete personal history.

Moving to human participants requires explaining the experiment, the limits of authorization, and what will be recorded, with an opportunity to withdraw where applicable. An appropriate research body reviews consent and participant-protection procedures in light of the institution and setting. These are proposed research-design requirements, not a determination of legal compliance in any particular country. Operational contact data are separated from analytical data, and merchant secrets or individual preferences are not published as open reproducibility materials.

Small de-identified instances or a generator matching aggregate properties may be released, but removing names must not be assumed always to prevent re-identification. If privacy techniques with mathematical guarantees are used, their model, unit of protection, parameters, and effect on quality must be specified; describing a system as private is insufficient. Verification queries that reveal sensitive information must be included in the policy's cost, even when they incur no monetary expenditure.

For models or services that cannot be redistributed, researchers retain instructions, call settings, execution dates, and outputs that may be shared, and identify what another researcher cannot reproduce exactly. The two computational experiments reported here use only the standard library and preserve all generated instances and results. Their numbers can therefore be recomputed without an external account, market data, or dependence on a changing model.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Sensitivity Analysis and a Taxonomy of Failure

A hypothesis is tested by varying its components, not simply rerunning the preferred configuration. Sensitivity analysis covers the query budget, computation cost, number of offers, resource scarcity, deadline length, and evidence age. Results are presented as performance-cost curves, since a method may excel with a large budget and fail with a small one. These curves must not be replaced by a single measurement selected at a point favorable to the proposed method.

Ablation studies isolate each component's effect: disabling decision-directed query selection, using all evidence without compression, removing resource reservations while retaining verification, or ignoring dependence between message arrivals. Each test changes one component while holding other resources as constant as possible. If an ablation changes available information or increases the effective budget, that change is recorded rather than interpreting the result as the component's isolated effect.

Failure classification begins with the nearest inspectable cause: no executable opportunity, evidence not arriving, misunderstanding a condition, value-estimation error, selecting an invalid composition, losing the resource before execution, or failing to fulfill a commitment. Several causes may occur along one trajectory, so the causal sequence is retained instead of forcing every failure into a single category. Classification should be reviewed independently of the model that produced the decision.

Legitimate adversarial cases are constructed within the simulator: linguistically similar offers differing in a decisive condition, an otherwise valid right whose exercise window has expired, and two messages that were correct when issued but cannot both remain valid after a state change. Unauthorized instructions embedded in an offer description are tested as content carrying no authority. The objective is to measure preservation of the boundary between information and authorization, not to allow a seller's descriptions to modify the buyer's policy.

Countervailing results warrant particular attention. Compression may reduce the number of options while increasing the cost of interpreting their representative, and evidence may arrive quickly but become stale before verification finishes. Reservations may improve completion for some agents by denying opportunities to others. Effects are therefore measured for all relevant parties and for actual resource use, not only through the success rate of the agent selected for observation.

Failures are re-examined using small examples solvable manually or by enumeration, so the simulator's size does not conceal a definitional error. If a state cannot be reconstructed, that is an auditability limitation that must be reported. The number of unclassified cases is published with examples that may be shared. Completing an attractive table does not justify inventing a cause unsupported by the available record.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Experimental Economics of Participation and the First Field Application

Computational decision quality alone does not establish the viability of an offer market. The buyer needs benefits exceeding learning and verification costs; the seller needs returns justifying the publication of terms and management of new commitments. The first field application should identify the operational problem it addresses, such as trading off service deadlines or available capacities, rather than testing a broad promise that making prices programmable will make every price more efficient.

A proposed starting point is a service with measurable terms and a small number of voluntarily participating providers. Offers, costs, and execution times are recorded in advance, and the agent initially operates in a recommendation mode reviewed by an operator before any actual commitment. This phase measures recommendation quality and cost without attributing execution success to transactions that were never executed. Limited authorization may follow if verification procedures demonstrate consistency and the parties accept the experimental scope.

The field test is prospective: comparators and metrics are specified before outcomes are collected, and inclusion and exclusion rules for requests are disclosed. Where the context permits, eligible requests can be randomly assigned to equivalent procedures, while cross-arm effects through shared resources are prevented or accounted for. If both arms compete for the same stock, interference changes outcomes; the experiment must be designed around that feature rather than ignoring it.

Buyer surplus and seller operating profit are measured using agreed definitions, including the costs of delays, returns, support, computation, and publishing offers. A lower price alone is not value creation: it may represent a transfer from seller to buyer or compensation for reduced service. The report examines whether an offer saves real resources or improves preference alignment, who receives the gains, and who bears the costs.

Participation indicators include continued offer publication by merchants, adoption of buyer-proposed offers, the number of revisions required, and repeat use after experimental incentives end. The effect of a temporary subsidy must be separated from that of the design itself. Evaluation also examines whether the rules disproportionately benefit merchants capable of technical integration: the system may reduce buyers' search costs while raising entry barriers for small suppliers.

Participants' incentives remain a separate issue from execution integrity. A participant may exaggerate scarcity, delay disclosure of a condition, or publish unattractive offers to occupy search results. The application therefore begins with declared behavioral assumptions, then observes and studies deviations without claiming that a constraint guard guarantees truthfulness. Moving to an open market requires separate analysis of manipulation, participation, equilibrium, and the distribution of benefits.

CHAPTER 6 RESEARCH PROGRAMME AND EVALUATION

Implementation Roadmap and Criteria for Scientific Contribution

A preliminary six-month programme is proposed for a small team, with its scope reviewed against actual resources. Month one establishes operational definitions, maps the literature, specifies the environment, and develops counterexamples. Its output is a defined problem whose assumptions and objective an independent expert can explain without referring to the project's commercial vision. Additional topics are not introduced at this stage merely because abstract economics is a broad title.

Months two and three focus on one algorithm concerning evidence sufficiency or offer compression, with a limited proof and tests in a finite world with known solutions. Deliverables include a strong comparison and cost report, not merely a demonstration interface. Month four introduces tests absent from development and independent review of counterexamples. If no practical benefit emerges, the class in which the method is useful is redefined, or its limitations are clearly reported.

Months five and six prepare the reproducibility package, paper draft, and a limited integration experiment. One extension is selected: resources that change over time, offers with two resources, or a natural-language interface. Combining these extensions at once makes outcomes harder to interpret and increases the likelihood of attributing errors to unexamined components. Each extension may become a later programme if it exposes a distinct research question.

A theoretical contribution is persuasive when it provides a bound, algorithm, or structural characterization beyond restating an established result, while explaining when its assumptions are necessary. An empirical contribution is persuasive when its design isolates a reproducible effect against serious comparators and explains failures. A tool or benchmark must demonstrate that it enables questions that could not previously be evaluated with comparable precision, rather than merely possessing a new name.

The two small experiments in this chapter do not make this manuscript a field-validated study; they check the implementation of selected definitions and bounds in computed instances. The feasibility of eliciting values, changing resources, merchant willingness, evidence privacy, strategic manipulation, and integration costs remain open questions. These limitations constrain current claims while also identifying where a substantive scientific contribution may emerge as the programme progresses.

Abstract economics can subsequently extend comparison to exchanges of computing services, data, and tasks among agents, and then to synthetic worlds with different resources, temporal structures, and transformation laws. Generality is assessed by whether a result retains its meaning across environments, with changes in assumptions explicitly identified. In this sense, the theory's breadth becomes a programme for transferring controlled results and testing their limits. Abstract pricing remains its initial domain for investigating the relationship between opportunities, information, commitment, and value.

SYNTHESIS AND RESEARCH FINDINGS

The theoretical structure of abstract pricing

The structure of the theory can now be expressed as relationships among distinct levels. The description of the environment and its elements determines what can happen. The perspective on value determines how particular outcomes are compared. The offer space determines which proposed agreements can be discovered or adopted. Evidence determines what an element can know about the state, while its policy and resources determine what it can choose and execute. These levels interact, but every transition between them requires an explicit assumption or mechanism.

In the economic specialization, a product with potential value is insufficient to determine a price or a transaction. Exchange requires a party able to make a commitment, another willing to accept it, resources that permit execution and information that arrives in time. Representing offers independently facilitates reuse and composition of terms across multiple situations. An actual commitment nevertheless remains attached to an entity that adopts it and to resources or rights that entity can control. Multiplying offers therefore does not multiply inventory or create new funding.

The representation-neutrality proposition supplies an important benchmark. If a change of representation preserves the option set, its semantics and the costs of access and solution, it does not raise the optimal value achieved by an ideal decision maker. This directs research towards what actually changes when a benefit appears: less search, fewer interpretation errors, better verification or newly available agreements. Commercial claims about improved pricing can then be decomposed and tested instead of being combined in an opaque indicator.

The kup concept serves a complementary epistemic function. An event may have occurred while evidence of it remains unavailable. Available evidence may also fail to distinguish two histories with different consequences. This motivates a separate study of record sufficiency relative to a particular question. What establishes that a party authorized payment may not establish eligibility for a discount or delivery of a service. Decision accuracy depends on which alternative states can be ruled out, rather than simply on the quantity of recorded text.

Value remains tied to a specified perspective throughout this structure. Seller profit can rise while buyer surplus falls without a change in real costs. Trading volume can increase because the same item changes hands repeatedly, without better satisfaction of a need. Any comparison must therefore declare its measure, purpose and counting unit. The theory requires no single measure valid in every space, but it does require consistent measures whenever they support an ordering or a claim of gain.

This structure supports a programme that can be pursued in separate parts. An offer-compression method can be tested before a complete market is built; evidence sufficiency can be examined in a finite world before being connected to a live protocol. Success in one component does not establish the remaining claims. Each contribution requires evidence appropriate to its level.

SYNTHESIS AND RESEARCH FINDINGS

Decomposing decision loss and exchange costs

We propose a simple accounting decomposition to clarify where a gain from abstract pricing might arise. Assume a fixed state, a finite offer set and a common scalar measure for comparison within the task. Let F be the set of feasible decisions and S the set of decisions that the agent has discovered and verified as valid, with S a subset of F. Let x be a decision selected from S. Include a no-commitment option so that these sets are nonempty.

$$W_F^* = \max_{y \in F} W(y), \qquad W_S^* = \max_{y \in S} W(y) \tag{S1}$$

Let c be the additional cost of search, verification and execution, converted into the same evaluation unit, nonnegative and not already included in W. Collecting the following terms then gives an exact identity rather than a statistical estimate.

$$W_F^* - [W(x) - c] = (W_F^* - W_S^*) + (W_S^* - W(x)) + c \tag{S2}$$

The first term is an availability or discovery loss, the second a selection loss within the available set, and the third the cost of the process. All are nonnegative under the stated assumptions. The proof uses the facts that maximizing over a larger set cannot reduce the optimum and that the optimum within S is at least the value of any element of S, followed by cancellation of the opposing terms on the right. If some costs are negative because of a subsidy, financial transfers must be separated from real costs instead of applying the nonnegativity condition mechanically.

The decomposition explains why an apparently strong linguistic improvement may fail. It may increase the number of discovered offers and reduce the first term while doubling verification costs. Alternatively, an agent may select the best offer from a short list whose representation omits conditions affecting validity. The assumption that S is contained in F then fails, and the identity cannot be interpreted in this way. A platform may also expand F itself by enabling a new agreement. In that case, two explicitly stated benchmarks must be compared, and the whole difference cannot be attributed to algorithm quality.

This decomposition is descriptive within a state known to the evaluator. It does not identify the causal contribution of each term, because the search policy selects S in response to costs and expectations. A causal interpretation requires changing a specified component in a controlled experiment or constructing a structural model of their interaction. The agent generally cannot calculate the first term because it does not know all offers. A simulator can provide it, or it can be bounded above and below.

The identity supports clearer reporting: present the quality of discovered opportunities, selection quality and tool costs separately, followed by net value. If the measures cannot be placed in a common unit, replace the sum with a multidimensional profile and state the trade-offs, rather than inventing a monetary price for every effect.

SYNTHESIS AND RESEARCH FINDINGS

Abstract pricing among pricing problems

The research subject is clarified by the question it seeks to answer. Pricing over time studies how terms change with state, demand or inventory. Bargaining studies how parties move towards an agreement, while mechanism design studies how selection and payment rules affect disclosure and behaviour. Abstract pricing here focuses on making exchange terms an independent object of representation, composition, access and verification, while retaining the economic constraints on those terms.

Problem	Central variable	Relationship to this framework
Price adjustment	A monetary amount or price schedule	A special case of offer modification
Bargaining	Messages, concessions and acceptance rules	A mechanism for selecting a point in the offer space
Allocation	Who receives which resources	A shared constraint on offer execution
Mechanism design	Incentives and decision and payment rules	Determines behaviour under private information
Abstract pricing	Representation, composition and availability of terms	Connects these problems to execution

These boundaries describe the scope of this monograph, not an exclusive classification of disciplines. An offer can simultaneously have a variable price, multiple attributes and negotiated terms. Separating offer representation from the product does not permit product quality or delivery to be ignored; those attributes remain part of the comparison. An item's name may change while the contract preserves the same economic meaning, or its name may remain fixed while a different warranty changes what is being exchanged.

The connection to information economics and bounded rationality is direct. Stigler examined the cost of searching for economic information, and Simon studied rational choice under limits on processing capacity. We use that perspective to characterize the costs of offer interfaces and verification, without attributing the idea of bounded search to the present framework (Stigler, 1961; Simon, 1955).

A scientific contribution should therefore be formulated as a specific result: an algorithm that improves performance under a declared budget, a bound linking representation size to decision loss, or a classification of situations in which a sufficient execution certificate is impossible. A broad name for an assemblage of existing techniques may supply an organizing language, but its explanatory or computational usefulness still requires separate evidence. A comparison is stronger when it identifies both the class in which a method works and the class in which it brings no improvement.

Offerato could provide an engineering application of these questions and, later, a source of evaluation data if appropriately controlled real use becomes available. The manuscript assumes no undocumented partnerships, customers or experimental transactions. Commercial product claims and research claims should be kept distinct so that each can be verified on its own terms.

SYNTHESIS AND RESEARCH FINDINGS

The first research problem and its boundaries

The proposed first project is decision making with sufficient evidence in a finite offer space. Preferences are initially fixed, while some eligibility conditions or resource requirements may be unknown. The agent receives queries of known cost that reveal truthful facts. It must choose a decision guaranteed to be valid across the states compatible with its evidence, or abstain, while reducing both value loss and query costs. This design confines uncertainty to knowledge relevant to execution, rather than simultaneously combining preference learning, bargaining and market change.

The initial model can use two kinds of condition: a right that cannot be used twice and an aggregate limit on a shared resource. The offer language specifies how conditions depend on records of earlier exchanges. Each query has a clear scope, such as reading a right, checking a balance or revealing a dependency between two offers. A freely phrased natural-language question is not assumed to correspond automatically to a trustworthy query. Interpreting questions is a separate layer that can be added once the basic model is fixed.

The theoretical question concerns the gap between the fewest queries sufficient in a particular instance and the queries used by an algorithm that does not know the instance in advance. A short certificate may exist but be difficult to find, or the worst case may require every condition to be checked. Attractive examples with fast verification are therefore insufficient. Hard families must be constructed and parameterized, and algorithm behaviour measured as the number of offers and the degree of resource sharing increase.

An algorithm might estimate each query's effect on eliminating decisions or narrowing the value gap. Computing that effect exactly may itself be expensive, so the cost of selecting a query belongs in the same budget. Evaluation should compare an approximate greedy policy with a strong reference policy feasible for small instances. A fully informed benchmark can show the upper limit of available opportunities, but it must not be presented as a competitor with equivalent information.

We expect the benefit of targeted queries to depend on the dependency structure among conditions. This is a hypothesis, not a result. If every field is necessary or full access is inexpensive, the advantage may disappear. Conversely, local certificates may save substantial cost when a decision's validity depends on only a few active constraints. Results should exhibit both patterns and explain improvements through separate measures of search, verification and execution.

The success criterion is a reproducible formulation and a comparable result, even if that result is an impossibility bound or a documented failure of an initial hypothesis. State changes, distributed authorization, untrustworthy evidence and private preferences should be added afterwards. Each extension requires revised guarantees and a new account of what the agent can know.

SYNTHESIS AND RESEARCH FINDINGS

Limits of generalization beyond economic spaces

The broad definitions of product and intelligence permit physical, biological and digital interactions to be described in one language, but the level of unification must remain explicit. The similarity between transferring a resource and delivering a service supports a shared analysis of constraints and events. It does not establish that a physical party has economic preferences or that interaction laws serve a contractual function. Where an application lacks intention or choice, we use the mechanism actually available to the element without adding hypothetical freedom of action.

To transfer a result from a market to a computing network, specify a map between states, operations and resources. If the map preserves the feasibility and costs of operations and the outcomes of tests, reasoning about that structure can transfer. If a resource that was scarce becomes copyable in the new application, funding, capability and compression conditions may change. The parts of the proof that relied on scarcity must then be identified and re-examined.

Designed universes with different laws are useful for constructing counterexamples. A model may allow information to be copied without consuming the original, permit at most one interaction per element per step, or make time a precedence order without numerical duration. Each choice gives capacity a different meaning. An exchange rate per second is undefined unless the model supplies seconds or a duration function, although completed exchanges within a window of ordered events can still be counted.

Nor is every abstract description guaranteed to have a consistent realization. Conditions may require and prohibit the same transition, require an unavailable resource, or permit unboundedly many exchanges without time advancing. These are properties worth detecting computationally and logically. The capacity theorems in this monograph carry assumptions that exclude or classify some such cases. They are not physical laws governing every possible universe.

A further difficulty arises when measuring intelligence over a family of spaces: the task distribution and environmental weights alter the ranking of elements. One element may succeed at transferring conserved resources but fail at composing information, while another displays the reverse pattern. A capability profile documents this difference more fully than an unjustified total ordering. If an application needs a single score, its weights are part of the application's purpose and must be disclosed.

The scientific value of generalization depends on a transferable result: one proof serving several classes, one simulator testing several laws, or a concept revealing a previously hidden error in an application. Declaring a unity of sciences without specifying the transfer map and its conditions supplies no testable prediction. This study therefore prioritizes the economic specialization, followed by controlled extensions from it.

SYNTHESIS AND RESEARCH FINDINGS

Established results and open questions

The manuscript proves results within specified models and states their limits. Representation neutrality requires preservation of meaning, information and costs. Offer compression requires replacement maps that preserve constraints. Monotonicity of the value of evidence assumes truthful evidence and a fixed state. Capacity bounds require a declared event unit and temporal structure. The arguments can be reconstructed from these assumptions without accepting the extended definition of intelligence or any commercial claim about the project.

Type of contribution	Status in this edition
Value-space and abstract-pricing language	Proposed conceptual and mathematical framework
Neutrality, compression, evidence and capacity	Conditional results with elementary proofs
Computations on finite worlds and instances	Limited synthetic checks already performed
Algorithmic superiority with real agents	Hypothesis requiring a separate experiment
Viability of a large offer market	Open economic and engineering question
An established general definition of intelligence	Not empirically validated by this work

The originality of individual results still requires a specialized literature review when each paper is prepared for publication. Relevant precedents exist in preference elicitation, mechanism design, compositional contracts, resource theory and agent systems. Something that appears new may be a direct application of an established result in different terminology. The contribution should then identify what the application adds in its model, guarantee, data or problem, while crediting the earlier tools.

Future empirical results require publication of instance descriptions, measurement rules and generation seeds, with data used to develop an algorithm separated from its test data. Changes in model versions, protocols and execution tools must be reported because they can alter the result. Real-market comparisons must also measure seller responses and the costs of integration and withdrawal, instead of relying solely on the success rate of purchase requests.

The criterion for continuing the programme is its ability to explain a difference, reduce a cost or establish a previously unclear bound. If independent offer representation brings no improvement after opportunities, constraints and resources are held fixed, that is an important result narrowing the hypothesis. If improvement appears, its mechanism and the cost of achieving it must be identified. The theory can develop through accumulated small results and explicit revisions.

Abstract Economics is the title of the overall project in this edition; abstract pricing is its specific research subject. The preceding pages provide initial material for several narrower studies, beginning with evidence sufficiency and offer-space compression, then extending to coordination, incentives and exchange capacity. The manuscript remains open to revision through peer review, recomputation and subsequent experiments.

REPRODUCIBILITY APPENDIX

Computing capacity in finite worlds

This appendix specifies the synthetic computation so that the results can be reimplemented without an external service or language model. A world state is a vector of four zeros and ones. In the transfer world, one means possession of a unit of a conserved resource; in the copying world it means possession of a copy of information. An undirected network specifies possible channels. Every operation is directed from an element with a one to an element with a zero in the state before the step.

Enumerate every set of operations in which no element participates twice within a step. Add a waiting operation that consumes one step and completes no exchange. Transfer changes each pair from 10 to 01; copying changes it from 10 to 11. Apply all operations in a set to the same previous state, so that a copy received during the step cannot be used in another simultaneous operation. The number of operations in the set is the transition reward.

$$C(s,0)=0, \qquad C(s,t)=\max_{(s',q)\in T(s)}[q+C(s',t-1)] \tag{S3}$$

T(s) comprises pairs consisting of a successor state and a completion count. Compute the recurrence for all sixteen states and horizons from zero to ten steps, storing intermediate results. No probabilistic approximation is required: every set is finite and can be enumerated. The transition reward must be retained with its successor state, because reaching the same state through different operations does not justify discarding what was completed along the way.

Implementation check	Required relationship
Every transfer transition	The number of ones is conserved
Every transition copying q units	The number of ones increases by q
Any world at horizon t	Capacity does not exceed 2t
Star network at horizon t	Capacity does not exceed t
Copying from k initial ones	Capacity does not exceed 4 minus k

The program performs these checks during enumeration and stops if any fail. It then compares the entire computed sequence for the selected initial states with the analytical formula at every horizon. This produces 528 state-horizon-world evaluations with bound checks, and 33 initial-state values matching the analytical formulas. Agreement verifies the example and its implementation over the enumerated domain; it does not validate another world generator or every possible exchange law.

Extending the model to five or six entities changes the state set and matching bounds and increases the search space. Nonuniform time requires a revised recurrence and an analysis of cycles, rather than merely changing the horizon. These details belong to the experimental definition and must be reported in any reproduction.

REPRODUCIBILITY APPENDIX

Generating offer instances and checking compression

The second experiment uses a deterministic sequence of instances once the random seed, 20261004, is fixed. Each instance has three buyers, each with eight generated offers and a zero-consumption, zero-value abstention offer. Each offer consists of integer resource consumption and integer value. Consumption is drawn uniformly from the integers 1-20, value from 0-40, and the shared budget from 5-30. These distributions are not estimates of a real market; they fully specify a synthetic test family.

Values are divided into buckets of width three. Within each bucket the algorithm retains the offer with the lowest resource consumption, breaking ties in favour of higher value. This preserves the abstention offer and gives every original offer a representative whose consumption is no greater and whose value loss is less than three units. Each buyer is treated separately: the algorithm transfers neither rights nor offers between buyers.

$$b(o) = \left\lfloor \frac{w(o)}{3} \right\rfloor, \qquad a(\phi(o)) \le a(o), \quad w(\phi(o)) \ge w(o) - 3 \tag{S4}$$

Compute the optimum by enumerating every triple that chooses one offer per buyer, discarding triples that exceed the budget. Before compression there are 729 possible combinations per instance. Repeat the enumeration after compression and store both menus, the budget, the two optimal values, the gap and the number of retained offers. Check that the gap is nonnegative and at most nine units, the sum of the individual loss bounds in this model.

Repeat generation and solution 200 times while continuing the same generator sequence. A seed alone is insufficient to reproduce the details: the generator type, order of random draws and tie-breaking rule are also required. The implementation uses Random from the Python standard library and generates the buyers' offers in order, followed by the budget. The resulting proportions do not generalize to a different generator or to distributions chosen deliberately to produce worst cases.

The results file retains all instances, including the fifteen with some value loss and the 185 that preserve the optimum. These records allow every numerical claim to be checked instead of relying on the mean alone. A code fingerprint also links the results to a specific implementation. A later sensitivity analysis could vary bucket width, buyer count and the dispersion of consumption, but it is not part of the computations reported here.

This test measures neither a language model's search costs nor the cost of building representatives when values are private or unknown. Its narrower purpose is to check that the implementation satisfies the proposition's conditions and to exhibit its behaviour on a declared set. Practical usefulness requires an algorithm operating with information that can actually be obtained, together with an accounting of acquisition costs.

REFERENCES

Theoretical and economic references

These references locate the framework and the tools on which it draws. Article details identify the original publication; links may lead to author or university copies.

Bob Coecke; Tobias Fritz; Robert W. Spekkens (2016). A mathematical theory of resources. Information and Computation, 250:59-86. DOI: 10.1016/j.ic.2016.02.008. https://arxiv.org/abs/1409.5531

On transformations, composition and resource orderings without assuming a universal monetary measure.

Jean-Charles Rochet; Jean Tirole (2003). Platform Competition in Two-Sided Markets. Journal of the European Economic Association 1(4), 990-1029. https://doi.org/10.1162/154247603322493212

On interaction between the sides of a platform; it does not imply automatic success for an offer market.

John F. Nash, Jr. (1950). The Bargaining Problem. Econometrica 18(2), 155-162. https://doi.org/10.2307/1907266

On axiomatic bargaining; the example derived in the text is limited by its stated assumptions.

Judea Pearl (1995). A Causal Calculus for Statistical Research. Pre-proceedings of the Fifth International Workshop on Artificial Intelligence and Statistics, PMLR R0:430-449; reissued by PMLR in 2022. https://proceedings.mlr.press/r0/pearl95a.html

On the distinction between observation and intervention. Causality in this manuscript is an explicit modelling assumption.

Leslie Lamport (1978). Time, Clocks, and the Ordering of Events in a Distributed System. Communications of the ACM, 21(7), 558-565. https://www.microsoft.com/en-us/research/publication/time-clocks-ordering-events-distributed-system/

On event ordering in distributed systems. The manuscript does not impose a shared physical clock.

Roger B. Myerson; Mark A. Satterthwaite (1983). Efficient Mechanisms for Bilateral Trading. Journal of Economic Theory 29(2), 265-281. https://doi.org/10.1016/0022-0531(83)90048-0

On limits to combining efficiency, incentives, participation and budget requirements in bilateral trade.

Simon, H. A. (1955). A Behavioral Model of Rational Choice. The Quarterly Journal of Economics, 69(1), 99-118. https://doi.org/10.2307/1884852

On bounded rationality and the relationship between choice and processing constraints.

Stigler, G. J. (1961). The Economics of Information. Journal of Political Economy, 69(3), 213-225. https://doi.org/10.1086/258464

A foundational account of the cost of searching for economic information.

REFERENCES

Offers contracts and agent capabilities

This group includes direct precedents in offer and contract representation, approximation and agent evaluation. Rephrasing their results in the terminology of this manuscript does not make those results new.

Alexander S. Klyubin; Daniel Polani; Chrystopher L. Nehaniv (2005). Empowerment: a universal agent-centric measure of control. Proceedings of the 2005 IEEE Congress on Evolutionary Computation, 1, 128-135. https://uhra.herts.ac.uk/id/eprint/282/

An information-theoretic measure of control, distinct from the value-exchange scores proposed here.

Christos H. Papadimitriou; Mihalis Yannakakis (2000). On the Approximability of Trade-offs and Optimal Access of Web Sources. Proceedings of the 41st Annual Symposium on Foundations of Computer Science, 86-92. DOI: 10.1109/SFCS.2000.892068. https://www.cs.purdue.edu/homes/yexiang/courses/18fall-cs590/papers/papadimitriou2000.pdf

On approximating trade-offs; the text distinguishes this approximation from exact dominance in resources.

Craig Boutilier; Holger H. Hoos (2001). Bidding Languages for Combinatorial Auctions. Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI-01). https://www.cs.toronto.edu/~cebly/Papers/bidlang-abs.html

On languages for specifying bids and bundles, where constraint semantics remain central to representation.

Jesper Andersen; Ebbe Elsborg; Fritz Henglein; Jakob Grue Simonsen; Christian Stefansen (2006). Compositional specification of commercial contracts. International Journal on Software Tools for Technology Transfer, 8, 485-516. DOI: 10.1007/s10009-006-0010-1. https://researchprofiles.ku.dk/en/publications/strongcompositional-specification-of-commercial-contractsstrong/

A direct precedent for contract composition. The manuscript does not claim to have invented compositional contracts.

Martin Bichler; Jayant Kalagnanam (2005). Configurable Offers and Winner Determination in Multi-Attribute Auctions. European Journal of Operational Research 160(2), 380-394. https://doi.org/10.1016/j.ejor.2003.07.014

Shane Legg; Marcus Hutter (2007). Universal Intelligence: A Definition of Machine Intelligence. Minds and Machines, 17(4), 391-444. https://arxiv.org/abs/0712.3329

A reference for comparing definitions of intelligence across environments. Value-exchange intelligence here is a proposed stipulation.

Stuart J. Russell; Devika Subramanian (1995). Provably Bounded-Optimal Agents. Journal of Artificial Intelligence Research, 2, 575-609. https://arxiv.org/abs/cs/9505103

On evaluating policies in relation to resources and computational constraints.

REFERENCES

Commerce protocols and recent evaluation

The technical review cutoff is 4 October 2026. This list distinguishes specifications and official announcements from peer-reviewed publications and preprints, and identifies specification versions where available.

David Huang; Edwin Lock; Francisco Marmolejo-Cossío; David C. Parkes (2026). Accelerated Preference Elicitation with LLM-Based Proxies. Web and Internet Economics (WINE 2025), LNCS 16266, pp. 631-648; DOI:10.1007/978-3-032-18660-7_34. https://link.springer.com/chapter/10.1007/978-3-032-18660-7_34

Conference paper.

Jiangyuan Wang; Kejun Xiao; Qi Sun; et al. (2025/2026). ShoppingBench: A Real-World Intent-Grounded Shopping Benchmark for LLM-based Agents. arXiv:2508.04266v4; accepted at AAAI 2026 (repository metadata). https://arxiv.org/abs/2508.04266

Updated repository version; the page reports acceptance at AAAI 2026.

Ke Jiang; Mohan Yu; Yuan Chang; et al. (2026). A Formal Analysis of Agent Payment Protocols. arXiv:2609.00060, submitted 30 August 2026. https://arxiv.org/abs/2609.00060

Preprint. Its results concern the stated consistency models and do not amount to a list of demonstrated production vulnerabilities.

OpenAI (2026). Powering Product Discovery in ChatGPT. OpenAI, 24 March 2026. https://openai.com/index/powering-product-discovery-in-chatgpt/

Primary technical source.

Shicheng Fan; Mingdai Yang; Duohao Wang; et al. (2026). Agentic Commerce World: An Auditable and Verifiable Environment for Vibe Commerce. arXiv:2608.02441v1, 3 August 2026. https://arxiv.org/abs/2608.02441

Research preprint.

Stavan Parikh; Rao Surapaneni (2025). Powering AI commerce with the new Agent Payments Protocol (AP2). Google Cloud Blog, 16 September 2025. https://cloud.google.com/blog/products/ai-machine-learning/announcing-agents-to-payments-ap2-protocol

Primary technical source.

Universal Commerce Protocol (2026). Discount Extension. UCP specification, version 2026-08-25. https://ucp.dev/2026-08-25/specification/shopping/extensions/discount

Primary technical source.

Wei Fan; Xinjie Shen; Xudong Guo; et al. (2026). E-Commerce Bench: Evaluating LLM Agents on Long-Horizon Autonomous Business Operation. arXiv:2608.30730v1, 31 August 2026. https://arxiv.org/abs/2608.30730

Research preprint.