State-Contingent Claim Infrastructure

One Payoff Function, Many Wrappers, and the Oracle–Collateral–Settlement Stack

Grant Stenger
Working paper

June 2026

Abstract

Prediction markets, insurance policies, derivatives, parametric risk products, catastrophe bonds, automated market-maker claims, and many tokenized instruments are institutionally different wrappers around a common mathematical object: a payoff function over future states of the world. This paper develops a theory of the infrastructure that implements that object. A state-contingent claim is not merely a function $g : \Omega \to \mathbb{R}$; it is an implemented pair (g, ρ), where ρ specifies the oracle, data permissions, legal form, collateral rule, margin process, participant restrictions, liquidity mechanism, dispute procedure, tax/reporting treatment, and settlement rail that deliver the payoff in practice. The paper's central distinction is therefore between the *payoff layer*, where prediction, insurance, and derivatives converge, and the *wrapper layer*, where purpose, law, liquidity, and admissibility remain different.

The paper treats payoff convergence as the starting observation, not as the main result: event contracts, indemnity insurance, derivatives, and parametric claims are restrictions of the same payoff-function technology, while their categories differ by institutional signatures. The central object is the delivery operator. Oracles, collateral, settlement, legal interpretation, dispute resolution, and access rules transform a target payoff into the payoff actually delivered. Infrastructure is modeled as a vector of reusable modules–specification, oracle, identity, collateral, margin, compliance, market making, settlement, dispute resolution, access, and reporting–so a module improvement lowers implementation costs across all claim categories that use the module and can shift category boundaries. The main theorem gives the timing duality of information. Data used for ex ante specification or ex post verification expands contractible payoff space and lowers dispute cost; the same data revealed publicly before pooling or trade can destroy risk-sharing value by eliminating the uncertainty that insurance exists to transfer. Additional results formalize the public-good price problem and show that payoff form alone cannot determine welfare or regulation: the same g can be hedge, insurance, gambling, speculation, governance, manipulation, or public-good information depending on holder, purpose, leverage, access, and externalities.

The paper deliberately does not re-solve the valuation, liquidity, admissibility, or agentic transaction-cost problems assigned to the companion papers. Instead it supplies the institutional substrate beneath them. The market span of Paper 1 becomes an implemented payoff; the liquidity regimes of Paper 2 become ways to support that payoff; the admissibility frontier of Paper 3 becomes a set of purpose and access constraints on that payoff; the agentic search problem of Paper 5 becomes the problem of discovering not only a valuable direction but a cheap, liquid, admissible wrapper. The conclusion is convergence without equivalence: the infrastructure of state-contingent claims will become increasingly shared, but welfare, liquidity, and legal treatment cannot be read from payoff shape alone.

Keywords: state-contingent claims; prediction markets; insurance; derivatives; parametric insurance; oracles; collateral; settlement; market design; financial infrastructure; public-good information.
JEL codes: D47, D52, D53, D61, D82, G13, G18, G22, L51.

Contents

1 Introduction

A prediction market, an insurance policy, and a derivative are usually taught as different institutions. One elicits beliefs, one transfers losses, and one prices a function of an underlying asset or index. They live in different legal categories, are distributed through different intermediaries, are collateralized in different ways, and are supervised by different regulators. The premise of this paper is that these distinctions are real but not primitive. At the payoff layer, all three are the same object: a transfer whose value depends on the realized state of the world.

An event contract pays $\mathbf{1}_A(\omega)$. An indemnity policy pays a function of realized loss $L(\omega)$, usually clipped by deductibles, limits, exclusions, and insurable-interest constraints. A derivative pays $h(P_T(\omega))$. A parametric disaster contract pays $A\mathbf{1}_{\{Z(\omega)\geq z^*\}}$. A catastrophe bond pays a coupon and principal schedule contingent on industry loss, modeled loss, or a physical trigger. A tokenized real-world-asset claim pays whatever its legal wrapper, oracle, and settlement system succeed in delivering. The common object is a payoff function $g : \Omega \to \mathbb{R}$. The difference is the wrapper that makes g observable, enforceable, collateralized, liquid, and permissible.

This observation is old in one sense. Arrow–Debreu theory has always represented commodities by state and date (Arrow and Debreu, 1954; Debreu, 1959); Radner studied sequential securities in time (Radner, 1972); derivatives theory prices functions of future states (Black and Scholes, 1973; Merton, 1973); insurance economics studies risk transfers under information and incentive constraints (Arrow, 1963; Rothschild and Stiglitz, 1976); prediction-market design studies securities over events and beliefs (Hanson, 2003, 2007; Wolfers and Zitzewitz, 2004). The new issue is institutional. If the cost of specifying, verifying, collateralizing, complying with, distributing, and settling a state-contingent payoff falls by orders of magnitude, then product boundaries that once looked fixed become endogenous. A claim that was too small for a derivative desk, too objective for indemnity insurance, too specialized for an exchange listing, and too regulated for an event venue may become feasible through a shared oracle–collateral–settlement stack.

The thesis can be stated in one sentence:

> *Prediction, insurance, derivatives, parametric products, and many tokenized instruments converge at the payoff layer, while remaining distinct at the wrapper layer; infrastructure determines which wrappers can cheaply implement which state-contingent claims.*

This paper is Paper 4 in the Trillion Markets research program. Paper 1 studies which residual payoff directions are valuable enough to create as markets. Paper 2 studies how long-tail claims can be quoted through factor hedging, residual diversification, or subsidy. Paper 3 studies which technically feasible claims should be allowed, restricted, subsidized, or prohibited. Paper 5 studies how agents and computation lower transaction costs and search over candidate claims. The role of Paper 4 is to explain why many familiar financial categories are special cases of a shared state-contingent claim infrastructure, and to show how that infrastructure changes the cost, timing, and institutional classification of claims.

The paper makes six contributions.

First, it states the payoff-layer convergence observation. Let $\mathcal{X} = L^2(\Omega, \mathcal{F}, \mathbb{P})$ be the payoff space. A claim category is not a different mathematical primitive; it is a restricted set of representations over the same base $\mathcal{X}$. Event contracts, indemnity insurance, derivatives, parametric products, and binary prediction markets are all fibers of a common implemented-claim map. What prevents broad representation in practice is not algebra. It is infrastructure, law, liquidity, collateral, information, and admissibility.

Second, it models the delivery and infrastructure stack as reusable modules. A representation ρ requires specification, oracle/event resolution, data permissioning, identity and eligibility checks, collateral, margining, compliance, market making, settlement, dispute resolution, participant access, and reporting. Infrastructure state I changes the cost and reliability of these modules. If an oracle standard, legal template, collateral rail, or automated compliance system improves, it lowers the cost not only of one product category but of every category that uses the module. This gives a spillover channel from improvements in one market family to entry in another.

Third, it makes category boundaries endogenous. The same payoff can be implemented as insurance, a derivative, an event contract, a parametric note, a reinsurance contract, a prediction-market share, or a tokenized claim, subject to legal and institutional restrictions. The observed category is often the least-cost admissible wrapper under the current infrastructure vector. When collateral, settlement, oracle, or compliance costs change, the least-cost wrapper can switch without the underlying payoff changing.

Fourth, it develops an information-timing theorem. "Better data" is not one thing. Data used to define a payoff before uncertainty realizes can expand the set of contractible claims. Data used after realization can improve settlement and reduce disputes. Data revealed publicly before agents can pool or trade can destroy insurance value by converting uncertainty into known heterogeneity. The same model can therefore expand parametric settlement and shrink insurance pools.

Fifth, it formalizes the public-good price result. Some information markets are valuable because their prices improve outside decisions: procurement, policy, journalism, risk management, emergency planning, or scientific funding. But once the price is visible, many users can benefit without paying the market maker or sponsor. The price is a nonrival information output. The result is under-entry or under-liquidity unless an institution can capture enough of the external decision value or subsidize the market.

Sixth, it separates payoff convergence from welfare equivalence. The same payoff g can be hedging for one holder, speculation for another, manipulation for a third, and socially valuable public information for a fourth. Any infrastructure capable of supporting trillions of state-contingent claims must carry purpose, access, leverage, oracle, collateral, and externality metadata. Claim form alone is not enough.

Reader's map

Section 2 positions the paper relative to complete markets, security design, insurance, prediction markets, derivatives, and infrastructure. Section 3 defines states, payoffs, wrappers, delivery operators, and category signatures. Section 4 states payoff-layer convergence and wrapper-fiber observations. Section 5 builds the modular cost stack and the infrastructure-spillover result. Section 6 gives boundary endogeneity and least-cost wrapper selection. Section 7 develops the information-timing duality, the paper's main theorem. Section 8 treats information prices as public goods while Paper 3 owns the policy admissibility version. Section 9 shows why purpose and access must be part of the wrapper layer. Section 10 connects the common payoff object to distinct liquidity mechanisms. Section 11 studies oracles, collateral, and settlement as delivery technologies rather than mere costs. Section 13 gives empirical predictions and research designs. Section 14 explains the paper's relation to the broader research program.

2 Related Literature and Positioning

Complete markets and state-contingent commodities. The deepest ancestor of this paper is Arrow–Debreu's state-contingent commodity framework (Arrow and Debreu, 1954; Debreu, 1959). In that framework, commodities can be indexed by states and dates, and complete markets permit allocations conditional on all relevant contingencies. Radner's sequential formulation studies securities traded over time when future markets and expectations matter (Radner, 1972). This paper takes the state-contingent payoff representation seriously but asks an institutional question: what infrastructure makes such payoffs implementable, and why do different legal wrappers persist over the same payoff space?

Security design and financial innovation. The security-design literature studies how payoffs are chosen under risk sharing, asymmetric information, incentives, and market incompleteness (Allen and Gale, 1994; Duffie and Rahi, 1995; DeMarzo and Duffie, 1999). Companion Paper 1 builds on this tradition by modeling costly basis selection: valuable residual payoff directions become markets only if a representation can implement them cheaply enough. The present paper moves one layer down. It studies the representation infrastructure itself: oracle, collateral, settlement, compliance, access, and dispute resolution.

Derivatives and dynamic replication. Derivatives theory treats securities as functions of underlying states or prices. Option-pricing theory shows how some contingent payoffs can be replicated dynamically under strong assumptions (Black and Scholes, 1973; Merton, 1973). More general arbitrage-pricing and asset-pricing theories identify factor structure and pricing kernels (Ross, 1976; Duffie, 2001). This paper is not primarily a pricing paper. It asks how payoff functions are institutionally specified and delivered, especially when dynamic replication is unavailable or too expensive.

Insurance, adverse selection, and information. Insurance economics studies risk transfer under moral hazard, adverse selection, and contract constraints (Arrow, 1963; Akerlof, 1970; Rothschild and Stiglitz, 1976). Hirshleifer's information-timing insight is central: information can reduce social risk-sharing opportunities if it arrives before trade (Hirshleifer, 1971). This paper generalizes that timing logic to state-contingent claim infrastructure. The same data feed may be market-creating when used as an ex post settlement oracle and market-destroying when used as pre-contract public classification.

Insurance-linked securities and capital-market convergence. The convergence of insurance, reinsurance, and capital-market risk transfer is already a substantial literature. Catastrophe bonds, industry-loss warranties, derivatives, and hybrid securitized solutions show that insurance and financial-market wrappers can compete to bear similar catastrophe exposures (Froot, 2001; Cummins and Weiss, 2009). This paper abstracts from the particular catastrophe-risk setting and asks what general infrastructure makes such wrapper migration possible.

Prediction markets and information aggregation. Prediction markets and scoring-rule market makers turn event-contingent payoffs into information aggregation mechanisms (Hanson, 2003, 2007; Wolfers and Zitzewitz, 2004). Grossman and Stiglitz's impossibility result warns that information acquisition cannot be costlessly competed away (Grossman and Stiglitz, 1980). The public-good price proposition below emphasizes a complementary problem: even when a market price is socially valuable, the value may leak to nonpayers, causing private underprovision of depth, subsidy, or market creation.

Market design and institutional engineering. Market design emphasizes that institutions matter when the price system alone is incomplete, fragile, or misaligned (Roth, 2015). State-contingent claim infrastructure is a market-design problem: the designer must specify resolution, collateral, margin, access, information rights, and default procedures. The point is not only to choose a payoff g, but to build the institutional machinery that makes g a reliable economic object.

Smart contracts, oracles, and programmable settlement. The smart-contract literature emphasizes executable agreements, automated settlement, and oracle dependence (Szabo, 1997). This paper uses that vocabulary without making a blockchain-specific claim. Programmable settlement is one possible implementation of the more general object: a state-contingent payoff whose delivery depends on data, authority, collateral, and enforcement.

What is new. The paper's contribution is not that contingent claims share mathematical structure; that is a standard insight. The contribution is to treat the wrapper layer as the object of analysis. A claim category is a bundle of purpose, law, collateral, oracle, liquidity, access, and settlement restrictions over a common payoff base. Infrastructure improvements therefore create cross-category spillovers and endogenous boundary shifts. The theory also clarifies why convergence of payoff infrastructure does not imply convergence of welfare, liquidity, or regulation.

Literature	Core object	This paper's additional object
Arrow–Debreu/Radner	state-contingent commodities and securities	implementation infrastructure for payoff functions
Security design	optimal or privately created payoff forms	wrapper modules and least-cost institutional signatures
Derivatives pricing	replication and pricing of contingent payoffs	oracle, collateral, settlement, and delivery errors
Insurance economics	risk transfer under information constraints	timing of data across classification, pooling, and verification
Prediction markets	event prices and information aggregation	public-good price output and subsidy need
Market design	engineered allocation institutions	claim-level resolution, access, collateral, and dispute procedures
Smart-contract/oracle design	executable agreements and external data	general infrastructure layer, not one technology family

3 Environment: Payoffs, Wrappers, and Delivery

3.1 States and payoff functions

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. For most formal results, Ω may be finite, but the notation is written in L^2 form. A state-contingent payoff is an element

$$g \in \mathcal{X} := L^2(\Omega, \mathcal{F}, \mathbb{P}).$$

The payoff $g(\omega)$ is the net transfer delivered in state ω. Constants can be separated through a risk-free asset, so companion papers often work in the centered space L^2_0. Paper 4 keeps constants because collateral, premiums, option prices, coupons, and principal schedules are part of the implemented wrapper.

A *target payoff* is the payoff the parties intend to write. An *implemented payoff* is the payoff actually delivered after oracle error, default, dispute, legal interpretation, margin shortfall, settlement delay, and other delivery imperfections.

Definition 3.1 (Representation and implemented claim). *A representation or wrapper is a tuple*

$$\rho = (\mathcal{F}^{spec}, \mathcal{F}^{orc}, \mathsf{law}, \mathsf{id}, \mathsf{data}, \mathsf{coll}, \mathsf{margin}, \mathsf{liq}, \mathsf{settle}, \mathsf{dispute}, \mathsf{access}, \mathsf{report}, \mathsf{purpose}),$$

where $\mathcal{F}^{spec}$ is the sigma-algebra of state distinctions available for specification, $\mathcal{F}^{orc}$ is the sigma-algebra resolved by the oracle or verification process, and the remaining components encode legal form, identity and eligibility, data permissions, collateral, margining, liquidity support, settlement, dispute resolution, participant access, reporting/tax treatment, and declared purpose. An implemented claim is a pair (g, ρ), where $g \in L^2(\mathcal{F}^{spec})$ is the target payoff.

The wrapper determines what is contractible, who may hold the claim, how payment is secured, how ambiguity is resolved, and whether the payoff is legally characterized as insurance, derivative, event contract, security, wager, reinsurance, loan, royalty, or something else.

3.2 Delivery operators

A contract label is not the payoff. A contract that references g may deliver something else. The paper represents this through a delivery operator.

Definition 3.2 (Delivery operator). *For each wrapper ρ, the delivery operator*

$$\mathsf{D}_\rho : \mathcal{X} \to \mathcal{X}$$

maps a target payoff g to the implemented payoff $\mathsf{D}_\rho g$. The primitive object is the operator, not a category label. When a wrapper separately records delivery channels, one may report channel losses

$$\delta_\rho^{orc}, \quad \delta_\rho^{def}, \quad \delta_\rho^{disp}, \quad \delta_\rho^{delay}, \quad \delta_\rho^{legal},$$

representing oracle error, default or collateral shortfall, dispute outcome, settlement delay, and legal interpretation risk. Additive formulas using these terms are modeling restrictions that require operational identification of each channel. The invariant payoff object is $\mathsf{D}_\rho g$.

Remark 3.3 (Value should be computed on delivered payoffs). *Paper 1's value object applies to payoff directions. Paper 4 adds that the relevant payoff is not the marketing label or intended exposure, but $\mathsf{D}_\rho g$. A weather derivative, parametric insurance policy, or event contract with a poor oracle may span a different payoff direction from the one it names.*

3.3 Category signatures

Let $\mathcal{K}$ be a finite set of institutional categories: prediction/event contract, indemnity insurance, parametric insurance, derivative, catastrophe bond, tokenized claim, revenue share, sportsbook wager, and so on. A category is a set of restrictions on wrappers and target payoffs.

Definition 3.4 (Category signature). *A category signature is a pair $(\mathcal{R}_k, \mathcal{G}_k)$, where $\mathcal{R}_k \subseteq \mathcal{R}$ is the set of wrappers admissible for category k, and $\mathcal{G}_k(\rho) \subseteq L^2(\mathcal{F}_\rho^{spec})$ is the set of target payoffs allowed under wrappers in that category. The implemented payoff set for category k is*

$$\mathsf{Pay}_k(I) = \{\mathsf{D}_\rho g : \rho \in \mathcal{R}_k(I),\ g \in \mathcal{G}_k(\rho)\}.$$

The dependence on infrastructure state I allows legal, oracle, collateral, and settlement improvements to expand or contract the category.

Examples:

- Event contracts restrict g to scaled indicators or finite partitions, $g = A\mathbf{1}_E + B\mathbf{1}_{E^c}$, with event-resolution rules.
- Indemnity insurance restricts g to functions of loss, insurable interest, deductibles, limits, exclusions, and indemnity principles.
- Parametric insurance restricts g to functions of observed triggers such as wind speed, rainfall, earthquake magnitude, grid outage, or satellite-measured flood depth.
- Derivatives restrict g to functions of reference prices, rates, indexes, baskets, credit events, volatility, or other underlyings, together with margin and clearing conventions.
- Catastrophe bonds combine debt-like cash flows with trigger-contingent principal impairment.
- Tokenized claims add transfer, custody, and settlement conventions, but still require an off-chain or on-chain rule for what payoff is being tokenized.

3.4 Infrastructure state

Let the infrastructure state be a vector

$$I = (I^{spec}, I^{orc}, I^{data}, I^{id}, I^{coll}, I^{margin}, I^{comp}, I^{liq}, I^{settle}, I^{disp}, I^{access}, I^{report}).$$

Each coordinate represents the quality, availability, cost, and legal reliability of a module. A higher coordinate means the relevant module is cheaper, more reliable, more standardized, or more legally accepted.

In the program notation, I is not a competing master state. Paper 5 defines the compact core state $S_t = (\mathcal{H}_t, \mathcal{Q}_t, \kappa_t, L_t, A_t)$ and the full program state $\mathsf{State}_t = (S_t, I_t, N_t)$. Paper 4 opens the infrastructure and delivery slice I_t: the machinery that determines the representation-cost envelope κ_t, the delivery operator D_ρ, and parts of the liquidity state L_t.

For a wrapper ρ and target payoff g, write

$$C(g, \rho; I)$$

for the implementation cost. This includes fixed and variable costs of specification, data, verification, identity, compliance, collateral, liquidity, settlement, dispute resolution, and reporting. It excludes, unless explicitly stated, the social externality and admissibility terms handled in Paper 3.

Assumption 3.5 (Module monotonicity)**.** *For every wrapper ρ, target payoff g, and module coordinate m, implementation cost is weakly decreasing in the quality of any module used by (g, ρ):*

$$\frac{\partial C(g, \rho; I)}{\partial I^m} \leq 0$$

when the derivative exists. If (g, ρ) does not use module m, the derivative is zero.

Remark 3.6 (Infrastructure is not only cost)**.** *Infrastructure can also change what is contractible by refining $\mathcal{F}^{spec}$ or $\mathcal{F}^{orc}$, not merely by lowering a scalar cost. A new satellite feed, court-recognized benchmark, or audited API can make a previously nonmeasurable payoff admissible to the contract.*

4 Payoff-Layer Convergence

This section formalizes convergence at the payoff layer. It does not claim institutional categories are unimportant. It claims they are wrappers over a common payoff base.

4.1 Canonical payoff forms

Definition 4.1 (Canonical state-contingent forms)**.** *The following common instruments can be written as state-contingent payoffs:*

$$\begin{aligned} \textit{event contract:} \quad & g(\omega) = a + b\mathbf{1}_A(\omega), \\ \textit{binary prediction share:} \quad & g(\omega) = \mathbf{1}_A(\omega), \\ \textit{derivative:} \quad & g(\omega) = h(P_T(\omega)), \\ \textit{indemnity insurance:} \quad & g(\omega) = \min\{\max(L(\omega) - d, 0), \ell\}, \\ \textit{parametric risk claim:} \quad & g(\omega) = A\mathbf{1}_{\{Z(\omega) \geq z^*\}}, \\ \textit{catastrophe bond impairment:} \quad & g(\omega) = c - P_0\,\lambda(T(\omega)), \\ \textit{revenue share:} \quad & g(\omega) = \alpha R_T(\omega) \wedge \ell, \\ \textit{royalty claim:} \quad & g(\omega) = \alpha \sum_{t \leq T} Y_t(\omega). \end{aligned}$$

Here L is loss, P_T a terminal price, Z a physical or index trigger, T a catastrophe trigger statistic, R_T revenue, and Y_t royalty income.

Observation 4.2 (One payoff function)**.** *Let Ω be finite. Every event contract, indemnity insurance policy, derivative, parametric risk claim, catastrophe-bond trigger, revenue share, or finite-outcome prediction-market share defines an element of $\mathcal{X} = \mathbb{R}^\Omega$. Conversely, if all singleton state indicators $\mathbf{1}_{\{\omega\}}$ are verifiable and admissible, every payoff $g \in \mathbb{R}^\Omega$ can be written as a finite portfolio of event claims:*

$$g = \sum_{\omega \in \Omega} g(\omega)\mathbf{1}_{\{\omega\}}.$$

Therefore the mathematical payoff layer is common. Institutional categories are restrictions on representable events, participant sets, purposes, collateral rules, legal forms, and delivery operators, not distinct payoff primitives.

Proof. The first statement follows by inspection: each listed instrument maps states ω to a real payment. For the converse, the singleton indicators form the canonical basis of $\mathbb{R}^\Omega$. Any vector $g \in \mathbb{R}^\Omega$ has coordinates $g(\omega)$ in this basis, giving the displayed representation. Verifiability and admissibility are required because the algebraic basis is economically meaningful only if the singleton events can be resolved and legally used in contracts. □

Remark 4.3 (Convergence does not require singleton claims)**.** *The singleton-basis representation is a benchmark. Real infrastructure usually verifies coarser partitions, such as weather-station readings, industry-loss indexes, price indexes, accounting data, or court-recognized events. In that case the attainable payoff space is $L^2(\mathcal{F}^{orc})$, not all of $L^2(\mathcal{F})$. Convergence still holds within the verifiable sigma-algebra.*

4.2 Wrappers as fibers over payoffs

Definition 4.4 (Forgetful map and wrapper fiber)**.** *Define the payoff-delivery map*

$$\mathsf{F} : (g, \rho) \mapsto \mathsf{D}_\rho g.$$

For a delivered payoff $x \in \mathcal{X}$, the wrapper fiber over x is

$$\mathsf{F}^{-1}(x) = \{(g, \rho) : \mathsf{D}_\rho g = x\}.$$

The category-specific fiber is $\mathsf{F}^{-1}(x) \cap (\mathcal{G}_k \times \mathcal{R}_k)$.

Proposition 4.5 (Same payoff, many wrappers)**.** *If two implemented claims (g, ρ) and (g', ρ') satisfy $\mathsf{D}_\rho g = \mathsf{D}_{\rho'} g'$, then they are payoff-equivalent but not necessarily institutionally equivalent. They can differ in cost, access, liquidity, collateral, tax treatment, legal enforceability, dispute risk, information production, and externality profile.*

Proof. Payoff equivalence is exactly equality under F. The other objects are coordinates of ρ and ρ', not functions of the delivered payoff alone. Therefore equality in $\mathcal{X}$ does not imply equality of wrapper properties. $\square$

Example 4.6 (Rainfall)**.** *A farmer's drought exposure can be represented by crop indemnity insurance, a parametric rainfall policy, a weather derivative, a futures hedge on crop prices, a catastrophe-linked note, or a local prediction market about rainfall. Some of these wrappers deliver similar payoffs in drought states. They differ in who may buy them, whether loss adjustment is needed, whether basis risk remains, whether collateral is posted, whether the product is regulated as insurance or derivative, and whether third parties may speculate.*

Example 4.7 (Election probability)**.** *A binary claim paying $\mathbf{1}_A$ if a candidate wins can be a prediction-market share, a sportsbook bet, a decision-market instrument, a news-organization forecasting subsidy, or a prohibited political wager. The payoff form is the same. The wrapper determines access, limits, disclosure, manipulation controls, and whether the price is produced for public information or private gambling.*

4.3 Wrapper dimensions

The wrapper layer has at least eight economically distinct dimensions:

Wrapper dimension	Economic role
Purpose	hedging, insurance, speculation, information production, entertainment, governance, financing
Participant access	who may hold, who may write, position limits, suitability, insurable interest, jurisdiction
Oracle/resolution	what facts decide the payoff and who verifies them
Collateral/margin	whether promised payments are credible under limited liability
Liquidity mechanism	dealer hedging, pooling, bookmaker residual flow, AMM/scoring-rule subsidy, bilateral OTC
Legal/accounting form	insurance, derivative, security, loan, wager, reinsurance, note, token, commodity interest
Settlement/dispute process	when and how money moves; how ambiguity is adjudicated
Reporting/tax/data rights	what must be reported; who owns and can reuse the data and price output

Remark 4.8 (The central discipline)**.** *Payoff convergence is not category erasure. Categories persist because the wrapper dimensions matter. Paper 4's object is the relation between the common payoff layer and the diverging wrapper layer.*

5 Infrastructure as Modular Cost Stack

5.1 Task technologies

Let $\mathcal{M}$ denote the set of infrastructure modules:

$$\mathcal{M} = \{spec, orc, data, id, coll, margin, comp, liq, settle, disp, access, report\}.$$

For each implemented claim (g, ρ), let $a_m(g, \rho) \geq 0$ measure its intensity of use of module m. Let $\chi_m(I^m)$ be the unit cost of module m, decreasing in the infrastructure state. A tractable cost specification is

$$C(g, \rho; I) = C_0(g, \rho) + \sum_{m \in \mathcal{M}} a_m(g, \rho)\chi_m(I^m) + \sum_{m,n \in \mathcal{M}} a_{mn}(g, \rho)\chi_{mn}(I^m, I^n). \tag{1}$$

The cross terms capture complementarities: identity may lower compliance cost; oracle quality may lower dispute cost; settlement automation may lower collateral cost; legal templates may lower specification and reporting cost.

Definition 5.1 (Infrastructure exposure vector)**.** *The infrastructure exposure vector of* (g, ρ) *is*

$$a(g, \rho) = (a_m(g, \rho))_{m \in \mathcal{M}}.$$

Claims with similar exposure vectors benefit from similar infrastructure shocks even if they belong to different legal categories.

Example 5.2 (Shared oracle module). *A standardized hurricane wind-speed oracle can support catastrophe bonds, parametric insurance, weather derivatives, reinsurance triggers, municipal disaster loans, and event contracts. The module is not owned by one category; it is a reusable state-resolution technology.*

5.2 Entry threshold

Let $V(D_\rho g \mid \mathcal{H})$ denote the gross value of the delivered payoff direction relative to the existing market span, as in Paper 1. Let $\phi_{g,\rho}$ be the fraction of gross surplus captured by the creator. Let $\Psi_{g,\rho}$ be the private admissibility, legal, or distribution burden. Private viability is

$$\phi_{g,\rho} V(\mathsf{D}_\rho g \mid \mathcal{H}) > C(g,\rho;I) + \Psi_{g,\rho}. \tag{2}$$

Social desirability additionally accounts for externality and public-good terms handled by Paper 3:

$$V(\mathsf{D}_\rho g \mid \mathcal{H}) + B^{info}_{g,\rho} + B^{public}_{g,\rho} - C(g,\rho;I) - E_{g,\rho} > 0. \tag{3}$$

Paper 4 treats C, D, and the infrastructure vector I as the central objects. It does not attempt to solve V or E from primitives.

5.3 Infrastructure spillovers

Proposition 5.3 (Infrastructure spillover across categories). *Suppose Assumption 3.5 holds. Let I' improve module m relative to I, with all other modules unchanged. Then for every implemented claim (g,ρ) that uses module m,*

$$C(g,\rho;I') \le C(g,\rho;I).$$

Therefore the privately viable set

$$\mathcal{A}^{priv}(I) = \{(g,\rho) : \phi_{g,\rho} V(\mathsf{D}_\rho g \mid \mathcal{H}) > C(g,\rho;I) + \Psi_{g,\rho}\}$$

weakly expands along all claims whose costs fall. Because module use is not category-specific, an infrastructure improvement originating in one category can create viable claims in other categories.

Proof. The cost inequality follows directly from module monotonicity. If a claim was privately viable under I, lowering its cost preserves viability. A previously nonviable claim may become viable if the cost reduction crosses the entry threshold. Since the set of claims using module m can include insurance, derivatives, prediction markets, parametric products, and tokenized claims, the expansion is cross-category. □

Corollary 5.4 (Category birth waves). *If many claims in different categories have entry gaps close to zero and share a module exposure $a_m > 0$, a discrete improvement in module m can produce a wave of market births across categories. The wave is governed by shared infrastructure exposure, not by category labels.*

Proof. Let the private entry gap be

$$\Delta_{g,\rho}(I) = \phi_{g,\rho} V(\mathsf{D}_\rho g \mid \mathcal{H}) - C(g,\rho;I) - \Psi_{g,\rho}.$$

A module improvement increases Δ for claims using the module. All claims whose gaps cross from negative to positive enter. The crossing set can span multiple categories. □

Example 5.5 (Stable settlement rail). *A settlement rail that can move collateralized value instantly and cheaply lowers settlement and counterparty-risk costs for derivatives, event contracts, tokenized securities, parametric insurance, revenue shares, and automated claims. The first visible application may be one category, but the cost shock is infrastructural.*

5.4 Infrastructure underprovision

Infrastructure modules can be public goods among claim creators. An oracle, benchmark, legal template, or settlement standard may be costly to build but cheap for later entrants to use.

Proposition 5.6 (Shared infrastructure under-entry)**.** *Let J be a set of potential implemented claims that all require a common infrastructure module with fixed cost F. Let ΔV_J be the gross value of the union of delivered payoff directions, counted without double-counting overlap. Suppose each individual first mover is privately nonviable if it must pay F:*

$$\phi_j V_j - c_j - F - \Psi_j < 0 \quad \textit{for all } j \in J,$$

but the cluster is socially valuable:

$$\Delta V_J + \Delta B_J^{info} - \sum_{j \in J} c_j - F - E_J > 0.$$

Then decentralized one-claim entry can fail even though building the module is socially efficient.

Proof. The first inequality rules out individual entry by any single claim creator. The second inequality states that the joint payoff and information value of the cluster exceeds claim-specific costs, the fixed infrastructure cost, and externalities. Because no individual captures the cluster surplus, the infrastructure can be undersupplied. □

Remark 5.7 (Why this belongs in Paper 4)**.** *Paper 1 contains an infrastructure-under-entry result for payoff spans. Paper 4 gives the institutional version: the underprovided object is often not the claim itself but the reusable oracle, template, benchmark, collateral convention, or settlement rail that makes many claims cheap.*

6 Boundary Endogeneity and Least-Cost Wrappers

A claim's institutional category is often treated as fixed. In this framework, it is partly a choice variable. If several wrappers can deliver similar payoffs, the observed wrapper is shaped by cost, law, distribution, liquidity, and admissibility.

6.1 Least-cost wrapper selection

For a target delivered payoff $x \in \mathcal{X}$, define the least-cost implementation in category k:

$$\kappa_k(x; I) = \inf_{(g,\rho):\rho \in \mathcal{R}_k(I),\ \mathsf{D}_\rho g = x} C(g, \rho; I) + \Psi_{g,\rho}.$$

If no category-k wrapper can deliver x, set $\kappa_k(x; I) = +\infty$. The least-cost category is

$$k^*(x; I) \in \arg\min_{k \in \mathcal{K}} \kappa_k(x; I).$$

Proposition 6.1 (Boundary endogeneity)**.** *Suppose two categories k and ℓ can both implement the same delivered payoff x, with finite costs $\kappa_k(x; I)$ and $\kappa_\ell(x; I)$. If there exist infrastructure states I_0, I_1 such that*

$$\kappa_k(x; I_0) < \kappa_\ell(x; I_0) \quad \textit{and} \quad \kappa_k(x; I_1) > \kappa_\ell(x; I_1),$$

then the least-cost category for x changes as infrastructure changes, even though the delivered payoff does not. Product boundaries are therefore endogenous to the infrastructure vector.

Proof. The conclusion follows immediately from the definition of $k^*(x;I)$. At I_0, category k is cheaper than ℓ; at I_1, category ℓ is cheaper than k. The payoff x is held fixed throughout. □

Example 6.2 (From indemnity to parametric)**.** *A flood exposure may first be insured by indemnity coverage because measurement technology is poor and regulators recognize indemnity forms. As satellite imagery, water-depth sensors, and standardized triggers improve, a parametric wrapper may become cheaper. The underlying exposure did not become a different mathematical object; the least-cost wrapper changed.*

Example 6.3 (From OTC derivative to event venue)**.** *A bespoke event exposure may first be written bilaterally as an OTC contract because only sophisticated parties can verify and collateralize it. If a recognized oracle and automated margin system later exist, the same payoff may migrate to an exchange or event-contract venue. The boundary moved because the wrapper cost changed.*

6.2 Legal categories as constraints and subsidies

Legal categories do not vanish in this model. They enter as constraints, taxes, subsidies, safe harbors, capital rules, participant restrictions, and enforcement probabilities. A legal category can make a wrapper infeasible, lower its cost through standardization, or raise its cost through compliance.

Definition 6.4 (Legal wedge)**.** *For category k, define the legal wedge for payoff x as*

$$\Lambda_k(x;I) = \kappa_k(x;I) - \kappa_k^{tech}(x;I),$$

where κ_k^{tech} is the least cost of technically implementing x under category-k technology absent legal and regulatory burdens. The wedge can be positive, zero, or negative if law provides a safe harbor, public guarantee, standardized form, or tax advantage.

Proposition 6.5 (Legal boundaries can distort basis choice)**.** *Let x and y be two delivered payoffs with gross values $V(x) > V(y)$. A legal wedge can cause the lower-value payoff y to be implemented before x if*

$$V(x) - V(y) < \kappa(x;I) - \kappa(y;I),$$

where κ includes legal wedges. Thus legal categories can change not only labels but the order in which payoff directions become markets.

Proof. Private or social implementation compares net values. If the cost disadvantage of x exceeds its gross-value advantage, y has higher net value and can be implemented first. □

6.3 Boundary invariance and boundary failure

A useful regulatory boundary should be invariant to payoff-preserving relabeling when the relevant risk, purpose, and externality are unchanged. A boundary fails when economically equivalent wrappers receive inconsistent treatment solely because of obsolete category labels.

Definition 6.6 (Payoff-purpose equivalence)**.** *Two implemented claims (g,ρ) and (g',ρ') are payoff-purpose equivalent for participant set A if*

$$\mathsf{D}_\rho g = \mathsf{D}_{\rho'} g',$$

and if their participant exposures, leverage, purpose, information rights, manipulation incentives, and externalities are the same for A.

Proposition 6.7 (Boundary invariance principle). *If two claims are payoff-purpose equivalent, then any difference in regulatory treatment must be justified by a difference in wrapper reliability, enforcement, collateral, disclosure, systemic linkage, or administrative cost. A pure label difference is not an economic distinction.*

Proof. By payoff-purpose equivalence, the objects that determine direct payoff consequences and participant motives are equal. If wrapper reliability and administrative objects are also equal, no modeled economic variable remains on which to base different treatment. Any remaining distinction is nominal. ☐

Remark 6.8 (Why labels still matter). *Labels are often proxies for real wrapper differences. Insurance law may bring solvency regulation and claims-handling duties. Derivative law may bring margin and clearing. Securities law may bring disclosure and antifraud rules. Gambling law may bring access limits. The point is not to ignore labels; it is to ask what wrapper function the label performs.*

7 Information Timing: The Dual Role of Data

The slogan "better data creates markets" is incomplete. Better data can create markets, destroy markets, or change their welfare interpretation depending on when and how it enters the claim lifecycle.

7.1 Five timings of information

A data system can enter at five distinct points:

Timing	Role	Market effect
Pre-contract classification	sorts risks before pooling or trade	can improve pricing but fragment pools and reduce insurance value
Ex ante specification	defines a payoff rule before uncertainty resolves	expands contractible payoff space
Interim public revelation	reveals state-relevant information before trade or pooling completes	can destroy risk-sharing value and create adverse selection
Ex post verification	resolves the payoff after state realization	lowers dispute cost, basis risk, and default ambiguity
Post-market price publication	disseminates the price as information	creates public-good decision value and capture problems

The same sensor, model, database, or oracle can appear in different rows. For example, a wildfire-risk model used to define a parametric trigger before the season can create a useful market. The same model used to publicly reveal which houses will burn before insurance is purchased can eliminate pooling.

7.2 Contractible sigma-algebras

Let $\mathcal{F}^C$ be the sigma-algebra of facts that can be specified in contracts ex ante, and $\mathcal{F}^V$ the sigma-algebra that can be verified at settlement. A payoff g is exactly implementable only if it is both specifiable and verifiable, abstracting from collateral and law:

$$g \in L^2(\mathcal{F}^C \cap \mathcal{F}^V).$$

In practice, specification and verification can be asymmetric: parties may specify an ideal payoff but only verify a coarser trigger, or they may verify rich data that law does not allow them to use.

Observation 7.1 (Best verifiable settlement approximation)**.** *Let $g \in L^2(\mathcal{F})$ be an ideal payoff and let $\mathcal{F}^V \subseteq \mathcal{F}$ be the settlement-verifiable sigma-algebra. Among all $\mathcal{F}^V$-measurable payoffs h, the unique L^2-best approximation is*

$$h^* = \mathbb{E}[g \mid \mathcal{F}^V].$$

The unavoidable basis risk from settlement verification is

$$g - \mathbb{E}[g \mid \mathcal{F}^V].$$

Proof. Conditional expectation is the orthogonal projection of g onto the closed subspace $L^2(\mathcal{F}^V)$. Orthogonal projection uniquely minimizes mean squared error. ☐

Remark 7.2 (Parametric basis risk)**.** *Paper 1 uses the same projection fact in the costly-basis setting. Paper 4 uses it only to interpret settlement infrastructure. Parametric insurance is not merely cheaper indemnity insurance. It replaces the ideal loss payoff with the best or most admissible payoff measurable by the trigger infrastructure. Basis risk is the component of loss not captured by the trigger sigma-algebra.*

7.3 Timing duality theorem

Consider a simple risk-sharing environment in which an agent faces random loss $L \in L^2(\mathcal{F})$. Under quadratic local utility, the value of actuarially fair full insurance is proportional to $\mathrm{Var}(L)$. Let $\mathcal{S} \subseteq \mathcal{F}$ be a public signal sigma-algebra.

Lemma 7.3 (Variance decomposition)**.** *For any $L \in L^2$ and signal sigma-algebra $\mathcal{S}$,*

$$\mathrm{Var}(L) = \mathrm{Var}(\mathbb{E}[L \mid \mathcal{S}]) + \mathbb{E}[\mathrm{Var}(L \mid \mathcal{S})].$$

Proof. This is the law of total variance. ☐

Theorem 7.4 (Information timing duality)**.** *Let $\mathcal{S}$ be information generated by a data system.*

(i) If $\mathcal{S}$ is used ex ante to specify or ex post to verify a payoff before the relevant uncertainty is resolved for contracting purposes, it weakly expands the contractible payoff space from $L^2(\mathcal{F}^V)$ to $L^2(\mathcal{F}^V \vee \mathcal{S})$ and weakly lowers best-approximation basis risk.

(ii) If $\mathcal{S}$ is publicly revealed before agents can pool or trade loss L, then the ex ante value of post-revelation insurance excludes the between-signal component $\mathrm{Var}(\mathbb{E}[L \mid \mathcal{S}])$. In the quadratic benchmark, the lost pooling value is proportional to

$$\mathrm{Var}(\mathbb{E}[L \mid \mathcal{S}]).$$

(iii) If L is $\mathcal{S}$-measurable, post-revelation insurance against L has no remaining uncertainty to pool at the $\mathcal{S}$ level.

Thus the same data system can create markets as verification infrastructure and destroy markets as pre-trade public revelation.

Proof. For (i), $\mathcal{F}^V \subseteq \mathcal{F}^V \vee \mathcal{S}$, so the projection space is larger. Orthogonal projection error weakly decreases in a larger subspace. For (ii), after $\mathcal{S}$ is publicly revealed, trades can insure only residual uncertainty within each signal atom. The remaining variance is $\mathbb{E}[\mathrm{Var}(L \mid \mathcal{S})]$. By Lemma 7.3, the eliminated cross-atom component is $\mathrm{Var}(\mathbb{E}[L \mid \mathcal{S}])$. Under the quadratic local benchmark, insurance value is proportional to variance reduced, giving the result. Part (iii) follows because $\mathrm{Var}(L \mid \mathcal{S}) = 0$ when L is $\mathcal{S}$-measurable. ☐

Remark 7.5 (AI underwriting)**.** *A model that predicts individual loss risk can improve pricing and reduce adverse selection, but if it publicly classifies risks before pooling, it can destroy the cross-type transfers that make insurance valuable. The infrastructure question is therefore not whether the model is accurate. It is when the model is used, who sees its output, and whether it supports pooling, settlement, or extraction.*

7.4 Adverse selection versus settlement accuracy

A common mistake is to treat all information asymmetry reduction as welfare-improving. If information reduces hidden type after contracts are written, it can reduce fraud and dispute costs. If it sorts buyers before coverage, it can unravel pooling.

Proposition 7.6 (Disclosure tradeoff)**.** *Let a data disclosure reduce expected settlement error by* $B^{settle} \geq 0$ *and reduce pooling value by* $B^{pool} \geq 0$ *through pre-contract classification. The disclosure is socially beneficial in the insurance-use case only if*

$$B^{settle} + B^{pricing} + B^{fraud} > B^{pool} + B^{privacy} + B^{exclusion},$$

where $B^{pricing}$ *is the value of improved risk pricing,* B^{fraud} *fraud reduction,* $B^{privacy}$ *privacy cost, and* $B^{exclusion}$ *the social cost of excluding or repricing high-risk types. The sign cannot be inferred from accuracy alone.*

Proof. This is an accounting identity over benefit and cost channels. Accuracy can enter with opposite signs depending on timing and use. The inequality states the condition for net positive social value. ☐

Remark 7.7. *The theorem does not imply that classification is bad. It implies that classification, specification, verification, and revelation are different economic operations. Paper 4's infrastructure stack must distinguish them.*

8 Public-Good Prices and Information-Market Underprovision

Some state-contingent claims are created mainly for price discovery. Their payoff is an instrument for producing an information output. That output can be a public good.

8.1 Decision value of prices

Let a market price process P_g for claim g produce an information signal used by outside decision makers $u \in U$. Let $\Delta_u(P_g) \geq 0$ be user u's improvement in decision value from observing the price. The total external information value is

$$B^{info}(g) = \sum_{u \in U} \Delta_u(P_g).$$

Let $R^{info}(g)$ be the revenue the market creator can capture from selling access, fees, data, sponsorship, or order flow. Typically

$$R^{info}(g) = \phi_g^{info} B^{info}(g), \qquad \phi_g^{info} < 1,$$

because price information is hard to exclude once observed.

Proposition 8.1 (Public-good price underprovision). *Suppose creating and subsidizing an information market for g costs K_g. The market is socially valuable if*

$$B^{info}(g) + V^{trade}(g) > K_g + E_g,$$

but privately created only if

$$R^{info}(g) + \phi_g V^{trade}(g) > K_g + \Psi_g.$$

If

$$R^{info}(g) + \phi_g V^{trade}(g) < K_g + \Psi_g \quad \text{and} \quad B^{info}(g) + V^{trade}(g) > K_g + E_g,$$

then the information market is privately undersupplied despite positive social value. Sponsorship, subsidy, mandate, club goods, data licensing, or public provision can close the gap.

Proof. The first inequality is the social net-benefit condition. The second is the private entry condition. If the social condition holds while the private condition fails, private actors do not create the market even though it is socially valuable. ☐

Remark 8.2 (Paper 3 owns the admissibility use). *This proposition is the infrastructure-side statement: price discovery can be a nonrival output. Paper 3 owns the policy version, including subsidy, access, manipulation, and harmful-flow tests. Paper 2 owns the liquidity version, where bounded-loss market makers or sponsors buy depth for information claims.*

Example 8.3 (Forecast as infrastructure). *A disaster-probability market may improve emergency procurement, insurance pricing, grid preparation, and public communication. Many beneficiaries need not trade. If the market sponsor cannot charge them for the improved decisions, the probability price is underprovided without subsidy.*

Corollary 8.4 (Why subsidized market makers are infrastructure). *For information markets with low direct hedging demand and high public decision value, a bounded-loss automated market maker or designated sponsor should be interpreted as information infrastructure, not as a failed trading venue. The subsidy buys a price.*

8.2 Price accuracy and endogenous information acquisition

Public-good prices also interact with incentives to acquire information. If informed traders cannot recover research costs because prices move too quickly or because participation is limited, the market may be liquid but uninformative. If informed traders recover too much through adverse selection against hedgers, the market may harm risk sharing.

Proposition 8.5 (Information compensation wedge). *Let $K^{research}$ be the private cost of producing information about event A. Let Π^{trade} be the expected trading profit available to an informed participant, and let $B^{decision}$ be the external decision value of the information incorporated into price. If*

$$\Pi^{trade} < K^{research} < \Pi^{trade} + B^{decision},$$

then the information is socially worth producing but privately underproduced through trading alone. If instead Π^{trade} is earned mainly by trading against uninformed hedgers, the same compensation mechanism can reduce hedging welfare. The institutional problem is to pay for information without destroying the risk-sharing use of the market.

Proof. The first claim is the comparison of private and social returns. The second follows because informed-trader profit is a transfer plus possible price-discovery benefit; when it is extracted from uninformed hedging flow, it can raise spreads and reduce insurance or hedging participation. ☐

Remark 8.6 (Prediction market design is not only liquidity design)**.** *A prediction market needs enough subsidy or informed trading profit to elicit information, but not so much adverse selection that hedgers and ordinary users exit. This is an infrastructure design problem involving access, subsidies, limits, and purpose.*

9 Purpose, Access, and Non-Payoff Metadata

The one-payoff-function observation is deliberately modest. It does not say that payoff-equivalent claims are welfare-equivalent. The same g can play different social roles depending on who holds it and why.

9.1 Purpose states

Let A be the set of participants. For participant a, let e_a denote their preexisting exposure and q_a their position in claim g. A participant is a hedger when $q_a g$ reduces the risk of e_a, an insurer or dealer when they warehouse compensated risk, an information trader when their position reflects costly signal acquisition, a speculator when it reflects disagreement without underlying exposure, and a manipulator when they can affect the event or resolution.

Definition 9.1 (Purpose profile)**.** *A purpose profile for a claim is*

$$\begin{aligned}
\pi &= (\pi_a)_{a\in A},\\
\pi_a &\in \{\text{hedge}, \text{insure}, \text{intermediate}, \text{inform},\\
&\qquad \text{speculate}, \text{consume}, \text{govern}, \text{manipulate}\}.
\end{aligned}$$

The purpose profile is part of the wrapper state when access, limits, disclosures, or duties depend on why participants trade.

Proposition 9.2 (No payoff-only welfare criterion)**.** *There is no payoff-only function $W : \mathcal{X} \to \mathbb{R}$ that can correctly rank all implemented claims when the same delivered payoff can be held under different exposure, purpose, leverage, information, or manipulation profiles with different externalities.*

Proof. Take a payoff g. In profile one, g is held by an exposed agent and reduces endowment risk, producing positive hedging value and no material externality. In profile two, the same g is held by an agent who can manipulate the underlying event or oracle, creating a negative externality exceeding any trading surplus. A payoff-only function assigns the same value to both because g is identical. Correct ranking requires non-payoff metadata. ☐

Remark 9.3 (Paper 3 owns the normative frontier)**.** *Paper 4's point is infrastructural: the wrapper must carry purpose and access information because payoff form is insufficient. The full normative rule for which profiles should be permitted belongs to Paper 3.*

9.2 Access controls as infrastructure

Participant restrictions are not merely paternalistic add-ons. They can be essential to making a claim admissible or even technically viable. A contract that is useful as a hedge for exposed parties can become manipulation-prone or gambling-like when made universally accessible with leverage.

Definition 9.4 (Access rule)**.** *An access rule is a map*

$$\alpha_\rho : A \to \{0,1\} \times \mathbb{R}_+ \times \mathcal{D},$$

assigning each participant eligibility, position limits, and disclosure obligations. Access rules may depend on exposure, sophistication, jurisdiction, insurable interest, fiduciary status, or manipulation risk.

Proposition 9.5 (Access can preserve a payoff market)**.** *Suppose a payoff g has positive hedging value for exposed participants but negative externality when held by manipulators or highly leveraged unexposed speculators. If an access rule excludes or limits the harmful participant set at cost C^{access}, then the claim is socially implementable whenever*

$$V^{hedge}(g) - C(g,\rho;I) - C^{access} - E^{residual} > 0,$$

even if the unrestricted version has negative social value. Access control can therefore be market-preserving rather than market-suppressing.

Proof. The restricted wrapper changes the participant profile and thus the externality term. If the remaining hedging value exceeds implementation, access, and residual externality costs, the restricted claim has positive social value even though the unrestricted profile may not. ☐

9.3 Purpose metadata and privacy

Purpose-aware access can require sensitive data: exposure, income, location, health, crop acreage, property characteristics, or business revenue. The infrastructure problem is to verify eligibility without unnecessary revelation.

Proposition 9.6 (Privacy-preserving eligibility as a module)**.** *Let eligibility be a predicate $E_a \in \{0,1\}$ depending on private data D_a. A privacy-preserving access module that proves $E_a = 1$ without revealing D_a can lower the joint cost*

$$C^{access} + E^{privacy} + E^{adverse\ selection}$$

relative to either unrestricted access or full disclosure, provided the proof is trusted by the venue and regulators.

Proof. Unrestricted access can raise externality or adverse-selection costs; full disclosure can raise privacy costs. A trusted predicate proof supplies eligibility information while reducing data revelation. The result follows from comparing the sum of access, privacy, and adverse-selection costs. ☐

Remark 9.7. *This is one reason identity and data-permissioning belong in Paper 4's infrastructure stack. The same payoff may become admissible only when access can be verified without turning private life into public collateral.*

10 Liquidity-Mechanism Convergence and Difference

The payoff object converges; liquidity mechanisms need not. A derivative, insurance pool, sportsbook, prediction market, and catastrophe bond can share a payoff form while relying on different sources of liquidity and capital.

10.1 Liquidity signatures

Definition 10.1 (Liquidity signature)**.** *A liquidity signature is a tuple*

$$\ell_\rho = (\mathsf{source}, \mathsf{capital}, \mathsf{inventory}, \mathsf{flow}, \mathsf{limits}, \mathsf{subsidy}),$$

where source *is factor hedging, pooling, residual diversification, AMM/scoring-rule subsidy, bilateral matching, or public backstop;* $\mathsf{capital}$ *is the balance sheet supporting payouts;* $\mathsf{inventory}$ *is the risk held by intermediaries;* flow *is the trader flow ecology;* limits *are position and exposure constraints; and* $\mathsf{subsidy}$ *is any explicit information or liquidity budget.*

Proposition 10.2 (Payoff convergence does not imply liquidity convergence)**.** *Two wrappers can deliver the same payoff x while having different liquidity signatures. Therefore a payoff-equivalent migration from one category to another can change spreads, depth, adverse selection, systemic exposure, and subsidy needs even when x is unchanged.*

Proof. Liquidity signatures are components of ρ, not functions solely of delivered payoff x. A dealer-hedged derivative, pooled insurance contract, and subsidized prediction market can all reference the same event but be supported by different capital and flow mechanisms. ☐

10.2 Four liquidity mechanisms

Mechanism	Common category	Infrastructure need	Failure mode
Factor-hedged dealer liquidity	derivatives, swaps, options	hedge basis, margin, collateral, market-maker capital	basis risk, hedge crowding, residual warehouses
Pooling/underwriting	insurance, reinsurance, cat bonds	underwriting data, solvency capital, claims verification	adverse selection, moral hazard, correlated losses
Residual-flow/bookmaker liquidity	sports/event venues, some binary markets	broad non-toxic flow, limits, fraud controls	sharp flow, addiction, manipulation, correlation
Subsidized information liquidity	prediction markets, decision markets, research forecasts	AMM/scoring rule, sponsor budget, oracle	subsidy exhaustion, ambiguous resolution, low information acquisition

Remark 10.3 (Paper 2 connection)**.** *Paper 2 models these mechanisms mechanically. Paper 4's contribution is to show that the same payoff infrastructure can feed multiple mechanisms. Oracle and settlement improvements lower costs for all of them, even though the liquidity source differs.*

10.3 Liquidity mechanism switching

Proposition 10.4 (Liquidity mechanism switching). *Let x be a delivered payoff that can be supported by two liquidity signatures ℓ and ℓ', with total support costs $L_\ell(x; I)$ and $L_{\ell'}(x; I)$. If infrastructure changes cause these costs to cross, then the same payoff can migrate from one liquidity mechanism to another. Such migration can change market access and welfare even when payoff and gross value remain fixed.*

Proof. The least-cost support mechanism is the minimizer of $L_\ell(x; I)$ over feasible signatures. A cost crossing changes the minimizer. Access and welfare can change because liquidity signatures include participant flow, leverage, limits, and capital structure. ☐

Example 10.5 (Catastrophe risk). *Catastrophe risk can be supported by insurer capital, reinsurance, catastrophe bonds, industry-loss warranties, parametric swaps, or government backstop. Improvements in catastrophe models, event reporting, collateral trusts, and securities settlement can shift risk from balance-sheet insurance to capital-market instruments without changing the underlying disaster states.*

11 Oracles, Collateral, and Settlement as Delivery Infrastructure

This section studies three modules that are often treated as back-office details but are in fact central to the payoff delivered: oracles, collateral, and settlement.

11.1 Oracles as sigma-algebras

An oracle is not merely a data feed. It is an institution that decides which state distinctions will be recognized by the contract.

Definition 11.1 (Oracle sigma-algebra). *An oracle O is a measurable map $O : \Omega \to \mathcal{Y}$. The oracle sigma-algebra is $\mathcal{F}^O = \sigma(O)$. A payoff is oracle-measurable if $g \in L^2(\mathcal{F}^O)$.*

Proposition 11.2 (Oracle refinement value). *If oracle O' refines oracle O, so $\mathcal{F}^O \subseteq \mathcal{F}^{O'}$, then for any ideal payoff g, the best oracle-measurable approximation error weakly falls:*

$$\mathbb{E}\left[(g - \mathbb{E}[g \mid \mathcal{F}^{O'}])^2\right] \le \mathbb{E}\left[(g - \mathbb{E}[g \mid \mathcal{F}^{O}])^2\right].$$

Proof. This is the monotonicity of orthogonal projection error as the projection subspace expands. ☐

Remark 11.3 (Oracle quality versus oracle granularity). *A more granular oracle is not always better if it raises privacy costs, manipulation incentives, dispute complexity, or pre-trade classification harms. Refinement lowers approximation error but can raise other wrapper costs.*

11.2 Oracle manipulation

If traders can affect the oracle output or the underlying event, the claim's payoff is endogenous. Paper 3 handles admissibility; Paper 4 records the infrastructure condition.

Definition 11.4 (Manipulation-stability condition). *Let participant a hold position q_a in payoff $g(O(\omega))$. Let $m_a(\Delta O)$ be the cost to induce oracle change ΔO, and let $P_a(\Delta O)$ be expected penalty. The claim is manipulation-stable for a if for all feasible manipulations,*

$$q_a \left[g(O + \Delta O) - g(O)\right] \leq m_a(\Delta O) + P_a(\Delta O).$$

It is manipulation-stable if the condition holds for all participants and relevant coalitions.

Proposition 11.5 (Position limits as oracle infrastructure). *If the maximum payoff gain from manipulation is linear in position size and manipulation cost plus penalty is bounded below by $M > 0$, then a position limit $\bar{q}$ satisfying*

$$\bar{q} \cdot \sup_{\Delta O} |g(O + \Delta O) - g(O)| \leq M$$

is sufficient for manipulation stability under Definition 11.4.

Proof. For any participant with $|q_a| \leq \bar{q}$, the payoff gain from any manipulation is bounded by the left side, which is at most M, and M is no greater than manipulation cost plus penalty. $\square$

Remark 11.6. *Position limits, resolution committees, cryptographic attestations, sensor redundancy, audit rights, and penalties are part of oracle infrastructure. They are not merely regulatory overlays.*

11.3 Collateral and limited liability

Promised state-contingent payoffs require credible payment. Collateral changes the delivered payoff.

Definition 11.7 (Collateralized delivery). *Let a short position promise to pay $g^+(\omega) = \max\{g(\omega), 0\}$ and post collateral K. Under limited liability and no recovery beyond collateral, the delivered payment is $\mathsf{D}^K g(\omega) = \min\{g^+(\omega), K(\omega)\} - g^-(\omega)$ for a two-sided convention, or the corresponding limited-liability payoff for the specific contract. Full delivery requires collateral sufficient in every payment state.*

Proposition 11.8 (Collateral sufficiency). *For a one-sided claim requiring payment $g^+(\omega) \geq 0$, finite static collateral K guarantees full settlement in all states if and only if g^+ is essentially bounded and*

$$K \geq \operatorname{ess\,sup}_\omega g^+(\omega).$$

If g^+ is unbounded, no finite static collateral level guarantees full settlement. If $K < \operatorname{ess\,sup} g^+$, then the delivered payoff differs from the promised payoff on a set of positive probability whenever $\mathbb{P}(g^+ > K) > 0$.

Proof. If $K \geq \operatorname{ess\,sup} g^+$, then $\min\{g^+, K\} = g^+$ almost surely. If g^+ is unbounded, every finite K is below the essential supremum. If $K < \operatorname{ess\,sup} g^+$, then by definition there is a positive-probability set on which $g^+ > K$, and delivery is capped at K on that set. $\square$

Remark 11.9 (Margin is dynamic collateral infrastructure). *Dynamic margin lowers the need for worst-case static collateral but introduces gap risk, liquidation risk, procyclicality, and settlement timing issues. It is a delivery technology, not a mere funding detail.*

11.4 Settlement finality

Settlement delay exposes participants to counterparty, liquidity, and information risk. Let τ be settlement lag and r_τ the cost of carrying settlement exposure.

Proposition 11.10 (Settlement lag as delivery haircut)**.** *If a payoff g is settled after lag τ with default intensity, funding cost, or liquidity discount summarized by haircut $h(\tau) \in [0, 1]$, then the present delivered payoff is*

$$\mathsf{D}^\tau g = (1 - h(\tau))g.$$

If h is increasing, reducing settlement lag increases delivered payoff value and lowers collateral/funding cost.

Proof. The displayed form defines the reduced-form haircut. Monotonicity of h implies that a lower lag weakly raises the multiplicative delivery factor. □

Remark 11.11. *Programmable settlement matters because it changes D_ρ, C, and collateral requirements simultaneously. It can make small claims economical and reduce default exposure, while also increasing automation risks if oracle or code failures become common-mode.*

12 Composability and Infrastructure Fragility

Reusable infrastructure creates spillovers. It also creates common-mode fragility. A shared oracle, settlement rail, collateral token, benchmark, or legal template can support many markets in normal time and fail many markets at once in stress.

Definition 12.1 (Infrastructure dependence graph)**.** *Let $G_I = (\mathcal{J}, \mathcal{M}, E)$ be a bipartite graph between implemented claims $j \in \mathcal{J}$ and infrastructure modules $m \in \mathcal{M}$. Edge (j, m) exists if claim j depends on module m. The module degree d_m is the number of claims depending on module m.*

Proposition 12.2 (Shared-module fragility)**.** *If module m fails and every dependent claim j suffers delivery loss $\ell_{jm} \geq 0$, total direct delivery loss is*

$$L_m = \sum_{j:(j,m)\in E} \ell_{jm}.$$

For fixed average loss per dependent claim, expected module-failure loss scales with module degree d_m. Infrastructure that creates large positive cost spillovers can also create large common-mode failure exposure.

Proof. The loss formula is additive by definition. If average loss is $\bar{\ell}_m = L_m/d_m$, then $L_m = d_m\bar{\ell}_m$, which scales linearly with degree for fixed $\bar{\ell}_m$. □

Remark 12.3 (Paper 3 connection)**.** *Whether this fragility makes a claim inadmissible is Paper 3's question. Paper 4 identifies the infrastructural source of the externality: reuse creates both economies of scale and correlated failure.*

Proposition 12.4 (Redundancy tradeoff)**.** *Suppose a claim can use one oracle at cost c with failure probability p and loss L, or two independent redundant oracles at cost $c + \Delta c$, failing only if both fail. Redundancy is privately or socially cost-justified when*

$$\Delta c < (p - p^2)L$$

for the relevant internalized or social loss L. If the claim creator internalizes only ϕL, redundancy is privately underchosen whenever

$$\phi(p - p^2)L < \Delta c < (p - p^2)L.$$

Proof. One oracle has expected failure loss pL. Two independent oracles have expected failure loss p^2L. The expected loss reduction is $(p - p^2)L$. Compare it to incremental cost Δc. The private underchoice condition follows by replacing L with ϕL in the private comparison. □

13 Empirical Predictions and Test Plan

The paper predicts that state-contingent claim markets appear and migrate when infrastructure modules improve. The empirical unit is not only a product category. It is a claim–wrapper–module triple (g, ρ, m).

13.1 Predictions

1. **Cross-category birth waves.** Improvements in shared modules–oracles, settlement rails, legal templates, identity, collateral, and automated compliance–should predict market births across multiple categories using the module.
2. **Boundary migration.** The same exposure should migrate across wrappers when relative wrapper costs cross: indemnity to parametric, OTC to exchange, bilateral to automated, insurance to capital-market risk transfer, or prediction venue to data-sponsored information market.
3. **Oracle refinement.** Better oracle granularity should reduce basis risk and dispute frequency, but may increase privacy, manipulation, or classification concerns.
4. **Settlement compression.** Faster and more final settlement should reduce collateral requirements and default haircuts, especially for small or high-frequency claims.
5. **Collateral standardization.** Standard collateral and margin conventions should increase transferability and secondary liquidity, but also create common liquidation channels.
6. **Information-timing asymmetry.** Data used as ex post verification should expand markets; the same data used as pre-contract public classification may shrink pooling or increase exclusion.
7. **Public-good underprovision.** Markets whose main value is outside decision improvement should require sponsorship, subsidy, data licensing, or public provision more often than markets with direct hedging flow.
8. **Purpose-dependent access.** Payoffs with strong hedging value but high manipulation or gambling risk should survive through access restrictions rather than unrestricted retail form.
9. **Infrastructure fragility.** Claims sharing an oracle, settlement rail, benchmark, collateral asset, or legal template should show correlated disruption when that module fails.

13.2 Observable module shocks

Module shock	Empirical examples
Oracle launch or upgrade	weather-station network, satellite feed, industry-loss index, audited API, benchmark methodology
Legal template or safe harbor	standardized contract, regulator no-action relief, model law, exchange rule approval
Settlement rail improvement	real-time payment system, clearing access, programmable escrow, custody standard
Collateral convention	margin methodology, eligible collateral expansion, central clearing, trust structure
Identity/access system	accredited status verification, insurable-interest proof, exposure attestation, jurisdiction controls
Reporting/tax clarity	tax classification, accounting treatment, transaction reporting rule
Dispute process	arbitration standard, resolution committee, court-recognized data source

13.3 Reduced-form market-birth specification

Let $Y_{jkt} = 1$ if claim j appears in category k at time t. Let M_{jmt} measure claim j's exposure to module m, and let I_{mt} measure module quality. A reduced-form hazard model is

$$\Pr(Y_{jkt} = 1) = \Lambda\left(\alpha_j + \mu_k + \tau_t + \beta V_{jt}^{res} - \gamma \Psi_{jkt} + \sum_{m \in \mathcal{M}} \theta_m M_{jmt} I_{mt} + \eta X_{jkt}\right),$$

where V_{jt}^{res} is the residual payoff value proxy from Paper 1 and X_{jkt} includes liquidity, access, and legal controls. The Paper 4 prediction is $\theta_m > 0$ for relevant module exposures, with cross-category effects.

13.4 Boundary-switch test

For exposure e, let $k_t(e)$ be the observed wrapper category. Estimate relative wrapper costs $\widehat{\kappa}_k(e;t)$. The model predicts

$$k_t(e) \in \arg\min_k \widehat{\kappa}_k(e;t)$$

subject to legal feasibility and distribution constraints. Boundary migration should occur when estimated cost rankings switch.

13.5 Information-timing tests

Data shocks should be classified by timing:

- Ex post verification shocks should reduce disputes, basis risk, and settlement delay, and increase market entry.
- Pre-contract public classification shocks should increase price discrimination and may reduce pooling, take-up by high-risk participants, or cross-subsidization.

- Post-price publication should generate measurable decision improvements outside the trading venue if the price has public-good value.

13.6 Candidate empirical surfaces

Candidate surfaces include catastrophe bonds, industry-loss warranties, parametric disaster insurance, weather derivatives, crop-risk products, energy and congestion claims, election and event markets, conditional prediction markets, tokenized real-world assets, revenue-share instruments, royalty markets, DeFi derivatives, automated market makers, stablecoin settlement rails, and oracle-supported smart contracts.

14 Relation to the Broader Research Program

This paper is the institutional bridge of the Trillion Markets program.

Paper 1: value and costly basis. Paper 1 studies $V(B \mid \mathcal{H})$, the value of a residual payoff subspace, and $\kappa(B, \mathcal{H}, T)$, the least-cost representation of that subspace. Paper 4 opens κ. It says that representation cost is the cost of a wrapper: oracle, data, collateral, compliance, settlement, dispute resolution, access, and reporting. It also says that the payoff direction actually added is $\mathsf{D}_\rho g$, not the advertised g.

Paper 2: liquidity. Paper 2 studies quote feasibility and liquidity regimes: factor-hedged dealer liquidity, residual-diversified bookmaker/insurer liquidity, and subsidized information-market liquidity. Paper 4 says those regimes are wrapper components. The same payoff can move from one liquidity mechanism to another as infrastructure changes.

Paper 3: admissibility. Paper 3 studies which claims should be allowed, restricted, subsidized, taxed, or prohibited. Paper 4 supplies the metadata admissibility needs: purpose, access, oracle, collateral, leverage, data rights, and settlement. It does not decide the normative frontier.

Paper 5: computational Coase theorem. Paper 5 studies agentic transaction costs and automated claim search. Paper 4 supplies the institutional search space. An agentic platform cannot merely discover a payoff vector; it must discover a wrapper ρ with cheap specification, reliable oracle, feasible collateral, suitable liquidity, compliant access, and admissible purpose.

Program state. The full program state can be written

$$\mathsf{State}_t = (S_t, I_t, N_t),$$

where $S_t = (\mathcal{H}_t, \mathsf{Q}_t, \kappa_t, L_t, A_t)$ is the compact core state, I_t is infrastructure and delivery, and N_t is the systemic-network slice used by Paper 3. Paper 4 is the theory of I_t and its relation to κ_t, L_t, A_t, and the actually delivered payoff.

Paper	Central object	Paper 4 connection
Paper 1	residual payoff value and costly basis	implemented payoff $\mathsf{D}_{\rho}g$ and representation cost modules
Paper 2	synthetic liquidity and residual risk	liquidity signature as wrapper component
Paper 3	claim-centric admissibility	purpose, access, oracle, collateral, and data metadata
Paper 4	shared claim infrastructure	payoff convergence plus wrapper divergence
Paper 5	agentic transaction costs and market search	search over (g, ρ), not just g

15 Non-Obvious Implications

1. The right primitive is not "prediction market," "insurance," or "derivative." It is an implemented payoff (g, ρ).
2. Product categories are fibers over a common payoff layer, not separate mathematical species.
3. Legal labels matter because they encode wrapper functions: solvency, disclosure, margin, access, tax, and enforcement.
4. A module improvement can create markets in categories far from the one that funded the improvement.
5. The same infrastructure that creates a parametric insurance market can create a derivative, event contract, reinsurance trigger, or tokenized claim.
6. Better data is market-creating when it supports specification or verification, and market-destroying when it reveals risk before pooling.
7. Oracle quality changes the payoff delivered, not only the cost of delivery.
8. Collateral is part of the payoff technology: undercollateralized claims deliver capped or default-contingent payoffs.
9. Settlement speed changes both implementation cost and delivered payoff through default and funding haircuts.
10. Public prediction prices are infrastructure outputs and may be underprovided because their users need not trade.
11. Purpose cannot be inferred from payoff. The same binary claim can be hedge, research subsidy, entertainment, manipulation, or gambling.
12. Access controls can preserve useful markets by preventing harmful purpose profiles.
13. Privacy-preserving eligibility is core infrastructure for personal, household, health, labor, and insurance-linked claims.
14. Payoff convergence can increase regulatory arbitrage unless regulation becomes wrapper-functional rather than label-formal.
15. Shared infrastructure creates both economies of scale and common-mode fragility.
16. A market-search algorithm must search over wrappers, not only payoff directions.
17. The convergence of state-contingent claim infrastructure is compatible with persistent legal and moral distinctions among claim uses.

16 Conclusion

The payoff layer of finance is simpler than its product taxonomy. Prediction markets, insurance policies, derivatives, parametric claims, catastrophe bonds, revenue shares, and many tokenized instruments are all state-contingent payoff functions. The complexity lies in the wrapper: how the payoff is specified, verified, collateralized, margined, distributed, traded, restricted, reported, disputed, and settled.

That distinction resolves a tension in the market-proliferation thesis. If one looks only at payoff functions, convergence seems to imply that all financial categories collapse into one. If one looks only at law and distribution, convergence seems overstated because insurance, derivatives, gambling, and securities remain different. The correct view is two-layered. At the payoff layer, they converge. At the wrapper layer, they remain different because purpose, liquidity, access, collateral, and externalities differ.

The infrastructure stack determines how far convergence goes. Oracles refine what can be resolved. Data permissions determine what can be used without violating privacy or destroying pools. Collateral and margin determine whether promised payoffs are credible. Settlement rails determine whether small and fast claims are economical. Legal templates and compliance determine which wrappers are admissible. Liquidity mechanisms determine whether a claim can be quoted. Dispute systems determine whether ambiguity destroys value after the fact. Improvements in any of these modules can spill across categories.

But the same infrastructure also carries limits. Data can arrive too early. Oracles can be manipulated. Collateral can transmit liquidations. Public prices can be underfunded. Purpose can transform a hedge into a weapon. Shared modules can fail systemically. The future of state-contingent claims is therefore not "everything becomes a prediction market" or "everything becomes insurance" or "everything becomes a derivative." It is a shared infrastructure for implementing payoff functions, with persistent differences in why people trade, who may participate, how liquidity is supplied, and which claims society admits.

The room is gone. The wrapper remains.

A Appendix A: A Finite-State Wrapper Example

Let $\Omega = \{1, 2, 3, 4\}$, with states representing combinations of rainfall $R \in \{low, high\}$ and crop price $P \in \{low, high\}$. A farmer's loss is

$$L = (100, 20, 80, 0),$$

where the states are ordered $(lowR, lowP), (lowR, highP), (highR, lowP), (highR, highP)$. Several wrappers can address this exposure:

1. Indemnity insurance pays $g_I = \min\{\max(L - d, 0), \ell\}$.
2. Parametric drought insurance pays $g_P = A\mathbf{1}_{\{R=low\}}$.
3. A crop-price derivative pays $g_D = B\mathbf{1}_{\{P=low\}}$.
4. A compound event contract pays $g_E = C\mathbf{1}_{\{R=low,P=low\}}$.

If only rainfall is verifiable, the best verifiable approximation to L is $\mathbb{E}[L \mid R]$, which is constant across price states within each rainfall atom. If both rainfall and price are verifiable, the space expands. If farm-level realized loss is verifiable and admissible, indemnity can approximate L directly. The different products are projections of the same target loss onto different verifiable sigma-algebras and wrapper constraints.

B Appendix B: Proof of Monotone Attainable Payoff Space

Let $\mathcal{A}(I)$ be the set of wrappers and target payoffs technically attainable under infrastructure state I, and define the attainable delivered payoff set

$$\overline{\mathcal{X}}(I) = \{\mathsf{D}_\rho g : (g, \rho) \in \mathcal{A}(I)\}.$$

Proposition B.1 (Attainable payoff monotonicity)**.** *If I' weakly expands feasible specification, oracle, collateral, legal, and settlement modules so that $\mathcal{A}(I) \subseteq \mathcal{A}(I')$ and $\mathsf{D}_\rho^{I'} g = \mathsf{D}_\rho^{I} g$ for old wrappers, then*

$$\overline{\mathcal{X}}(I) \subseteq \overline{\mathcal{X}}(I').$$

If new infrastructure also improves delivery for old wrappers, the set inclusion may fail literally because old payoffs are replaced by improved payoffs; in that case the correct monotone object is the feasible graph of target-wrapper pairs, not the delivered payoff set.

Proof. Under the stated assumptions, every old attainable pair remains attainable and delivers the same payoff under I'. Therefore every old delivered payoff belongs to the new delivered payoff set. If delivery changes, then the same pair maps to a different payoff, so the graph of feasible pairs expands but the image can move. ☐

Remark B.2. *This subtlety matters. Better collateral or settlement can improve a claim by changing $\mathsf{D}_\rho g$, not only by adding new payoffs. Image-set monotonicity is not the right statement when delivery quality changes.*

C Appendix C: Binary Event Claims Span Finite Partitions

Let $\mathcal{P} = \{A_1, \dots, A_n\}$ be a finite verifiable partition of Ω. The space of $\sigma(\mathcal{P})$-measurable payoffs is

$$L^2(\sigma(\mathcal{P})) = \left\{ \sum_{i=1}^{n} x_i \mathbf{1}_{A_i} : x_i \in \mathbb{R} \right\}.$$

Thus a complete set of Arrow securities for the partition spans every payoff measurable with respect to the partition. A binary-event infrastructure that can create and settle $\mathbf{1}_{A_i}$ for each atom can implement any payoff on the partition by portfolio. In practice, portfolio construction can be blocked by transaction costs, margin, access limits, indivisibility, legal restrictions, or manipulation risk.

D Appendix D: LMSR as Information Infrastructure

For K mutually exclusive outcomes, the logarithmic market scoring rule uses cost function

$$C(q) = b \log \left(\sum_{k=1}^{K} e^{q_k/b} \right),$$

with prices

$$p_k(q) = \frac{e^{q_k/b}}{\sum_{\ell=1}^{K} e^{q_\ell/b}}.$$

Starting from $q = 0$, worst-case loss is bounded by $b \log K$. The parameter b is therefore an explicit infrastructure budget for information-market depth. The sponsor is buying a probability vector over the outcome partition. Oracle ambiguity, nonexhaustive outcomes, and manipulation risk are outside the bound and must be handled by wrapper design.

E Appendix E: Bibliographic Notes for Development

A polished journal version should deepen three literatures beyond the formal spine. First, insurance law and economics: insurable interest, indemnity, adverse selection, moral hazard, and solvency. Second, market infrastructure: central clearing, margin, settlement finality, custody, and payment systems. Third, information-market design: scoring rules, subsidized liquidity, manipulation, and decision markets. The present draft includes canonical anchors rather than an exhaustive survey.

References

Akerlof, George A. 1970. “The Market for ‘Lemons’: Quality Uncertainty and the Market Mechanism.” *Quarterly Journal of Economics* 84(3): 488–500.

Allen, Franklin, and Douglas Gale. 1994. *Financial Innovation and Risk Sharing.* Cambridge, MA: MIT Press.

Arrow, Kenneth J. 1963. “Uncertainty and the Welfare Economics of Medical Care.” *American Economic Review* 53(5): 941–973.

Arrow, Kenneth J., and Gerard Debreu. 1954. “Existence of an Equilibrium for a Competitive Economy.” *Econometrica* 22(3): 265–290.

Black, Fischer, and Myron Scholes. 1973. “The Pricing of Options and Corporate Liabilities.” *Journal of Political Economy* 81(3): 637–654.

Coase, Ronald H. 1937. “The Nature of the Firm.” *Economica* 4(16): 386–405.

Cummins, J. David, and Mary A. Weiss. 2009. “Convergence of Insurance and Financial Markets: Hybrid and Securitized Risk-Transfer Solutions.” *Journal of Risk and Insurance* 76(3): 493–545.

Debreu, Gerard. 1959. *Theory of Value: An Axiomatic Analysis of Economic Equilibrium.* New Haven: Yale University Press.

DeMarzo, Peter, and Darrell Duffie. 1999. “A Liquidity-Based Model of Security Design.” *Econometrica* 67(1): 65–99.

Duffie, Darrell. 2001. *Dynamic Asset Pricing Theory.* 3rd ed. Princeton: Princeton University Press.

Duffie, Darrell, and Rohit Rahi. 1995. “Financial Market Innovation and Security Design: An Introduction.” *Journal of Economic Theory* 65(1): 1–42.

Froot, Kenneth A. 2001. “The Market for Catastrophe Risk: A Clinical Examination.” *Journal of Financial Economics* 60(2–3): 529–571.

Grossman, Sanford J., and Joseph E. Stiglitz. 1980. “On the Impossibility of Informationally Efficient Markets.” *American Economic Review* 70(3): 393–408.

Hanson, Robin. 2003. “Combinatorial Information Market Design.” *Information Systems Frontiers* 5(1): 107–119.

Hanson, Robin. 2007. “Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation.” *Journal of Prediction Markets* 1(1): 3–15.

Hayek, Friedrich A. 1945. “The Use of Knowledge in Society.” *American Economic Review* 35(4): 519–530.

Hirshleifer, Jack. 1971. “The Private and Social Value of Information and the Reward to Inventive Activity.” *American Economic Review* 61(4): 561–574.

Merton, Robert C. 1973. “Theory of Rational Option Pricing.” *Bell Journal of Economics and Management Science* 4(1): 141–183.

Radner, Roy. 1972. “Existence of Equilibrium of Plans, Prices, and Price Expectations in a Sequence of Markets.” *Econometrica* 40(2): 289–303.

Ross, Stephen A. 1976. “The Arbitrage Theory of Capital Asset Pricing.” *Journal of Economic Theory* 13(3): 341–360.

Roth, Alvin E. 2015. “The Theory and Practice of Market Design.” Nobel Prize Lecture.

Rothschild, Michael, and Joseph Stiglitz. 1976. “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information.” *Quarterly Journal of Economics* 90(4): 629–649.

Savage, Leonard J. 1954. *The Foundations of Statistics.* New York: Wiley.

Szabo, Nick. 1997. “Formalizing and Securing Relationships on Public Networks.” *First Monday* 2(9).

Wolfers, Justin, and Eric Zitzewitz. 2004. “Prediction Markets.” *Journal of Economic Perspectives* 18(2): 107–126.