The Costly Basis of Incomplete Markets

Verifiable Information, Residual Spanning Value, and Costly Basis Selection

Grant Stenger

Working paper

June 2026

Abstract

This paper studies market incompleteness as an endogenous consequence of costly claim creation. In a finite-state economy with common beliefs, mean-variance agents, zero-net-supply risky claims, and a risk-free asset, the gross risk-sharing value of a tradable claim system depends only on the residual payoff subspace it adds relative to the existing market span. This value is a coordinate-free quadratic form in a valuation-dispersion operator. In the zero-cost benchmark, the optimal residual basis is given by the leading eigenvectors of this operator, recovering the optimal-security logic of frictionless risk-sharing design. Actual markets do not choose arbitrary payoff directions. They choose verifiable, admissible contractual representations–indexes, swaps, futures, insurance triggers, revenue shares, and other payoff rules–with representation-dependent costs, distortions, and constraints. The paper's central result is an endogenous-resolution threshold: in an orthogonal residual dictionary, and in particular in the residual eigenbasis, a payoff direction becomes tradable exactly when its marginal valuation-dispersion eigenvalue exceeds its marginal representation and admissibility cost. This turns market creation into a costly-fidelity problem: concentrated spectra justify high-resolution bespoke claims, while flat or low-value tails favor coarse indexes, baskets, and standardized benchmarks. Technology is modeled as a refinement of a verifiable information structure, and the mean-variance geometry is shown to be the constant-curvature case of a local welfare theorem for general expected utility. The core thesis remains: value lives on residual payoff subspaces; cost, admissibility, and implementation failure live on contractual representations. Approximate indexing, dictionary selection, infrastructure under-entry, welfare wedges, and empirical market-birth predictions are consequences of this spine.

Keywords: incomplete markets; financial innovation; security design; market design; transaction costs; state-contingent claims; endogenous asset creation; market completeness; sparse approximation.
JEL codes: D47, D52, D53, D61, G12, G13, G18, L51.

Contents

1 Introduction

Complete-markets theory asks what allocations would be possible if sufficiently many state-contingent payoff directions were tradable. The benchmark is foundational, but it treats the market span as part of the environment. In actual economies, payoff directions are not automatically tradable. They must be specified, measured, verified, priced, collateralized, legally enforced, distributed, cleared, and settled. A state-contingent claim becomes a market only when some intermediary, platform, exchange, protocol, insurer, dealer, market maker, or public institution finds it privately or socially worthwhile to instantiate that payoff rule.

This paper turns the set of tradable claims into the object of analysis. The core thesis is:

> *The economic value of a new market is span invariant, but the cost of implementing that span is representation dependent. Markets are not merely missing claims; they are missing low-cost bases for valuable residual payoff directions.*

The distinction matters. A housing-risk exposure can be represented by an individual-home derivative, a metropolitan price-index future, a real-estate ETF, an insurance contract, a mortgage-linked security, or a swap on a regional benchmark. These instruments may span related payoff directions, but their verification, legal, liquidity, and settlement costs differ radically. The relevant economic object is not the label "future," "ETF," "insurance," "swap," or "event contract." It is the residual payoff subspace the instrument adds relative to already traded claims, together with the cost of implementing that subspace.

The first contribution is a geometry of market-span value. Let $\mathcal{X} = L^2_0(\mathbb{P})$ be the centered payoff space and let $\mathcal{H} \subseteq \mathcal{X}$ be the existing span of risky traded claims. Agents' risk-sharing motives generate a positive semidefinite valuation-dispersion operator Q. The welfare value of a span is

$$W(\mathcal{H}) = \frac{1}{2}\operatorname{tr}(\Pi_{\mathcal{H}}\mathsf{Q}).$$

A candidate claim system G contributes only through its residual subspace

$$B(G \mid \mathcal{H}) = \operatorname{col}((I - \Pi_{\mathcal{H}})G) \subseteq \mathcal{H}^{\perp}.$$

The incremental surplus is

$$V(B \mid \mathcal{H}) = \frac{1}{2}\operatorname{tr}(\Pi_B\mathsf{Q}).$$

This representation makes three facts immediate. First, a candidate system already spanned by existing markets adds no risk-sharing value. Second, a nonzero residual payoff direction can still be welfare-irrelevant if no agent has trade-generating valuation along it. Third, if implementation were free and an institution could choose any m new payoff directions, the optimal missing-market basis would be the leading eigenvectors of the residualized operator $(I - \Pi_{\mathcal{H}})\mathsf{Q}(I - \Pi_{\mathcal{H}})$. This frictionless eigenbasis result should be read as a benchmark, not as the paper's main novelty. The paper's contribution begins when those high-value directions must be implemented through costly, verifiable, admissible, institutionally feasible claim forms.

The second contribution moves one layer upstream from payoff vectors to verifiable information. A fact may be economically meaningful without being contractible. A technology/legal regime T determines a sigma-algebra $\mathcal{F}_T$ of state distinctions that can be observed, verified, enforced, and admitted into contracts. The attainable payoff space is

$$\mathcal{X}_T = L^2_0(\mathcal{F}_T).$$

Technology is therefore not merely a scalar that lowers cost. It can refine the contractible state space. If $\mathcal{F}_T \subseteq \mathcal{F}_{T'}$, the gross value of the refinement is the value of the newly verifiable residual subspace. This gives a formal version of the claim that sensors, APIs, audited data feeds, legal templates, indexes, oracles, and settlement rails make new markets possible by making new state distinctions contractible.

The third contribution is the costly-basis problem. Gross value is a property of subspaces; implementation cost is a property of representations. Let ρ denote a contractual representation–a payoff rule, legal form, data source, oracle, collateral rule, settlement convention, and venue–and let $\pi(\rho)$ be its payoff. For a residual subspace B, define the least-cost exact implementation

$$\kappa(B, \mathcal{H}, T) = \inf_{\rho:\ \mathrm{col}((I-\Pi_{\mathcal{H}})\pi(\rho))=B} C(\rho, \mathcal{H}, T),$$

where T is a technology vector. Private entry requires

$$\phi_B V(B \mid \mathcal{H}) > \kappa(B, \mathcal{H}, T) + \Psi_B,$$

while social entry requires

$$V(B \mid \mathcal{H}) - \kappa(B, \mathcal{H}, T) - E_B > 0.$$

The inequality is the entry comparison. The paper's substantive objects are the two sides of that comparison: a residual projection functional that depends on the current market span, and a least-cost envelope over admissible contractual bases.

The fourth contribution is a disciplined model of endogenous market resolution. A finite dictionary of candidate representations generates a market set M and therefore a span $\mathcal{H}(M)$. Adding one claim changes the residual value of other claims. Claims are payoff substitutes when they span overlapping residual directions. They are cost complements when they share oracles, data standards, legal templates, collateral conventions, clearing systems, or distribution. For an orthogonal dictionary, the listing rule is explicit: include a direction when its marginal valuation-dispersion term exceeds its marginal representation and admissibility cost. Applied to the residual eigenbasis, the rule says that the resolution at which an economy prices itself is set where the residual eigenvalue tail crosses marginal implementation cost. Concentrated spectra support bespoke, high-fidelity claims; flat or low spectra make coarse indexes and standardized baskets privately and socially optimal. The general market set is a fixed point or, for a platform or planner, the solution to a penalized projection problem:

$$\max_{M \subseteq \mathcal{J}} \ \phi V(\mathrm{span}\, D_M \mid \mathcal{H}) - \sum_{j \in M} c_j - \Psi(M).$$

This nests market listing, index construction, and sparse basis selection.

The fifth contribution is institutional. Many missing markets are missing not because the individual residual directions have no value, but because the first entrant must build infrastructure that later claims use. A corrected infrastructure-under-entry result uses cluster surplus

$$W(\mathcal{H} + \mathrm{span}\{g_j : j \in J\}) - W(\mathcal{H})$$

rather than summing overlapping one-claim surpluses. This avoids double-counting and explains why public data, legal templates, benchmark construction, or platform provision can be socially valuable even when no single claim can privately pay the fixed cost.

The paper deliberately avoids the claim that more markets are always good. Falling claim-creation costs can increase the dimension of the tradable span while reducing welfare if high-externality claims cross the private threshold first. The policy object is therefore not simply "market versus no market," but claim-level governance: allow, standardize, subsidize, tax, margin, restrict, or prohibit.

Reader's map and hierarchy

The paper is organized around a hierarchy of results. The *core theorem spine* is deliberately short: theorem 4.1 values a market span using the operator Q, theorem 4.3 gives the frictionless optimal residual basis, theorem 10.9 states the endogenous-resolution threshold, and theorem 7.1 gives the local expected-utility extension. Around those the paper records several useful projection facts—proposition 4.6 and lemma 4.13—whose role is to keep the geometry coordinate-free, not to carry the novelty by themselves. The substantive contribution is proposition 10.5: gross surplus lives on residual payoff subspaces, while implementation cost, admissibility, and failure modes live on contractual representations. The upstream verifiability layer, propositions 5.1 and 5.3, explains how technology changes which state distinctions are contractible.

The *main economic mechanisms* are approximate bases and indexing (propositions 9.3 and 9.4 and example 9.5), endogenous market resolution (corollary 10.8 and theorem 10.9), finite-dictionary market design (propositions 10.7 and 10.11), and infrastructure under-entry (proposition 12.1). These show how the abstract geometry becomes finance: benchmarks, templates, oracles, legal forms, collateral conventions, and distribution rails determine which residual payoff directions are actually implemented and at what fidelity.

The *consequences* are private-social wedges, admissible completion frontiers, and empirical market-birth predictions. Fat-tailed proliferation is kept as an appendix-level companion result, while synthetic liquidity is left to Paper 2, so that the first paper remains a theory of costly residual basis selection rather than a general manifesto about financial innovation. Heterogeneous beliefs are treated as a scope boundary for the welfare claims and as a companion-paper frontier, because disagreement changes the interpretation of trading gains.

2 Related Literature and Positioning

The paper is closest to several literatures. The novelty is not that markets are incomplete, not that financial innovation exists, and not that security design matters. The novelty is the combination of residual span value, representation-dependent implementation costs, technology-specific cost compression, and private-social admissibility wedges.

Closest-literature wedge. A skeptical reader may ask whether this is simply security design with transaction costs. The answer is narrower and more precise. Prior work studies which assets are created and how securities expand risk sharing. This paper isolates an invariant that is often implicit: gross market-completion value lives on residual payoff subspaces, while implementation cost lives on contractual representations. The gap between those two objects is the costly-basis problem.

Complete and incomplete markets. Arrow and Debreu's complete-markets framework gives a benchmark in which commodities can be indexed by state and date, and competitive equilibrium can be studied in a specified commodity space (Arrow and Debreu, 1954; Debreu, 1959). Radner's sequential equilibrium formulation studies plans, prices, and price expectations when markets are sequential and securities are traded over time (Radner, 1972). The incomplete-markets literature studies equilibrium and welfare once the asset structure is given, including the possibility that incomplete markets can be constrained inefficient (Hart, 1975; Geanakoplos et al., 1990; Magill and Quinzii, 1996). This paper asks a prior institutional question: which payoff directions become tradable at all?

Transaction costs and institutional boundaries. Coase explained the boundary between markets and firms by the cost of using the price system (Coase, 1937). This paper moves the same logic inside finance. Some exposures remain inside firms, households, insurers, bilateral contracts, collateralized lending relationships, government programs, and informal norms because transforming them into tradable state-contingent claims is too costly.

Information and price discovery. Hayek emphasized the price system as a mechanism for aggregating dispersed knowledge (Hayek, 1945). Grossman and Stiglitz showed that perfectly informationally efficient markets are inconsistent with costly information acquisition because informed traders must be compensated (Grossman and Stiglitz, 1980). Hirshleifer showed that information can reduce social risk-sharing opportunities when it arrives before trade (Hirshleifer, 1971). The present paper asks an even earlier question: before asking how informative a price is, one must ask whether that price exists. Costly information production affects not only price accuracy but market existence, and the timing of information affects whether better data creates markets or destroys insurance value.

Financial innovation and security design. This is the closest prior literature. Allen and Gale study financial innovation and risk sharing (Allen and Gale, 1994); Duffie and Rahi survey financial market innovation and security design in incomplete markets (Duffie and Rahi, 1995); Pesendorfer, Demange and Laroque, Ohashi, Hara, and Bisin study endogenous innovation, incomplete asset structures, and security design in general equilibrium environments (Pesendorfer, 1995; Demange and Laroque, 1995; Ohashi, 1995; Hara, 1995; Bisin, 1998). Athanasoulis and Shiller derive optimal risk-sharing securities using eigenvectors of transformed income-variance objects (Athanasoulis and Shiller, 2000, 2001). The eigenbasis result below is therefore best understood as the frictionless zero-cost benchmark within this paper's notation. The contribution is not the broad claim that assets are endogenous, nor the claim that eigenvector-based optimal securities are new. It is the pairing of a residual-subspace value functional, invariant to contractual basis, with a representation-dependent implementation-cost envelope.

Financial innovation under disagreement and private incentives. The common-beliefs baseline below gives the surplus formulas a welfare interpretation. With heterogeneous beliefs, new assets may support risk sharing, information aggregation, or speculation. Simsek shows that financial innovation can increase speculative variance under belief disagreement and that a profit-seeking market maker may introduce speculative assets that raise portfolio risk (Simsek, 2013). Carvajal, Rostek, and Weretka show that private innovation incentives can leave markets incomplete even when innovation is essentially costless (Carvajal et al., 2012). These papers discipline the interpretation of the private-entry condition: cheap market creation is not automatically welfare-improving, and surplus capture is itself an institutional object.

Macro markets, factor structure, market design, and sparse approximation. Shiller's macro-markets agenda argues that societies lack markets for large aggregate risks and proposes institutions to manage them (Shiller, 1993). Ross's arbitrage pricing theory motivates a low-dimensional factor view of risky payoffs (Ross, 1976). Market design studies how institutions can be engineered when prices alone do not perform all allocation tasks (Roth, 2015). The costly-basis problem also has a mathematical analogy to sparse approximation and dictionary selection (Natarajan, 1995): market creation can be written as choosing a costly dictionary of payoff directions to approximate high-value residual valuation factors.

Recent security design and algorithmic selection. Recent surveys emphasize that security design spans classic corporate finance, intermediation, fintech, sustainability, and health-care finance (Allen and Bărbălău, 2024). Recent mechanism-design work derives optimal securities for risk-averse investors from primitive constraints (Gershkov et al., 2025). This paper differs by asking when designed payoff directions become market-wide tradable subspaces under verifiability and representation costs. On the algorithmic side, the dictionary-selection problem connects to approximate submodularity and sparse approximation: when the residual payoff dictionary is sufficiently well conditioned, greedy listing rules can have approximation guarantees (Das and Kempe, 2018).

What is new. Prior work endogenizes securities, studies security design, and shows that incomplete markets matter. This paper abstracts from many general-equilibrium complications to isolate a canonical geometric decomposition. The gross value of a candidate market is a function of the residual payoff subspace it spans. The market form that appears is the cheapest admissible representation of that subspace. This distinction makes it possible to say when a new instrument adds market-completion value, when it is merely a re-basing of existing span, why indexes often arrive before bespoke claims, why infrastructure is underprovided, and how technology-specific cost shocks create predictable market-birth waves.

Literature	Core object	This paper's additional object
Arrow-Debreu	allocations with complete contingent commodity markets	creation of the claim span itself
Radner and incomplete markets	equilibrium for a given asset structure	endogenous movement of the asset structure
Coase	firm-market boundary under transaction costs	tradable-claim boundary under claim-creation costs
Financial innovation/security design	endogenous securities and risk sharing	residual-subspace value plus representation-cost envelope
Shiller macro markets	institutions for large untraded risks	selection criterion for which risks get private or public markets
Market design	rules for allocation and exchange	payoff rules as admissible, costly market objects
Sparse approximation	costly basis/dictionary selection	economic interpretation as claim creation

3 Finite-State Environment

3.1 States, payoffs, and covariance geometry

Let $\Omega = \{1, \dots, K\}$ with common probability $\mathbb{P}(\omega) > 0$ for all states. A payoff is a vector $x \in \mathbb{R}^K$. Define

$$\mathbb{E}[x] = \sum_{\omega \in \Omega} \mathbb{P}(\omega) x(\omega), \qquad \widetilde{x} = x - \mathbb{E}[x]\mathbf{1}.$$

The risk-free payoff $\mathbf{1}$ is available. Since constants can be transferred through the risk-free asset, risky payoff directions live in the centered Hilbert space

$$\mathcal{X} = L_0^2(\mathbb{P}) = \left\{x \in \mathbb{R}^K : \mathbb{E}[x] = 0\right\}$$

with inner product and norm

$$\langle x, y\rangle = \mathbb{E}[xy], \qquad \|x\|^2 = \mathbb{E}[x^2].$$

For a subspace $\mathcal{H} \subseteq \mathcal{X}$, let $\Pi_{\mathcal{H}}$ be the orthogonal projection and $\mathcal{H}^{\perp}$ its orthogonal complement. In finite states, all subspaces are closed and all projections are well-defined.

3.2 Matrix and normalization conventions

The paper uses the covariance inner product throughout. If $R = (r_1, \ldots, r_m)$ is a residual claim system, write $R(\omega) = (r_1(\omega), \ldots, r_m(\omega))^{\top}$. For $x \in \mathcal{X}$,

$$R^{\top}x \equiv \mathbb{E}[R(\omega)x(\omega)] \in \mathbb{R}^m, \qquad \Sigma_R \equiv R^{\top}R \equiv \mathbb{E}[R(\omega)R(\omega)^{\top}] \in \mathbb{R}^{m\times m}.$$

Equivalently, if the columns of the $K \times m$ matrix R are payoff vectors and $D = \operatorname{diag}(\mathbb{P}(1), \ldots, \mathbb{P}(K))$, then $R^{\top}x$ means $R'Dx$ and $R^{\top}R$ means $R'DR$. A residual basis need not be orthonormal. Normalization is handled by Σ_R in coordinates or by the projection $\Pi_{\mathrm{col}(R)}$ in coordinate-free statements. Span value therefore does not depend on units, labels, leverage, or invertible changes of basis.

Assumption 3.1 (Baseline maintained assumptions)**.** *Unless otherwise stated, the model uses common beliefs $\mathbb{P}$, centered risky payoff directions with a risk-free asset available, zero-net-supply risky claims, mean-variance certainty equivalents, frictionless trading within the existing span $\mathcal{H}$, and finite states. Heterogeneous beliefs and positive net supply are treated as extensions in the appendix.*

Residual terminology. This paper uses *payoff residual* for the component of a candidate claim not spanned by existing markets:

$$r_g = (I - \Pi_{\mathcal{H}})g.$$

It uses *valuation-residual surplus* for the risk-sharing value of the residual payoff subspace, weighted by valuation dispersion:

$$V(\operatorname{span}\{r_g\} \mid \mathcal{H}) = \frac{1}{2}\langle u_g, \mathsf{Q}u_g\rangle, \qquad u_g = r_g/\|r_g\|.$$

when $r_g \neq 0$. Thus payoff novelty and market-completion value are different objects. A claim can be payoff-residual relative to $\mathcal{H}$ but have little value if agents do not have dispersed marginal valuations along that direction. The companion synthetic-liquidity paper uses *hedging residual* for the risk left after a dealer hedges factor exposures. The shared projection language is useful, but the objects should not be conflated.

3.3 Agents and existing markets

There are agents $a = 1, \ldots, A$. Agent a has endowment $W_a \in \mathbb{R}^K$, risk aversion $\gamma_a > 0$, and mean-variance certainty equivalent

$$U_a(c) = \mathbb{E}[c] - \frac{\gamma_a}{2}\operatorname{Var}(c).$$

Let $\theta_a = 1/\gamma_a$ denote risk tolerance and $\Theta = \sum_a \theta_a$. Let $\mathcal{H} \subseteq \mathcal{X}$ denote the span of existing risky traded claims. Given $\mathcal{H}$, agent a can choose a centered payoff transfer $z_a \in \mathcal{H}$. Feasibility requires

$$\sum_{a=1}^{A} z_a = 0.$$

Consumption under existing markets is

$$c_a^{\mathcal{H}} = W_a + z_a^{\mathcal{H}},$$

where $(z_a^{\mathcal{H}})_a$ is the competitive risk-sharing allocation over $\mathcal{H}$ under zero net supply. Under common beliefs and the quadratic/CARA-normal environment, the competitive allocation over a given span coincides with the planner's allocation subject to $\sum_a z_a = 0$. This is why the surplus formulas below have a welfare interpretation in the baseline.

Define the Riesz representation of agent a's marginal valuation over risky payoff directions:

$$Y_a = -\gamma_a \widetilde{W}_a \in \mathcal{X}, \qquad \bar{Y} = \frac{\sum_a \theta_a Y_a}{\Theta}, \qquad Z_a = Y_a - \bar{Y}.$$

For $x \in \mathcal{X}$, the risk part of the directional derivative of $U_a(W_a + tx)$ at $t = 0$ is $\langle Y_a, x \rangle$. Thus Y_a is not an asset payoff; it is the Riesz vector representing the agent's marginal willingness to receive payoff in each state. The risk-tolerance-weighted common component $\bar{Y}$ affects prices but not zero-net-supply risk-sharing gains. The trade-generating vector is Z_a.

3.4 The valuation-dispersion operator

Define the positive semidefinite operator $\mathsf{Q} : \mathcal{X} \to \mathcal{X}$ by

$$\mathsf{Q} = \sum_{a=1}^{A} \theta_a Z_a \otimes Z_a, \qquad (Z_a \otimes Z_a)x = Z_a \langle Z_a, x \rangle .$$

This is the paper's canonical value-side object: it records the directions in which agents' marginal valuations differ after risk-tolerance weighting. Then, for any subspace $B \subseteq \mathcal{X}$,

$$\operatorname{tr}(\Pi_B \mathsf{Q}) = \sum_a \theta_a \left\| \Pi_B Z_a \right\|^2 .$$

The operator Q is the sufficient statistic for risk-sharing value in the common-beliefs mean-variance baseline. It compresses all endowments and risk tolerances into a finite-dimensional object whose eigenvectors are the highest-value directions of missing risk-sharing.

Remark 3.2 (When Q is fixed)**.** *The comparative statics below treat Q as fixed while the market span changes. In the finite-state mean-variance baseline with common beliefs, fixed endowments, and zero-net-supply risky claims, this is the relevant local value-side sufficient statistic after frictionless trade in the inherited span $\mathcal{H}$. In richer environments, new markets can change information, endowments, participation, beliefs, or risk tolerances; then Q should be interpreted as a local operator evaluated at the pre-entry allocation, or replaced by the appropriate post-entry valuation-dispersion object.*

Lemma 3.3 (Rank of welfare-relevant risk)**.** *The valuation-dispersion operator satisfies*

$$\operatorname{rank}(\mathsf{Q}) \leq \min\{K - 1, A - 1\}.$$

Proof. The operator's range is contained in $\operatorname{span}\{Z_1, \dots, Z_A\} \subseteq \mathcal{X}$, so its rank is at most $K - 1$ and at most the dimension of the span of the Z_a's. Moreover,

$$\sum_a \theta_a Z_a = \sum_a \theta_a (Y_a - \bar{Y}) = 0,$$

so the A vectors are linearly dependent when all $\theta_a > 0$. Hence the rank is at most $A - 1$. □

Remark 3.4. *The Arrow-Debreu state space may have $K-1$ risky directions, but only the directions in the range of Q have risk-sharing value in this baseline. A complete set of state claims is sufficient for full risk sharing, but it is not necessary when the economically relevant valuation dispersion is low-dimensional.*

4 Risk-Sharing Value of a Market Span

4.1 Global value of an existing span

Theorem 4.1 (Risk-sharing value of a span)**.** *For any existing market span $\mathcal{H} \subseteq \mathcal{X}$, the maximum aggregate certainty-equivalent gain from allowing zero-net-supply trades in $\mathcal{H}$, relative to no risky trade, is*

$$W(\mathcal{H}) = \frac{1}{2}\sum_{a=1}^{A} \theta_a \left\|\Pi_{\mathcal{H}} Z_a\right\|^2 = \frac{1}{2}\operatorname{tr}(\Pi_{\mathcal{H}}\mathsf{Q}).$$

The equilibrium risky transfer to agent a is

$$z_a^{\mathcal{H}} = \theta_a \Pi_{\mathcal{H}} Z_a.$$

Proof. Since expected aggregate consumption is unchanged by zero-net-supply centered trades, aggregate expected utility changes only through covariance and variance terms. For transfers $z_a \in \mathcal{H}$ with $\sum_a z_a = 0$, the aggregate certainty-equivalent gain is

$$\sum_a \left\{ \langle Y_a, z_a \rangle - \frac{\gamma_a}{2}\left\|z_a\right\|^2 \right\}.$$

Form the Lagrangian in $\mathcal{H}$, with multiplier $\lambda \in \mathcal{H}$ for $\sum_a z_a = 0$:

$$\mathcal{L} = \sum_a \left\{ \langle \Pi_{\mathcal{H}} Y_a, z_a \rangle - \frac{\gamma_a}{2}\left\|z_a\right\|^2 \right\} - \left\langle \lambda, \sum_a z_a \right\rangle.$$

The first-order condition is $\Pi_{\mathcal{H}} Y_a - \gamma_a z_a - \lambda = 0$, so $z_a = \theta_a(\Pi_{\mathcal{H}} Y_a - \lambda)$. Imposing $\sum_a z_a = 0$ gives $\lambda = \Pi_{\mathcal{H}} \bar{Y}$. Therefore $z_a^{\mathcal{H}} = \theta_a \Pi_{\mathcal{H}}(Y_a - \bar{Y}) = \theta_a \Pi_{\mathcal{H}} Z_a$. Substitution into the concave quadratic objective yields

$$W(\mathcal{H}) = \frac{1}{2}\sum_a \gamma_a \left\|z_a^{\mathcal{H}}\right\|^2 = \frac{1}{2}\sum_a \theta_a \left\|\Pi_{\mathcal{H}} Z_a\right\|^2 = \frac{1}{2}\operatorname{tr}(\Pi_{\mathcal{H}}\mathsf{Q}).$$

□

Remark 4.2 (Prices versus surplus)**.** *A common component of valuation moves prices but not zero-net-supply trading surplus. Surplus comes from dispersion in residual marginal valuations, weighted by risk tolerance.*

4.2 The frictionless optimal missing-market basis

If claim creation were free and an institution could choose any m additional payoff directions, the optimal span has a closed-form characterization. This is the paper's benchmark result. It recovers, in the present residual-payoff notation, the optimal-security/eigenvector logic associated with Athanasoulis and Shiller's risk-sharing design work (Athanasoulis and Shiller, 2000, 2001). The costly-basis contribution starts after this benchmark: actual institutions cannot list arbitrary eigenvectors of valuation dispersion; they must choose contract forms, data, oracles, legal wrappers, collateral rules, and venues.

Theorem 4.3 (Frictionless optimal residual basis)**.** *Fix an existing span $\mathcal{H} \subseteq \mathcal{X}$ and let*

$$\mathsf{Q}_{\mathcal{H}} = (I - \Pi_{\mathcal{H}})\mathsf{Q}(I - \Pi_{\mathcal{H}})$$

be the residual valuation-dispersion operator on $\mathcal{H}^{\perp}$. Let its eigenvalues on $\mathcal{H}^{\perp}$ be ordered

$$\lambda_1 \geq \lambda_2 \geq \cdots \geq 0.$$

Among all m-dimensional residual subspaces $B \subseteq \mathcal{H}^{\perp}$, the maximum incremental surplus is

$$\max_{\substack{B \subseteq \mathcal{H}^{\perp} \\ \dim B = m}} V(B \mid \mathcal{H}) = \frac{1}{2} \sum_{\ell=1}^{m} \lambda_{\ell},$$

where

$$V(B \mid \mathcal{H}) = W(\mathcal{H} + B) - W(\mathcal{H}) = \frac{1}{2} \operatorname{tr}(\Pi_B \mathsf{Q}).$$

Any span of m eigenvectors associated with the m largest residual eigenvalues is optimal.

Proof. Because $B \subseteq \mathcal{H}^{\perp}$, $\Pi_B \mathsf{Q}$ has the same trace as $\Pi_B \mathsf{Q}_{\mathcal{H}}$. Therefore the problem is

$$\max_{\dim B = m} \operatorname{tr}(\Pi_B \mathsf{Q}_{\mathcal{H}}).$$

This is the Ky Fan maximum principle for a symmetric positive semidefinite operator. The maximum is the sum of the largest m eigenvalues and is attained by the corresponding eigenspace. □

Corollary 4.4 (Value gap of a realized span)**.** *If the realized market expansion has residual dimension m and residual span B, its basis inefficiency relative to the frictionless best m-dimensional expansion is*

$$\mathcal{G}_m(B \mid \mathcal{H}) = \frac{1}{2} \sum_{\ell=1}^{m} \lambda_{\ell} - \frac{1}{2} \operatorname{tr}(\Pi_B \mathsf{Q}) \geq 0.$$

This gap measures the loss from implementing an institutionally cheap basis rather than the frictionless high-value basis.

Remark 4.5 (Why this theorem matters)**.** *The theorem gives a canonical benchmark. It says what an omniscient, zero-cost market designer would span first. The result is a market-completion analogue of principal components analysis: the frictionless missing-market basis is spectral, while the actual implemented basis is chosen from costly, admissible contractual representations. The theorem itself is a benchmark, not the main claim of novelty; the economic contribution is the wedge between this benchmark and the costly basis selected in real institutions.*

4.3 Incremental value of a candidate residual subspace

Let G be a $K \times m$ matrix of candidate risky claim payoffs, with columns centered in $\mathcal{X}$. Relative to an existing span $\mathcal{H}$, define the residual claim matrix

$$G^{\perp} = (I - \Pi_{\mathcal{H}})G$$

and its residual subspace

$$B(G \mid \mathcal{H}) = \operatorname{col}(G^{\perp}) \subseteq \mathcal{H}^{\perp}.$$

The gross incremental value of adding G is

$$S_G(\mathcal{H}) = W(\mathcal{H} + B(G \mid \mathcal{H})) - W(\mathcal{H}).$$

Because $B(G \mid \mathcal{H}) \perp \mathcal{H}$, $\Pi_{\mathcal{H}+B} = \Pi_{\mathcal{H}} + \Pi_B$.

Proposition 4.6 (Coordinate-free incremental surplus)**.** *The gross incremental risk-sharing surplus from adding candidate claim system G to existing span $\mathcal{H}$ is*

$$S_G(\mathcal{H}) = \frac{1}{2}\sum_{a=1}^{A} \theta_a \left\|\Pi_{B(G|\mathcal{H})} Z_a\right\|^2 = \frac{1}{2}\operatorname{tr}(\Pi_{B(G|\mathcal{H})}\mathsf{Q}).$$

In particular, $S_G(\mathcal{H}) = 0$ if and only if $\Pi_{B(G|\mathcal{H})} Z_a = 0$ for all agents with positive risk tolerance. A sufficient condition is $B(G \mid \mathcal{H}) = \{0\}$, i.e. every candidate payoff is already spanned by existing markets.

Proof. By theorem 4.1,

$$S_G(\mathcal{H}) = \frac{1}{2}\sum_a \theta_a \left(\|\Pi_{\mathcal{H}+B} Z_a\|^2 - \|\Pi_{\mathcal{H}} Z_a\|^2\right).$$

Since $B \subseteq \mathcal{H}^\perp$, the projections onto $\mathcal{H}$ and B are orthogonal and $\Pi_{\mathcal{H}+B} = \Pi_{\mathcal{H}} + \Pi_B$. The Pythagorean theorem gives the first expression. The trace expression follows from the definition of Q. □

Remark 4.7 (Zero residual versus zero value)**.** *If $B(G \mid \mathcal{H}) = \{0\}$, the candidate system is redundant and adds no value. The converse is not logically necessary: a nonzero residual direction can have zero risk-sharing value if no agent has residual valuation for it. Market incompleteness is therefore not the same as welfare-relevant incompleteness.*

4.4 Equilibrium pricing and trading of a residual claim system

Let $R = (r_1, \ldots, r_m)$ be a full-rank basis for a residual subspace $B \subseteq \mathcal{H}^\perp$. Let

$$\Sigma_R = R^\top R = \mathbb{E}[R(\omega)R(\omega)^\top]$$

be its covariance matrix under the convention in Section 3.2. Full rank means Σ_R is positive definite. If agent a buys $q_a \in \mathbb{R}^m$ units of R at price vector $p \in \mathbb{R}^m$, the incremental final payoff is $q_a^\top R - q_a^\top p\,\mathbf{1}$. Conditional on the existing allocation, define the residual valuation vector

$$b_a(R \mid \mathcal{H}) = \mathbb{E}[R] - \gamma_a \operatorname{Cov}(c_a^{\mathcal{H}}, R).$$

In the centered common-beliefs baseline, $\mathbb{E}[R] = 0$. Because $R \perp \mathcal{H}$ and $z_a^{\mathcal{H}} \in \mathcal{H}$, this can be written as $b_a = R^\top Y_a$.

The next lemma is the load-bearing step behind the exact additive decomposition. Existing trade moves each agent's allocation inside $\mathcal{H}$; residual claims lie in $\mathcal{H}^\perp$. Under the constant-curvature baseline, those two facts imply that marginal valuations in residual directions can be measured using the original valuation vectors Y_a even after agents have already traded in $\mathcal{H}$.

Lemma 4.8 (Residual valuations after existing trade)**.** *Let $R \subseteq \mathcal{H}^\perp$ be centered. Then*

$$b_a(R \mid \mathcal{H}) = R^\top Y_a, \qquad p^* = R^\top \bar{Y}, \qquad b_a - p^* = R^\top Z_a.$$

Thus the local price/demand formula below is the coordinate version of the global span-value formula.

Proof. Since R is centered, $\mathbb{E}[R] = 0$. Since $z_a^{\mathcal{H}} \in \mathcal{H}$ and $R \perp \mathcal{H}$, $\operatorname{Cov}(z_a^{\mathcal{H}}, R) = 0$. Hence

$$b_a = -\gamma_a \operatorname{Cov}(W_a, R) = R^\top(-\gamma_a \widetilde{W}_a) = R^\top Y_a.$$

The expression for p^* follows by risk-tolerance averaging, and $b_a - p^* = R^\top(Y_a - \bar{Y}) = R^\top Z_a$. □

Remark 4.9 (Why the residual-valuation lemma matters)**.** *Lemma 4.8 is why $W(\mathcal{H}+B)-W(\mathcal{H})$ can be computed by projecting the same operator Q onto $B \subseteq \mathcal{H}^{\perp}$. The existing market allocation changes agents' consumptions, but in the quadratic baseline it changes them only along directions orthogonal to the new residual subspace. This is the technical heart of the clean additive formula; it is also the step that becomes allocation- and metric-dependent once state-dependent curvature is allowed.*

Lemma 4.10 (Projection formula for a residual basis)**.** *Let R be full rank and let $B = \mathrm{col}(R)$. For any $z \in \mathcal{X}$,*

$$\Pi_B z = R\Sigma_R^{-1}R^\top z, \qquad \|\Pi_B z\|^2 = (R^\top z)^\top \Sigma_R^{-1}(R^\top z).$$

Proof. The projection $\Pi_B z$ has the form $R\alpha$. Orthogonality of $z - R\alpha$ to the columns of R gives $R^\top(z - R\alpha) = 0$. Since $R^\top R = \Sigma_R$, $\alpha = \Sigma_R^{-1}R^\top z$. Substitution gives both displayed formulas. □

The incremental certainty-equivalent gain is

$$\Delta_a(q_a; p) = q_a^\top(b_a - p) - \frac{\gamma_a}{2} q_a^\top \Sigma_R q_a.$$

Proposition 4.11 (Incremental equilibrium and surplus)**.** *Suppose R is full rank and introduced in zero net supply. Then the unique linear-demand equilibrium has*

$$q_a(p) = \theta_a \Sigma_R^{-1}(b_a - p), \qquad p^* = \frac{\sum_a \theta_a b_a}{\Theta},$$

with equilibrium positions

$$q_a^* = \theta_a \Sigma_R^{-1}(b_a - p^*).$$

Gross incremental surplus is

$$S_R(\mathcal{H}) = \frac{1}{2}\sum_a \theta_a (b_a - p^*)^\top \Sigma_R^{-1}(b_a - p^*).$$

This matrix formula equals the coordinate-free expression in proposition 4.6.

Proof. Agent a's objective is a strictly concave quadratic in q_a. The first-order condition gives

$$b_a - p - \gamma_a \Sigma_R q_a = 0,$$

so $q_a(p) = \theta_a \Sigma_R^{-1}(b_a - p)$. Market clearing requires $\sum_a q_a(p) = 0$, hence

$$\Sigma_R^{-1}\left(\sum_a \theta_a b_a - \Theta p\right) = 0,$$

which implies $p = p^*$. Substituting the optimized value of a quadratic objective gives

$$\max_q \left\{ q^\top(b_a - p) - \frac{\gamma_a}{2} q^\top \Sigma_R q \right\} = \frac{\theta_a}{2}(b_a - p)^\top \Sigma_R^{-1}(b_a - p).$$

Summing over agents at $p = p^*$ gives the surplus formula. Finally, by lemma 4.8, $b_a - p^* = R^\top Z_a$. Applying lemma 4.10 with $z = Z_a$ gives

$$\|\Pi_B Z_a\|^2 = (b_a - p^*)^\top \Sigma_R^{-1}(b_a - p^*),$$

which proves equivalence with proposition 4.6. □

Corollary 4.12 (One-claim formula and corrected two-agent sign)**.** *For a single residual claim $h \in \mathcal{H}^\perp$, let $\sigma_h^2 = \mathrm{Var}(h)$, $\tau_a = (\gamma_a \sigma_h^2)^{-1}$, and $b_a = b_a(h \mid \mathcal{H})$. Then*

$$p^* = \frac{\sum_a \tau_a b_a}{\sum_a \tau_a}, \qquad S_h(\mathcal{H}) = \frac{1}{2} \sum_a \tau_a (b_a - p^*)^2.$$

For two agents, define $\Delta b = b_1 - b_2$. Then

$$q_1^* = -q_2^* = \frac{\Delta b}{(\gamma_1 + \gamma_2)\,\mathrm{Var}(h)}, \qquad S_h(\mathcal{H}) = \frac{(\Delta b)^2}{2(\gamma_1 + \gamma_2)\,\mathrm{Var}(h)}.$$

Thus the higher-valuation agent buys the residual claim.

4.5 Residual and basis invariance

Lemma 4.13 (Residual span and basis invariance)**.** *Let G be any candidate claim matrix and let $R = (I - \Pi_{\mathcal{H}})G$. Then*

$$S_G(\mathcal{H}) = S_R(\mathcal{H}).$$

If $\mathrm{rank}(R) = 0$*, then $S_G(\mathcal{H}) = 0$. If R has full column rank and $A \in \mathbb{R}^{m \times m}$ is invertible, then*

$$S_{RA}(\mathcal{H}) = S_R(\mathcal{H}).$$

Thus gross risk-sharing surplus depends only on the residual subspace $\mathrm{col}((I - \Pi_{\mathcal{H}})G)$*, not on the contractual basis used to represent it.*

Proof. First, $\mathcal{H} + \mathrm{col}(G) = \mathcal{H} + \mathrm{col}((I - \Pi_{\mathcal{H}})G)$, so adding G and adding R generate the same expanded payoff span. Hence they have the same W-increment. If $\mathrm{rank}(R) = 0$, then $\mathcal{H} + \mathrm{col}(G) = \mathcal{H}$, so the increment is zero. If A is invertible, $\mathrm{col}(RA) = \mathrm{col}(R)$, so proposition 4.6 gives $S_{RA} = S_R$. In coordinates, $\Sigma_{RA} = A^\top \Sigma_R A$, $b_a(RA) = A^\top b_a(R)$, and the quadratic form is unchanged. ☐

Remark 4.14 (The quotient-space object)**.** *The invariant object is the equivalence class of a payoff system modulo the existing span, i.e. its image in $\mathcal{X}/\mathcal{H}$. The orthogonal representative is $(I - \Pi_{\mathcal{H}})G \in \mathcal{H}^\perp$. This is the standard marketed-subspace equivalence in incomplete-markets theory (Magill and Quinzii, 1996): asset structures with the same span have the same frictionless allocation possibilities. The paper's contribution is not this invariance by itself. It is the wedge created when gross value descends to the residual subspace while implementation cost remains attached to the contractual representation.*

5 Verifiable Information and Contractible Payoffs

The preceding section treats the payoff space $\mathcal{X}$ as if all state distinctions were, in principle, contractible. That is the right geometry once payoffs are available, but it starts one layer too late. A contract cannot condition on a fact merely because the fact is economically meaningful. It can condition only on facts that are observable, verifiable, enforceable, and admissible.

In finite states, a sigma-algebra is equivalent to a partition of Ω. If two states lie in the same atom of a verification partition, no admissible contract under that technology regime can distinguish them. Let $\mathcal{F}_T$ denote the sigma-algebra of facts contractible under technology/legal regime T and define

$$\mathcal{X}_T = L_0^2(\mathcal{F}_T) = \{x \in \mathcal{X} : x \text{ is } \mathcal{F}_T\text{-measurable}\}.$$

Thus T is not just a scalar that lowers costs. It determines which state distinctions can enter payoff rules.

Proposition 5.1 (Value of a verification refinement)**.** *Fix an inherited market span $\mathcal{H} \subseteq \mathcal{X}$. Suppose $\mathcal{F}_T \subseteq \mathcal{F}_{T'}$, so that regime T' weakly refines the verifiable information structure. Let $\mathcal{X}_T = L_0^2(\mathcal{F}_T)$ and $\mathcal{X}_{T'} = L_0^2(\mathcal{F}_{T'})$. The gross value of the refinement, before representation costs, is*

$$\Delta W(T \to T' \mid \mathcal{H}) = \frac{1}{2} \operatorname{tr}\left(\left[\Pi_{\mathcal{H}+\mathcal{X}_{T'}} - \Pi_{\mathcal{H}+\mathcal{X}_T}\right]\mathsf{Q}\right) \geq 0.$$

It is strictly positive if and only if the newly verifiable residual subspace

$$(\mathcal{H} + \mathcal{X}_{T'}) \ominus (\mathcal{H} + \mathcal{X}_T)$$

has nonzero projection of some trade-generating valuation vector Z_a.

Proof. Because $\mathcal{F}_T \subseteq \mathcal{F}_{T'}$, $\mathcal{X}_T \subseteq \mathcal{X}_{T'}$ and hence $\mathcal{H} + \mathcal{X}_T \subseteq \mathcal{H} + \mathcal{X}_{T'}$. Apply theorem 4.1 to the two spans and subtract. For nested subspaces, the difference of the projections is the projection onto the orthogonal complement of the smaller subspace inside the larger one, so the trace with the positive semidefinite operator Q is nonnegative. Strict positivity is exactly the condition that some Z_a has a nonzero component in the newly available subspace. □

Remark 5.2 (Financial innovation as verifiable-state refinement)**.** *Indexes, parametric triggers, benchmark methodologies, credit scores, oracles, audited data feeds, public registries, and settlement rails can all be understood as ways of refining $\mathcal{F}_T$. They do not merely lower a generic cost; they make new state distinctions contractible. The slogan becomes: financial innovation is costly refinement of the verifiable state space.*

Proposition 5.3 (Optimal measurable approximation)**.** *Let $x \in L^2(\mathbb{P})$ be an ideal target payoff and let $\mathcal{F}_T$ be the verification sigma-algebra. Among all $\mathcal{F}_T$-measurable payoffs $y \in L^2(\mathcal{F}_T)$, the unique L^2 projection of x is the conditional expectation:*

$$\arg\min\nolimits_{y \in L^2(\mathcal{F}_T)} \mathbb{E}[(x-y)^2] = \mathbb{E}[x \mid \mathcal{F}_T].$$

The unavoidable basis risk under regime T is

$$x - \mathbb{E}[x \mid \mathcal{F}_T].$$

Proof. $L^2(\mathcal{F}_T)$ is a closed linear subspace of $L^2(\mathbb{P})$. The orthogonal projection of x onto this subspace is characterized by $\mathbb{E}[(x-y)z] = 0$ for all $z \in L^2(\mathcal{F}_T)$, which is precisely the defining property of $y = \mathbb{E}[x \mid \mathcal{F}_T]$. □

Example 5.4 (Indexes as verification partitions)**.** *Suppose an individual-home payoff x depends on house-level states, but courts, data vendors, and exchanges can verify only a metropolitan index partition $\mathcal{F}_T$. Then the best contractible index payoff is $\mathbb{E}[x \mid \mathcal{F}_T]$. The residual $x - \mathbb{E}[x \mid \mathcal{F}_T]$ is not merely statistical error; it is noncontractible state variation under regime T. A later technology that verifies house-level transactions refines $\mathcal{F}_T$ and can convert some of that basis risk into tradable span.*

6 Information Timing and the Limits of Market Creation

The verification-refinement result is not a monotonic claim that more information always creates more markets or more welfare. Timing matters. Information that becomes verifiable *after* agents can contract can improve settlement and reduce basis risk. Information that is publicly revealed *before* agents can trade can destroy insurance value by eliminating the uncertainty that made risk sharing valuable in the first place. This is the Hirshleifer force: information can have private value while reducing the social value of risk-sharing opportunities (Hirshleifer, 1971).

Proposition 6.1 (Verification versus pre-trade revelation). *Let $x \in L^2(\mathbb{P})$ be a risky payoff direction and let $\mathcal{F}$ be a public signal sigma-algebra. If agents can contract ex ante on an $\mathcal{F}$-measurable settlement rule before $\mathcal{F}$ is realized, the refinement can expand the attainable payoff space as in proposition 5.1. If instead $\mathcal{F}$ is publicly revealed before agents can trade, the ex ante risk-sharing value of claims on x is limited to the value of uncertainty remaining within the atoms of $\mathcal{F}$. In particular, if $\mathcal{F}$ fully reveals x, then no post-revelation zero-net-supply trade in x can insure the revealed variation across atoms.*

Proof. Before $\mathcal{F}$ is realized, an $\mathcal{F}$-measurable contract transfers resources across future signal atoms. After $\mathcal{F}$ is publicly realized, all agents condition on the realized atom and trades can only share residual uncertainty within that atom. If x is fully $\mathcal{F}$-measurable, its value is known after the signal, so a claim on x is locally riskless conditional on the atom and cannot transfer resources across the now-known states. Thus the cross-atom insurance value is available only to ex ante contracts. □

Remark 6.2 (The AI prediction limit). *AI and data infrastructure can create markets by making settlement cheaper and state distinctions verifiable. They can also eliminate markets by revealing idiosyncratic risk before pooling or trade occurs. Perfect pre-contract prediction of a house fire does not deepen fire insurance; it turns the premium into the known loss. The market-creating version of information is verifiability at settlement and better contract design. The market-destroying version is public revelation before risk sharing.*

7 Local Welfare Geometry Beyond Mean-Variance

The mean-variance baseline is deliberately transparent. It is not meant to claim that quadratic utility is the only environment in which the geometry appears. The same projection logic is the local quadratic approximation to a broad class of expected-utility economies.

Let $c_a^{\mathcal{H}}$ be the existing incomplete-market allocation and let u_a be twice continuously differentiable, increasing, and concave. Define the marginal utility process and local curvature process

$$m_a(\omega) = u_a'(c_a^{\mathcal{H}}(\omega)), \qquad h_a(\omega) = -u_a''(c_a^{\mathcal{H}}(\omega)) > 0.$$

For a candidate residual subspace B, define the local quadratic welfare problem

$$\mathcal{W}^{loc}(B) = \sup_{z_a \in B,\ \sum_a z_a = 0} \sum_a \left\{ \mathbb{E}[m_a z_a] - \frac{1}{2}\mathbb{E}[h_a z_a^2] \right\}.$$

Theorem 7.1 (Local expected-utility spanning value). *For small feasible transfers $z_a \in B$ around $c^{\mathcal{H}}$, aggregate expected utility admits the expansion*

$$\sum_a \mathbb{E}[u_a(c_a^{\mathcal{H}} + z_a)] - \sum_a \mathbb{E}[u_a(c_a^{\mathcal{H}})] = \sum_a \left\{ \mathbb{E}[m_a z_a] - \frac{1}{2}\mathbb{E}[h_a z_a^2] \right\} + o\left(\sum_a \|z_a\|^2 \right).$$

Consequently, the local value of adding B is the concave quadratic program $\mathcal{W}^{loc}(B)$. If $h_a(\omega) = \gamma_a$ is constant across states for each agent, then

$$\mathcal{W}^{loc}(B) = \frac{1}{2}\sum_a \theta_a \|\Pi_B(m_a - \bar{m})\|^2 = \frac{1}{2}\operatorname{tr}(\Pi_B \mathsf{Q}^{loc}),$$

where $\theta_a = 1/\gamma_a$, $\bar{m} = \sum_a \theta_a m_a / \sum_a \theta_a$, and

$$\mathsf{Q}^{loc} = \sum_a \theta_a (m_a - \bar{m}) \otimes (m_a - \bar{m}).$$

Proof. The first display is the second-order Taylor expansion of each u_a around $c_a^{\mathcal{H}}$, integrated over states. The remainder is uniform on finite Ω. The feasible local value is therefore the stated concave quadratic program. When curvature is constant by agent, the proof of theorem 4.1 applies with Y_a replaced by m_a and γ_a by the local curvature coefficient. □

Remark 7.2 (What constant curvature buys). *The local theorem does not imply that the single-operator geometry survives unchanged under arbitrary expected utility. With state-dependent curvature $h_a(\omega)$, each agent's welfare loss is measured in the curvature-weighted norm $\mathbb{E}[h_a z_a^2]$. The relevant orthogonality is then allocation-, agent-, and curvature-dependent rather than the common $L^2(\mathbb{P})$ orthogonality used above. In general there need not be a single global operator Q, a single residual subspace independent of the traders, or an exact additive identity $W(\mathcal{H}+B) = W(\mathcal{H})+V(B \mid \mathcal{H})$. Constant curvature is therefore not merely the case in which the local approximation is globally exact; it is the case in which the common residual-subspace and single-Q representation exist in the simple form used throughout the paper.*

8 Spectral Geometry of Effective Incompleteness

The frictionless optimal-basis theorem gives more than a benchmark. It gives a spectral measure of how incomplete an economy is in welfare-relevant directions.

Definition 8.1 (Effective incompleteness tail). *Let $\lambda_1 \geq \lambda_2 \geq \cdots \geq 0$ be the eigenvalues of the residual operator $\mathsf{Q}_{\mathcal{H}}$ on $\mathcal{H}^{\perp}$. Define*

$$\mathcal{I}_m(\mathcal{H}) = \frac{1}{2}\sum_{\ell>m}\lambda_\ell, \qquad d_{\text{eff}}(\varepsilon \mid \mathcal{H}) = \min\left\{m : \mathcal{I}_m(\mathcal{H}) \leq \varepsilon\right\}.$$

Proposition 8.2 (Effective completeness). *After the best frictionless m-dimensional market expansion, the maximum remaining risk-sharing value is $\mathcal{I}_m(\mathcal{H})$. Hence an economy can be Arrow-Debreu incomplete but effectively complete at tolerance ε if $d_{\text{eff}}(\varepsilon \mid \mathcal{H})$ is small.*

Proof. By theorem 4.3, the best m-dimensional residual expansion captures $\frac{1}{2}\sum_{\ell=1}^{m}\lambda_\ell$. The total residual value is $\frac{1}{2}\sum_{\ell\geq 1}\lambda_\ell$. Subtract. □

Remark 8.3 (Why the spectrum matters). *The spectral tail is not an additional piece of geometric machinery. It is the economic object that measures how much risk-sharing value remains after the best low-dimensional residual expansion. The empirical stability of an estimated residual basis is addressed in section 19, where perturbation bounds matter for interpreting an estimated $\widehat{\mathsf{Q}}$.*

9 Approximate Bases, Indexing, and a Canonical Example

Real markets often do not implement exact individual risks. They implement cheap approximations: indexes, baskets, triggers, benchmarks, standardized forms, and parametric rules. The theory therefore distinguishes exact spanning from approximate spanning.

Definition 9.1 (Projection distance between subspaces). *For subspaces $B, \tilde{B} \subseteq \mathcal{H}^{\perp}$, define*

$$d(B, \tilde{B}) = \left\|\Pi_B - \Pi_{\tilde{B}}\right\|_{op}.$$

This is the sine of the largest principal angle between the two subspaces.

Definition 9.2 (ε-implementation)**.** *Let ρ denote a claim representation system with centered payoff matrix $\pi(\rho)$ and implementation cost $C(\rho, \mathcal{H}, T)$. A representation system ρ ε-implements a target residual subspace $B \subseteq \mathcal{H}^\perp$ if*

$$d\big(B, \operatorname{col}((I - \Pi_{\mathcal{H}})\pi(\rho))\big) \leq \varepsilon.$$

The least cost of ε-implementing B is

$$\kappa_\varepsilon(B, \mathcal{H}, T) = \inf_{\rho\colon d(B, \operatorname{col}((I-\Pi_{\mathcal{H}})\pi(\rho)))\leq\varepsilon} C(\rho, \mathcal{H}, T).$$

Exact implementation is $\varepsilon = 0$.

Proposition 9.3 (Continuity of residual surplus)**.** *Let $B, \tilde{B} \subseteq \mathcal{H}^\perp$. Then*

$$\Big|V(B \mid \mathcal{H}) - V(\tilde{B} \mid \mathcal{H})\Big| \leq \frac{1}{2} d(B, \tilde{B}) \operatorname{tr}(\mathsf{Q}),$$

where $V(B \mid \mathcal{H}) = \frac{1}{2}\operatorname{tr}(\Pi_B \mathsf{Q})$.

Proof. For each Z_a,

$$\|\Pi_B Z_a\|^2 - \big\|\Pi_{\tilde{B}} Z_a\big\|^2 = \big\langle Z_a, (\Pi_B - \Pi_{\tilde{B}}) Z_a \big\rangle\,.$$

Taking absolute values and using the operator norm gives

$$\Big|\|\Pi_B Z_a\|^2 - \big\|\Pi_{\tilde{B}} Z_a\big\|^2\Big| \leq d(B, \tilde{B})\, \|Z_a\|^2\,.$$

Multiply by $\theta_a/2$ and sum. Since $\operatorname{tr}(\mathsf{Q}) = \sum_a \theta_a \|Z_a\|^2$, the result follows. □

Proposition 9.4 (When a cheap approximate basis dominates)**.** *Let B be a target residual payoff subspace and $\tilde{B}$ an approximate implementation. Suppose private capture is ϕ, admissibility burdens are $\Psi_B, \Psi_{\tilde{B}}$, and exact implementation costs are $\kappa_0(B, \mathcal{H}, T), \kappa_0(\tilde{B}, \mathcal{H}, T)$. A private creator prefers $\tilde{B}$ to exact implementation of B whenever*

$$\kappa_0(B, \mathcal{H}, T) - \kappa_0(\tilde{B}, \mathcal{H}, T) > \phi\big(V(B \mid \mathcal{H}) - V(\tilde{B} \mid \mathcal{H})\big) + \Psi_{\tilde{B}} - \Psi_B.$$

Proof. Compare private net values

$$\phi V(\tilde{B} \mid \mathcal{H}) - \kappa_0(\tilde{B}, \mathcal{H}, T) - \Psi_{\tilde{B}}$$

and

$$\phi V(B \mid \mathcal{H}) - \kappa_0(B, \mathcal{H}, T) - \Psi_B.$$

Rearrangement gives the stated inequality. □

Example 9.5 (A four-state index-versus-bespoke example)**.** *Let $\Omega = \{1, 2, 3, 4\}$ with uniform probabilities. Let*

$$h = (1, 1, -1, -1), \qquad n = (1, -1, 1, -1).$$

Then $h, n \in \mathcal{X}$, $\|h\| = \|n\| = 1$, and $\langle h, n\rangle = 0$. Think of h as a metro housing shock and n as idiosyncratic neighborhood variation. There are two agents with risk tolerances $\theta_1 = \theta_2 = 1$ and trade-generating valuation vectors

$$Z_1 = L(h + \varepsilon n), \qquad Z_2 = -L(h + \varepsilon n),$$

where $L > 0$ and $0 < \varepsilon < 1$. The exact bespoke direction is $B = \mathrm{span}\{h + \varepsilon n\}$; the cheap index direction is $\tilde{B} = \mathrm{span}\{h\}$. Since $\mathsf{Q} = 2L^2(h + \varepsilon n) \otimes (h + \varepsilon n)$,

$$V(B) = L^2(1 + \varepsilon^2), \qquad V(\tilde{B}) = L^2.$$

The index loses only $L^2\varepsilon^2$ of surplus. If the exact representation costs c_E and the index costs c_I, the index is privately preferred when

$$c_E - c_I > L^2\varepsilon^2 + \Psi_I - \Psi_E$$

for $\phi = 1$. Thus the first market may be a low-cost approximate basis, not the most precise claim.

Remark 9.6 (Why indexes arrive before bespoke claims)**.** *A metropolitan housing index may lose some residual surplus relative to individual-home derivatives but save large verification, appraisal, legal, liquidity, and dispute-resolution costs. Proposition 9.4 and example 9.5 give a precise condition under which the index appears first. Equivalently, private creators choose over a frontier of $(\varepsilon, \kappa_\varepsilon)$ tradeoffs rather than over exact Arrow securities alone.*

10 Costly Basis Selection

10.1 Representations, payoffs, and implementation costs

A payoff subspace can be implemented by many contractual representations. Let ρ denote a representation: a legal form, payoff rule, oracle, settlement convention, collateral rule, venue, and distribution technology. Let $\pi(\rho) \in \mathcal{X}$ denote the centered payoff generated by representation ρ. A system of representations $\rho = (\rho_1, \dots, \rho_n)$ generates payoff matrix

$$\pi(\rho) = \big(\pi(\rho_1), \dots, \pi(\rho_n)\big).$$

Let $T = (T_1, \dots, T_R)$ be a technology vector. Implementation cost is decomposed as

$$\begin{aligned} C(\rho, \mathcal{H}, T) = {} & C_{\text{spec}} + C_{\text{data}} + C_{\text{verify}} + C_{\text{legal}} + C_{\text{comply}} \\ & + C_{\text{custody}} + C_{\text{collateral}} + C_{\text{price}} + C_{\text{liq}} + C_{\text{settle}} + C_{\text{dispute}} + C_{\text{dist}}. \end{aligned}$$

Each component may depend on the existing market span $\mathcal{H}$, because existing claims can provide hedge instruments, legal templates, collateral conventions, data standards, market-maker inventory offsets, or distribution channels.

A minimal microfoundation for this reduced form is a task technology. Each representation ρ requires a finite set of tasks $\mathcal{E}(\rho)$: writing the payoff rule, observing the underlying state variable, verifying the trigger, enforcing the contract, computing margin, routing orders, clearing, and resolving disputes. Each task e has cost $\chi_e(T) \geq 0$, with $\partial \chi_e / \partial T_r \leq 0$ for technologies that automate or standardize that task. Some tasks are shared across a cluster of claims and therefore have fixed infrastructure costs. The reduced-form cost $C(\rho, \mathcal{H}, T)$ should be read as the lower envelope of these task costs, net of reusable infrastructure already present in $\mathcal{H}$.

Remark 10.1 (Liquidity is a cost-side object here)**.** *The span-value theorems characterize the welfare value of making payoff directions tradable. They do not by themselves prove that a market will be deep or liquid. Liquidity provision, market-maker capital, inventory risk, collateral, and hedgeability enter through C_{liq}, $C_{\text{collateral}}$, and C_{price}, and hence through the implementation envelope κ.*

Definition 10.2 (Least-cost implementation of a residual subspace)**.** *For $B \subseteq \mathcal{H}^{\perp}$, define*

$$\kappa(B, \mathcal{H}, T) = \inf_{\rho:\, \mathrm{col}((I-\Pi_{\mathcal{H}})\pi(\rho))=B} C(\rho, \mathcal{H}, T),$$

with $\kappa(B, \mathcal{H}, T) = +\infty$ if no admissible representation implements B.

Definition 10.3 (Representation fibers)**.** *Let $\mathcal{R}$ be the set of contractual representations and define*

$$\pi_{\mathcal{H}} : \mathcal{R} \to \mathrm{Sub}(\mathcal{X}/\mathcal{H}), \qquad \pi_{\mathcal{H}}(\rho) = \mathrm{col}((I - \Pi_{\mathcal{H}})\pi(\rho)).$$

Two representations are payoff-equivalent relative to $\mathcal{H}$ if they lie in the same fiber $\pi_{\mathcal{H}}^{-1}(B)$. Span value descends to the base space of residual subspaces; costs live on the representation fiber. The implementation envelope can be written

$$\kappa(B, \mathcal{H}, T) = \inf_{\rho \in \pi_{\mathcal{H}}^{-1}(B)} C(\rho, \mathcal{H}, T).$$

Remark 10.4 (Delivered payoffs)**.** *A representation may distort the intended payoff. Let x be a target payoff and let D_ρ be a delivery operator capturing default, oracle error, settlement delay, dispute risk, legal uncertainty, margin calls, collateral haircuts, or manipulation. Then the relevant residual span is generated by $D_\rho x$, not by the label on x, and the creator solves*

$$V(\mathrm{col}((I - \Pi_{\mathcal{H}})D_\rho x) \mid \mathcal{H}) - C(\rho, \mathcal{H}, T).$$

This distinguishes contracts that reference the same exposure but deliver different state-contingent payoffs.

Proposition 10.5 (Costly basis selection)**.** *For a residual subspace B, gross surplus $V(B \mid \mathcal{H})$ is invariant across all representations implementing B, while implementation cost is generally not. If a private creator can capture fraction $\phi_B \in [0, 1]$ of gross surplus and faces admissibility burden Ψ_B, then B is privately viable if and only if*

$$\phi_B V(B \mid \mathcal{H}) > \kappa(B, \mathcal{H}, T) + \Psi_B.$$

A social planner facing externality cost E_B finds B socially desirable if and only if

$$V(B \mid \mathcal{H}) - \kappa(B, \mathcal{H}, T) - E_B > 0.$$

Proof. The surplus side follows from lemma 4.13: all representations with the same residual span generate the same gross risk-sharing value. The cost side is minimized by definition 10.2. Private and social viability are direct comparisons of benefits and costs. □

Example 10.6 (Same span, different representation cost)**.** *Let two representations ρ_1 and ρ_2 implement payoff systems with the same residual span B:*

$$\mathrm{col}((I - \Pi_{\mathcal{H}})\pi(\rho_1)) = \mathrm{col}((I - \Pi_{\mathcal{H}})\pi(\rho_2)) = B.$$

Then $V(B \mid \mathcal{H})$ is identical for both by lemma 4.13. If $C(\rho_1, \mathcal{H}, T) < C(\rho_2, \mathcal{H}, T)$ and admissibility burdens are the same, the private and social rankings both prefer ρ_1. This is the simplest form of the costly-basis wedge: two contracts can be economically equivalent in payoff space but very different institutional objects.

Missing risk direction	Low-cost basis likely first	High-cost basis delayed
Housing risk	metro index future or ETF	house-specific derivative
Weather risk	parametric trigger	indemnity contract requiring loss adjustment
Creator income risk	platform royalty share	bespoke lifetime income swap
Small-business revenue	standardized revenue-share note	customized cash-flow derivative
Labor-income risk	occupation-region index	individual wage-income claim
Climate risk	satellite-triggered claim	litigated loss-contingent insurance

10.2 Dictionary selection and sparse market design

Let $\mathcal{J} = \{1, \dots, N\}$ be a finite dictionary of admissible representations. Candidate j has payoff $d_j \in \mathcal{X}$, cost $c_j \geq 0$, admissibility burden $\psi_j \geq 0$, and externality $e_j \geq 0$. For $M \subseteq \mathcal{J}$, write

$$D_M = \{(I - \Pi_{\mathcal{H}})d_j : j \in M\}, \qquad B_M = \operatorname{span} D_M \subseteq \mathcal{H}^{\perp}.$$

Proposition 10.7 (Market creation as penalized projection)**.** *If a platform can choose any subset of the finite dictionary and captures fraction ϕ, its private market-creation problem is*

$$M^{priv} \in \arg\max_{M \subseteq \mathcal{J}} \left\{ \phi \, \frac{1}{2} \operatorname{tr}(\Pi_{B_M} \mathsf{Q}) - \sum_{j \in M} c_j - \Psi(M) \right\}.$$

A social planner's analogue is

$$M^{soc} \in \arg\max_{M \subseteq \mathcal{J}} \left\{ \frac{1}{2} \operatorname{tr}(\Pi_{B_M} \mathsf{Q}) - \sum_{j \in M} c_j - E(M) \right\}.$$

Thus endogenous market creation is a cost-weighted basis-selection problem over residual payoff directions.

Proof. For any chosen set M, proposition 4.6 gives gross value $\frac{1}{2} \operatorname{tr}(\Pi_{B_M} \mathsf{Q})$. Additive claim costs and set-level admissibility or externality terms yield the stated objectives. □

Corollary 10.8 (Orthogonal dictionary threshold)**.** *Suppose the residualized dictionary vectors $d_j^{\perp} = (I - \Pi_{\mathcal{H}})d_j$ are nonzero and mutually orthogonal, and define $u_j = d_j^{\perp} / \left\| d_j^{\perp} \right\|$. If costs and admissibility burdens are additive, then a private platform includes j independently if and only if*

$$\frac{\phi}{2} \langle u_j, \mathsf{Q} u_j \rangle > c_j + \psi_j.$$

A social planner includes j independently if and only if

$$\frac{1}{2} \langle u_j, \mathsf{Q} u_j \rangle > c_j + e_j.$$

Proof. With an orthonormal residual dictionary, $\Pi_{B_M} = \sum_{j \in M} u_j \otimes u_j$. Therefore

$$\frac{1}{2}\operatorname{tr}(\Pi_{B_M}\mathsf{Q}) = \sum_{j \in M} \frac{1}{2}\langle u_j, \mathsf{Q}u_j\rangle .$$

The platform and planner objectives separate across j. □

Theorem 10.9 (Endogenous market-resolution threshold)**.** *Let $e_1, e_2, \dots$ be an orthonormal eigenbasis of the residual valuation-dispersion operator $\mathsf{Q}_{\mathcal{H}}$ on $\mathcal{H}^{\perp}$, with eigenvalues $\lambda_1 \geq \lambda_2 \geq \cdots \geq 0$. Suppose an admissible representation dictionary contains exact normalized implementations of these residual directions, with additive private implementation costs c_ℓ, admissibility burdens ψ_ℓ, and social externality costs χ_ℓ. A private platform that captures fraction ϕ implements eigendirection ℓ independently if and only if*

$$\frac{\phi}{2}\lambda_\ell > c_\ell + \psi_\ell.$$

A social planner implements eigendirection ℓ independently if and only if

$$\frac{1}{2}\lambda_\ell > c_\ell + \chi_\ell.$$

Thus the resolution at which an economy prices residual risk is set by the crossing of the valuation-dispersion spectrum and the marginal cost of implementing admissible representations.

More generally, if an exact target residual subspace B and a cheaper approximation $\widetilde{B}$ are both available, with private implementation costs $\kappa(B)$ and $\kappa(\widetilde{B})$, the approximation is privately preferred whenever

$$\kappa(B) - \kappa(\widetilde{B}) > \phi\left[V(B \mid \mathcal{H}) - V(\widetilde{B} \mid \mathcal{H})\right],$$

with the analogous social comparison replacing ϕ by one and adding externality costs. The value-loss term is controlled by proposition 9.3; hence coarse indexes, baskets, and benchmarks dominate bespoke claims when their cost savings exceed the local spectral value they discard.

Proof. For the exact eigenbasis dictionary, corollary 10.8 applies with $u_\ell = e_\ell$ and

$$\langle e_\ell, \mathsf{Q}e_\ell\rangle = \lambda_\ell.$$

The private and social threshold inequalities follow immediately. For the approximation comparison, subtract the private objective value of $\widetilde{B}$ from that of B and rearrange:

$$\phi V(\widetilde{B} \mid \mathcal{H}) - \kappa(\widetilde{B}) > \phi V(B \mid \mathcal{H}) - \kappa(B)$$

is equivalent to the displayed condition, suppressing common burdens. The continuity bound supplies an upper bound on the value loss from replacing B by $\widetilde{B}$. □

Remark 10.10 (Resolution, not just listing)**.** *Theorem 10.9 is the rigorous counterpart of the resolution metaphor. The zero-cost benchmark says which residual directions an omniscient market designer would span first. The costly-basis threshold says how far down that spectrum actual institutions go, and whether they implement a high-fidelity bespoke claim or a lower-cost approximation. A concentrated residual spectrum justifies paying for precise state distinctions. A flat or low spectrum makes coarse representations optimal even when the exact missing claim is imaginable.*

Proposition 10.11 (One-step marginal value of a nonorthogonal claim)**.** *Let M be an existing dictionary set with span $\mathcal{H}(M) = \mathcal{H} + B_M$. For a new candidate $j \notin M$, define the residual*

$$r_j(M) = (I - \Pi_{\mathcal{H}(M)})d_j.$$

If $r_j(M) \neq 0$, the gross one-step surplus from adding j is

$$S_j(M) = \frac{1}{2}\left\langle \frac{r_j(M)}{\|r_j(M)\|}, \mathsf{Q}\frac{r_j(M)}{\|r_j(M)\|} \right\rangle.$$

If $r_j(M) = 0$, then $S_j(M) = 0$.

Proof. Adding a single claim expands the span by the one-dimensional residual subspace generated by $r_j(M)$. Apply proposition 4.6 with that one-dimensional subspace. □

Remark 10.12 (Scale normalization)**.** *The normalized residual direction $r_j(M)/\|r_j(M)\|$ appears because a one-claim market adds a subspace, not a unit of measurement. Multiplying a payoff by a nonzero scalar changes its units and price, but not the residual span it adds. Margin, collateral, tick-size, or contract-size effects belong in representation cost, not span value.*

Remark 10.13 (No unconditional submodularity claim)**.** *A larger span mechanically lowers a candidate's residual norm, but marginal surplus also depends on the direction of the remaining residual relative to Q. Global submodularity of market creation is therefore not a general property of arbitrary nonorthogonal dictionaries. Complementarity can enter through costs: shared oracles, benchmark status, legal templates, collateral conventions, clearing infrastructure, distribution channels, and market-maker balance sheets.*

Definition 10.14 (Submodularity ratio)**.** *For a normalized, monotone set function $F : 2^{\mathcal{J}} \to \mathbb{R}_+$, define the submodularity ratio at scale (L, k) by*

$$\gamma_{L,k}(F) = \min_{S,T\subseteq\mathcal{J},\ |S|\leq L,\ |T|\leq k,\ S\cap T=\emptyset} \frac{\sum_{j\in T}\big(F(S\cup\{j\}) - F(S)\big)}{F(S\cup T) - F(S)},$$

with the ratio interpreted as one when the denominator is zero.

Remark 10.15 (Algorithmic interpretation and companion result)**.** *A platform, exchange, or regulator need not solve the combinatorial basis-selection problem exactly. Under weak-submodularity conditions, standard greedy guarantees imply that sequentially listing the largest marginal residual-value claims is approximately disciplined (Das and Kempe, 2018). Without those conditions, greedy listing is only a heuristic. Paper 5 develops the algorithmic-search version of this statement, including hardness and greedy approximation guarantees; Paper 1 uses the submodularity ratio only to clarify when the static residual-value objective is compatible with simple listing heuristics.*

11 Endogenous Market-Set Equilibrium

11.1 Candidate claims and pairwise stability

Let $\mathcal{J} = \{1, \ldots, N\}$ be a finite set of candidate representations. Candidate j has payoff $g_j \in \mathcal{X}$, capture share ϕ_j, admissibility burden Ψ_j, and cost $C_j(M, T)$ if created when the existing market set is $M \subseteq \mathcal{J}$. Define

$$\mathcal{H}(M) = \mathcal{H}_0 + \operatorname{span}\{g_j : j \in M\},$$

where $\mathcal{H}_0$ is the inherited span. For $j \notin M$, define its marginal residual surplus

$$S_j(M) = W(\mathcal{H}(M) + \mathrm{span}\{g_j\}) - W(\mathcal{H}(M)).$$

The private net value of adding j to M is

$$\Pi_j(M; T) = \phi_j S_j(M) - C_j(M, T) - \Psi_j.$$

Definition 11.1 (Pairwise-stable market set)**.** *A market set $M^* \subseteq \mathcal{J}$ is pairwise stable at technology T if:*
(i) for every $j \notin M^*$, $\Pi_j(M^*; T) \leq 0$;
(ii) for every $j \in M^*$,

$$\phi_j S_j(M^* \setminus \{j\}) - C_j(M^* \setminus \{j\}, T) - \Psi_j \geq 0.$$

The first condition rules out profitable entry. The second rules out markets that would not be viable if reconsidered without themselves. This is a local stability notion, not a universal existence theorem.

Proposition 11.2 (Pure stability under an ordinal potential)**.** *Suppose there exists a function $\Phi : 2^{\mathcal{J}} \to \mathbb{R}$ such that for every $M \subseteq \mathcal{J}$ and $j \notin M$,*

$$\mathrm{sign}\left(\Phi(M \cup \{j\}) - \Phi(M)\right) = \mathrm{sign}\left(\Pi_j(M; T)\right),$$

and for every $j \in M$, the sign of deleting j is the negative of the sign of adding j to $M \setminus \{j\}$. Then any global maximizer of Φ is pairwise stable.

Proof. Let M^* maximize Φ. If some $j \notin M^*$ had $\Pi_j(M^*; T) > 0$, then $\Phi(M^* \cup \{j\}) > \Phi(M^*)$, contradicting maximality. If some $j \in M^*$ violated the viability condition, then adding j to $M^* \setminus \{j\}$ would have negative private net value, so deleting j would increase Φ, again contradicting maximality. □

Remark 11.3 (Mixed equilibria)**.** *Under arbitrary strategic interactions among multiple creators, pure stable sets need not exist. A finite entry game in which creators choose binary listing decisions and receive finite payoffs has a mixed Nash equilibrium. The paper uses pairwise stability for comparative statics but does not rely on universal pure-existence claims.*

Proposition 11.4 (Residual norm weakly falls with a larger span)**.** *If $\mathcal{H} \subseteq \mathcal{H}' \subseteq \mathcal{X}$, then for every candidate claim $g \in \mathcal{X}$,*

$$\|(I - \Pi_{\mathcal{H}'})g\| \leq \|(I - \Pi_{\mathcal{H}})g\|.$$

Proof. $\Pi_{\mathcal{H}'} g$ is the best approximation to g in the larger subspace $\mathcal{H}'$, while $\Pi_{\mathcal{H}} g \in \mathcal{H}'$ is feasible in that approximation problem. Hence the residual distance to $\mathcal{H}'$ is weakly smaller. □

Remark 11.5 (Payoff substitutes, cost complements)**.** *Two claims are payoff substitutes when one spans part of the other's residual direction. They are cost complements when one lowers the other's implementation cost through shared data, legal templates, settlement rails, liquidity, collateral, or distribution. A theory of financial innovation must keep these two channels separate.*

12 Private Under-Entry of Market Infrastructure

Some implementation costs are shared infrastructure costs. Let $J \subseteq \mathcal{J}$ be a cluster of claims requiring common infrastructure I: an oracle, data standard, index methodology, legal template, clearing rule, collateral convention, or settlement rail. Let $F > 0$ be the fixed cost of building I, and let c_j be claim-specific cost after I exists. Define the cluster residual subspace and cluster gross surplus by

$$B_J = \operatorname{span}\{(I - \Pi_{\mathcal{H}})g_j : j \in J\}, \qquad \Delta V_J = V(B_J \mid \mathcal{H}) = W(\mathcal{H} + B_J) - W(\mathcal{H}).$$

This cluster definition avoids double-counting overlapping one-claim surplus.

Proposition 12.1 (Infrastructure under-entry without double-counting)**.** *Suppose that for every $j \in J$, a first mover that must pay the full infrastructure cost is not privately viable:*

$$\phi_j V(\operatorname{span}\{(I - \Pi_{\mathcal{H}})g_j\} \mid \mathcal{H}) - c_j - F - \Psi_j < 0.$$

Suppose the cluster is socially valuable:

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

Then decentralized one-claim entry can fail even though infrastructure creation is socially efficient. A platform that internalizes cluster surplus creates I whenever

$$\phi_J \Delta V_J - \sum_{j \in J} c_j - F - \Psi_J > 0,$$

where ϕ_J is the platform's captured share of cluster surplus and Ψ_J is the platform-level admissibility burden.

Proof. The first inequality rules out every individual first mover. The second inequality is the planner's net surplus from creating the common infrastructure and the whole payoff span generated by the cluster. Because ΔV_J is defined as the value of the union span, overlapping payoff directions are counted only once. The platform condition is the corresponding private net value for an actor that can capture surplus across the cluster. □

Example 12.2 (Why cluster surplus is not the sum of one-claim surpluses)**.** *Let $u, v \in \mathcal{H}^\perp$ be orthonormal and let*

$$\mathsf{Q} = \lambda_u u \otimes u + \lambda_v v \otimes v, \qquad \lambda_u \geq \lambda_v > 0.$$

Consider two infrastructure-dependent claims $g_1 = u$ and $g_2 = u + \varepsilon v$. The individual one-claim values are

$$V(\operatorname{span}\{g_1\} \mid \mathcal{H}) = \frac{\lambda_u}{2}, \qquad V(\operatorname{span}\{g_2\} \mid \mathcal{H}) = \frac{1}{2}\frac{\lambda_u + \varepsilon^2 \lambda_v}{1 + \varepsilon^2}.$$

Adding these values double-counts the common u direction. The cluster value is instead the value of the union span. For $\varepsilon \neq 0$,

$$V(\operatorname{span}\{g_1, g_2\} \mid \mathcal{H}) = \frac{1}{2}(\lambda_u + \lambda_v),$$

which is the correct object for deciding whether to build shared infrastructure.

Remark 12.3 (Constructive policy role)**.** *The result explains why socially valuable markets may be absent even when many individual claims would be viable after standardization. Public data, legal safe harbors, benchmark construction, and open settlement standards can be welfare-improving infrastructure rather than subsidies to a single product.*

13 Dynamic Cost Complementarities and Tipping

Payoff overlap creates substitution in residual value, while shared infrastructure creates complementarity in costs. When cost complementarities are strong, the market frontier can exhibit tipping and hysteresis.

Let

$$C_j(M,T) = c_j(T) - \eta_j(M),$$

where $\eta_j(M)$ is the cost reduction from templates, data standards, settlement rails, liquidity, distribution, or collateral conventions already created by M.

Proposition 13.1 (Two-market tipping)**.** *Consider two candidate markets with identical captured surplus ϕS and standalone cost c. Suppose each market reduces the other's cost by $\eta > 0$ if it exists. If*

$$\phi S - c < 0 \qquad \text{and} \qquad \phi S - c + \eta > 0,$$

then both the empty set and the two-market set are pairwise stable. Early infrastructure or coordinated platform entry can move the economy from the low-market equilibrium to the high-market equilibrium.

Proof. Starting from the empty set, either market would have negative private net value $\phi S - c$, so no single entrant enters. If both markets exist, deleting either removes a market whose net value conditional on the other is $\phi S - c + \eta > 0$, so neither is deleted. Thus both sets satisfy pairwise stability. □

Remark 13.2 (Hysteresis)**.** *A temporary subsidy, legal safe harbor, benchmark mandate, or public data buildout can permanently change the market set if it creates infrastructure that makes later private maintenance viable. This is distinct from the payoff-spanning channel: it is a cost-complementarity channel.*

14 Technology, Cost Compression, and Entry Waves

Let the cost of claim j have components indexed by $r = 1, \dots, R$:

$$C_j(M,T) = \sum_{r=1}^{R} C_{jr}(M, T_r), \qquad \frac{\partial C_{jr}}{\partial T_r} \le 0.$$

Define the private entry gap

$$\Delta_j(M,T) = \phi_j S_j(M) - C_j(M,T) - \Psi_j.$$

Proposition 14.1 (Component-specific cost compression)**.** *Holding residual surplus and admissibility fixed,*

$$\frac{\partial \Delta_j}{\partial T_r} = -\frac{\partial C_{jr}}{\partial T_r} \ge 0.$$

Claims most likely to enter after an improvement in technology component T_r are those with high residual surplus, low admissibility burden, and high cost exposure to component r.

Proof. Immediate from differentiating the entry gap. □

Technology improvement	Main costs lowered	First likely markets
Sensors, APIs, remote measurement	data, verification	parametric insurance, supply-chain claims
Legal templates and safe harbors	specification, legal, dispute	revenue shares, royalties, standardized notes
Risk engines and market-making	pricing, liquidity, collateral	factor-hedgeable long-tail derivatives
Settlement and custody rails	custody, settlement	transferable, fractional, tokenized claims
Agentic distribution	distribution, attention, specification	micro-hedges, household/SMB contracts
Regulatory clarity	admissibility, legal, compliance	event contracts, new derivatives

Proposition 14.2 (Entry time under exponential cost decline)**.** *Suppose a claim has fixed captured surplus $\phi_j S_j$, fixed admissibility burden Ψ_j, and cost $C_j(t) = C_j(0)e^{-\lambda_j t}$. If $\phi_j S_j \le \Psi_j$, the claim never becomes privately viable. If $\phi_j S_j > \Psi_j$, the entry time is*

$$\tau_j = \max\left\{0, \frac{1}{\lambda_j}\log\left(\frac{C_j(0)}{\phi_j S_j - \Psi_j}\right)\right\}.$$

Proof. Entry requires $C_j(0)e^{-\lambda_j t} < \phi_j S_j - \Psi_j$. Solve for t, imposing immediate entry if the inequality already holds at $t = 0$. □

Remark 14.3 (Non-monotonicities)**.** *Individual viability is monotone in lower own cost, holding other objects fixed. Equilibrium market count need not be monotone: a broad low-cost benchmark can replace many narrow claims. Private welfare and social welfare need not be monotone either, because entry can redistribute rents or create externalities. A safe monotone result requires cumulative entry or no-delisting assumptions.*

15 Welfare and Admissibility

For a residual subspace B, define private net value

$$P_B = \phi_B V(B \mid \mathcal{H}) - \kappa(B, \mathcal{H}, T) - \Psi_B^{priv}$$

and social net value

$$W_B^{soc} = V(B \mid \mathcal{H}) - \kappa(B, \mathcal{H}, T) - E_B.$$

Here E_B includes externalities such as privacy loss, coercion, manipulation, moral hazard, adverse selection, addictive speculation, inequality, oracle fragility, and systemic-risk contribution.

Proposition 15.1 (Private-social wedge classification)**.** *Claims fall into four regions:*

Region	*Condition*	*Policy implication*
Privately viable, socially valuable	$P_B > 0,\ W_B^{soc} > 0$	*allow; standardize; perhaps encourage*
Privately viable, socially harmful	$P_B > 0,\ W_B^{soc} < 0$	*tax, margin, restrict, or prohibit*
Privately nonviable, socially valuable	$P_B < 0,\ W_B^{soc} > 0$	*subsidize, public provision, mandate data standards*
Privately nonviable, socially harmful	$P_B < 0,\ W_B^{soc} < 0$	*ignore or prohibit*

Falling creation costs can increase market completeness while reducing welfare if claims with $P_B > 0$ and $W_B^{soc} < 0$ cross the private entry threshold.

Proof. The classification is the sign pattern of private and social net values. If a harmful claim has P_B cross from negative to positive as κ falls while $W_B^{soc} < 0$, private entry occurs and social welfare changes by $W_B^{soc} < 0$. □

Remark 15.2 (Claim-level regulation)**.** *The policy object is not a venue or technology in the abstract. It is the claim's payoff rule, externality profile, manipulability, privacy cost, collateral requirement, and systemic-risk contribution.*

Definition 15.3 (Admissible representation set)**.** *Let*

$$\mathcal{R}^{ad}(T) = \mathcal{R}^{ver}(T) \cap \mathcal{R}^{IC}(T) \cap \mathcal{R}^{privacy}(T) \cap \mathcal{R}^{systemic}(T) \cap \mathcal{R}^{legal}(T)$$

be the set of representations satisfying verifiability, incentive, privacy, systemic-risk, and legal constraints. The constrained social design problem is

$$\max_{\rho \in \mathcal{R}^{ad}(T)} \{V(\pi_{\mathcal{H}}(\rho) \mid \mathcal{H}) - C(\rho, \mathcal{H}, T)\}.$$

Proposition 15.4 (Scalar externalities as a reduced-form relaxation)**.** *Suppose a soft admissibility screen is represented by violation functions $v_k(\rho) \geq 0$, where $v_k(\rho) = 0$ means that representation ρ satisfies constraint k. If a planner attaches shadow costs $\mu_k \geq 0$, then*

$$E(\rho) = \sum_k \mu_k v_k(\rho)$$

is a reduced-form penalty for relaxing the hard constrained problem in definition 15.3. The scalar term E_B in the wedge classification should therefore be read as a shorthand for externality and constraint costs over representations that implement B, not as the primitive admissibility object.

Proof. Replacing hard membership in $\mathcal{R}^{ad}(T)$ by nonnegative violation penalties yields the displayed penalized objective. Taking the infimum of penalties over representations in a fiber $\pi_{\mathcal{H}}^{-1}(B)$ gives the corresponding reduced-form penalty for the residual subspace B. □

16 Technology-Driven Completion and Its Limits

Let $\mathcal{A}(T)$ be the set of admissible representations at technology level T:

$$\mathcal{A}(T) = \{\rho :\ C(\rho, T) < \infty,\ \Psi_\rho(T) < \infty,\ \pi(\rho) \text{ is verifiable and enforceable}\}.$$

Define the attainable span

$$\overline{\mathcal{H}}(T) = \mathrm{span}\{\pi(\rho) : \rho \in \mathcal{A}(T)\} \subseteq \mathcal{X}.$$

Proposition 16.1 (Admissible completion frontier)**.** *If $T' \succeq T$ implies $\mathcal{A}(T) \subseteq \mathcal{A}(T')$, then*

$$\overline{\mathcal{H}}(T) \subseteq \overline{\mathcal{H}}(T') \quad \textit{and} \quad \dim \overline{\mathcal{H}}(T) \leq \dim \overline{\mathcal{H}}(T').$$

Full Arrow-Debreu risk completeness in this finite-state setting requires $\overline{\mathcal{H}}(T) = \mathcal{X}$, so $\dim \overline{\mathcal{H}}(T) = K - 1$. If privacy, manipulation, moral hazard, legal prohibition, unverifiability, or systemic-risk constraints permanently exclude some payoff directions, then $\dim \overline{\mathcal{H}}(T) < K - 1$ for all feasible T.

Proof. Set inclusion of admissible representations implies span inclusion and hence weakly increasing dimension. Full risk completeness requires spanning all centered state-contingent payoff directions. Permanent exclusions imply the attainable span is a strict subspace. □

Proposition 16.2 (Actual span monotonicity under cumulative entry)**.** *Suppose markets, once created, remain tradable; admissibility is stable; and the entry process at higher T can add claims but does not delete existing payoff directions. Then the realized span $\mathcal{H}(\mathcal{M}(T))$ is weakly increasing in T:*

$$T' \succeq T \quad \Rightarrow \quad \mathcal{H}(\mathcal{M}(T)) \subseteq \mathcal{H}(\mathcal{M}(T')),$$

and therefore $\dim \mathcal{H}(\mathcal{M}(T')) \geq \dim \mathcal{H}(\mathcal{M}(T))$.

Proof. Under cumulative entry, $\mathcal{M}(T) \subseteq \mathcal{M}(T')$. Taking spans preserves set inclusion. Dimension is monotone with subspace inclusion. □

Remark 16.3. *Without cumulative entry, market count can fall and even the realized span can change non-monotonically. A broad benchmark can displace narrow claims; regulation can prohibit a previously traded direction; a settlement technology can favor transferable but low-welfare claims.*

17 Endogenous Exposures and Reflexive Claims

The baseline comparative statics hold the valuation-dispersion operator Q fixed while the market span changes. That is the right local object for the theorem spine. In a dynamic economy, however, new markets can change endowments, production choices, leverage, occupational choice, location choice, information acquisition, and risk-taking. Then the risk-sharing operator itself becomes endogenous.

Proposition 17.1 (Completion effect versus behavioral effect)**.** *Let $W(\mathcal{H}, \mathsf{Q}) = \frac{1}{2}\mathrm{tr}(\Pi_{\mathcal{H}}\mathsf{Q})$. Suppose a market innovation changes both the span from $\mathcal{H}$ to $\mathcal{H}'$ and the valuation-dispersion operator from Q to Q'. Then*

$$W(\mathcal{H}', \mathsf{Q}') - W(\mathcal{H}, \mathsf{Q}) = \frac{1}{2}\mathrm{tr}\left((\Pi_{\mathcal{H}'} - \Pi_{\mathcal{H}})\mathsf{Q}\right) + \frac{1}{2}\mathrm{tr}\left(\Pi_{\mathcal{H}'}(\mathsf{Q}' - \mathsf{Q})\right).$$

The first term is the market-completion effect holding exposures fixed. The second term is the behavioral or endogenous-exposure effect. More complete markets can reduce welfare if the second term is sufficiently negative after accounting for moral hazard, leverage, adverse selection, or socially excessive risk-taking.

Proof. Add and subtract $\frac{1}{2}\mathrm{tr}(\Pi_{\mathcal{H}'}\mathsf{Q})$. □

Remark 17.2 (Markets can change the risk they price). *Insurance can induce moral hazard; hedging can encourage concentration; liquid credit protection can change lending behavior; prediction markets can affect the decision or event being predicted. A full dynamic theory should therefore track* $(\mathcal{H}_t, \mathsf{Q}_t, \kappa_t)$*, not only* $\mathcal{H}_t$*. The fixed-*Q *results are the local geometry; the endogenous-*Q *term is where behavioral and general-equilibrium feedback enters.*

Remark 17.3 (Reflexive underlyings). *The paper usually treats a claim as a payoff function* $g : \Omega \to \mathbb{R}$ *on an exogenous state space. Some markets are reflexive: the price or existence of the market changes the underlying event. Credit spreads can affect financing, governance markets can affect decisions, and public prediction markets can alter incentives. For such claims, admissibility requires a stable price-to-state-to-payoff mapping, not merely a measurable payoff rule.*

18 Common-Beliefs Scope

The baseline is intentionally common-beliefs. Under common beliefs, the surplus formulas above measure risk-sharing gains. With heterogeneous beliefs, the same algebra gives perceived trading surplus, but that object mixes risk sharing, speculation, information aggregation, and redistribution. This distinction is not cosmetic. Under disagreement, cheaper claim creation can open valuable information markets, but it can also open high-volume speculative directions whose private trading gains are not social risk-sharing gains (Simsek, 2013). The common-beliefs baseline is therefore the correct environment for the span-value welfare theorem. The heterogeneous-beliefs extension in appendix A should be developed in the broader research program as a theory of when belief-driven market creation is welfare-improving, merely redistributive, or harmful.

19 Empirical Predictions and Test Plan

The theory predicts market births when valuation-residual surplus is high and implementation costs fall. The natural empirical unit is (g, d, t): claim form g, domain or exposure class d, and time t. Let $Y_{gdt} = 1$ if a market is born: first listing, exchange approval, sustained volume, ETF launch, liquidity pool, standardized contract, insurance product, or legally enforceable traded claim.

The main empirical object should match the theory. Let

$$r_{gdt} = (I - \Pi_{\mathcal{H}_{dt}})g_{dt}$$

be the candidate payoff residual relative to already traded claims in domain d at time t, and let $u_{gdt} = r_{gdt}/\left\|r_{gdt}\right\|$ when $r_{gdt} \neq 0$. A theory-consistent one-dimensional value proxy is

$$\widehat{V}^{res}_{gdt} = \frac{1}{2}\left\langle u_{gdt}, \widehat{\mathsf{Q}}_{dt} u_{gdt}\right\rangle.$$

Raw residual payoff novelty, such as $1 - R^2$ from projecting g_{dt} on existing traded factors, is not by itself residual span value. It is a screening measure for whether the candidate adds a payoff direction. The value proxy must also measure valuation dispersion along that residual direction.

A reduced-form hazard specification is

$$\begin{aligned}\Pr(Y_{gdt} = 1) = \Lambda\Big(&\beta_1 \widehat{V}^{res}_{gdt} + \beta_2 \text{PayoffResidualNovelty}_{gdt} - \beta_3 \text{VerificationCost}_{gdt} \\ &- \beta_4 \text{LegalCost}_{gdt} - \beta_5 \text{SettlementCost}_{gdt} - \beta_6 \text{LiquidityCost}_{gdt} + \beta_7 \text{Template}_{gdt} \\ &+ \beta_8 \text{DataAvailability}_{gdt} - \beta_9 \text{AdmissibilityBurden}_{gdt} + \mu_d + \tau_t\Big).\end{aligned}$$

The most distinctive prediction is an interaction:

$$\widehat{V}^{res}_{gdt} \times \text{VerificationShock}_{dt}.$$

A verification shock should create markets where valuation-residual surplus was high but measurement was previously expensive. It should not create markets for already-spanned exposures, and it should not create many markets for residual payoff directions with little valuation dispersion.

Theoretical object	Empirical proxy
Payoff residual novelty	$1 - R^2$ or residual norm from projecting candidate exposure onto existing traded factors
Valuation-residual surplus	dispersion in marginal valuations, hedging demand, forecasts, regional risk, or balance-sheet sensitivity projected onto the normalized residual direction; estimate $\frac{1}{2}\left\langle u_g, \widehat{\mathsf{Q}}u_g \right\rangle$ where possible
Verification cost	sensor coverage, data APIs, auditability, weather-station density, satellite observability
Legal cost	standard contract availability, regulatory clarity, approval route, jurisdictional enforceability
Settlement cost	clearing access, custody infrastructure, settlement lag, programmable rails
Liquidity cost	factor hedgeability, inventory cost, bid-ask proxies, correlated open interest
Infrastructure complementarity	shared oracle, benchmark index, legal template, collateral standard
Admissibility burden	prohibitions, retail-access limits, privacy exposure, manipulation sensitivity

Candidate empirical surfaces include:

1. futures and options listings, contract specifications, and exchange product filings;
2. ETF and thematic-product launches, index methodology changes, and fund closures;
3. private secondary-market contracts and tender-offer programs;
4. revenue-based finance and small-business cash-flow claims;
5. music royalty and creator-revenue markets;
6. insurance-linked securities and parametric insurance;
7. prediction markets and event contracts;
8. tokenized real-world assets and programmable settlement systems;
9. compute, energy, transmission, and demand-response markets;
10. DeFi and automated-market-maker pool creation.

Identification strategies include difference-in-differences around legal-template adoption, oracle launches, satellite or sensor coverage expansions, exchange rule changes, settlement-rail introductions, and regulatory clarity events; event studies around data-standard releases; and survival models of claim birth.

19.1 Revealed spectral estimation

The reduced-form hazard model tests whether market births follow valuation-residual surplus and cost shocks. A more structural empirical version treats $\mathsf{Q}_{\mathcal{H}}$ as an object to be estimated. Market births reveal inequalities of the form

$$Y_{jt} = 1 \Rightarrow \phi_j \frac{1}{2} \operatorname{tr}(\Pi_{B_j(M_t)} \mathsf{Q}) - C_{jt} - \Psi_j > 0,$$

while nonbirths reveal the opposite inequality, subject to observability, strategic delay, and omitted admissibility constraints. This suggests a revealed-preference estimator

$$\widehat{\mathsf{Q}} \in \arg\min_{\mathsf{Q} \succeq 0} \sum_{j,t} \ell \left(Y_{jt}, \phi_j \frac{1}{2} \operatorname{tr}(\Pi_{B_j(M_t)} \mathsf{Q}) - C_{jt} - \Psi_j \right) + \lambda \left\| \mathsf{Q} \right\|_* ,$$

where $\|\mathsf{Q}\|_*$ is the nuclear norm. Under cost identification, the target is not only to predict listings but to recover the hidden residual spectrum of missing markets. Proposition 19.1 gives the stability condition under which estimated missing-market bases are precise enough to interpret.

This structural interpretation requires cost identification. If C_{jt} and Ψ_j are unobserved, a market birth cannot by itself distinguish a high-value residual direction from a low-cost representation. Recovering Q therefore requires either measured cost and admissibility proxies, or exogenous variation in cost terms–for example an oracle launch, legal-template adoption, reporting-rule change, exchange approval path, settlement-rail introduction, or collateral-standard shock–that moves implementation cost while leaving the underlying valuation-dispersion object plausibly fixed. Without such variation, the empirical claim should be stated more modestly: births test whether listing is correlated with residual-value proxies and cost shocks, not that the hidden spectrum has been structurally recovered.

Proposition 19.1 (Stability of an estimated residual basis)**.** *Let $\widehat{\mathsf{Q}}_{\mathcal{H}}$ be an estimated residual valuation-dispersion operator and suppose*

$$\left\| \widehat{\mathsf{Q}}_{\mathcal{H}} - \mathsf{Q}_{\mathcal{H}} \right\|_{op} \leq \delta.$$

If $\Delta_m = \lambda_m - \lambda_{m+1} > 2\delta$, and $\widehat{P}_m$ and P_m are the top-m spectral projections of $\widehat{\mathsf{Q}}_{\mathcal{H}}$ and $\mathsf{Q}_{\mathcal{H}}$, then

$$\left\| \widehat{P}_m - P_m \right\|_{op} \leq \frac{2\delta}{\Delta_m}.$$

Thus an estimated missing-market basis is empirically interpretable only when the residual spectrum has a gap large relative to estimation error.

Proof. This is the Davis–Kahan sin-theta theorem applied to the two symmetric operators. The stated constant is a conservative version sufficient for the paper's empirical interpretation (Davis and Kahan, 1970). □

20 Non-Obvious Implications

The theory yields implications that are not captured by the slogan "markets exist when benefits exceed costs."

1. The frictionless missing-market basis is an eigenbasis of residual valuation dispersion; the implemented basis follows a cost threshold along that spectrum.

2. The economy's market resolution is set where marginal valuation-dispersion eigenvalues cross marginal representation and admissibility costs.
3. Technology can create markets by refining the verifiable state space, not merely by lowering a scalar implementation cost.
4. Better information can also destroy markets when it reveals risk before agents can contract to share it.
5. The best coarse index for an unverifiable bespoke exposure is a conditional expectation with respect to the verification partition.
6. An economy can be Arrow-Debreu incomplete but effectively complete if the residual eigenvalue tail is small.
7. The same risk may become tradable under the cheapest basis, not the most natural description.
8. Market count and market span can diverge: many new products may add little span, while one benchmark can add a lot.
9. Some financial innovation is re-basing: it repackages existing span without adding residual risk-sharing value.
10. Standardization is basis compression: it lowers the cost of representing many adjacent payoff directions.
11. Verification breakthroughs create discontinuous waves of market creation.
12. Broad indexes can displace bespoke claims while increasing welfare and span.
13. The largest raw exposures are not necessarily the most important missing markets; residual exposure after projection onto existing claims is what matters.
14. Some socially valuable markets require subsidies because price discovery and risk-sharing benefits are diffuse public goods.
15. Settlement technology alone favors transferable claims, not necessarily high-welfare claims.
16. Privacy and manipulation constraints can cap attainable completeness permanently.
17. Claims can be payoff substitutes and cost complements at the same time.
18. Liquidity fragmentation is less binding when many claims share hedgeable factor structure.
19. Regulation by venue is coarse; claim-level payoff rules and externalities are the welfare objects.
20. Falling transaction costs can create excessive markets as well as missing markets.
21. New markets can change the valuation-dispersion operator itself by changing behavior, leverage, exposure choice, or information acquisition.
22. Reflexive claims require stable price-to-state-to-payoff mappings, not only measurable payoff rules.
23. The market frontier is path dependent because early claims change both residual value and implementation costs of later claims.
24. Market births and nonbirths can be read as noisy revealed-preference inequalities about the hidden valuation-dispersion spectrum only when implementation costs are measured or shifted by identifiable shocks.

21 Conclusion

The central object of this paper is not a security label but a payoff subspace. Complete-markets theory asks what happens when the relevant state-contingent payoff directions are tradable. This paper asks which state distinctions become verifiable, which payoff directions become tradable, which contractual bases implement them, and how technology moves the boundary. In the frictionless benchmark, the most valuable missing directions are the leading eigenvectors of residual valuation dispersion. In real economies, only costly, verifiable, admissible representations of those directions can become markets. The dimensionality of tradable risk is endogenous to the economics of claim creation, and the resolution of the market system is set by the crossing of the residual value spectrum and marginal representation cost.

The same framework also gives the limits. Information can arrive too early and eliminate insurance value. Markets can change the exposures they price. Reflexive claims can alter their own underlying states. Cheap market creation is therefore neither a monotone good nor a mere proliferation story. It is a theory of when latent payoff directions become verifiable, valuable, liquid, and admissible enough to be made explicit.

A Heterogeneous Beliefs and Speculative Surplus

The baseline uses common beliefs so that S is a welfare-relevant risk-sharing surplus. With heterogeneous beliefs, trading gains mix risk sharing with disagreement.

Suppose agent a has probability $\mathbb{P}_a$, mean $\mathbb{E}_a$, and covariance matrix $\Sigma_{a,R} = \mathrm{Var}_a(R)$. Define

$$b_a(R \mid \mathcal{H}) = \mathbb{E}_a[R] - \gamma_a \,\mathrm{Cov}_a(c_a^{\mathcal{H}}, R), \qquad A_a = (\gamma_a \Sigma_{a,R})^{-1}.$$

Demand is

$$q_a(p) = A_a(b_a - p).$$

For a zero-net-supply derivative, the clearing price is

$$p^* = \left(\sum_a A_a\right)^{-1} \left(\sum_a A_a b_a\right),$$

and perceived gross trading surplus is

$$S_R^{perceived}(\mathcal{H}) = \frac{1}{2}\sum_a (b_a - p^*)^\top A_a (b_a - p^*).$$

This object need not equal social surplus under any planner's probability measure. A paper that begins with heterogeneous beliefs must specify whether disagreement gains are welfare, speculation, information aggregation, or redistribution. The baseline therefore uses common beliefs and treats heterogeneous beliefs as an extension.

B Additional Derivations

B.1 Positive net supply

If a new claim system R has exogenous net supply $\bar{q} \in \mathbb{R}^m$, market clearing is $\sum_a q_a(p) = \bar{q}$. Under common covariance Σ_R, demand remains $q_a(p) = \theta_a \Sigma_R^{-1}(b_a - p)$. Clearing gives

$$\Sigma_R^{-1}\left(\sum_a \theta_a b_a - \Theta p\right) = \bar{q},$$

so

$$p^* = \frac{\sum_a \theta_a b_a}{\Theta} - \frac{1}{\Theta}\Sigma_R \bar{q}.$$

Zero net supply is the clean derivative case; positive net supply adds an issuance or capital-raising component.

B.2 A planner's exact first-order conditions

For a residual matrix R, the planner's incremental problem is

$$\max_{q_1,\ldots,q_A} \sum_a \left\{ q_a^\top b_a - \frac{\gamma_a}{2} q_a^\top \Sigma_R q_a \right\} \quad \text{s.t.} \quad \sum_a q_a = 0.$$

The Lagrangian is

$$\mathcal{L} = \sum_a \left\{ q_a^\top b_a - \frac{\gamma_a}{2} q_a^\top \Sigma_R q_a \right\} - \lambda^\top \sum_a q_a.$$

FOCs give $q_a = \theta_a \Sigma_R^{-1}(b_a - \lambda)$. Clearing gives $\lambda = p^*$. Hence competitive equilibrium implements the planner's incremental risk-sharing allocation under common beliefs and quasi-linear risk-free transfers.

B.3 A note on CARA-normal preferences

If terminal payoffs are jointly normal and agents have CARA utility $u_a(c) = -\exp(-\gamma_a c)$, certainty equivalents are

$$CE_a(c) = \mathbb{E}[c] - \frac{\gamma_a}{2}\,\mathrm{Var}(c),$$

so all finite-dimensional formulas in the paper apply exactly. In finite non-normal states, mean-variance utility should be read as a tractable quadratic approximation.

C Fat-Tailed Proliferation as a Companion Result

Rank candidate residual subspaces by captured private value $V_j = \phi_j S_j = Aj^{-\alpha}$ with $\alpha > 0$. Suppose a common creation/admissibility threshold is $B_t = C_t + \Psi_t$. Claim j is viable when $Aj^{-\alpha} > B_t$. Therefore the number of viable markets is

$$M_t = \#\{j \leq n : Aj^{-\alpha} > B_t\} \approx \min\left\{n, \left(\frac{A}{B_t}\right)^{1/\alpha}\right\}.$$

If $\Psi_t = 0$ or is small relative to C_t, and

$$C_t = C_0 e^{-\lambda t},$$

then before saturation,

$$M_t \approx \left(\frac{A}{C_0}\right)^{1/\alpha} e^{\lambda t/\alpha}.$$

If there is an irreducible admissibility burden $\Psi > 0$, then

$$M_t \approx \left(\frac{A}{C_0 e^{-\lambda t} + \Psi}\right)^{1/\alpha},$$

which plateaus at

$$M_\infty = \left(\frac{A}{\Psi}\right)^{1/\alpha}.$$

This result belongs in the research program, not the core theorem paper: it explains market-count growth conditional on a distribution of viable residual values, while the main text explains where those residual values come from.

References

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

Allen, Franklin, and Adelina Bărbălău. 2024. "Security Design: A Review." *Journal of Financial Intermediation* 60: 101113.

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

Athanasoulis, Stefano G., and Robert J. Shiller. 2000. "The Significance of the Market Portfolio." *Review of Financial Studies* 13(2): 301–329.

Athanasoulis, Stefano G., and Robert J. Shiller. 2001. "World Income Components: Measuring and Exploiting Risk-Sharing Opportunities." *American Economic Review* 91(4): 1031–1054.

Bisin, Alberto. 1998. "General Equilibrium with Endogenously Incomplete Financial Markets." *Journal of Economic Theory* 82(1): 19–45.

Carvajal, Andrés, Marzena Rostek, and Marek Weretka. 2012. "Competition in Financial Innovation." *Econometrica* 80(5): 1895–1936.

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

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

Das, Abhimanyu, and David Kempe. 2018. "Approximate Submodularity and Its Applications: Subset Selection, Sparse Approximation and Dictionary Selection." *Journal of Machine Learning Research* 19(3): 1–34.

Davis, Chandler, and W. M. Kahan. 1970. "The Rotation of Eigenvectors by a Perturbation. III." *SIAM Journal on Numerical Analysis* 7(1): 1–46.

Demange, Gabrielle, and Guy Laroque. 1995. "Private Information and the Design of Securities." *Journal of Economic Theory* 65(1): 233–257.

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

Fan, Ky. 1949. "On a Theorem of Weyl Concerning Eigenvalues of Linear Transformations. I." *Proceedings of the National Academy of Sciences* 35(11): 652–655.

Geanakoplos, John, Michael Magill, Martine Quinzii, and Jacques Drèze. 1990. "Generic Inefficiency of Stock Market Equilibrium when Markets are Incomplete." *Journal of Mathematical Economics* 19(1–2): 113–151.

Gershkov, Alex, Benny Moldovanu, Philipp Strack, and Mengxi Zhang. 2025. "Optimal Security Design for Risk-Averse Investors." *American Economic Review* 115(6): 2050–2092.

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

Hara, Chiaki. 1995. "Commission-Revenue Maximization in a General Equilibrium Model of Asset Creation." *Journal of Economic Theory* 65(1): 258–298.

Hart, Oliver D. 1975. "On the Optimality of Equilibrium when the Market Structure is Incomplete." *Journal of Economic Theory* 11(3): 418–443.

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.

Magill, Michael, and Martine Quinzii. 1996. *Theory of Incomplete Markets.* Cambridge, MA: MIT Press.

Natarajan, Balas K. 1995. "Sparse Approximate Solutions to Linear Systems." *SIAM Journal on Computing* 24(2): 227–234.

Ohashi, Kazuhiko. 1995. "Endogenous Determination of the Degree of Market-Incompleteness in Futures Innovation." *Journal of Economic Theory* 65(1): 198–217.

Pesendorfer, Wolfgang. 1995. "Financial Innovation in a General Equilibrium Model." *Journal of Economic Theory* 65(1): 79–116.

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.

Shiller, Robert J. 1993. *Macro Markets: Creating Institutions for Managing Society's Largest Economic Risks.* Oxford: Oxford University Press.

Simsek, Alp. 2013. "Speculation and Risk Sharing with New Financial Assets." *Quarterly Journal of Economics* 128(3): 1365–1396.