Admissible Market Design

Claim-Centric Regulation When Market Creation Is Cheap

Grant Stenger

Working paper

June 2026

Abstract

When the cost of creating state-contingent claims is high, many harmful markets are never born because they cannot recover their fixed costs. As specification, verification, settlement, and algorithmic market-making costs fall, that accidental screen disappears. The central normative question becomes not whether a claim can trade, but under what conditions it may trade. This paper develops a claim-centric theory of admissible market design for an economy in which markets can be embedded in exchanges, apps, wallets, protocols, devices, personal agents, and legal entities. The object of regulation is not a venue label but a representation: payoff rule, oracle, data permissions, participant set, access rule, collateral, leverage, liquidity mechanism, governance, and externality profile.

The paper builds on a costly-basis theory of market creation and a synthetic-liquidity theory of long-tail claims. It separates technical feasibility, legal permissibility, and normative admissibility. Private entry is governed by captured surplus net of implementation, liquidity, and private compliance costs; social admissibility is governed by risk-sharing, immediacy, information, and legitimate consumption benefits net of privacy, coercion, manipulation, moral hazard, addiction, dignity, norm-erosion, and systemic externalities. The main results show: private market creation and social desirability coincide only under knife-edge capture and externality assumptions; venue-based rules are generally dominated by claim-level rules when venues host heterogeneous claims; liquidity is not welfare because the same payoff can be supported by hedging, information, entertainment, exploitation, or manipulation; belief-driven markets are admissible only when their information or hedging value dominates speculative risk amplification; better information can create markets through ex post verification while destroying insurance through pre-trade revelation; human-capital claims require a no-control boundary; manipulable and reflexive claims require incentive-compatibility constraints; and systemic admissibility is non-additive because a claim harmless in isolation can be dangerous through shared hedges, collateral reuse, cross-margining, and forced liquidation.

The paper's practical output is an admissibility frontier and policy operator: allow, standardize, subsidize, tax, margin, restrict access, impose oracle governance, or prohibit. The framework rejects both naive market optimism and blanket anti-financialization. Cheap market creation is valuable when it moves risk to better bearers, supplies useful immediacy, produces decision-relevant public information, or enables legitimate consumption; it is harmful when private volume is driven by extraction, compulsion, manipulation, addictive speculation, privacy leakage, or hidden systemic leverage.

Keywords: market design; financial innovation; state-contingent claims; incomplete markets; prediction markets; derivatives; insurance; regulation; welfare economics; speculation; privacy; systemic risk.

JEL codes: D47, D53, D61, D62, D82, G13, G18, K23, L51.

Contents

1 Introduction

A market can be technically possible, liquid, legal, and still bad. A different market can be thin, subsidized, or institutionally awkward and still be socially valuable. The distinction is easy to miss in a world where market creation is expensive. High fixed costs perform a crude screening function: many claims that would be exploitative, trivial, manipulative, or systemically dangerous are never created because they cannot amortize specification, verification, settlement, liquidity, and compliance costs. But that screening function is accidental. If those costs collapse, feasibility ceases to be permission.

This paper is the normative and institutional layer of the Trillion Markets research program. Paper 1 studies when a residual payoff direction is valuable enough to justify claim creation. Paper 2 studies when a long-tail claim can be quoted by mapping common risk into hedgeable factors, diversifying residual flow, or subsidizing information-market liquidity. Paper 4 studies the shared infrastructure beneath prediction, insurance, derivatives, and parametric claims. Paper 5 studies how agents lower the minimum efficient transaction size and search over implemented claims. This paper asks the question that the others deliberately postpone: *which technically feasible markets should exist?*

The answer cannot be given at the level of venues. In the old regulatory picture, markets lived in identifiable rooms, exchanges, broker-dealers, clearinghouses, betting shops, insurance carriers, and banks. Those institutions still matter, but they no longer exhaust the design space. A state-contingent claim can be offered through an exchange, a bilateral template, a protocol, a wallet, an app, an employer platform, a creator-financing marketplace, a parametric insurance product, a prediction-market automated market maker, or an agent-to-agent contract. The venue is a container. The welfare object is the claim.

A claim-centric approach asks a different set of questions. What is the payoff rule? What facts does it condition on? Who observes them? Who resolves disputes? Who may trade? Why do they trade? Is the dominant flow hedging, immediacy, information, entertainment, disagreement, addiction, manipulation, or coercion? Does the claim transfer an existing risk to a better bearer, or does it manufacture a risk no one needed to hold? Does its price improve decisions outside the market, or merely arm informed traders against uninformed ones? Does it invade privacy? Does it give investors control over a person's life choices? Does it change the underlying event? Does it create moral hazard? Does it load on crowded hedges or reusable collateral? Does it push tail risk onto the same residual warehouse as thousands of other claims?

These questions are not refinements after the real economics is done. They are the economics of admissibility. Market design has traditionally asked how to make a desired market work. In a cheap-creation economy, market design must also ask when the desired object is a market at all, when it should be a public statistic, when it should be insurance, when it should be a restricted professional contract, when it should be subsidized, and when it should not be allowed to exist.

The paper's central claim is:

> *When market creation becomes cheap, regulation must move from venue-centric oversight to claim-centric admissibility. The admissibility of a claim is a function of its payoff rule, representation, participants, flow ecology, information timing, control rights, manipulation incentives, privacy burden, and systemic externality.*

This claim is not an argument for suppressing financial innovation. Many missing markets are missing for bad reasons: verification is expensive, risk pools are coarse, prices are public goods, retail access is bundled with predation, or no single intermediary can capture enough of the social surplus. A claim-centric theory can justify subsidies, public data, safe harbors, standardized templates, designated market makers, and public provision. Nor is the claim an argument for financializing everything. Some markets are socially harmful precisely because cheap creation makes them profitable. A claim-centric theory can justify access limits, margin, taxation, oracle rules, fiduciary duties, privacy constraints, hard prohibitions, and public floors.

The constructive task is to define an *admissibility frontier*. Claims on one side may trade, perhaps with design conditions. Claims on the other side should be restricted or prohibited. Some claims should be subsidized because their prices are public goods or their hedging value is diffuse. Some claims should be taxed or margined because private trading creates externalities. Some claims should be allowed only for hedgers or sophisticated counterparties. Some claims should never attach to intimate life choices, personhood, political rights, medical decisions, reproductive decisions, religious practice, or coercive labor.

Contributions

The paper makes seven contributions.

First, it separates *technical feasibility*, *legal permissibility*, and *normative admissibility*. A claim is technically feasible when it can be specified, verified, settled, and quoted. It is legally permissible when it can be offered under a given statute, venue rule, or regulatory interpretation. It is normatively admissible when it should be allowed, restricted, subsidized, taxed, margined, or prohibited after accounting for welfare, externalities, rights, and systemic effects. The paper is about the third concept.

Second, it gives a welfare decomposition for state-contingent claims. In the common-beliefs benchmark, the risk-sharing component is the residual-span value from Paper 1. Under heterogeneous beliefs, the paper separates risk-sharing dispersion from pure-disagreement dispersion in the spirit of Simsek (2013) and evaluates pure disagreement using the belief-neutral welfare discipline of Brunnermeier et al. (2014). More generally, trading in a claim can produce hedging value, immediacy value, information value, legitimate consumption value, and private speculative gains, while also producing privacy loss, coercion, manipulation, moral hazard, addictive speculation, dignity harms, norm erosion, and systemic risk. The sign of social welfare cannot be inferred from volume, liquidity, or payoff shape.

Third, it states a private-social divergence principle. Private creation and social desirability coincide only when captured private surplus equals social surplus, implementation and liquidity costs are internalized, and externalities are zero or perfectly priced. Otherwise, falling costs can increase the number of markets while reducing welfare if the first claims crossing the private threshold are socially harmful.

Fourth, it formalizes the claim-centric critique of venue-based regulation. If a venue hosts two claims with the same venue label but opposite welfare signs, a venue-level rule must either admit both or restrict both. A claim-level rule weakly dominates whenever the two claims require different treatment.

Fifth, it gives admissibility tests for the most important boundary cases: disagreement and speculation, information timing, human-capital claims, oracle and event manipulation, reflexive underlyings, and systemic externality. These tests are stated as constraints, not slogans.

Sixth, it integrates Paper 2's liquidity regimes into welfare analysis. Factor-hedged dealer liquidity, bookmaker or insurer residual-diversification liquidity, and subsidized information-market liquidity have different welfare interpretations. The same payoff can be supported by hedging, entertainment, addiction, manipulation, or public-good price discovery. Liquidity source is an admissibility input.

Seventh, it gives a policy operator. The planner's action set is not simply allow or ban. The relevant instruments include standardization, public data, subsidy, access rules, suitability, disclosure, margin, leverage limits, position limits, oracle governance, dispute procedures, privacy-preserving verification, fiduciary duties for personal agents, taxes, and hard prohibitions.

Reader's map and formal spine

The core formal spine is as follows. Section 3 defines claims, representations, technical feasibility, legal permissibility, and normative admissibility. Section 4 defines private and social value. Principle 4.3 states the private-social divergence principle and Observation 4.5 gives the four-region policy table. Proposition 5.1 gives the venue-insufficiency result. Section 6 defines flow ecology, proves that liquidity is not welfare, and gives a measurable harmful-cell diagnostic. Section 7 gives the disagreement/speculation boundary, including the disagreement-risk-sharing separation theorem. Section 8 gives the admissibility implication of the information-timing theorem, whose infrastructure version belongs to Paper 4. Section 9 gives the human-capital and no-control boundary. Section 10 gives oracle, manipulation, and reflexivity constraints. Section 11 gives the paper's systemic theorem: non-additive systemic admissibility. Section 13 turns the theory into a policy operator. Section 15 gives empirical tests and policy surfaces.

2 Related Literature and Positioning

The paper sits between incomplete-markets theory, security design, market design, welfare economics, information economics, privacy economics, microstructure, systemic risk, and moral limits of markets. Its novelty is not that markets can fail, not that externalities matter, and not that some transactions are repugnant. The novelty is to put these facts into the same claim-level framework as endogenous state-contingent market creation.

Complete and incomplete markets. Arrow-Debreu theory defines the complete-contingent-commodity benchmark (Arrow and Debreu, 1954; Debreu, 1959). Radner and the incomplete-markets literature study equilibrium and welfare for a given asset structure (Radner, 1972; Hart, 1975; Magill and Quinzii, 1996). Paper 1 of this program asks which payoff directions become tradable at all; this paper asks which technically tradable directions should be permitted.

Financial innovation and security design. The endogenous-security literature studies financial innovation and the design of securities under incomplete markets (Allen and Gale, 1994; Duffie and Rahi, 1995; Pesendorfer, 1995; Demange and Laroque, 1995; Hara, 1995; Bisin, 1998). Recent security-design work continues to derive optimal securities from primitives and constraints (Gershkov et al., 2025). This paper treats security design as only one layer. A claim can be well-designed as a payoff and still fail admissibility because of flow, access, manipulation, privacy, or systemic externality.

Transaction costs and the boundary of the price system. Coase explains institutional boundaries by the cost of using the price system (Coase, 1937). Paper 5 moves this logic to the minimum efficient transaction. This paper adds that the efficient boundary is not merely $v - c - T > 0$ but $v - c - T - E > 0$, with hard rights constraints when some harms are not commensurable with price.

Information, disagreement, and speculation. Hayek emphasizes the price system as a mechanism for aggregating dispersed knowledge (Hayek, 1945). Grossman and Stiglitz show that costly information prevents perfect informational efficiency (Grossman and Stiglitz, 1980). Hirshleifer shows that information can reduce social risk-sharing value when it arrives before trade (Hirshleifer, 1971). Simsek shows that financial innovation under belief disagreement can increase speculative variance even as it expands risk-sharing opportunities (Simsek, 2013). Brunnermeier et al. (2014) supply the welfare criterion this paper needs: when beliefs are heterogeneously distorted, mutually inconsistent perceived gains should not be counted automatically as social surplus. These results are central: belief-driven volume is not automatically welfare, and better information is not monotonically market-creating.

Market design, repugnance, and moral limits. Market design studies allocation rules and institutions, including settings where prices alone are insufficient (Roth, 2015). Roth's analysis of repugnance shows that social constraints can determine which markets exist even when trade would be mutually beneficial in narrow terms (Roth, 2007). Titmuss, Satz, Sandel, and Bowles emphasize that markets can affect norms, personhood, civic relations, and preferences (Titmuss, 1970; Satz, 2010; Sandel, 2012; Bowles, 1998). This paper translates that tradition into claim-level constraints: some payoff rules are inadmissible not because they cannot be priced but because pricing changes or violates the object priced.

Privacy, addiction, and vulnerable agents. Privacy economics studies the value and consequences of personal-information disclosure (Acquisti et al., 2016). Behavioral models of hyperbolic discounting and self-control show why voluntary trades can be exploitative when counterparties design products around present bias or naivete (Laibson, 1997; O'Donoghue and Rabin, 1999). These literatures discipline the paper's claim that consent is necessary but not always sufficient.

Microstructure, liquidity, and systemic risk. Informed-trading and dealer models show that spreads reflect adverse selection, inventory, and immediacy costs (Glosten and Milgrom, 1985; Kyle, 1985; Garman, 1976; Ho and Stoll, 1981). Paper 2 turns those ideas into a synthetic-liquidity model for long-tail claims. Systemic-risk research emphasizes amplification, interconnectedness, forced liquidation, clearing networks, and phase transitions in financial networks (Eisenberg and Noe, 2001; Elliott et al., 2014; Acemoglu et al., 2015; Brunnermeier and Oehmke, 2013). This paper uses those mechanisms as admissibility constraints: a liquid market can be socially harmful if its liquidity source is extractive or systemically fragile.

Reflexivity and real effects. Prices can affect the real actions that determine the payoffs being priced. The real-effects literature studies feedback from market prices to corporate and economic decisions (Bond et al., 2012); the performativity literature emphasizes that financial models and market devices can reshape the markets they describe (MacKenzie, 2006). This paper does not attempt a full model of real effects. It imports the lesson as an admissibility burden: reflexive claims require a stable price-to-action-to-state fixed point.

Prediction markets and public information. Prediction-market and scoring-rule mechanisms can produce information prices, sometimes with bounded subsidy requirements (Hanson, 2003, 2007; Wolfers and Zitzewitz, 2004). Paper 4 treats prediction, insurance, derivatives, and parametric products as wrappers around state-contingent payoff functions. This paper adds purpose: the same event payoff can be hedging, gambling, manipulation, governance, or public information production.

What is new. Prior work gives the pieces: incomplete markets, financial innovation, speculation, privacy, repugnance, microstructure, and systemic risk. This paper's contribution is to unify them at the level of the claim. The state variable is not simply "market exists" or "venue allowed." It is a claim representation embedded in a market state. The admissibility decision is a constrained social-design problem over those representations.

3 Claims, Market State, and Three Kinds of Admissibility

3.1 State-contingent claims and representations

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a finite probability space unless otherwise stated. A payoff is a measurable function $g : \Omega \to \mathbb{R}$. A tradable claim is not merely g. It is an implemented representation.

Definition 3.1 (Claim representation)**.** *A claim representation is a tuple*

$$\rho = (g_\rho, \mathcal{O}_\rho, \mathcal{D}_\rho, \mathcal{U}_\rho, \mathcal{L}_\rho, \mathcal{M}_\rho, \mathcal{G}_\rho, \mathcal{A}_\rho),$$

where g_ρ is the delivered payoff rule, $\mathcal{O}_\rho$ is the oracle or resolution process, $\mathcal{D}_\rho$ is the data and permissioning regime, $\mathcal{U}_\rho$ is the participant and access set, $\mathcal{L}_\rho$ is the liquidity mechanism, $\mathcal{M}_\rho$ is the margin, collateral, and leverage rule, $\mathcal{G}_\rho$ is governance and dispute resolution, and $\mathcal{A}_\rho$ is the advertised purpose or permitted-use class.

The same payoff rule can have different representations. A weather-contingent payoff can be an exchange-traded derivative, a farmer's parametric insurance policy, a prediction-market event contract, a reinsurance sidecar, or a private bilateral hedge. These representations can share g but differ in access, margin, oracle, dispute procedure, leverage, legal treatment, liquidity source, and welfare consequences.

Let $\mathcal{R}$ be the universe of possible representations. Let $\pi(\rho) = g_\rho$ denote the delivered payoff. Delivered payoff matters: if default, oracle error, manipulation, or legal uncertainty changes settlement, the payoff is not the label on the term sheet but the state-contingent cash flow actually delivered.

3.2 Market state

A claim's value and admissibility depend on the market state in which it is introduced. Let

$$\Xi = (S, N), \qquad S = (\mathcal{H}, \mathsf{Q}, \kappa, L, A)$$

be the Paper 3 admissibility state. Paper 5 defines the program's compact core state $S = (\mathcal{H}, \mathsf{Q}, \kappa, L, A)$ and full state $\mathsf{State} = (S, I, N)$. Paper 3 opens the network slice N because systemic admissibility often requires collateral, hedges, exposures, oracle commonality, leverage, and cross-market dependencies to be represented explicitly. Paper 4 opens the infrastructure slice I. Paper 5 treats the dynamic update

$$(S_t, N_t) \to (S_{t+1}, N_{t+1})$$

as the computational problem of autonomous market proposal. This paper treats Ξ as the state at which admissibility is assessed.

3.3 Technical, legal, and normative admissibility

Definition 3.2 (Technical feasibility)**.** *A representation ρ is technically feasible at technology state T if its payoff can be specified, the relevant state variables can be observed or inferred, the oracle can resolve them with bounded error, settlement can occur, and at least one liquidity mechanism can quote or clear the claim at finite cost. The set of technically feasible representations is denoted $\mathcal{R}^{tech}(T)$.*

Technical feasibility is the domain of Papers 1, 2, 4, and 5: specification, verification, costly basis selection, synthetic liquidity, infrastructure, and transaction costs.

Definition 3.3 (Legal permissibility)**.** *A representation ρ is legally permissible in jurisdiction or rule system j if it can be offered under the applicable statutes, licenses, exemptions, venue rules, disclosure rules, and regulatory interpretations. The set is denoted $\mathcal{R}_j^{leg}$.*

Legal permissibility is partly exogenous to this paper. The paper is not jurisdiction-specific legal drafting. It studies the welfare logic that should inform legal categories.

Definition 3.4 (Normative admissibility)**.** *A representation ρ is normatively admissible at market state Ξ if it satisfies hard admissibility constraints and its marginal social value, under the appropriate policy action, is nonnegative. The unconstrained social net value is*

$$W_\rho(\Xi) = B_\rho^{soc}(\Xi) - K_\rho^{soc}(\Xi) - E_\rho(\Xi),$$

where B_ρ^{soc} is social benefit, K_ρ^{soc} is social resource cost including implementation and liquidity cost, and E_ρ is externality and constraint burden. The hard-constraint set is denoted $\mathcal{R}^{hard}(\Xi)$. Thus the simplest allow condition is

$$\rho \in \mathcal{R}^{tech}(T) \cap \mathcal{R}_j^{leg} \cap \mathcal{R}^{hard}(\Xi) \quad \text{and} \quad W_\rho(\Xi) \geq 0.$$

Hard constraints represent boundaries that are not safely priced by Pigouvian taxes: coercive control over persons, sale of political rights, claims on minors' intimate futures, nonconsensual biometric surveillance, or payoff rules whose manipulation would create unacceptable physical harm. Soft externalities can be priced, margined, disclosed, insured, or restricted. Hard constraints define the outside boundary.

Remark 3.5 (Admissibility is representation-level)**.** *Admissibility is not a property of g alone. The same payoff shape can be admissible for one participant set and inadmissible for another; admissible with low leverage and inadmissible with hidden leverage; admissible with privacy-preserving verification and inadmissible with invasive data extraction; admissible with a robust oracle and inadmissible with manipulable resolution.*

3.4 Residual payoff value from Paper 1

To connect the normative layer to the costly-basis layer, suppose the risky payoff space is the centered Hilbert space $\mathcal{X} = L_0^2(\mathbb{P})$ and existing markets span $\mathcal{H} \subseteq \mathcal{X}$. For a claim system G, Paper 1 defines the residual subspace

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

In the common-beliefs mean-variance benchmark, the risk-sharing value of a residual subspace B is

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

where Q is the valuation-dispersion operator. This paper uses $V(B \mid \mathcal{H})$ as the hedging and risk-sharing component of social benefit, not as total welfare.

Definition 3.6 (Representation cost and liquidity cost)**.** *Let $\kappa(\rho, \Xi)$ denote the implementation cost of representation ρ: specification, data, verification, legal form, compliance, custody, collateral, pricing, settlement, dispute resolution, and distribution. Let $\ell(\rho, \Xi)$ denote the liquidity cost stack from Paper 2: factor hedge cost, residual capital, adverse-selection loss, model and oracle risk, operating cost, funding, collateral, and capital constraints. The social resource cost is*

$$K_\rho^{soc}(\Xi) = \kappa(\rho, \Xi) + \ell^{soc}(\rho, \Xi),$$

where ℓ^{soc} includes real resource and risk-bearing costs rather than pure transfers.

4 The Welfare Object

4.1 Benefit components

A claim can generate value through several channels. Let

$$B_\rho^{soc} = S_\rho^{hedge} + S_\rho^{imm} + S_\rho^{info} + S_\rho^{cons} + S_\rho^{other} + S_\rho^{dist},$$

where:

- S_ρ^{hedge} is risk-sharing, insurance, or financing value. In the common-beliefs residual-span benchmark, $S_\rho^{hedge} = V(B(\rho \mid \mathcal{H}) \mid \mathcal{H})$ or the marginal value of the residual subspace introduced by ρ.
- S_ρ^{imm} is immediacy, execution, convenience, or balance-sheet transformation value.
- S_ρ^{info} is the social value of decision-relevant price discovery.
- S_ρ^{cons} is legitimate consumption or participation value, such as entertainment, identity, or engagement, counted only when participation is voluntary, non-exploitative, and not driven by manipulation or addiction.
- S_ρ^{other} captures other positive spillovers, such as standardization, benchmark creation, or infrastructure learning.
- S_ρ^{dist} is a distributional welfare term when the planner uses non-unit welfare weights.

The decomposition is not meant to imply that each term is easy to measure. It is meant to prevent the mistake of treating all volume as risk sharing or all speculation as zero value.

4.2 Externality and constraint burden

Let

$$E_\rho = E_\rho^{priv} + E_\rho^{coer} + E_\rho^{manip} + E_\rho^{moral} + E_\rho^{addict} + E_\rho^{ineq} + E_\rho^{oracle} + E_\rho^{sys} + E_\rho^{dignity} + E_\rho^{norm} + E_\rho^{legal-risk}.$$

These terms include privacy loss, coercion, manipulation of the underlying event or oracle, moral hazard, addictive speculation, inequality and exploitation of vulnerable participants, oracle fragility, systemic risk, dignity harms, norm erosion, and legal uncertainty not already counted in resource cost.

Some terms are continuous and priceable. Others are constraints. Formally, let hard violations be $h_k(\rho, \Xi) \leq 0$ for $k = 1, \ldots, m$. The hard-admissible set is

$$\mathcal{R}^{hard}(\Xi) = \{\rho \in \mathcal{R} : h_k(\rho, \Xi) \leq 0 \text{ for all } k\}.$$

Equivalently, one may write an extended-valued externality

$$\bar{E}_\rho(\Xi) = E_\rho(\Xi) + \sum_{k=1}^{m} \mu_k \max\{h_k(\rho, \Xi), 0\},$$

with $\mu_k = +\infty$ for non-priceable rights constraints.

4.3 Private creation

A market creator does not capture all social benefit. Let B_ρ^{priv} be the private monetizable benefit from fees, spreads, listing revenue, market-maker rents, data revenue, index intellectual property, collateral revenue, distribution control, and other captured sources. Let $\phi_\rho \in [0, \infty)$ be the capture parameter mapping underlying activity into creator rents. In many information markets $\phi_\rho < 1$ because prices are public goods. In extractive distribution channels ϕ_ρ can be high even when social welfare is negative.

Private net value is

$$P_\rho(\Xi) = \phi_\rho B_\rho^{priv}(\Xi) - K_\rho^{priv}(\Xi) - \Psi_\rho^{priv}(\Xi), \tag{1}$$

where K_ρ^{priv} is private implementation and liquidity cost and Ψ_ρ^{priv} is private legal, compliance, capital, and reputational burden. Private entry occurs when $P_\rho > 0$.

Remark 4.1 (The capture parameter is endogenous)**.** *The parameter ϕ_ρ is not a nuisance. It is shaped by market architecture: exchange fees, spreads, data fees, benchmark intellectual property, exclusive distribution, custody, settlement, compliance, public-good leakage, competition, and regulation. A policy that changes who can monetize a price changes which claims are privately created, even when social value is unchanged.*

4.4 Social net value

Social net value is

$$W_\rho(\Xi) = B_\rho^{soc}(\Xi) - K_\rho^{soc}(\Xi) - E_\rho(\Xi). \tag{2}$$

A claim can be privately viable but socially harmful, privately nonviable but socially valuable, both, or neither. This is the central wedge.

Assumption 4.2 (Baseline comparison)**.** *When stating algebraic threshold results, assume a single marginal claim ρ, fixed market state Ξ, no hard-constraint violation, and no interaction with other candidate claims except through the displayed benefit, cost, and externality terms. Later sections relax additivity for systemic risk and infrastructure.*

Principle 4.3 (Private-social divergence)**.** *Private creation and social desirability have the same sign for every candidate claim only under restrictive alignment conditions: private captured benefit net of private burdens must be an increasing affine transformation of social benefit net of externality, and implementation/liquidity costs must be internalized in the same way. In the one-claim threshold model*

$$P_\rho = \phi B^{priv} - C - \Psi, \qquad W_\rho = B^{soc} - C - E,$$

there is a privately viable but socially harmful interval of costs whenever

$$B^{soc} - E < C < \phi B^{priv} - \Psi. \tag{3}$$

There is a privately nonviable but socially valuable interval whenever

$$\phi B^{priv} - \Psi < C < B^{soc} - E. \tag{4}$$

Consequently, a fall in creation cost can increase the number of markets while reducing social welfare if the claims crossing the private threshold satisfy (3)*.*

Proof. Private entry is $P_\rho > 0$, equivalently $C < \phi B^{priv} - \Psi$. Social harm is $W_\rho < 0$, equivalently $C > B^{soc} - E$. Both hold exactly on (3). Private nonviability and social value are $C > \phi B^{priv} - \Psi$ and $C < B^{soc} - E$, which gives (4). If costs fall from above the private threshold into the harmful interval, private entry occurs and social welfare changes by $W_\rho < 0$. Coincidence of signs for every claim requires the two thresholds to coincide claim by claim, which is the stated alignment condition. □

Corollary 4.4 (Cheap creation is not monotone welfare)**.** *If there exists a claim with* $\phi B^{priv} - \Psi > B^{soc} - E$*, then lowering creation cost can move that claim from absent to privately created while it remains socially harmful. Thus cheaper market creation is not, by itself, a welfare theorem.*

Observation 4.5 (Four-region classification)**.** *For a candidate claim* ρ *with private net value* P_ρ *and social net value* W_ρ*, the policy classification is:*

Region	*Condition*	*Default policy implication*
Privately viable, socially valuable	$P_\rho > 0$, $W_\rho > 0$	*allow; standardize; reduce unnecessary frictions*
Privately viable, socially harmful	$P_\rho > 0$, $W_\rho < 0$	*tax, margin, restrict access, redesign, or prohibit*
Privately nonviable, socially valuable	$P_\rho < 0$, $W_\rho > 0$	*subsidize, public provision, mandate data standards, or create safe harbor*
Privately nonviable, socially harmful	$P_\rho < 0$, $W_\rho < 0$	*ignore, monitor, or prohibit if circumvention risk is high*

Proof. The table enumerates the four possible sign pairs of (P_ρ, W_ρ). The policy implication follows from whether private entry will occur and whether that entry is socially desirable. □

5 Why Venue-Based Regulation Is Not Enough

Venue labels are useful for licensing, supervision, custody, clearing, disclosure, and enforcement. They are not sufficient welfare objects. The same venue can list an income hedge, an addictive zero-sum event contract, a manipulable governance claim, a socially valuable public forecast, and a leveraged derivative whose systemic risk depends on cross-margining. A single venue label cannot encode these differences.

Proposition 5.1 (Venue insufficiency)**.** *Let a regulator observe only a venue attribute* $v(\rho)$ *and choose a venue-level action* $a(v)$ *that applies to all claims with that attribute. Suppose there exist two claims* ρ_1, ρ_2 *with* $v(\rho_1) = v(\rho_2)$*, hard constraints satisfied, and opposite welfare signs:*

$$W_{\rho_1} > 0, \qquad W_{\rho_2} < 0.$$

If at least one action set contains separate treatments that would allow ρ_1 and restrict ρ_2, then any venue-level rule is weakly dominated by a claim-level rule and strictly dominated whenever allowing both or restricting both is welfare inferior to separate treatment.

Proof. A venue-level rule must assign the same action to ρ_1 and ρ_2 because it observes only v. If it allows both, it incurs the negative welfare of ρ_2. If it restricts both, it forgoes the positive welfare of ρ_1. A claim-level rule can allow ρ_1 and restrict ρ_2, weakly improving welfare and strictly improving it whenever the lost value or avoided harm is nonzero relative to the common action. □

Remark 5.2 (Venue regulation remains necessary)**.** *The theorem does not say venues are irrelevant. Venue-level supervision supplies capital, custody, reporting, dispute, governance, and enforcement infrastructure. The claim is that venue-level permission is too coarse as the final admissibility screen.*

Definition 5.3 (Claim-level policy operator)**.** *A policy operator is a mapping*

$$\Pi : \mathcal{R} \times \Xi \to \mathcal{P}(\mathcal{Y}),$$

where $\mathcal{Y}$ is the action set

$\mathcal{Y} = \{\textit{allow}, \textit{standardize}, \textit{subsidize}, \textit{tax}, \textit{margin}, \textit{restrict access}, \textit{disclose}, \textit{govern oracle}, \textit{cap}, \textit{prohibit}\}$.

The output may be a bundle of actions rather than a single action.

A claim-level operator can still use venue information as an input. It simply refuses to make the venue the entire object.

6 Flow Ecology: Why People Trade

A claim's welfare depends on why people trade it. The payoff rule describes the cash flow. The flow ecology describes the demand sustaining the market.

Definition 6.1 (Flow mix)**.** *For a claim ρ, let*

$$\lambda_\rho = (\lambda^H, \lambda^I, \lambda^F, \lambda^C, \lambda^S, \lambda^M, \lambda^X)$$

denote the local intensity or share of hedging/insurance flow (H), immediacy/convenience flow (I), information flow (F), consumption/entertainment flow (C), disagreement/speculative flow (S), manipulative flow (M), and extractive or coercive flow (X).

Flow type	Liquidity effect	Welfare interpretation
Hedging / insurance	supplies natural counterparties and risk transfer	generally positive when pricing is fair and pools are not destroyed
Immediacy / convenience	pays dealers for execution and balance-sheet service	positive if voluntary, disclosed, and non-extractive
Information	creates adverse selection and price discovery	positive when prices improve decisions outside the market
Consumption / entertainment	can support broad non-toxic flow	positive only if non-addictive and non-exploitative
Disagreement / speculation	creates volume, volatility, and private perceived gains	ambiguous; can amplify risk without social benefit
Manipulative / coercive	can generate high volume	negative or inadmissible

Proposition 6.2 (Liquidity is not welfare). *There exist two claims with identical payoff rule g, technical feasibility, and quoted liquidity, but opposite social welfare signs. Therefore payoff shape and liquidity are insufficient statistics for admissibility.*

Proof. Fix a payoff g and implementation cost K. Construct representation ρ^+ in which the dominant flow is hedging and the claim reduces agents' uninsurable income risk, yielding $S^{hedge} > K+E$. Then $W_{\rho^+} > 0$. Construct representation ρ^- with the same g and same mechanical liquidity but with flow generated by addiction, manipulation, or coercive distribution, so $S^{hedge} = S^{info} = 0$ and $E > K$ or $E > S^{cons} - K$. Then $W_{\rho^-} < 0$. Because payoff and liquidity are the same by construction, they cannot determine admissibility. □

Remark 6.3 (The sportsbook lesson). *Bookmaker liquidity can be technically excellent when flow is broad, independent, and non-toxic. That fact says little by itself about welfare. If flow is insurance-like or legitimate entertainment, welfare may be positive. If flow is addiction, cognitive bias, or predation, the same liquidity technology is a harm amplifier.*

Definition 6.4 (Extractive liquidity). *A claim has extractive liquidity if a material part of spread revenue, fees, or creator rents is generated by participants whose demand is produced by addiction, deception, hidden leverage, coercion, dark-pattern distribution, or systematic misunderstanding of the payoff. Extractive liquidity enters E_ρ rather than B_ρ^{soc}.*

Proposition 6.5 (Harmful-cell diagnostic). *Let the local flow mix be statistically estimable, and write a reduced-form private revenue approximation and social welfare approximation as*

$$P_\rho^{flow} \approx \sum_k r_k \lambda_\rho^k - K_\rho - \Psi_\rho, \qquad W_\rho^{flow} \approx \sum_k (w_k - e_k)\lambda_\rho^k - K_\rho - E_\rho^{nonflow}.$$

Here r_k is the creator's monetized revenue per unit of flow type k, w_k is the social benefit of that flow type, e_k is its flow-specific externality, and $E_\rho^{nonflow}$ includes manipulation, privacy, rights, and systemic costs not captured by flow shares. A candidate is in the privately viable/socially harmful cell whenever

$$\sum_k r_k \lambda_\rho^k > K_\rho + \Psi_\rho \quad \textit{and} \quad \sum_k (w_k - e_k)\lambda_\rho^k < K_\rho + E_\rho^{nonflow}.$$

Thus high liquidity supplied by components with large r_k but negative $w_k - e_k$, such as extractive or manipulative flow, is an observable warning sign rather than evidence of welfare.

Proof. The first inequality is the private-entry condition under the local revenue approximation. The second is the negative-social-welfare condition under the local welfare approximation. Their intersection is exactly the harmful cell in Observation 4.5. The diagnostic is reduced form: it identifies the empirical primitives to estimate rather than claiming those primitives are structural constants. □

Proposition 6.6 (Flow-source admissibility). *Let λ_ρ be observable or statistically estimable. If the marginal social value of flow component k is w_k and the marginal externality is e_k, then a local flow-based welfare approximation is*

$$W_\rho \approx \sum_k (w_k - e_k)\lambda_\rho^k - K_\rho.$$

A shift in flow mix toward components with $w_k - e_k < 0$ can make a technically unchanged market inadmissible.

Proof. The expression is the first-order approximation of social welfare in the flow shares. Holding payoff, technology, and cost fixed, the sign changes when the weighted flow mix crosses zero. □

7 Disagreement, Speculation, and Information

Common beliefs give risk-sharing formulas a clean welfare interpretation. Heterogeneous beliefs complicate admissibility. A new claim may permit useful hedging, reveal information, or merely create a new object of disagreement. The same trading volume can be welfare-improving, redistributive, or harmful.

7.1 Perceived surplus versus social surplus

Consider a zero-net-supply residual claim $u \in \mathcal{H}^{\perp}$. Agent a has subjective mean

$$\mu_a = \mathbb{E}_a[u],$$

common second moments, local risk aversion $\gamma_a > 0$, and background exposure W_a. Let

$$\kappa_a = \operatorname{Cov}(W_a, u)$$

be the covariance of the candidate payoff with the agent's background risk. In a local CARA-Gaussian or quadratic approximation, the relevant reservation value is not the subjective mean alone. It is the mean net of the risk penalty from background exposure:

$$b_a = \mu_a - \gamma_a \kappa_a.$$

Thus cross-sectional valuation dispersion can come from hedging demand, belief disagreement, or their interaction. That interaction is the object lost when one simply assumes a clean additive split between a risk-sharing operator and a disagreement operator.

Theorem 7.1 (Primitive disagreement-risk-sharing separation). *Let $u \in \mathcal{H}^{\perp}$ be a one-dimensional residual claim with common variance $\sigma_u^2 > 0$. Agent a's local demand is generated by a mean-variance objective with subjective mean $\mu_a = \mathbb{E}_a[u]$, background covariance $\kappa_a = \operatorname{Cov}(W_a, u)$, and risk aversion $\gamma_a > 0$. Define*

$$b_a = \mu_a - \gamma_a \kappa_a, \qquad \tau_a = \frac{1}{\gamma_a \sigma_u^2}, \qquad T = \sum_a \tau_a,$$

and let Var_τ and Cov_τ denote cross-sectional variance and covariance under weights τ_a/T. Then the competitive one-claim price is

$$p^* = \frac{\sum_a \tau_a b_a}{T},$$

and perceived gross trading surplus is

$$S_u^{perceived} = \frac{1}{2} T \operatorname{Var}_\tau(b) = \frac{1}{2} T \left[\operatorname{Var}_\tau(\gamma\kappa) + \operatorname{Var}_\tau(\mu) - 2 \operatorname{Cov}_\tau(\mu, \gamma\kappa) \right].$$

Under a belief-neutral social evaluation that does not credit pure mutually inconsistent subjective gains unless they produce separately modeled information, consumption, or institutional value, the common-belief risk-sharing component is

$$S_u^R = \frac{1}{2} T \operatorname{Var}_\tau(\gamma\kappa).$$

If a private creator captures fraction ϕ of perceived surplus and pays private burden $C + \Psi$, while the social evaluator pays $C + E$, then u is privately viable but socially harmful exactly when

$$\frac{\phi T}{2} \left[\operatorname{Var}_\tau(\gamma\kappa) + \operatorname{Var}_\tau(\mu) - 2 \operatorname{Cov}_\tau(\mu, \gamma\kappa) \right] > C + \Psi$$

and

$$\frac{T}{2}\operatorname{Var}_\tau(\gamma\kappa) < C + E.$$

Equivalently, the disagreement-and-alignment term must satisfy

$$\operatorname{Var}_\tau(\mu) - 2\operatorname{Cov}_\tau(\mu,\gamma\kappa) > \frac{2(C+\Psi)}{\phi T} - \operatorname{Var}_\tau(\gamma\kappa),$$

while the hedging-dispersion term remains below the social cost threshold.

Proof. Agent a's local certainty-equivalent gain from choosing quantity q_a at price p is, up to constants independent of q_a,

$$q_a(\mu_a - p) - \frac{\gamma_a}{2}\operatorname{Var}(W_a + q_a u).$$

Using the common second moment $\sigma_u^2 = \operatorname{Var}(u)$ and $\kappa_a = \operatorname{Cov}(W_a, u)$, the first-order condition is

$$\mu_a - p - \gamma_a\kappa_a - \gamma_a q_a \sigma_u^2 = 0.$$

Hence

$$q_a(p) = \tau_a(b_a - p), \qquad \tau_a = \frac{1}{\gamma_a\sigma_u^2}.$$

Zero net supply gives $\sum_a q_a(p^*) = 0$, so $p^* = \sum_a \tau_a b_a/T$. Standard quadratic surplus aggregation gives

$$S_u^{perceived} = \frac{1}{2}\sum_a \tau_a(b_a - p^*)^2 = \frac{1}{2}T\operatorname{Var}_\tau(b).$$

Since $b = \mu - \gamma\kappa$,

$$\operatorname{Var}_\tau(b) = \operatorname{Var}_\tau(\mu - \gamma\kappa) = \operatorname{Var}_\tau(\gamma\kappa) + \operatorname{Var}_\tau(\mu) - 2\operatorname{Cov}_\tau(\mu,\gamma\kappa).$$

If beliefs are common, $\mu_a \equiv \bar{\mu}$, so $\operatorname{Var}_\tau(\mu) = \operatorname{Cov}_\tau(\mu,\gamma\kappa) = 0$ and perceived surplus collapses to $\frac{1}{2}T\operatorname{Var}_\tau(\gamma\kappa)$, the risk-sharing component. The private and social inequalities are then the definitions of private viability and social harm under the maintained belief-neutral evaluation. □

Remark 7.2 (Relation to Simsek and belief-neutral welfare)**.** *Simsek (2013) decomposes the effects of financial innovation under belief disagreement into risk-sharing and speculative components. Brunnermeier et al. (2014) supply the welfare discipline used here: pure gains from mutually inconsistent distorted beliefs are not automatically counted as social surplus. The theorem makes the local diagnostic observable. The relevant axes are belief dispersion* $\operatorname{Var}_\tau(\mu)$*, hedging dispersion* $\operatorname{Var}_\tau(\gamma\kappa)$*, and their covariance. If optimistic beliefs are concentrated among natural buyers of the hedge, then* $\operatorname{Cov}_\tau(\mu,\gamma\kappa) < 0$ *and disagreement reinforces hedging demand. If optimistic beliefs sit with agents whose background exposure makes them natural sellers, then the cross term reduces perceived surplus. The old additive notation* $Q^{tot} = Q^R + Q^S$ *is therefore only a special case with a zero alignment term.*

Proposition 7.3 (Pure-disagreement boundary)**.** *Suppose a zero-net-supply claim has no hedging value, produces no decision-relevant public information, creates no legitimate consumption value, and is traded only because agents disagree about its payoff. Under a belief-neutral social evaluation in the sense of Brunnermeier et al. (2014), which does not count mutually inconsistent subjective gains as aggregate welfare by default, the claim's social value is nonpositive after implementation, liquidity, and externality costs. If trading increases portfolio risk or creates externalities, its social value is strictly negative.*

Proof. Zero-net-supply payoffs are transfers across agents. With no hedging value, no information value, and no consumption value, aggregate social benefit from the transfer is zero under the maintained social evaluation. Implementation and liquidity costs are real resource costs, and externalities are weakly positive costs. Therefore social net value is weakly negative. If the claim induces additional risk-bearing or externality, the inequality is strict. □

The proposition is narrow by construction. It does not condemn all speculation. It identifies the case in which speculation has no social channel beyond inconsistent private beliefs.

7.2 Information value

A price can be a public good. Let a decision maker choose action $d \in \mathcal{D}$ before state ω is realized and suffer loss $L(d, \omega)$. Without the claim price, the decision maker chooses d_0. With price signal p_ρ, the decision maker chooses $d(p_\rho)$. Define decision value

$$S_\rho^{info} = \mathbb{E}\left[L(d_0, \omega) - L(d(p_\rho), \omega)\right].$$

This value can accrue to many actors who do not pay the market creator.

Proposition 7.4 (Public-good information subsidy)**.** *Suppose a claim's only social benefit is information value $S^{info} > 0$, implementation and subsidy cost is C, and the creator captures fraction $\phi < 1$ of the information value. If*

$$\phi S^{info} < C < S^{info},$$

then the information market is privately nonviable but socially valuable. A subsidy, procurement contract, data license, or public/platform provision can implement the market whenever it transfers enough of the nonrival decision value to cover the creator's deficit without inducing larger admissibility costs.

Proof. Private net value is $\phi S^{info} - C < 0$. Social net value is $S^{info} - C > 0$. □

Remark 7.5 (Subsidy is not a failure)**.** *Some prediction markets require subsidy not because they are fake markets but because their output is a public statistic. Paper 4 owns the infrastructure-side public-good price result. Paper 3 uses the same inequality as a policy test: subsidized information liquidity is justified when decision value exceeds subsidy and externality costs.*

7.3 Admissibility of belief-driven markets

Definition 7.6 (Belief-market admissibility)**.** *A belief-driven claim is admissible only if at least one of the following channels is positive and dominates costs: (i) hedging or risk-sharing value; (ii) decision-relevant information value; (iii) legitimate non-exploitative consumption value; or (iv) an institutional purpose such as governance or resource allocation. Pure disagreement volume is not by itself a social benefit.*

Proposition 7.7 (Speculation tax or access rule)**.** *If a class of claims has high private speculative surplus, low information value, and externalities proportional to volume or leverage, then a volume tax, leverage margin, access restriction, or position limit can improve welfare by reducing the speculative component while preserving hedging or information flow when those flows have lower elasticity or can be separately identified.*

Proof. Let welfare be $W(q_H, q_I, q_S) = B_H(q_H) + B_I(q_I) - E_S(q_S) - K(q_H, q_I, q_S)$, where q_S is speculative volume. If an instrument reduces q_S with smaller effect on q_H and q_I, and $\partial E_S/\partial q_S$ exceeds the marginal social benefit of the reduced speculative flow, welfare rises locally. □

8 Information Timing: The Hirshleifer Boundary

Better information can create markets or destroy them. It creates markets when it makes state distinctions contractible at specification or settlement. It can destroy risk-sharing markets when it reveals risk before agents can pool or trade.

Let $\mathcal{F}_T$ be the sigma-algebra of facts verifiable under technology T. Paper 1 shows that a refinement $\mathcal{F}_T \subseteq \mathcal{F}_{T'}$ can expand the contractible payoff space. But a signal can arrive at different times: before contract formation, during the life of a contract, at settlement, or after price publication.

Paper 4 gives the infrastructure theorem for this timing duality using the law of total variance and the oracle–settlement stack. Paper 3 uses the same logic only as an admissibility input: a claim review must ask not just how accurate a data system is, but when it is observed, by whom, and for what institutional purpose.

Definition 8.1 (Information role)**.** *A data process Z has role:*

- pre-contract classification *if it is observed before agents choose whether and on what terms to contract;*
- ex ante specification *if it lets parties write a better payoff rule before the state is known;*
- interim public revelation *if it reveals the state before risk-sharing trades are made;*
- ex post verification *if it resolves the payoff after the contract is already in force;*
- post-market publication *if it produces a public decision-relevant price.*

Proposition 8.2 (Information-timing admissibility)**.** *Let $\mathcal{F}$ be a public signal sigma-algebra and let $x \in L^2(\mathbb{P})$ be a risky payoff. If agents can contract ex ante on $\mathcal{F}$-measurable settlement before $\mathcal{F}$ is realized, the refinement can expand the attainable payoff space and create risk-sharing value. If instead $\mathcal{F}$ is publicly revealed before trade, then trade can insure only residual uncertainty within the realized atom of $\mathcal{F}$; cross-atom insurance value is lost. If x is $\mathcal{F}$-measurable and $\mathcal{F}$ is revealed before trade, post-revelation trade in x has no insurance value for the revealed variation.*

Proof. Before $\mathcal{F}$ is realized, a contract measurable with respect to $\mathcal{F}$ transfers resources across future signal atoms. After $\mathcal{F}$ is publicly realized, all agents condition on the realized atom and can trade only claims on remaining within-atom uncertainty. If x is fully $\mathcal{F}$-measurable, then conditional on the realized atom x is known. A known payoff can be priced as a riskless transfer and cannot insure the variation across atoms that has already been revealed. □

Corollary 8.3 (AI prediction is not monotonically market-creating)**.** *An AI system that improves ex post verification can increase admissible markets by reducing settlement and dispute costs. The same system, if used for pre-contract classification or public pre-trade revelation, can fragment pools and reduce insurance value. Admissibility must specify the timing and use of the information, not merely its accuracy.*

Example 8.4 (Fire risk)**.** *A building-level fire model used after contract formation to verify parametric triggers can reduce adjustment cost. The same model used before contract formation to predict each building's loss with near certainty can eliminate pooling: the premium converges to the known loss and the insurance value disappears.*

Proposition 8.5 (Disclosure tradeoff)**.** *Suppose disclosure reduces adverse selection by $A(d)$ but reduces pooling value by $P(d)$, where d indexes disclosure intensity. The optimal disclosure rule solves*

$$\max_d \; A(d) - P(d) - C(d) - E(d),$$

subject to privacy and rights constraints. Full disclosure is optimal only if its marginal reduction in adverse selection and dispute cost exceeds its marginal pooling, privacy, and externality costs.

Proof. The objective is the social net value of disclosure intensity. The first-order comparison follows from marginal benefits and costs. Hard privacy or rights constraints restrict the feasible set of d. □

9 The Human-Capital and Personhood Boundary

Cheap claim creation makes it tempting to trade claims on individuals: future income, creator royalties, athlete prize streams, medical outcomes, attention, reputation, romantic choices, political choices, or life decisions. Some of these claims can be welfare-improving if they finance human capital, diversify income risk, or monetize specified cash flows without control. Others cross a boundary.

The boundary is not "never contract with a person." Employment, insurance, royalties, loans, equity-like revenue shares, and income-contingent repayment all involve persons. The boundary is: *claim, not person.* A permissible claim can attach to specified cash flows under caps, duration limits, hardship protections, and privacy-preserving verification. It must not give holders control over personhood, intimate choices, political rights, medical decisions, reproductive choices, religious practice, romantic association, coercive labor, or basic mobility.

Definition 9.1 (Protected personal domain)**.** *Let $\mathcal{D}^{prot}$ be a set of protected personal decisions: medical, reproductive, romantic, religious, political, intimate, bodily, family, and other decisions treated as non-alienable for purposes of market design. A representation ρ violates the no-control boundary if holders of the claim receive direct or indirect control rights over decisions in $\mathcal{D}^{prot}$ or if payoff incentives are designed to coerce those decisions.*

Principle 9.2 (No-control human-capital boundary)**.** *If a representation ρ gives claim holders enforceable control rights over protected personal decisions, then $\rho \notin \mathcal{R}^{hard}(\Xi)$ under a rights-respecting admissibility frontier. No positive financial surplus can make such a representation admissible while the corresponding hard constraint has infinite shadow cost.*

Proof. By definition, hard admissibility constraints enter the extended externality $\bar{E}_\rho$ with infinite penalty when violated. A representation that gives control over protected personal decisions violates such a constraint. Therefore its net admissibility value is $-\infty$ regardless of finite financial surplus. □

Remark 9.3 (Why not just price it?)**.** *Some harms are not safely treated as ordinary externalities. If a market lets investors buy control over whom someone dates, what medical treatment they accept, how they vote, or whether they carry a pregnancy, the problem is not that the price is too low. The problem is that the object is outside the admissible commodity space.*

Definition 9.4 (Permissible personal cash-flow claim)**.** *A personal cash-flow claim is presumptively admissible only if it satisfies: (i) no control over protected decisions; (ii) informed consent; (iii) cap on share of income or cash flow; (iv) cap on duration; (v) hardship, bankruptcy, and exit rules; (vi) privacy-preserving verification; (vii) special protection or prohibition for minors and vulnerable persons; (viii) no assignment of intimate data beyond what is necessary for settlement; and (ix) no hidden leverage or compounding penalty that creates coercive labor incentives.*

Proposition 9.5 (Consent is necessary but not sufficient)**.** *Suppose participation in a personal claim is increasing in vulnerability v and the claim imposes uninternalized expected harm $E(v)$ on the participant or third parties. If the seller's outside option is sufficiently poor, voluntary acceptance can occur even when social welfare is negative. Therefore consent alone cannot be a sufficient admissibility test for personal claims.*

Proof. Let the participant accept when price p exceeds private reservation value $r(v)$, where $r(v)$ falls as vulnerability worsens. Social value is $p - r(v) - E(v)$ plus any external benefits. If $E(v) > p - r(v)$ for vulnerable participants, acceptance occurs while social value is negative. Thus observed consent does not imply admissibility. □

10 Oracle, Manipulation, Moral Hazard, and Reflexivity

A state-contingent claim is dangerous when traders can profitably change the state, corrupt the oracle, or alter behavior because the claim exists. This is not merely fraud. It is endogenous payoff formation.

10.1 Oracle manipulation

Let $y(\omega)$ be the true underlying event and $\hat{y} = \mathcal{O}_\rho(\omega, a_o)$ be the oracle output, where a_o is manipulation or corruption effort. The delivered payoff is $g(\hat{y})$. A trader with position q can gain

$$G_o(a_o; q) = q\,[g(\mathcal{O}_\rho(\omega, a_o)) - g(\mathcal{O}_\rho(\omega, 0))]$$

from oracle manipulation, at private cost $c_o(a_o)$ and expected penalty $\pi_o(a_o)F_o$.

Proposition 10.1 (Oracle incentive-compatibility condition)**.** *A sufficient condition for oracle manipulation to be privately unprofitable for positions in admissible set $\mathcal{Q}_\rho$ is*

$$\sup_{q \in \mathcal{Q}_\rho} \sup_{a_o \neq 0} \{\mathbb{E}[G_o(a_o; q)] - c_o(a_o) - \pi_o(a_o)F_o\} \leq 0.$$

If this condition fails, admissibility requires position limits, oracle redesign, delayed settlement, dispute procedures, penalties, or prohibition.

Proof. The expression is the maximum expected private net gain from oracle manipulation over admissible positions and manipulation actions. If it is nonpositive, no such manipulation is privately profitable. If positive, some trader-position pair has an incentive to manipulate unless design changes remove the incentive. □

10.2 Underlying-event manipulation

Now let the true event itself depend on trader action a_y: $y = y(\omega, a_y)$. The social harm of event manipulation is $H_y(a_y)$. The private trading gain is $G_y(a_y; q)$.

Proposition 10.2 (Underlying manipulation condition)**.** *If*

$$\sup_{q \in \mathcal{Q}_\rho} \sup_{a_y} \{\mathbb{E}[G_y(a_y; q)] - c_y(a_y) - \pi_y(a_y)F_y\} > 0,$$

then the claim creates a manipulation incentive. Its social externality includes

$$E_\rho^{manip} \geq \mathbb{E}[H_y(a_y^*)],$$

where a_y^* is the induced manipulation action. A claim whose dominant traders can cheaply alter a harmful underlying event is inadmissible unless position limits, access restrictions, monitoring, penalties, or payoff redesign make the manipulation constraint bind at zero harmful action.

Proof. If the supremum is positive, some action maximizes private net trading gains away from zero. The resulting action changes the underlying event and creates social harm $H_y(a_y^*)$ not internalized by the trader unless penalties or design fully internalize it. That harm enters externality cost. □

Example 10.3 (Thin event markets)**.** *A market on whether a small firm misses a financing milestone can be informative if traders passively aggregate beliefs. It can be manipulative if a trader can influence the milestone by threatening counterparties, spreading false information, or withholding financing while holding a position. The payoff rule is not enough; admissibility depends on who can trade and what they can affect.*

10.3 Moral hazard and adverse selection

Insurance and hedging claims can change behavior. A claim that pays on a loss may reduce prevention effort. A claim that lets a borrower sell revenue risk may change project choice. A market that reveals a price may affect the event being predicted.

Definition 10.4 (Behavioral response operator)**.** *Let a denote actions that affect the payoff distribution. Without the claim, agents choose a^0. With representation ρ, positions and prices induce action a^ρ. The behavioral externality is*

$$E_\rho^{beh} = \mathbb{E}\left[H(a^\rho,\omega) - H(a^0,\omega)\right] - \textit{internalized payments},$$

where H is social harm or resource cost.

Proposition 10.5 (Completion effect versus behavioral effect)**.** *Let $W(\mathcal{H},\mathsf{Q}) = \frac{1}{2}\operatorname{tr}(\Pi_{\mathcal{H}}\mathsf{Q})$ be the local risk-sharing value from Paper 1. If a new market changes the span from $\mathcal{H}$ to $\mathcal{H}'$ and changes behavior so that the valuation-dispersion operator changes from Q to Q', then*

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

The first term is the completion effect. The second is the behavioral effect. Admissibility requires the completion effect plus other benefits to dominate behavioral and externality costs.

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

10.4 Reflexive claims

A claim is reflexive when its price or existence changes the underlying state. Credit spreads can affect refinancing, governance markets can affect decisions, prediction markets can change behavior, and public ranking claims can affect reputation.

Definition 10.6 (Reflexive stability)**.** *Let the state distribution induced by price p and claim existence be $\mathbb{P}_{\rho,p}$. A reflexive claim has a stable admissible equilibrium if there exists a price p^* such that (i) p^* prices g_ρ under $\mathbb{P}_{\rho,p^*}$ and the permitted trading protocol; (ii) induced actions satisfy manipulation, rights, and moral-hazard constraints; and (iii) the resulting social value is nonnegative.*

Proposition 10.7 (Reflexivity is an admissibility burden)**.** *A non-reflexive claim requires a payoff rule and a settlement oracle. A reflexive claim additionally requires existence and stability of the price-to-action-to-state-to-payoff fixed point. Failure of this fixed point, or existence only through harmful manipulation, makes the claim inadmissible or requires restricted access.*

Proof. If the claim's price affects the state, payoff distribution cannot be evaluated independently of the market. Admissibility must therefore be assessed at a fixed point of induced behavior and pricing. If no acceptable fixed point exists, the representation has no stable welfare object satisfying the constraints. □

Proposition 10.8 (Contraction sufficient condition for reflexive stability)**.** *Let $p \in P$ be a price in a complete metric space and let $\Phi_\rho : P \to P$ be the induced price-to-action-to-state-to-price map for a reflexive representation ρ, after imposing access, position, oracle, and manipulation constraints. If Φ_ρ is a contraction with modulus $L < 1$, then ρ has a unique fixed point p^*. If, in addition, the induced actions at p^* satisfy hard constraints and $W_\rho(p^*) \geq 0$, then the reflexive fixed-point burden is satisfied.*

Proof. Banach's fixed-point theorem gives existence and uniqueness of p^*. The remaining two conditions are the admissibility requirements evaluated at the fixed point. The proposition is only sufficient: non-contractive reflexive markets may still have acceptable equilibria, but they require separate existence, selection, and stability analysis. □

11 Systemic Admissibility

A claim can be harmless in isolation and dangerous in a system. The danger comes from leverage, crowded hedges, cross-margining, collateral reuse, common oracles, common market makers, and forced liquidation. Paper 2 shows that synthetic liquidity can consolidate risk at the level of factor hedges and residual warehouses. This section treats that consolidation as an admissibility input.

11.1 Network state

Let a market set M generate a network state $N(M)$. The network includes:

- factor hedge exposures β_ρ;
- residual covariance and tail dependence Σ_ε;
- leverage and margin rules ℓ_ρ;
- collateral reuse and rehypothecation links;
- shared market makers and residual warehouses;
- common oracle dependencies;
- cross-margin and liquidation rules.

Let $E^{sys}(M)$ be the systemic externality of market set M. It is generally non-additive. The marginal systemic cost of adding ρ is

$$\Delta E^{sys}_\rho(M) = E^{sys}(M \cup \{\rho\}) - E^{sys}(M).$$

Definition 11.1 (Systemic admissibility)**.** *A claim ρ is systemically admissible relative to market set M if*

$$\Delta W_\rho(M) = \mathcal{W}(M \cup \{\rho\}) - \mathcal{W}(M) \geq 0$$

when $\Delta E^{sys}_\rho(M)$ is included, and if all hard systemic constraints, such as leverage caps or collateral segregation rules, are satisfied.

11.2 Cascade threshold

A simple reduced form captures non-additivity. Let $R(M)$ be a matrix of cross-market amplification: entry R_{ij} is the fraction of a stress loss in market j that forces liquidation, margin calls, or loss propagation into market i. If the spectral radius $r(R(M)) < 1$, shocks decay in the linear approximation. If $r(R(M)) \geq 1$, shocks can amplify.

Theorem 11.2 (Non-additive systemic risk). *Suppose systemic loss from a unit shock is proportional to*

$$\left\|(I - R(M))^{-1} s\right\|$$

when $r(R(M)) < 1$. Consider a sequence of nonnegative irreducible amplification matrices R_n with $r(R_n) \uparrow 1$. If the stress vector s has nonzero projection onto the Perron direction along the sequence, then amplified loss becomes unbounded or discontinuously large. Consequently, a claim ρ with small standalone risk can have arbitrarily large marginal systemic cost if adding it moves $r(R(M \cup \{\rho\}))$ close to or above one. Therefore standalone claim safety is not sufficient for systemic admissibility.

Proof. For $r(R) < 1$, the Neumann series gives $(I - R)^{-1} = \sum_{n \geq 0} R^n$. For nonnegative irreducible matrices, Perron–Frobenius theory identifies a positive leading eigenvector. When the shock has nonzero overlap with that direction, the component of $(I - R)^{-1}s$ in the leading eigenspace scales like $(1 - r(R))^{-1}$. A claim that changes R only slightly in entries can nevertheless move the spectral radius close to the threshold. The resulting marginal change in amplified loss can be arbitrarily large relative to the claim's standalone loss. □

Corollary 11.3 (Cross-margining and shared hedges). *A claim that uses the same collateral pool, liquidation rule, or crowded hedge as many existing claims must be evaluated by its marginal effect on the network, not by its own payoff variance alone.*

11.3 Residual warehouses

Proposition 11.4 (Residual warehouse externality). *Suppose n claims are quoted by the same residual-risk warehouse. Each claim is individually quote-feasible under normal covariance, but stress covariance increases joint residual loss and forces warehouse liquidation across all n claims. If individual creators do not pay for the increased probability of warehouse failure, each claim's private entry condition omits a positive systemic externality. A capital charge based only on standalone variance underprices the marginal claim.*

Proof. The marginal claim increases the warehouse's stress loss distribution and the probability or severity of joint liquidation. If creators pay only standalone normal-time quote costs, they do not internalize the increased tail dependence and liquidation spillover imposed on other markets. The omitted term is a systemic externality. □

Remark 11.5 (Normal-time liquidity is not robustness). *A market can be quote-feasible at normal covariance and inadmissible at system scale if its liquidity relies on a common hedge or residual warehouse that fails in stress. Admissibility requires stress transparency, capital, collateral segregation, and resolution design.*

12 Liquidity-Source Admissibility

Paper 2 identifies three liquidity mechanisms: factor-hedged dealer liquidity, bookmaker or insurer residual-diversification liquidity, and subsidized information-market liquidity. Paper 3 asks whether the source of liquidity is socially acceptable.

Liquidity source	Technical logic	Welfare-positive use	Admissibility risk
Factor-hedged dealer	common risk hedged in deep markets; residual warehoused	hedging, financing, useful price discovery	crowded hedges, tail concentration, hidden leverage
Bookmaker / insurer	broad independent non-toxic flow diversifies	insurance pools; legitimate entertainment	addiction, exploitation, sharp/insider flow, pool collapse
Subsidized information	sponsor pays bounded loss for public price	decision-relevant public information	manipulation, biased sponsor, ambiguous resolution
Extractive liquidity	volume from bias, coercion, opacity, dark patterns	none or limited private entertainment value	predation, addiction, privacy loss, wealth extraction

Proposition 12.1 (Same liquidity, different admissibility)**.** *Two claims can have identical quoted spreads and depth but different admissibility because their liquidity is supported by different flow ecologies or different systemic backstops. A liquidity test is therefore necessary for feasibility but not sufficient for admissibility.*

Proof. Quoted spread and depth summarize execution terms, not the welfare source of order flow or the externality of the balance sheet supporting the quote. Holding spread and depth fixed, changing flow from hedging to addiction changes B^{soc} and E. Changing the backstop from diversified capital to a fragile residual warehouse changes E^{sys}. Thus admissibility can differ. □

Definition 12.2 (Liquidity transparency requirement)**.** *A claim satisfies liquidity transparency if participants and supervisors can identify the dominant liquidity mechanism, material hedge dependencies, residual-risk warehouses, subsidy budgets, flow-toxicity controls, and conditions under which quotes may be withdrawn.*

Proposition 12.3 (Liquidity transparency as a constraint)**.** *If a claim's social value depends on continued liquidity or on the absence of hidden leverage, then opacity about the liquidity source increases E_ρ and can make an otherwise valuable claim inadmissible. Disclosure, stress testing, designated market-maker obligations, and collateral segregation reduce this externality.*

Proof. Opacity causes participants and regulators to underestimate execution risk, liquidation risk, and systemic dependence. This misallocation creates expected harm in stress. Information and design requirements reduce the probability or magnitude of that harm, lowering E_ρ. □

13 Admissibility as Constrained Market Design

The planner does not merely choose whether a claim exists. The planner can choose a representation, access rule, margin rule, oracle, subsidy, disclosure, tax, or prohibition.

13.1 The constrained design problem

Let $M \subseteq \mathcal{R}$ be a set of claim representations. Aggregate welfare is

$$\mathcal{W}(M;\Xi) = B^{soc}(M;\Xi) - K^{soc}(M;\Xi) - E(M;\Xi),$$

where each term can be non-additive. The constrained social design problem is

$$\max_{M \subseteq \mathcal{R}^{tech}(T) \cap \mathcal{R}_j^{leg}} \mathcal{W}(M; \Xi) \quad \text{s.t.} \quad M \subseteq \mathcal{R}^{hard}(\Xi). \tag{5}$$

The marginal admissibility of a claim depends on the existing set:

$$\Delta \mathcal{W}_\rho(M) = \mathcal{W}(M \cup \{\rho\}) - \mathcal{W}(M).$$

A claim is marginally admissible relative to M if $\Delta \mathcal{W}_\rho(M) \geq 0$ and hard constraints remain satisfied.

Proposition 13.1 (Marginal, not standalone, admissibility)**.** *If benefits or externalities are non-additive, standalone social value $W_\rho(\emptyset)$ does not determine admissibility relative to an existing market set M. The correct object is $\Delta \mathcal{W}_\rho(M)$.*

Proof. By definition, non-additivity means $\mathcal{W}(M \cup \{\rho\}) - \mathcal{W}(M)$ need not equal $\mathcal{W}(\{\rho\}) - \mathcal{W}(\emptyset)$. Thus standalone value can differ from marginal value. □

13.2 Policy instruments as parameter changes

Policy actions change private entry, social welfare, or both:

- *Standardization* lowers κ and dispute cost for socially valuable classes.
- *Public data and oracle infrastructure* lower verification cost and adverse selection.
- *Subsidy or public provision* raises private viability for public-good prices or diffuse hedging value.
- *Taxes and fees* reduce entry of privately viable harmful claims.
- *Margin, leverage limits, and position limits* reduce systemic, manipulation, and addiction externalities.
- *Access restrictions* separate hedging and professional information uses from retail exploitation.
- *Suitability and fiduciary duties* reduce extraction through agents and platforms.
- *Privacy-preserving verification* lowers privacy cost without destroying settlement.
- *Oracle governance and dispute rules* reduce manipulation and ambiguity.
- *Hard prohibition* removes claims violating non-priceable constraints.

Proposition 13.2 (Pigouvian alignment in the special case)**.** *Suppose private benefit equals social benefit, private cost equals social resource cost, there are no hard constraints, and externality E_ρ is observable and additive. A tax $t_\rho = E_\rho$ aligns private entry with social desirability.*

Proof. With the tax, private net value is $B_\rho - K_\rho - t_\rho = B_\rho - K_\rho - E_\rho = W_\rho$. Hence private entry occurs exactly when social net value is positive. □

Remark 13.3 (Why Pigou is not enough)**.** *The proposition requires strong assumptions. When prices are public goods, $\phi < 1$ and subsidy may be needed. When benefits are private but harms are hard constraints, taxes are insufficient. When systemic risk is non-additive, the tax depends on the market set. When participants are vulnerable or present-biased, disclosure may not cure extraction.*

Proposition 13.4 (Subsidy for public-good prices)**.** *If a claim has social information value S^{info}, creator capture ϕS^{info}, and no offsetting externalities, then a subsidy*

$$s_\rho \in [C - \phi S^{info},\ S^{info} - C]$$

can make private creation viable without making a socially negative market viable, provided the interval is nonempty.

Proof. Private viability with subsidy requires $\phi S^{info} + s_\rho - C \geq 0$, so $s_\rho \geq C - \phi S^{info}$. Social net value including subsidy as a transfer-financed resource cost is nonnegative if the total cost does not exceed S^{info}; in the simplest case this requires $s_\rho \leq S^{info} - C$ after accounting for financing costs. The stated interval implements both. □

13.3 The admissibility checklist

A practical claim review should answer the following questions in order:

A1. Define the representation. What is the payoff rule, oracle, data source, access set, margin rule, liquidity mechanism, governance, and purpose?

A2. Check technical feasibility. Can the claim be specified, verified, settled, and quoted under stress-relevant assumptions?

A3. Check legal permissibility. What legal wrapper, license, exemption, or prohibition applies?

A4. Identify the flow ecology. What share of expected flow is hedging, immediacy, information, consumption, disagreement, manipulation, or extraction?

A5. Measure benefits. Estimate hedging value, immediacy value, information value, legitimate consumption value, infrastructure spillovers, and distributional weights.

A6. Measure costs and externalities. Estimate implementation, liquidity, privacy, coercion, manipulation, moral hazard, addiction, oracle, dignity, norm, and systemic costs.

A7. Apply hard constraints. Does the claim touch protected personal domains, political rights, minors, nonconsensual data, coercive labor, or physical harm manipulation?

A8. Stress the system. How does the claim affect shared hedges, residual warehouses, collateral reuse, cross-margining, and liquidation cascades?

A9. Choose the least restrictive welfare-improving policy. Allow, standardize, subsidize, tax, margin, restrict access, redesign oracle, impose fiduciary duties, or prohibit.

14 Claim Categories and Boundary Rules

The framework is abstract. This section applies it to recurring categories.

14.1 Hedging and insurance claims

A claim designed to transfer a pre-existing risk from a high-cost bearer to a lower-cost bearer is presumptively valuable when pricing is fair, adverse selection is controlled, and moral hazard is limited. The main risks are pool destruction from excessive classification, privacy invasion, moral hazard, and capital fragility.

Proposition 14.1 (Insurance admissibility)**.** *An insurance-like claim is admissible when*

$$S^{hedge} + S^{imm} + S^{info}_{settlement} > K + E^{privacy} + E^{moral} + E^{selection} + E^{sys},$$

and hard constraints are satisfied. Classification data should be allowed only to the extent that its reduction in adverse selection, moral hazard, and settlement cost exceeds its pooling, privacy, and exclusion costs.

14.2 Prediction and event claims

An event claim can be a public forecast, a hedge, entertainment, or manipulation. The key tests are information value, insider access, manipulation of the event, oracle clarity, and whether trading affects the underlying decision.

Proposition 14.2 (Prediction-market admissibility)**.** *A prediction or event claim is presumptively admissible when its expected decision value plus legitimate hedging or consumption value exceeds implementation, subsidy, manipulation, insider-trading, oracle, and externality costs. It is presumptively inadmissible when the dominant private value comes from manipulating or corrupting the event, exploiting nonpublic control, or generating addictive zero-sum flow with no offsetting information value.*

14.3 Revenue shares and creator/athlete finance

Revenue shares can finance human capital and diversify risk. They can also become disguised control over a person. The line is cash-flow claim versus personhood claim.

Proposition 14.3 (Revenue-share admissibility)**.** *A revenue-share claim on a specified project, catalog, firm, or prize stream is more admissible when it is capped, time-limited, non-controlling, bankruptcy-aware, privacy-preserving, and tied to verifiable cash flows. It becomes less admissible as it expands toward uncapped lifetime income, invasive monitoring, transfer restrictions that bind labor mobility, or investor control over personal choices.*

14.4 Micro-markets and agentic transactions

Agentic micro-markets can allocate compute, bandwidth, energy flexibility, delivery windows, charging priority, reservations, or congestion rights. Their risks are privacy leakage, stratification, norm erosion, attention capture, and platform extraction.

Proposition 14.4 (Agentic market fiduciary condition)**.** *If a personal agent negotiates claims on behalf of a user, admissibility requires the agent to respect user-defined hard constraints and fiduciary duties. A platform-controlled agent that optimizes venue revenue while controlling user transaction flow increases $E^{extract}$ and can make otherwise efficient micro-markets inadmissible.*

Proof. The user's welfare depends on preferences and hard constraints that may conflict with platform revenue. If the agent optimizes platform revenue, it can steer the user into trades with positive platform capture but negative user surplus or rights violations. Fiduciary constraints reduce this agency externality. □

15 Empirical Predictions and Test Plan

The empirical object is a claim representation ρ in market state Ξ_t, not merely a venue or product label. The theory predicts both market births and regulatory interventions.

15.1 Market birth predictions

Privately created claims should appear first where captured private value is high relative to implementation, liquidity, and private compliance costs:

$$P_{\rho t} = \phi_{\rho t} B^{priv}_{\rho t} - K^{priv}_{\rho t} - \Psi^{priv}_{\rho t} > 0.$$

Socially valuable but privately undersupplied claims should have high estimated $W_{\rho t}$ but low $P_{\rho t}$ because of public-good prices, diffuse hedging value, or infrastructure fixed costs. Socially harmful private markets should have high $P_{\rho t}$ and low or negative $W_{\rho t}$ due to extraction, addiction, manipulation, or systemic externality.

A reduced-form hazard for private birth is

$$\Pr(Y_{\rho t}^{birth} = 1) = \Lambda\left(\beta_1 \widehat{B}_{\rho t}^{priv} - \beta_2 \widehat{K}_{\rho t}^{priv} - \beta_3 \widehat{\Psi}_{\rho t}^{priv} + \beta_4 \widehat{\phi}_{\rho t} + \mu_{d(\rho)} + \tau_t\right).$$

A social-desirability measure should instead use

$$\widehat{W}_{\rho t} = \widehat{B}_{\rho t}^{soc} - \widehat{K}_{\rho t}^{soc} - \widehat{E}_{\rho t}.$$

The distinctive prediction is divergence: high private-birth probability need not imply high social desirability.

15.2 Regulatory intervention predictions

Regulatory restrictions should be more likely when claims score high on manipulation, retail vulnerability, hidden leverage, privacy invasion, or systemic connectivity:

$$\Pr(Y_{\rho t}^{restrict} = 1) = \Lambda\left(\gamma_1 \widehat{E}_{\rho t}^{manip} + \gamma_2 \widehat{E}_{\rho t}^{privacy} + \gamma_3 \widehat{E}_{\rho t}^{addict} + \gamma_4 \widehat{E}_{\rho t}^{sys} - \gamma_5 \widehat{S}_{\rho t}^{hedge/info} + \eta_j + \tau_t\right).$$

Subsidy or public provision should be more likely when information value or hedging value is high but capture is low.

15.3 Empirical proxies

Theoretical object	Empirical proxy
Residual hedging value	projected exposure not spanned by existing markets; demand from natural hedgers; variance reduction
Immediacy value	willingness to pay for execution speed, financing, or balance-sheet service
Information value	forecast use by decision makers; price impact on procurement, policy, insurance, or investment decisions
Consumption value	entertainment demand net of problem-use indicators and complaint rates
Speculative/disagreement flow	high turnover unrelated to hedging need; leverage; disagreement proxies; post-trade risk increase
Manipulation risk	trader ability to affect event or oracle; insider/control relationships; abnormal event-linked activity
Privacy cost	sensitivity of data; identifiability; secondary use; breach harm; consent quality
Coercion/vulnerability	participant distress, income, age, bargaining power, outside option, concentration of marketing
Systemic risk	shared hedges, collateral reuse, margin links, market-maker concentration, stress correlations
Oracle fragility	dispute frequency, data latency, governance concentration, override history
Capture parameter	fee revenue, spreads, data monetization, exclusive distribution, benchmark IP, public-good leakage

15.4 Policy surfaces

Promising empirical and institutional surfaces include:

1. prediction-market insider and control-person rules;
2. event-contract prohibited categories and purpose-based access;
3. parametric insurance and classification-data regulation;
4. tokenized securities and real-world-asset claims;
5. income-share agreements and human-capital financing;
6. athlete, creator, and royalty finance;
7. social-token and personal-token failures;
8. DeFi oracle manipulation and governance attacks;
9. maximum extractable value, latency predation, and batch auctions;

10. retail suitability, leverage, and disclosure rules;
11. market-maker concentration and residual warehouse stress tests;
12. public forecasting markets for procurement, disaster response, and policy.

15.5 Identification strategies

Cost shocks. Use verification, legal-template, oracle, settlement, or regulatory clarity shocks to observe which claims enter. Compare private entry with estimated social value.

Access-rule changes. Study outcomes when claims move from professional-only to retail access, or when leverage and margin rules change. Measure flow mix, complaint rates, post-trade losses, and hedging use.

Oracle events. Analyze disputes, oracle failures, manipulation attacks, and resolution changes. Estimate how oracle governance affects spreads, participation, and harm.

Systemic stress. Identify claims linked by common hedges, collateral, or market makers. Test whether stress widening and quote withdrawal propagate through those links beyond payoff fundamentals.

Information public goods. Measure whether published probabilities are used by nontrading decision makers. Estimate under-entry by comparing social decision value to fee capture.

16 Relation to the Broader Research Program

Paper 3 is the permission layer. It should inherit the program's vocabulary without blurring it:

- **Paper 1 value object.** $V(B \mid \mathcal{H})$ is the valuation-dispersion value of adding residual payoff span relative to existing markets. Paper 3 treats this as S^{hedge}, one component of social benefit.
- **Paper 1 cost object.** $\kappa(B, \mathcal{H}, T)$ is representation cost. Paper 3 asks whether the cheapest representation is admissible and whether a more expensive representation is required to satisfy constraints.
- **Paper 2 liquidity object.** The quote-feasibility stack determines technical liquidity. Paper 3 asks whether the liquidity source is welfare-positive, extractive, or systemically fragile.
- **Paper 4 infrastructure object.** Prediction, insurance, derivatives, and parametric claims are all implemented payoff functions. Paper 3 asks who may hold them, why they trade, and what externalities they create.
- **Paper 5 computational object.** An agentic platform searches over $(\mathcal{H}, \mathsf{Q}, \kappa, L, A, N)$. Paper 3 supplies A: the admissibility frontier and the constraints that prevent market search from becoming extraction.

The program's central discipline is layered. A payoff direction can be valuable but too costly to implement; implementable but illiquid; liquid but socially harmful; or socially valuable but privately undersupplied. Paper 3 owns the last two distinctions.

17 Non-Obvious Implications

1. The correct regulatory object is a representation, not a payoff label and not a venue label.
2. Technical feasibility, legal permissibility, and normative admissibility are distinct.
3. Liquidity is not welfare; flow ecology matters.
4. A market can be socially valuable because it is subsidized, not despite being subsidized.
5. A market can be socially harmful because it is liquid, if liquidity comes from addiction, manipulation, or hidden leverage.
6. The same payoff can be insurance, hedging, gambling, governance, manipulation, or public information production.
7. Better information can expand markets through verification and shrink markets through pre-trade classification.
8. Consent is necessary but insufficient when claims target vulnerable participants or protected personal domains.
9. Human-capital markets require a no-control boundary.
10. Privacy is not merely a disclosure cost; it can be a hard design constraint.
11. Prediction markets require rules for control persons, event manipulation, and oracle governance.
12. Systemic admissibility is marginal and network-dependent; standalone safety is insufficient.
13. Pigouvian taxes align private and social incentives only in a special case.
14. Public data and templates can be pro-market admissibility tools.
15. Access restrictions can preserve hedging and information value while reducing extractive flow.
16. Personal agents must be fiduciary or they become distribution channels for exploitation.
17. The admissibility frontier is dynamic: each market birth changes spans, costs, liquidity, behavior, and systemic links.
18. The goal is not fewer markets or more markets. The goal is the right claims under the right representations.

18 Conclusion

The coming problem is not that markets will be impossible. It is that too many claims will be possible for possibility to carry normative weight. When the cost of claim creation falls, the economy loses a crude but important screen. Markets that once would have been too expensive to build can appear inside protocols, wallets, apps, platforms, personal agents, and legal templates. Some will complete missing risk-sharing opportunities, finance useful projects, generate public information, or allocate scarce flexibility. Others will monetize addiction, invade privacy, manipulate events, sell pieces of personhood, or wire small claims into large cascades.

The answer is not to financialize everything. It is not to ban financialization. It is to make admissibility the central object of market design. A claim-centric regulator asks what payoff is being traded, why participants trade it, how it is verified, who may access it, what data it uses, what incentives it creates, what liquidity supports it, what rights it touches, and how it connects to the system. Then it chooses the least restrictive instrument that makes the claim socially valuable: allow, standardize, subsidize, tax, margin, restrict, redesign, or prohibit.

The market was never the room. It was the state-contingent transfer. As the room dissolves into infrastructure, the old venue boundary cannot do the whole job. The claim becomes the unit of permission.

A Appendix A: A Two-Claim Example of Venue Insufficiency

Consider a venue that can list two binary event claims, both technically feasible and equally liquid. Claim 1 pays if regional rainfall falls below a threshold and is bought by farmers hedging crop risk. Its social value is

$$W_1 = S_1^{hedge} - K_1 - E_1 = 10 - 2 - 1 = 7.$$

Claim 2 pays if a thinly traded small-firm milestone fails and can be influenced by dominant traders. It has no hedging or decision value and manipulation externality $E_2 = 8$:

$$W_2 = 0 - 2 - 8 = -10.$$

A venue-level rule allowing the venue admits both and yields -3. A venue-level ban yields 0 but forgoes 7. A claim-level rule allows Claim 1 and prohibits Claim 2, yielding 7. The example illustrates Proposition 5.1.

B Appendix B: Belief-Driven Trading in a One-Claim Quadratic Model

Two agents trade a zero-net-supply claim g with variance σ^2. Agent a has risk aversion γ_a and subjective expected payoff m_a. Demand at price p is

$$q_a(p) = \frac{m_a - p}{\gamma_a \sigma^2}.$$

Market clearing gives

$$p^* = \frac{\tau_1 m_1 + \tau_2 m_2}{\tau_1 + \tau_2}, \qquad \tau_a = \frac{1}{\gamma_a \sigma^2}.$$

Perceived surplus is

$$S^{perceived} = \frac{1}{2} \sum_a \tau_a (m_a - p^*)^2.$$

If $m_1 \neq m_2$ only because of pure disagreement and there is no hedging, information, or consumption value, this perceived surplus is not social surplus under a common social evaluation. It is a transfer motivated by inconsistent beliefs, net of real costs and any induced risk.

C Appendix C: Manipulation Bound with Position Limits

Suppose manipulation action a changes event probability by $\Delta p(a)$ and a trader holds position q in a unit-payoff event claim. Expected trading gain is $q\Delta p(a)$. Manipulation cost is $c(a)$ and expected penalty is $\pi(a)F$. If positions are capped by $\bar{q}$, a sufficient no-manipulation condition is

$$\bar{q} \sup_a \Delta p(a) \leq \inf_a \{c(a) + \pi(a)F\} \quad \text{for all profitable manipulations } a.$$

More generally, admissible position limit $\bar{q}$ must satisfy

$$\bar{q} \leq \inf_{a:\Delta p(a)>0} \frac{c(a) + \pi(a)F}{\Delta p(a)}.$$

Thus position limits are not arbitrary paternalism; they are an incentive-compatibility instrument.

D Appendix D: Policy Operator in Pseudocode

```
Input: representation rho, market state Xi
1. If rho not technically feasible: reject or route to infrastructure design.
2. If rho not legally permissible: reject or route to legal reform analysis.
3. If rho violates hard protected-domain constraints: prohibit.
4. Estimate flow mix lambda_rho.
5. Estimate social benefits: hedge, immediacy, info, consumption, spillover.
6. Estimate resource costs: implementation, liquidity, settlement, dispute.
7. Estimate externalities: privacy, coercion, manipulation, moral hazard,
   addiction, inequality, oracle, dignity, norms, systemic.
8. Compute marginal social value relative to current market set.
9. If W >= 0 and private entry likely: allow or standardize.
10. If W >= 0 and private entry unlikely: subsidize, public provision,
    or safe harbor.
11. If W < 0 and private entry likely: tax, margin, restrict, redesign,
    or prohibit.
12. If W < 0 and private entry unlikely: monitor; prohibit if
    circumvention risk is high.
Output: policy bundle Pi(rho, Xi).
```

References

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman. 2016. "The Economics of Privacy." *Journal of Economic Literature* 54(2): 442–492.

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2015. "Systemic Risk and Stability in Financial Networks." *American Economic Review* 105(2): 564–608.

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

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

Arrow, Kenneth J., 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.

Bond, Philip, Alex Edmans, and Itay Goldstein. 2012. "The Real Effects of Financial Markets." *Annual Review of Financial Economics* 4: 339–360.

Bowles, Samuel. 1998. "Endogenous Preferences: The Cultural Consequences of Markets and Other Economic Institutions." *Journal of Economic Literature* 36(1): 75–111.

Brunnermeier, Markus K., and Martin Oehmke. 2013. "Bubbles, Financial Crises, and Systemic Risk." In *Handbook of the Economics of Finance*, vol. 2B, 1221–1288. Elsevier.

Brunnermeier, Markus K., Alp Simsek, and Wei Xiong. 2014. "A Welfare Criterion for Models with Distorted Beliefs." *Quarterly Journal of Economics* 129(4): 1753–1797.

Carvajal, Andres, 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.

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.

Eisenberg, Larry, and Thomas H. Noe. 2001. "Systemic Risk in Financial Systems." *Management Science* 47(2): 236–249.

Elliott, Matthew, Benjamin Golub, and Matthew O. Jackson. 2014. "Financial Networks and Contagion." *American Economic Review* 104(10): 3115–3153.

Garman, Mark B. 1976. "Market Microstructure." *Journal of Financial Economics* 3(3): 257–275.

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.

Glosten, Lawrence R., and Paul R. Milgrom. 1985. "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics* 14(1): 71–100.

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

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

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

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.

Ho, Thomas, and Hans R. Stoll. 1981. "Optimal Dealer Pricing under Transactions and Return Uncertainty." *Journal of Financial Economics* 9(1): 47–73.

Kyle, Albert S. 1985. "Continuous Auctions and Insider Trading." *Econometrica* 53(6): 1315–1335.

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112(2): 443–477.

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

MacKenzie, Donald. 2006. *An Engine, Not a Camera: How Financial Models Shape Markets.* Cambridge, MA: MIT Press.

O'Donoghue, Ted, and Matthew Rabin. 1999. "Doing It Now or Later." *American Economic Review* 89(1): 103–124.

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.

Roth, Alvin E. 2007. "Repugnance as a Constraint on Markets." *Journal of Economic Perspectives* 21(3): 37–58.

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

Sandel, Michael J. 2012. *What Money Can't Buy: The Moral Limits of Markets*. New York: Farrar, Straus and Giroux.

Satz, Debra. 2010. *Why Some Things Should Not Be for Sale: The Moral Limits of Markets*. Oxford: Oxford University Press.

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.

Titmuss, Richard M. 1970. *The Gift Relationship: From Human Blood to Social Policy*. London: Allen and Unwin.

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