

# The Missing-System Pricing Law: Decision-Target Recovery from Price-Formation Inputs and Realized Prices

**Erik Postnieks**

Founder and Principal Researcher

Center for Decision Accounting

Salt Lake City, Utah, USA

Corresponding author: [erik@decisionaccounting.org](mailto:erik@decisionaccounting.org)

Working Paper v2.4 — July 2026

**Research program:** <https://decisionaccounting.org/>

**Companion theorem release archive:** <https://decisionaccounting.org/research/mst/mst-v1.3.2-release-20260728.zip>

**Companion theorem landing page:** <https://decisionaccounting.org/papers/mst/>

**Companion formal proof page:** <https://decisionaccounting.org/proof/>

**Publication and version index:** <https://decisionaccounting.org/publication-roadmap/>

**Paper-specific materials:** a stable page for this paper, its correction record, and any released replication materials will be added upon public release. No replication archive for this paper is public at the date of this version.

This manuscript supersedes the earlier eponymously titled working paper on the same result (SAPM Working Paper No. 73). The formal lineage is governed by *The Missing System Theorem: The Limits of Welfare Recovery from Payoff-Space Information* (v1.3.2). The direct release archive above freezes that source, and every result required for this paper's conclusions is also stated and proved here.

## Abstract

A market price is produced from bids, asks, reports, observed attributes, and a rule that maps those inputs into a realized price. This paper asks when that architecture carries the system-cost information needed for a declared decision: a marginal external cost, a compliance threshold, a resource-stock trigger, a benchmark-integrity adjustment, or another target specified before the audit. The information answer has two levels. The complete inputs recover the decision target exactly when the target is constant on every input fiber. The realized price recovers it exactly when the target is constant on every price fiber; price recovery therefore requires input recovery and a pricing rule that retains the needed distinctions. Exact recovery of total system welfare is a stronger optional diagnostic and is not required when the decision needs only a marginal or threshold statistic. The incentive question asks what public disclosure changes when the disclosed coordinate leaves payoff primitives, feasible actions, and beliefs about exogenous payoff-relevant states unchanged. Established public-randomization and sunspot logic implies that continuation best-response correspondences are unchanged for fixed conjectures while equilibrium beliefs about others' actions and equilibrium selection may change. Four diagnostic institutional routes and two scoped case audits then show how measurement, private valuation or cost, rules, and anticipated consequences change a price-forming architecture. The contribution is an applied synthesis and audit protocol; the accounting and return studies remain prospective.

## Highlights.

- Equal complete inputs imply equal realized prices.
- Decision-target recovery has distinct input- and price-fiber tests.
- Payoff-orthogonal public signals can select existing equilibria.
- Four routes diagnose changes in price-forming architecture.
- Accounting and return studies remain prospective protocols.

**Keywords:** system welfare, price formation, externalities, disclosure, information and market efficiency, institutional design, welfare accounting

*JEL classification:* D62, D82, D83, G14, H23

## Plain-language summary

Every day, markets set prices for coal, opioids, data, and firearms. A pricing decision may need a particular piece of system-cost information: the additional harm from one more unit, a threshold indicating that an order should stop, or a verified integrity adjustment. This paper asks whether the complete input record and the realized price retain that declared decision target. A target absent from the complete inputs cannot appear in the price. A target present in the inputs can still disappear when the pricing rule compresses them. Recovering total system welfare is a stronger question and should be asked only when a named decision requires it. Public disclosure that leaves payoff and feasibility primitives unchanged may still coordinate participants by changing strategic beliefs and selecting among equilibria supported by the underlying game.

Two scoped case audits and four diagnostic routes show how an application can identify the target, input owner, pricing rule, response path, and evidence needed for a defensible conclusion.

## 1. The problem in plain language

In transactions where no carbon price, liability rule, contractual term, or buyer preference carries marginal climate damage into private costs, the fuel price may omit a decision-relevant system cost. Similar gaps can arise when the shared object is an institution, such as a benchmark's integrity, or a population's health. Section 7 provides scoped audits of benchmark reform and pharmaceutical distribution. The benchmark audit is supported by the Wheatley Review's documented reform of LIBOR inputs and governance. The pharmaceutical audit is a prospective architecture test grounded in federal suspicious-order and ARCOS requirements; it does not claim an observed equal-price witness without transaction-level price data.

The familiar name for this pattern is externality, and a century of economics has studied it. This paper asks a narrower and more exacting question, and it asks the question twice. First: as a matter of information, when can a price-formation process (an order book, a negotiation, a posted-price algorithm) carry the condition of a shared system at all? Second: as a matter of incentives, what determines whether anyone feeds that condition into the process, or acts on it once fed? The two questions have different answers with different logical standing, and most of what has gone wrong in earlier discussions of this topic, including an earlier version of this paper, came from merging them.

The information answer is a two-level result. A pricing process computes its output from what it receives. Before applying the result, the analyst declares the information required for the decision, written  $T$ : for example, marginal external cost for one more unit, a diversion-risk stop threshold, or a benchmark-integrity adjustment. If two admissible situations present the process with identical complete inputs and different  $T$ , neither the inputs nor the realized price recover the decision target on a domain containing the pair. Even when the inputs distinguish the situations, the pricing rule may map them to the same price. Input recovery is therefore weaker than price recovery. Exact recovery of total system welfare  $W$  uses the same fiber test but is a stronger diagnostic whose relevance must be justified by a named decision. Section 4 states both levels and explains why decision-target recovery can matter even when total welfare is not recoverable.

The incentive answer applies established public-randomization and sunspot logic. Suppose the system coordinate changes no utility or feasible action and reveals nothing about an exogenous payoff-relevant state. Publishing it leaves those primitives and the best-response operator unchanged. The public signal can still change equilibrium beliefs about others' actions and select a different equilibrium of the underlying game. Disclosure can therefore change conduct and expected system welfare through coordination while adding no payoff or feasibility consequence. Section 5 states that result using continuation strategies in a finite Bayesian game and distinguishes primitive invariance, strategic beliefs, realized-state welfare, and expected welfare.

The formal results leave a practical question: how does an application change its price-forming architecture? Section 6 presents four recurring routes: an informative measurement entering a price input; a direct private

valuation or cost; a rule-imposed consequence such as a liability, tax, standard, or contractual term; and an anticipated future valuation or consequence. These labels are diagnostic categories for a fully described application. An analyst must enumerate the actual paths before using the taxonomy; no necessity or exhaustiveness claim attaches to the four labels.

The primary result can now be stated precisely. On a domain containing an equal-input, unequal-target pair, the declared inputs cannot recover the decision target and the realized price cannot recover it. Removing that witness requires inputs that preserve every distinction in  $T$  needed for the decision. Price recovery additionally requires the pricing rule to retain those distinctions. If  $T$  is marginal system cost and a maintained decision rule sets the corrective component of price equal to  $T$ , recovery supplies the information needed for that rule. It still does not establish correct measurement, behavioral response, distributional acceptability, compensation, or implementation quality. Exact recovery of the welfare level  $W$  is separately labeled throughout.

Three qualifications bound the contribution. First, the existence of witness situations in a real market is an empirical premise, established by construction inside a specified model or investigated at a stated field measurement resolution; the companion theorem paper supplies the input-witness protocol. Second, the formal components have close predecessors: externalities as missing markets, structural nonconvexity, sufficient-statistic and price-revelation results, public randomization, and sunspot equilibrium. The contribution claimed here is an applied synthesis of decision-target and stronger welfare-level fiber tests, a payoff-orthogonal disclosure application, a diagnostic institutional taxonomy, and an application-audit protocol. Third, where parties' payoffs already internalize the system cost, the application must model that path directly.

The paper proceeds as follows. Section 2 positions the boundary against its nearest literatures and states what each established. Section 3 defines the objects. Section 4 states and proves the input- and price-recovery results and works through an eight-situation pricing example. Section 5 states the payoff-orthogonal public-disclosure result. Section 6 presents the diagnostic taxonomy. Section 7 reads existing instruments and objections through it. Section 8 reports the evidence protocols and their status. Section 9 separates formal failure conditions from application audits. Sections 10 and 11 state the design implication, limitations, and conclusion.

## 2. What this boundary adds to established results

The intellectual ancestry determines the contribution's scope. The defensible contribution is the assembled diagnostic framework.

K. William Kapp documented, across air pollution, water pollution, occupational disease, and resource depletion, how private production shifts costs onto society without compensation (Kapp 1950). Pigou formalized the wedge between private and social marginal cost and the corrective tax that closes it (Pigou 1920); the corrective program requires damage measurement and an institution able to impose the charge.

Coase analyzed the reciprocal character of harmful effects and compared alternative institutional arrangements (Coase 1960). Bargaining applications require defined rights and affected parties or representatives able to negotiate at tolerable information and transaction costs. Arrow described externalities as missing markets (Arrow 1969), and Starrett proved that fundamental nonconvexities can defeat the existence of a price-supported optimum with detrimental externalities (Starrett 1972). Samuelson's public-goods result (Samuelson 1954), the constrained-inefficiency results of Greenwald and Stiglitz (1986) and Geanakoplos and Polemarchakis (1986), and the welfarist-invariance results of Sen (1977) and d'Aspremont and Gevers (1977) each bound what prices, markets, or welfare functionals can carry under their stated conditions.

Several information and equilibrium results sit close to the two formal layers. Grossman and Stiglitz show that costly acquisition and fully revealing prices are incompatible in their information-market model (Grossman and Stiglitz 1980). Their state is payoff-relevant to traders. Grossman and Milgrom show how verifiable disclosure can unravel when product information bears on market payoffs (Grossman 1981; Milgrom 1981). Chetty's sufficient-statistics program recovers welfare objects from selected observables under maintained economic structure (Chetty 2009). This paper uses a simpler exact factorization question: whether the complete input or the realized price preserves every distinction required to recover a declared decision target, with total system welfare retained as a stronger diagnostic.

The direct predecessors to Section 5 are public randomization and sunspot equilibrium. Aumann (1974) studies subjective random devices and correlated strategies, including information that need not be publicly identical across players. The result in Section 5 uses the narrower case of one common public signal indexing Bayes–Nash equilibria. Cass and Shell (1983) show that payoff-irrelevant public uncertainty can affect rational-expectations equilibrium allocations in an overlapping-generations economy with restricted market participation, while also identifying a complete static Arrow–Debreu setting in which it does not. Duffy and Fisher (2005) report laboratory evidence that extrinsic public signals can coordinate market outcomes and that the effect depends on market institution and information flow. Section 5 applies this established logic to disclosure of a payoff-orthogonal system coordinate: the public signal leaves payoff and feasibility primitives fixed while permitting equilibrium selection.

The disclosure literature also gives direct evidence and boundary conditions for sustainability applications. Ben-Shahar and Schneider (2014) document the practical limits of mandated disclosure across legal settings. Christensen, Floyd, Liu, and Maffett (2017) study a mandate that placed mine-safety records in financial reports and report changes in safety, productivity, market reactions, and ownership. Christensen, Hail, and Leuz (2021) review the economics of mandatory corporate-social-responsibility and sustainability reporting, including capital-market effects, real effects, standardization benefits, externalities, and implementation costs. These studies show why a domain application must identify the response path instead of treating disclosure as a single intervention.

On the commons side, Gordon and Hardin describe rivalrous stocks depleted by uncoordinated appropriation (Gordon 1954; Hardin 1968). Ostrom documents varied cases in which communities

developed institutions for governing common-pool resources (Ostrom 1990). In the diagnostic taxonomy, monitoring can create a measurement path and graduated sanctions can create a rule-imposed consequence. A domain study must establish how the particular institution changes  $q$ ,  $p$ , and the relevant outcomes.

On the measurement side, Muller, Mendelsohn, and Nordhaus constructed industry-level gross external damages and compared them to value added, finding several industries whose air-pollution damages exceed their value added (Muller, Mendelsohn, and Nordhaus 2011). That study is a predecessor for the measurement protocol described in Section 8. In finance, Bolton and Kacperczyk (2021) report higher returns for high-emissions firms in their sample and institutional screening in salient industries. Pástor, Stambaugh, and Taylor (2021) model sustainable investing in equilibrium. These works show that environmental exposure can become relevant to investor decisions under particular preferences and institutions. They do not establish the event-concentrated repricing hypothesis reserved for the companion return-side study.

**Table 1. How this paper fits the literature.**

| Author/work                     | What it established                                                                                                                                | Inputs and assumptions                                                      | What remains unresolved                                                                          | What this paper adds                                                                                     |
|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------|
| Pigou (1920)                    | A corrective tax can align private and social marginal cost under the analysis's conditions.                                                       | A damage estimate and an institution able to impose the tax.                | Whether a named pricing process's inputs and realized price recover the declared damage measure. | Separate input- and price-recovery tests.                                                                |
| Kapp (1950)                     | Private production shifts social costs across the documented industries and institutions.                                                          | Institutionalist and historical documentation.                              | A common audit for the information retained by a named pricing process.                          | The architecture audit and diagnostic route taxonomy.                                                    |
| Coase (1960)                    | Institutional arrangements, rights, and transaction costs shape how harmful effects are resolved.                                                  | Defined institutional setting and comparative institutional analysis.       | Whether a named bargain's input and price maps recover a declared system measure.                | Treats bargaining rights and enforcement as application-specific causal paths.                           |
| Arrow (1969); Starrett (1972)   | Missing markets and fundamental nonconvexities create structural limits under the stated general-equilibrium conditions.                           | Competitive general equilibrium and the respective externality assumptions. | An operational test on a fixed observed pricing process.                                         | Separate fiber tests on declared inputs and realized prices; no priority claim on structural incapacity. |
| Samuelson (1954)                | Derived an optimality condition for collective-consumption goods and identified the difficulty of decentralized preference revelation and pricing. | Collective-consumption technology and stated preference information.        | Whether a named price architecture carries a declared decision target.                           | Separates target recovery from the provision and revelation problems.                                    |
| Grossman–Stiglitz (1980)        | Costly acquisition and fully revealing prices cannot coexist in the model's informational equilibrium.                                             | Costly acquisition; payoff-relevant private information.                    | Exact recovery of a declared decision target from a specified input or price map.                | Separates target-input recovery from target-price recovery.                                              |
| Grossman (1981); Milgrom (1981) | Verifiable disclosure can unravel under their product-information and market-inference conditions.                                                 | Payoff-relevant, verifiable information and specified market beliefs.       | Disclosure of a payoff-orthogonal system coordinate.                                             | Applies equilibrium-selection logic to that coordinate.                                                  |

Table 1 (continued).

| Author/work                                           | What it established                                                                                                                                                                                           | Inputs and assumptions                                                                       | What remains unresolved                                                                                                   | What this paper adds                                                                          |
|-------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------|
| Aumann (1974); Cass–Shell (1983); Duffy–Fisher (2005) | Shared randomization can correlate strategic behavior; extrinsic uncertainty can affect equilibrium allocations in specified economies; laboratory market effects depend on institution and information flow. | Multiple equilibria or correlation opportunities; model- or institution-specific conditions. | Application to disclosure of a declared system coordinate.                                                                | Applies established equilibrium-selection logic and states its welfare scope.                 |
| Ben-Shahar–Schneider (2014)                           | Disclosure mandates often fail to change behavior as intended across the settings examined.                                                                                                                   | Legal and institutional evidence.                                                            | A domain-specific causal classification of response paths.                                                                | Uses the four-route audit plus coordination as diagnostic questions.                          |
| Christensen et al. (2017, 2021)                       | Mine-safety reporting had measured real and market effects; sustainability-reporting economics includes capital-market, real, externality, standardization, and cost channels.                                | A specific U.S. mandate; a broad economics review.                                           | An exact input/price recovery test for a declared system measure.                                                         | Connects disclosure mechanisms to the architecture audit without claiming a universal effect. |
| Sen (1977); d’Aspremont–Gevers (1977)                 | Sen analyzes informational restrictions in social-welfare analysis; d’Aspremont and Gevers characterize collective-choice principles under specified informational and equity conditions.                     | Social-choice frameworks with declared informational bases and invariance conditions.        | Application to a fixed price architecture with an operational witness test.                                               | Uses factorization on declared decision targets and, separately, welfare levels.              |
| Greenwald–Stiglitz (1986)                             | In their imperfect-information and incomplete-market framework, equilibria are rarely constrained Pareto optima, and Pareto-improving tax interventions generally exist.                                      | Model-specific incomplete markets, imperfect information, and intervention conditions.       | Whether the target needed by a named intervention is recoverable from a specified institutional record or realized price. | A signal-specific factorization audit; no new constrained-inefficiency theorem.               |

Table 1 (continued).

| Author/work                                    | What it established                                                                                                                            | Inputs and assumptions                                                                                                      | What remains unresolved                                                              | What this paper adds                                                                                                                             |
|------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------|
| Geanakoplos–Polemarchakis (1986)               | With incomplete asset markets, competitive-equilibrium allocations are generically constrained suboptimal under the paper’s stated conditions. | An incomplete asset market, the paper’s numeraire and regularity conditions, and feasible reallocations of existing assets. | Whether a named decision target is computable from a specified input or price map.   | A signal-specific factorization audit; no new incomplete-markets theorem.                                                                        |
| Chetty (2009)                                  | Selected empirical elasticities can support welfare analysis under maintained economic structure.                                              | Model-specific structural assumptions and sufficient statistics.                                                            | An exact factorization test for a proposed statistic.                                | States separate fiber conditions for complete inputs and realized prices.                                                                        |
| Ostrom (1990)                                  | Communities in the studied cases developed varied institutions for governing common-pool resources.                                            | Case-specific monitoring, rules, participation, and sanctions.                                                              | How those institutions alter a named pricing architecture.                           | Classifies monitoring and sanctions as recurring measurement and rule paths.                                                                     |
| Muller–Mendelsohn–Nordhaus (2011)              | Estimated gross external air-pollution damages by U.S. industry and compared them with value added.                                            | Emissions data, dispersion modeling, exposure, health and other damage valuation.                                           | Whether a named pricing process’s inputs or price recover the measured damage.       | Supplies an input/price audit that a domain study can apply to such an estimate.                                                                 |
| Bolton–Kacperczyk (2021); Pástor et al. (2021) | Carbon exposure can affect returns and investor demand under specified preferences and institutions.                                           | Empirical return and holdings evidence; an equilibrium model of sustainable investing.                                      | The causal route and timing by which a particular institutional event affects price. | A protocol for testing an anticipated-consequence route without attributing that result to these papers.                                         |
| Companion theorem paper (Postnieks 2026a)      | The input-recoverability characterization, witness protocol, and path results.                                                                 | Declared signal; specified domain.                                                                                          | Recovery from the realized price and the incentive layer.                            | Adds the P-fiber test, the disclosure application, and the diagnostic taxonomy. The cited v1.3.2 source is frozen in the direct release archive. |

The disclosure result applies established public-randomization and sunspot logic to a system-welfare coordinate. The recoverability results apply elementary factorization conditions separately to complete price inputs and realized prices. The paper’s prospective contribution is their integration with an application-level institutional taxonomy, an audit protocol, and prospective measurement and return-study designs. The contribution is an applied synthesis and diagnostic framework whose components have close formal and empirical predecessors.

### 3. Objects and assumptions

A **transaction situation**  $\omega$  is a complete description of an admissible state of affairs when a trade may occur: the parties, their actions and reports, the environment, and the state of any shared system the analyst has declared.  $\Omega$  is the set of admissible situations. Every claim below is relative to a stated domain.

A **decision target**  $T: \Omega \rightarrow \mathbb{R}$  is the state-dependent quantity required by a named pricing or governance decision. Examples include marginal external cost per additional unit, a diversion-risk stop indicator, a resource-stock threshold, and a benchmark-integrity adjustment. The analyst must declare  $T$ 's unit, baseline, activity boundary, and the decision rule that uses it before testing recovery.  $\mathbb{R}$  may be any finite or measurable codomain; the exact factorization results do not require a monetary target.

**System welfare**  $W: \Omega \rightarrow \mathbb{R}$  is a declared global or total measure of the shared system's condition, with a stated boundary, baseline, and unit. Exact recovery of  $W$  is a stronger diagnostic. Its relevance must be justified when the named decision depends on a marginal, threshold, or adjustment statistic rather than on the welfare level itself.

A **price-formation process** is a pair  $(q, p)$ . The **complete declared input**  $q: \Omega \rightarrow D$  collects everything the process receives in situation  $\omega$ : bids, asks, reports, observed attributes, relevant history, and the chosen actions or order flow that the process conditions on. The **rule**  $p: D \rightarrow \mathbb{R}$  maps the input to a price. The realized price function is  $P = p \circ q$ .

**Axiom M (Mediation).** The price depends on the transaction situation only through the complete declared input:  $P = p \circ q$ . An application must state what  $q$  contains, who or what defines the fixed map  $p$ , and every direct and indirect causal path by which the system coordinate can change  $q$ . If the operative pricing rule varies across situations, the rule state must be included in an expanded input and mapped to price by a fixed higher-level rule.

Axiom M is an application-level modeling commitment. A price difference across a claimed equal-input pair normally diagnoses an incomplete input enumeration or an incorrectly measured input; it does not by itself refute the mathematical identity after the true input has been specified.

For the disclosure result,  $\Gamma$  is a finite Bayesian game: players  $i \in N$ , finite action sets  $A_i$ , a finite payoff-relevant state  $\theta \in \Theta$ , finite private signals  $t_i \in T_i$ , utility functions  $u_i(a, \theta)$  over action profiles  $a \in A = \times_i A_i$ , and a common prior. A behavioral strategy maps  $t_i$  into  $\Delta(A_i)$ . Let  $E(\Gamma)$  denote the set of Bayes–Nash equilibrium strategy profiles. The finite system coordinate is  $x \in X$ , and the system-welfare mapping is  $W(a, x)$ .

### 4. The information boundary

The intuition comes first. A target can be reconstructed from the complete input exactly when situations sharing that input also share the target. It can be reconstructed from the realized price exactly when situations sharing that price also share the target. Because the price is computed from the input, price recovery is the

stricter requirement. The same tests can be applied to total system welfare, but a pricing decision may need only a marginal cost or threshold.

**Proposition 1 (Input Variation).** Let  $(q, p)$  satisfy Axiom M. If  $P(\omega') \neq P(\omega'')$ , then  $q(\omega') \neq q(\omega'')$ , and  $p$  is nonconstant on the relevant input variation. Equivalently, if  $q(\omega') = q(\omega'')$ , then  $P(\omega') = P(\omega'')$ .

*Proof.* By Axiom M,  $P(\omega) = p(q(\omega))$ . Equal arguments of  $p$  have equal values. The contrapositive gives the first statement. ■

**Proposition 2 (Decision-target recovery and stronger welfare-level diagnostic).** Let  $(q, p)$  satisfy Axiom M, let  $P = p \circ q$ , and let  $T$  be declared on  $\Omega_0 \subseteq \Omega$ .

- (a) **Input recovery.** There exists a function  $g$  with  $T = g \circ q$  on  $\Omega_0$  if and only if  $T$  is constant on every  $q$ -fiber in  $\Omega_0$ .
- (b) **Price recovery.** There exists a function  $h$  with  $T = h \circ P$  on  $\Omega_0$  if and only if  $T$  is constant on every  $P$ -fiber in  $\Omega_0$ .
- (c) **Relationship.** Price recovery implies input recovery. Input recovery implies price recovery only when  $p$  preserves every input distinction needed to recover  $T$ .
- (d) **Stronger welfare-level diagnostic.** Parts (a)–(c) also hold after replacing  $T$  by  $W$ . Recovery of  $W$  requires  $W$  to be constant on the corresponding fibers and is distinct from recovery of a decision target.

A **target q-witness** is a pair  $\omega', \omega'' \in \Omega_0$  with  $q(\omega') = q(\omega'')$  and  $T(\omega') \neq T(\omega'')$ . It rules out both target-input and target-price recovery. A **target P-witness** is a pair with  $P(\omega') = P(\omega'')$  and  $T(\omega') \neq T(\omega'')$ . It rules out target-price recovery. A target  $P$ -witness can exist even when  $q(\omega') \neq q(\omega'')$  and target-input recovery holds. A **welfare q-witness** or **welfare P-witness** is defined by substituting  $W$  for  $T$  and tests the stronger welfare-level claim.

*Proof.* For (a), necessity follows because a composition with  $q$  is constant on  $q$ -fibers. For sufficiency, assign to each realized  $q$ -value the common  $T$ -value of its fiber and extend arbitrarily outside  $q(\Omega_0)$ . The same argument with  $P$  proves (b). For (c), if  $T = h \circ P$  and  $P = p \circ q$ , then  $T = (h \circ p) \circ q$ . The reverse implication requires  $T$  to be constant on  $P$ -fibers, as part (b) states. Substituting  $W$  for  $T$  proves (d). ■

**Meaning of “recover.”** Target-input recovery means  $T = g \circ q$ . Target-price recovery means  $T = h \circ P$ . Welfare-level recovery uses  $W$  in the same equations. These are informational conditions. They do not establish correct measurement, efficiency, compensation of every loss, or complete insurance.

**Corollary 3 (Decision sufficiency for a declared target).** Suppose a full-information pricing or governance decision has the maintained form  $dF(\omega) = r(T(\omega))$ . If  $T = g \circ q$ , the decision is implementable from the complete input as  $dF = r \circ g \circ q$ . If  $T = h \circ P$ , an observer of the realized price can reproduce the decision as

$dF = r \circ h \circ P$ . When  $r$  is injective on  $T(\Omega_0)$ , failure of target recovery also rules out exact recovery of the decision from the same signal.

*Proof.* Substitute the recovery equation into  $dF(\omega) = r(T(\omega))$ . Under injectivity, exact recovery of  $dF$  would imply exact recovery of  $T$  by applying the inverse of  $r$  on its realized range. ■

**Worked example: a marginal target can be recovered even when the welfare level is not.** Let  $\Omega_0$  contain eight situations indexed by a private cost state  $z \in \{0, 1\}$ , a current marginal-damage state  $x \in \{0, 1\}$ , and an accumulated system-stock state  $y \in \{0, 1\}$ . Declare the target  $T = 5x$ , the charge needed by a maintained corrective rule  $dF = T$ . Define the stronger welfare-level diagnostic as  $W = 100 - 20x - 30y$ . Initially the posted-price process receives  $q(z, x, y) = z$  and sets  $P = 100 + 10z$ . For either value of  $z$ , situations that differ in  $x$  form a target  $q$ -witness: they have the same input and price but different  $T$ .

Now add an audited marginal-damage measure  $m = x$  to obtain  $q_1(z, x, y) = (z, m)$ . The target is recoverable from  $q_1$  because  $T = 5m$ . If the posted-price rule remains  $P_1 = 100 + 10z$ , equal- $z$  situations with different  $x$  remain target  $P$ -witnesses. The complete input carries  $T$  while the realized price discards it. Finally, let the maintained rule impose the declared charge:  $P_2 = 100 + 10z + 5m$ . The realized price now identifies  $x$  and therefore  $T$  on the domain, so the corrective decision is both implementable from  $q_1$  and auditable from  $P_2$ . The welfare level  $W$  remains unrecoverable from  $q_1$  and  $P_2$  because  $y$  is absent from both. This example isolates the paper's primary claim: a price architecture can carry the information needed for a declared marginal decision while leaving the stronger total-welfare question unresolved.

The results are signal-specific and domain-specific. An audited emissions measure, a benchmark-integrity measure, or a resource-stock measure changes  $q$  and creates a different application. Neither proposition asserts that a witness exists in every market. A witness must be constructed in a specified model or investigated under a stated field measurement resolution. The companion theorem paper supplies the input-witness protocol; the additional  $P$ -fiber audit checks whether the pricing map retains the recovered distinction.

## 5. Payoff-orthogonal disclosure and continuation-equilibrium selection

A public announcement can affect conduct in two conceptually separate ways. It can change payoff or feasibility primitives, or it can coordinate expectations about which equilibrium participants will play. This section isolates the second channel by defining a signal that is independent of exogenous payoff-relevant information and enters no utility or action set.

**Definition 1 (payoff-orthogonal system coordinate).** In  $\Gamma$ ,  $x$  is *payoff-orthogonal* if (i) no utility function  $u_i$  depends on  $x$ ; (ii) no action set  $A_i$  depends on  $x$ ; and (iii)  $x$  is independent of the payoff-relevant state and all private signals  $(\theta, t_1, \dots, t_n)$  under the common prior.

$x$  changes no payoff or feasible action and reveals nothing about an exogenous payoff-relevant matter. If an aquifer condition predicts future input costs, for example, disclosure changes beliefs about a private cost and

does not satisfy this definition. Intrinsic, professional, or institutional motives belong in  $u_i$  if an application relies on them. Beliefs about other players' actions remain endogenous and can respond to a public signal.

**Definition 2 (pure disclosure).** A *pure disclosure* is a public signal  $\tau = f(x)$  announced before play, with utilities, action sets, and the prior over  $(\theta, t)$  otherwise unchanged.

**Proposition 4 (Payoff-orthogonal public disclosure and equilibrium selection).** Let  $\Gamma$  have a payoff-orthogonal system coordinate  $x$ , and let  $\Gamma^\tau$  be the game after pure disclosure of  $\tau = f(x)$ .

- (a) **Primitive invariance.** Disclosure leaves utilities, feasible actions, beliefs about exogenous payoff-relevant states, and each player's best-response operator as a function of beliefs unchanged. It may change equilibrium beliefs about opponents' actions and thereby change the selected best response. A  $\tau$ -ignoring strategy profile is a Bayes–Nash equilibrium of  $\Gamma^\tau$  if and only if it belongs to  $E(\Gamma)$ .
- (b) **Continuation-equilibrium selection.** In every Bayes–Nash equilibrium of  $\Gamma^\tau$ , the continuation strategy profile conditional on each positive-probability realization  $\tau_0$  belongs to  $E(\Gamma)$ . Conversely, every  $\tau$ -indexed selection of strategy profiles from  $E(\Gamma)$  is a Bayes–Nash equilibrium of  $\Gamma^\tau$ . Pure disclosure can therefore select among pre-existing Bayes–Nash equilibria, including selections contingent on the disclosed state.
- (c) **State-conditional welfare scope.** For  $\sigma \in E(\Gamma)$ , define  $w(\sigma, x)$  as expected system welfare in state  $x$  under  $\sigma$ , integrating over  $\theta$ , private signals, and mixed actions. Let  $W_0$  be a baseline. If  $w(\sigma, x) < W_0$  for every  $\sigma \in E(\Gamma)$  and every  $x \in X$  with positive probability, then every equilibrium of  $\Gamma^\tau$  remains below  $W_0$  state by state. State-contingent disclosure can still improve welfare when it selects higher-w equilibria in some states, including cases in which every selected equilibrium remains below  $W_0$ .

*Proof.* For (a),  $\tau$  is independent of  $(\theta, t)$  by Definition 1(iii), and neither utilities nor feasible actions depend on  $x$  by (i)–(ii). The best-response operator for any fixed conjecture about opponents is therefore unchanged. Opponents' equilibrium strategies may vary with  $\tau$ , so the conjecture supplied to that operator may vary. A  $\tau$ -ignoring equilibrium consequently has the same deviation inequalities in  $\Gamma$  and  $\Gamma^\tau$ . For (b), fix a positive-probability realization  $\tau_0$ . If the continuation strategy profile were outside  $E(\Gamma)$ , some player would have a profitable type-contingent deviation after  $\tau_0$ ; using that deviation only after  $\tau_0$  would contradict equilibrium in  $\Gamma^\tau$ . Conversely, selecting an element of  $E(\Gamma)$  after every  $\tau$  satisfies every continuation deviation inequality. Part (c) follows by applying its premise to the selected  $\sigma(\tau) \in E(\Gamma)$  at every realized  $x$ . ■

The result is the finite-Bayesian-game form of established public-randomization and sunspot logic. A public signal can choose among equilibria supported by the underlying payoff and feasibility primitives. Strategic beliefs and chosen actions can change even though the primitive best-response operator does not. A disclosure claim must therefore specify whether it concerns payoff or feasibility primitives, strategic beliefs, continuation-equilibrium selection, state-conditional welfare, or ex ante welfare.

## 6. Four diagnostic routes into a price-forming process

The following causal taxonomy is an application audit. It carries no necessity, sufficiency, mutual-exclusivity, or exhaustiveness claim. An application must identify the price setter, the fixed rule  $p$ , the inputs  $q$ , and every path by which  $x$  changes  $q$  before classifying a route. If a pricing policy changes across situations, its rule state belongs in the expanded input described after Axiom M.

**D1 — Measurement or input route.** An  $x$ -informative measurement, whether or not it recovers  $T$  or  $W$ , enters  $q$ . It may affect price through a formula, through participants' use of the information, or through public coordination among existing equilibria.

**D2 — Direct private valuation or cost.**  $x$  changes a participant's current private benefit, cost, or feasible action without a new public rule. A buyer who values verified clean production, or a firm whose resource condition changes its operating cost, illustrates this route.

**D3 — Rule-imposed consequence.** A tax, liability, standard, contractual term, procurement rule, or enforceable right makes  $x$  consequential to a participant's payoff or feasible action.

**D4 — Anticipated future valuation or consequence.** A credible expectation that D2 or D3 will operate later changes current continuation values and present choices.

The routes can combine and overlap. D1 supplies measurement; D2 or D3 often supplies a current response path; D4 can capitalize an expected future response into a current price. A price can vary with  $x$  under Axiom M whenever  $q$  varies with  $x$  and  $p$  responds to that variation. Input recovery and price recovery remain governed by Proposition 2's separate fiber conditions. Classification under one or more routes establishes neither condition by itself.

The boundary can now be stated with its scope attached. **On a documented target  $q$ -witness domain, the declared inputs and the realized price fail to recover the decision target. Repairing target-input recovery requires an input architecture whose fibers preserve  $T$ . Repairing target-price recovery additionally requires  $p$  to preserve the relevant distinctions.** The same statement applies to  $W$  when exact welfare-level recovery is the declared objective. Measurement, direct valuation or cost, rule-imposed consequence, and anticipated future valuation or consequence are recurring routes an analyst can test while explaining how the architecture changes.

## 7. Reading instruments and objections through the taxonomy

The taxonomy organizes instruments without ranking them. Each application requires a causal account of the price setter, inputs, and response mechanism.

### 7.1 Benchmark manipulation: a documented input and governance repair

The Wheatley Review supplies a concrete architecture audit for LIBOR reform. The relevant price-forming object was a benchmark assembled from contributing banks' submissions under an administrator's calculation

and governance rules. For an integrity decision, a defensible target  $T$  could be a declared transaction-support or exception indicator used to accept, challenge, or adjust a submission. The review found weaknesses in the prevailing submission and oversight structure and recommended that submissions be supported by relevant transaction data, with stronger administration, governance, regulation, and sanctions (HM Treasury 2012).

In the paper's notation, the pre-reform input  $q$  included bank submissions and the administrator's calculation inputs. The reform agenda enlarged and verified that input architecture through transaction evidence and strengthened the rule state through governance and enforcement. Those changes instantiate D1 and D3. The public report supports the architecture claim; it does not by itself establish an historical target  $q$ -witness or target  $P$ -witness at a specified observation resolution. Completing that test requires archived submissions, underlying eligible transactions, the calculation rule in force, exception or integrity classifications, and the published benchmark for matched fixing dates. The decision target must also be fixed before comparing fibers.

This case shows why target recovery and realized-price recovery are separate. Transaction support can make an integrity target computable from the administrator's record. A trimmed or averaged benchmark can still compress distinctions in that record. The administrator may therefore be able to make an integrity decision even when an outside observer cannot reconstruct the target from the published benchmark alone.

## 7.2 Pharmaceutical distribution: a prospective suspicious-order pricing audit

Federal controlled-substance distribution rules provide a second, prospective application. Registered distributors must maintain systems designed to identify suspicious orders and inform the Drug Enforcement Administration of suspicious orders when discovered; the regulation lists unusual size, pattern, and frequency as examples (21 C.F.R. § 1301.74(b)). The Drug Enforcement Administration also describes ARCOS as a system that collects and summarizes controlled-substance transaction information and has released enhanced aggregate data intended to help registrants identify, report, and stop suspicious opioid orders (Drug Enforcement Administration 2019).

For a named release, hold, or surcharge decision,  $T$  could be a declared suspicious-order stop indicator or a quantified marginal diversion-risk cost. The complete input  $q$  must distinguish the commercial pricing record from compliance information such as customer history, order size, frequency, pattern, alerts, and the operative review rule. If the compliance coordinate is absent from the price setter's input, a target  $q$ -witness can arise when otherwise identical commercial inputs receive different  $T$  classifications. If the coordinate enters  $q$  but the realized contract price is unchanged, the pair is a target  $P$ -witness even though the hold decision may be implementable from the complete record.

This manuscript does not claim that either witness has been observed. A defensible field test requires transaction-level prices and contract terms, orders and customer identifiers, compliance alerts, review dispositions, shipment or hold decisions, the versioned suspicious-order rule, and a predeclared matching resolution. ARCOS records can help establish distribution quantities and flows; they do not supply the complete contract-price and decision record needed for the  $P$ -fiber audit. The likely routes are D1 through

monitoring and D3 through regulatory duties, with possible D4 effects when distributors anticipate enforcement consequences.

**Pigouvian taxation** is D3: an institution measures a damage estimate and attaches a charge. It makes the estimated harm a private cost. Its performance depends on measurement, rate design, attribution, enforcement, and political durability.

**Property rights and Coasean bargains** can make an affected interest consequential through D3 or, where a resource condition directly changes a party's costs, through D2. Coasean bargaining requires enforceable rights and affected parties or representatives able to bargain at tolerable information and transaction costs. Atmospheric, benchmark, and population harms may be represented by people, firms, governments, trustees, or statutory bodies. The practical question is whether representation, standing, information, and enforcement are adequate in the specified case.

**Ostrom's cases** demonstrate that communities can build monitoring and sanctions. Local observability, identifiable users, repeated interaction, and enforceable community rules are frequent features emphasized in this application of the framework; they are not asserted here as universal necessary-and-sufficient conditions. In taxonomy terms, monitoring can be D1 and sanctions D3.

**Disclosure mandates** can create D1. Under Proposition 4's assumptions, a pure disclosure changes no payoff or feasibility primitive and reveals nothing about an exogenous payoff state. It can still alter strategic beliefs and coordinate selection among existing equilibria. Where disclosure changes conduct, the analyst should test whether it changed direct private valuation or cost, made a rule enforceable, altered an expected future consequence, or coordinated equilibrium selection. The environmental-finance and sustainability-disclosure literatures motivate these questions; they do not identify a single universal response route.

**Whistleblower awards** may create a D3 consequence for a person who can supply information. A case-specific analysis must establish the award's terms, the information channel, and the enforcement pathway. **Benchmark reform** can change  $q$  by replacing a manipulable reported input with a transaction-based or otherwise independently verified input. Both claims require application-specific evidence rather than a general inference from the taxonomy.

Several familiar objections become questions for an application. Reputation can operate through D1 plus D2 when customers or counterparties observe a signal and alter demand. Liability is a D3 institution whose measurement, attribution, and enforcement costs must be specified. Competition selects on the payoffs firms face, so an application must show which route links the system coordinate to those payoffs. Externality theory describes wedges and corrective institutions; this paper adds separate input- and price-recovery tests and a payoff-orthogonal disclosure application. Aggregation and Vickrey–Clarke–Groves mechanisms can be defined over broad sets of agents. Their feasibility in a particular system-harm setting turns on representation, identification, measurement, and mechanism design; the present paper does not claim a general impossibility result for them.

## 8. Evidence, with its status stated

This section describes protocols and pending tests. The formal propositions do not depend on any empirical result reported here.

**Cross-domain accounting protocol.** The companion measurement program specifies how a later domain study can estimate an average ratio  $\beta_W = \Delta W / \Pi$ , where  $\Pi$  is annual industry revenue and  $\Delta W$  is a monetized system-welfare loss defined at a stated activity boundary. It requires source records, uncertainty distributions, overlap review, dependence assumptions, and a documented activity boundary. A domain-specific study must define its own applicable channels and cannot infer its result from a screened panel of other domains.

The 61-domain measurement layer remains under its July 2026 release hold. Its values, simulations, propagated intervals, and aggregate results are outside this manuscript. This paper reports neither a sign pattern nor a magnitude. The protocol identifies the evidence a later application would need before it draws an architecture or route inference from a welfare estimate.

**A pending return-side test.** A companion public-data study is designed to examine whether a system-welfare exposure is associated with ordinary-period returns and with return movements at events that may change expected future valuation or rule-imposed consequence. Its data assembly, event classifications, analysis, and replication run remain under independent verification. This paper draws no empirical conclusion from that study. Once released, the study must distinguish the anticipated-consequence explanation from ordinary industry, regulatory, volatility, and event-risk explanations.

**What empirical evidence could and could not do.** A positive measured welfare loss does not establish that a market excludes the welfare term; partial pricing with large residual harm can produce the same observation. No accounting result substitutes for the input enumeration, target-witness test, and P-fiber audit required for an application of Proposition 2.

## 9. Formal failure conditions and application audits

The mathematical statements have counterexample conditions. Empirical applications also have architecture and scope audits. They answer different questions.

**T1 (mediation and input enumeration).** Within an application that fixes  $q$  and  $p$  in advance, an equal- $q$  pair with different realized prices contradicts  $P = p \circ q$ . In practice, such a finding normally identifies an omitted, mismeasured, or bypassed input and defeats that application's enumeration.

**T2 (claimed target recovery).** An equal- $q$ , unequal- $T$  pair refutes target-input and target-price recovery on a domain containing the pair. An equal- $P$ , unequal- $T$  pair refutes target-price recovery. These tests implement Proposition 2's primary fiber conditions.

**T3 (claimed welfare-level recovery).** An equal- $q$ , unequal- $W$  pair refutes welfare-level input recovery. An equal- $P$ , unequal- $W$  pair refutes welfare-level price recovery. These tests address Proposition 2(d), separately from the declared target.

**T4 (continuation-equilibrium selection).** Under Definitions 1–2, exhibit an equilibrium of  $\Gamma^\tau$  whose continuation strategy profile after a positive-probability  $\tau$  lies outside  $E(I)$ . This would refute Proposition 4(b).

**T5 (state-conditional welfare).** Under Proposition 4(c)’s premise, exhibit  $x$  and the selected  $\sigma(\tau) \in E(I)$  with  $w(\sigma(\tau), x) \geq W_0$ . This would refute the proposition’s welfare conclusion.

**A1 (taxonomy coverage).** Identify a causal  $x$ -to- $q$  path in a fully specified application that the four diagnostic routes classify poorly. The finding expands or revises the taxonomy; it does not affect Propositions 1–4.

**A2 (practical reach).** Identify a domain in which the realized bilateral price recovers the declared  $T$  and the named pricing or governance decision performs as intended under a careful outcome study. The finding narrows the framework’s practical motivation in that domain. Target recovery and decision performance must be tested separately because Proposition 2 establishes an informational condition.

## 10. The design implication

The architecture audit yields six practical questions. What decision is being made? What target  $T$  does that decision require? What information about  $T$  enters the complete input? Which target distinctions does the pricing map retain? Whose present valuation or cost changes with the condition? Which current or expected rules make it consequential? A domain study should identify the evidence for each path, its maintenance costs, its failure modes, and the conditions under which it ceases to operate. The companion design and measurement work develops the implementation protocol.

An institution that creates a route can later be weakened or removed. Applications should therefore evaluate measurement, rules, and expectations as maintained infrastructure, with identifiable operating costs and failure modes.

## 11. Limitations and conclusion

The limitations follow from the assumptions. The input- and price-fiber results are domain-specific and only as good as the declared decision, target, input enumeration, and pricing map. Target recovery can be operationally useful while the target is misspecified or poorly measured, so outcome validation remains essential. Exact welfare-level recovery additionally depends on the stated welfare boundary, baseline, and unit. Payoff-orthogonality is demanding: many real system coordinates correlate with private costs, revenues, or beliefs about exogenous states, and Proposition 4 does not apply there. The proposition covers finite Bayesian games; richer strategy spaces, endogenous information acquisition, and dynamic trading require

separate treatment. The taxonomy is a diagnostic organizing device that requires an application-specific model of the price setter and input paths. The accounting and return-side work remains under separate release review and supplies no empirical conclusion here. The paper’s contribution is an applied synthesis of established factorization, public-randomization, sunspot, externality, common-pool, and disclosure ideas.

The conclusion is scope-limited. Equal complete inputs imply equal realized prices. Recovery of a declared decision target requires  $T$  to be constant on  $q$ -fibers; recovery from the price requires  $T$  to be constant on  $P$ -fibers. Price recovery therefore requires both an adequate input architecture and a pricing map that preserves the distinctions the decision needs. Exact recovery of total system welfare applies the same test to  $W$  and is a stronger diagnostic. Payoff-orthogonal public disclosure leaves payoff and feasibility primitives fixed while permitting strategic-belief changes and selection among existing Bayes–Nash equilibria. The four diagnostic routes organize recurring changes in price-forming architecture: measurement, direct private valuation or cost, rule-imposed consequence, and anticipated future valuation or consequence. Each application must establish its decision, target, inputs, pricing map, causal paths, and evidence.

## Statements and declarations

**Funding.** This research received no external funding.

**Competing interests.** The author is founder of the Center for Decision Accounting and has developed software that applies Decision Accounting. The formal results in this paper do not evaluate that software. The author declares no other competing interests.

**Data availability.** The formal propositions require no data. The exact companion theorem source is frozen at <https://decisionaccounting.org/research/mst/mst-v1.3.2-release-20260728.zip>. Related research-program materials are available at <https://decisionaccounting.org/papers/mst/>, <https://decisionaccounting.org/proof/>, and <https://decisionaccounting.org/publication-roadmap/>. The cross-domain and return-side studies discussed in Section 8 remain under separate release review, and this manuscript reports no result from either. The research program is available at <https://decisionaccounting.org/>. A stable page for this paper and any released materials will be added there upon release; no paper-specific replication archive is public at the date of this version.

**Generative AI disclosure.** During the preparation of this manuscript, the author used Anthropic Claude, OpenAI ChatGPT and Codex, and Google Gemini to assist with literature discovery, adversarial review of earlier drafts, formal stress-testing, consistency checking, and editorial revision. The author directed the research program, selected the claims, reviewed the cited sources and derivations, and takes full responsibility for the content.

## References

- Arrow KJ (1969) The organization of economic activity: issues pertinent to the choice of market versus nonmarket allocation. In: *The analysis and evaluation of public expenditure: the PPB system*, vol 1. US Government Printing Office, Washington, DC, pp 47–64
- Aumann RJ (1974) Subjectivity and correlation in randomized strategies. *J Math Econ* 1(1):67–96.  
[https://doi.org/10.1016/0304-4068\(74\)90037-8](https://doi.org/10.1016/0304-4068(74)90037-8)
- Ben-Shahar O, Schneider CE (2014) *More than you wanted to know: the failure of mandated disclosure*. Princeton University Press, Princeton
- Bolton P, Kacperczyk M (2021) Do investors care about carbon risk? *J Financ Econ* 142(2):517–549.  
<https://doi.org/10.1016/j.jfineco.2021.05.008>
- Cass D, Shell K (1983) Do sunspots matter? *J Polit Econ* 91(2):193–227. <https://doi.org/10.1086/261139>
- Chetty R (2009) Sufficient statistics for welfare analysis: a bridge between structural and reduced-form methods. *Annu Rev Econom* 1:451–488. <https://doi.org/10.1146/annurev.economics.050708.142910>
- Christensen HB, Floyd E, Liu LY, Maffett M (2017) The real effects of mandated information on social responsibility in financial reports: evidence from mine-safety records. *J Account Econ* 64(2–3):284–304.  
<https://doi.org/10.1016/j.jacceco.2017.08.001>
- Christensen HB, Hail L, Leuz C (2021) Mandatory CSR and sustainability reporting: economic analysis and literature review. *Rev Account Stud* 26:1176–1248. <https://doi.org/10.1007/s11142-021-09609-5>
- Coase RH (1960) The problem of social cost. *J Law Econ* 3:1–44. <https://doi.org/10.1086/466560>
- d’Aspremont C, Gevers L (1977) Equity and the informational basis of collective choice. *Rev Econ Stud* 44(2):199–209.  
<https://doi.org/10.2307/2297061>
- Drug Enforcement Administration (2019) DEA announces enhanced tool for registered drug manufacturers and distributors to help combat opioid crisis. <https://www.dea.gov/press-releases/2019/02/26/dea-announces-enhanced-tool-registered-drug-manufacturers-and>
- Duffy J, Fisher EO’N (2005) Sunspots in the laboratory. *Am Econ Rev* 95(3):510–529.  
<https://doi.org/10.1257/0002828054201350>
- Geanakoplos J, Polemarchakis H (1986) Existence, regularity, and constrained suboptimality of competitive allocations when the asset market is incomplete. In: Heller WP, Starr RM, Starrett DA (eds) *Uncertainty, information and communication: essays in honor of Kenneth J. Arrow*, vol 3. Cambridge University Press, Cambridge, pp 65–96
- Gordon HS (1954) The economic theory of a common-property resource: the fishery. *J Polit Econ* 62(2):124–142.  
<https://doi.org/10.1086/257497>
- Greenwald BC, Stiglitz JE (1986) Externalities in economies with imperfect information and incomplete markets. *Q J Econ* 101(2):229–264. <https://doi.org/10.2307/1891114>
- Grossman SJ (1981) The informational role of warranties and private disclosure about product quality. *J Law Econ* 24(3):461–483. <https://doi.org/10.1086/466995>
- Grossman SJ, Stiglitz JE (1980) On the impossibility of informationally efficient markets. *Am Econ Rev* 70(3):393–408
- Hardin G (1968) The tragedy of the commons. *Science* 162(3859):1243–1248.  
<https://doi.org/10.1126/science.162.3859.1243>

- HM Treasury (2012) *The Wheatley Review of LIBOR: final report*.  
[https://assets.publishing.service.gov.uk/media/5a7b3fe2e5274a319e77e076/wheatley\\_review\\_libor\\_finalreport\\_280912.pdf](https://assets.publishing.service.gov.uk/media/5a7b3fe2e5274a319e77e076/wheatley_review_libor_finalreport_280912.pdf)
- Kapp KW (1950) *The social costs of private enterprise*. Harvard University Press, Cambridge, MA
- Milgrom PR (1981) Good news and bad news: representation theorems and applications. *Bell J Econ* 12(2):380–391.  
<https://doi.org/10.2307/3003562>
- Muller NZ, Mendelsohn R, Nordhaus W (2011) Environmental accounting for pollution in the United States economy. *Am Econ Rev* 101(5):1649–1675. <https://doi.org/10.1257/aer.101.5.1649>
- Ostrom E (1990) *Governing the commons: the evolution of institutions for collective action*. Cambridge University Press, Cambridge
- Pástor L, Stambaugh RF, Taylor LA (2021) Sustainable investing in equilibrium. *J Financ Econ* 142(2):550–571.  
<https://doi.org/10.1016/j.jfineco.2020.12.011>
- Pigou AC (1920) *The economics of welfare*. Macmillan, London
- Postnieks E (2026a) The Missing System Theorem: the limits of welfare recovery from payoff-space information. Center for Decision Accounting working paper, v1.3.2. Release archive: <https://decisionaccounting.org/research/mst/mst-v1.3.2-release-20260728.zip>. Landing page: <https://decisionaccounting.org/papers/mst/>. Formal proof: <https://decisionaccounting.org/proof/>
- Samuelson PA (1954) The pure theory of public expenditure. *Rev Econ Stat* 36(4):387–389.  
<https://doi.org/10.2307/1925895>
- Sen AK (1977) On weights and measures: informational constraints in social welfare analysis. *Econometrica* 45(7):1539–1572.  
<https://doi.org/10.2307/1913949>
- Starrett DA (1972) Fundamental nonconvexities in the theory of externalities. *J Econ Theory* 4(2):180–199.  
[https://doi.org/10.1016/0022-0531\(72\)90148-2](https://doi.org/10.1016/0022-0531(72)90148-2)
- Electronic Code of Federal Regulations (2026) 21 C.F.R. § 1301.74(b), in § 1301.74, “Other security controls for non-practitioners; narcotic treatment programs and compounders for narcotic treatment programs; mobile narcotic treatment programs.” Current through 24 July 2026.  
<https://www.ecfr.gov/current/title-21/chapter-II/part-1301/subpart-D/section-1301.74>