

Shadow Futures: Contribution Uncertainty and the Self-Reinforcing Market

Martin Erlic

First posted December 2025; revised July 2026

Abstract

Can a market observe productive inputs perfectly and still be unable to learn what those inputs contributed to its rewards? This paper studies adaptive allocation systems in which verified work, effort, measurable quality, cost, foregone income, and capital at risk directly affect reward probabilities, while past reward changes future exposure. I define the market's comparison budget as the cumulative probability mass remaining outside the currently dominant alternative. For a broad class of locally equivalent allocation rules with a common predictable design, finite total comparison makes the complete single-history laws generated by distinct contribution parameters mutually absolutely continuous. No estimator based on one realized market can consistently recover every nonconstant contribution functional, and no test can separate two contribution parameters with vanishing total error. A finite-horizon bound shows that attribution precision is limited by the comparison budget rather than transaction count. Strong reinforcement is a sharp corollary because it exhausts the budget and produces eventual allocation monopoly. With latent position, contribution and position can be exactly observationally equivalent. The unrealized paths needed to separate these explanations are shadow futures. The result does not deny that work or risk matters. It shows that a real causal effect can remain unrecoverable from the path that rewarded it, with implications for platforms, competition policy, entrepreneurship, and merit-sensitive taxation.

Keywords: increasing returns; path dependence; identification; cumulative advantage; monopoly; antitrust; taxation; redistribution

JEL classifications: C13; C18; D31; D43; D83; D85; H21; L41

1. Introduction

A successful firm can arrive with a warehouse of evidence. It can produce payroll records, patents, code commits, invoices, hours worked, capital committed, and risks survived. Every document can be genuine. The missing exhibit is the same firm run again after a different first customer, a different ranking, a different distribution channel, or a different early shock. That firm exists only as an unrealized history.

An increasing share of economic life is organized through systems that turn early outcomes into later access. A freelancer's first rating affects the next client. A creator's first burst of attention affects the next recommendation. A startup's first customer affects financing, recruitment, and distribution. A paper's first placement affects the citations later treated as evidence of scientific importance. People perform real work and accept real risk, but they do so inside reinforced tournaments in which the result of one round changes the odds in the next.

Economics has long understood that increasing returns and cumulative advantage can create concentration, lock-in, and historical selection (Arthur, 1989; David, 1985; Merton, 1968). This paper asks a prior question. Once reward compounds through position, can the realized market history reveal how much of the outcome was caused by work, quality, effort, judgment, capital, or risk?

The distinction is between causal effect and recoverability. A better product may raise the probability of winning every customer. Greater effort may improve every submission. Capital placed at risk may support experimentation and command an ex ante premium. None of this guarantees that one realized history contains enough independent comparison to reveal the magnitude of those effects. A market can keep paying after it has stopped learning.

The information problem is related to, but different from, Akerlof's (1970) market for lemons. In a lemons market, relevant quality exists but is hidden from one side of the trade. Here productive inputs may be public, measured without error, and written into the allocation rule. The missing information is historical rather than private. It consists of the unrealized allocation paths needed to learn how the same inputs would have performed under different early shocks and positions. Better disclosure cannot reveal an experiment that was never run.

I call the resulting problem contribution uncertainty: uncertainty about the causal share of observed productive inputs in a realized reward. Contribution uncertainty is not uncertainty about whether work occurred. A time sheet, audit, proof, or ledger may establish that fact perfectly. It is uncertainty about the counterfactual difference the input made once it passed through a self-reinforcing allocation process.

The paper makes three claims. First, transaction volume is not the relevant sample size for attribution. The relevant object is a comparison budget that records how much probability remained available for the allocation to go another way. Second, when a market consumes only a finite total amount of comparison, no method can consistently recover a nonconstant contribution quantity from its single realized history. Third, competition and policy institutions matter partly because they create independent comparison paths. This makes concentration an epistemic problem even without an appeal to unfairness, deliberate exclusion, or conventional monopoly pricing.

The argument connects several literatures without claiming that their individual ingredients are new. Increasing-returns theory explains lock-in. Reinforced-process theory explains dominance under nonlinear feedback (Pemantle, 2007; Oliveira, 2009). Statistical work links absolute continuity to failures of consistent estimation, and Le Goff and Soulier (2017) establish such a

failure for a reinforced urn. Hayek (2002) treats competition as discovery. The contribution here is the economic bridge: productive inputs can genuinely affect reward while the endogenous allocation process destroys the comparisons needed to recover their contribution, and the same missing comparison constrains competition policy and merit-sensitive taxation.

2. The Self-Reinforcing Market and Its Comparison Budget

There are n alternatives. Before reward $t + 1$ is allocated, alternative i has a verified input profile x_{it} . The profile may include output, hours, costly effort, measurable quality, capital committed, foregone income, or exposure to loss. The profiles may change with the observed history. The contribution parameter β gives the direct loading of those inputs on reward odds. A predictable state index s_{it} records accumulated position, including prior sales, reputation, citations, financing, ranking, or platform placement.

For every β , the experiment is defined on the same raw filtered space, with terminal sigma-field $F_\infty = \sigma(\cup_t F_t)$. The initial law is common, and the input profiles and state indices are common predictable functions of the realized allocation history. All parameter-dependent randomness in the next observed allocation enters through the recipient kernel below. If additional observed variables have parameter-dependent laws, their information must be included separately.

$$\Pr_{\beta}(J_{t+1} = i \mid \mathcal{F}_t) = p_{it}(\beta) = \frac{\exp(x_{it}^T \beta + s_{it})}{\sum_{j=1}^n \exp(x_{jt}^T \beta + s_{jt})}. \quad (1)$$

Verified inputs therefore affect every reward probability directly. The conclusion is not obtained by assuming that work, quality, or risk is irrelevant. The state term permits past reward to change future access.

At date t , let the current leader be any alternative with the largest conditional allocation probability. Define residual contestability as the probability that the next reward goes elsewhere:

$$\varepsilon_t(\beta) = 1 - p_{wt}(\beta) = 1 - \max_i p_{it}(\beta). \quad (2)$$

The comparison budget through date T is the cumulative residual contestability:

$$B_T(\beta) = \sum_{t=0}^{T-1} \varepsilon_t(\beta), \quad B_\infty(\beta) = \sum_{t=0}^{\infty} \varepsilon_t(\beta). \quad (3)$$

Transaction count is T. Comparison budget is B(T). They need not grow together. If the leader receives the next allocation with probability 0.999, another thousand transactions are not equivalent to a thousand open comparisons. The market may be commercially busy while the statistical experiment is nearly exhausted.

The paper studies contribution as a counterfactual functional of β . For example, one may ask how much additional cumulative reward an increase in an agent's verified quality or effort would cause, including the later feedback generated by the initial advantage. Realized profit, market share, citation, or audience size is one draw from that counterfactual system. It is not itself the contribution.

3. The Comparison-Budget Impossibility

In the conditional-logit model, the one-period score is the chosen input profile minus its conditional mean. Its conditional Fisher information is the cross-sectional covariance of the profiles. Let D_x bound the distance between any two profiles. The following inequality gives the paper's basic information accounting rule.

Proposition 1 - Comparison controls information. For every date t,

$$\text{tr} \mathcal{I}_t(\beta) \leq D_x^2 \varepsilon_t(\beta). \quad (4)$$

The conditional mean minimizes expected squared distance. Comparing it with the current leader's profile gives the bound immediately. Summing over dates shows that cumulative

information is bounded by D_x^2 times $B(T)$. An allocation produces attribution evidence only to the extent that it remains capable of going somewhere else.

For two conditional recipient laws, let h_t^2 denote their one-period squared Hellinger distance:

$$h_t^2(\beta, \beta') = \sum_{i=1}^n \left(\sqrt{p_{it}(\beta)} - \sqrt{p_{it}(\beta')} \right)^2.$$

Assume local equivalence, meaning that the one-period laws have the same support after every history. Also assume comparison-dominated separation: for each β and β' there is a finite constant such that

$$h_t^2(\beta, \beta') \leq K_{\beta, \beta', \varepsilon_t}(\beta)$$

for every date and history, with the same condition after reversing β and β' . Bounded conditional logit satisfies this requirement because Hellinger distance is bounded by Kullback-Leibler divergence, and the latter is bounded by residual contestability through Proposition 1.

Theorem 1 - Finite comparison-budget impossibility. Suppose the design is common and predictable, local equivalence and comparison-dominated separation hold, and

$$B_\infty(\beta) < \infty \quad \mathbb{P}_\beta\text{-almost surely} \quad \text{for every } \beta \in \Theta. \quad (5)$$

1. For every $\beta, \beta' \in \Theta$, the complete single-history laws are mutually absolutely continuous on F_∞ .
2. No sequence of estimators based on one realized history can be consistent at every β for any nonconstant contribution functional.
3. No sequence of tests can distinguish two contribution parameters with both type-I and type-II errors converging to zero.
4. No confidence procedure can have asymptotic coverage one at every β while shrinking to a point at every β .

Proof. Finite comparison makes the cumulative conditional Hellinger process finite. The filtered-space Hellinger criterion therefore makes the complete-history law under β absolutely continuous with respect to the law under β' . Reversing the parameters gives equivalence (Jacod and Shiryaev, 2003; Gabrielyan, 2011). If an estimator converges to a constant under one law, equivalence forces the same convergence under the other. It cannot also converge to a different contribution value. The testing and confidence-set claims follow from the same change-of-measure argument. The result is not a complaint about noise or computational cost. It rules out every method on the stated information set. The Technical Appendix gives the full likelihood-ratio, testing, and confidence-set arguments.

The theorem is a failure of universal learning from one endogenous history, not a claim that different parameters generate identical probability distributions. Independent replications can identify the parameter because statistical separation adds across markets. Time extends the inherited path. Replication reopens the laboratory.

The same accounting gives a finite-market result. For the history through date T , conditional logit satisfies

$$D_{\text{KL}}(\mathbb{P}_{\beta}^T \parallel \mathbb{P}_{\beta'}^T) \leq K_{X,\theta} \parallel \beta - \beta' \parallel^2 \mathbb{E}_{\beta} B_T(\beta). \quad (6)$$

A Le Cam two-point bound then implies that local attribution precision cannot improve faster than the inverse square root of the expected comparison budget. If $B(T)$ grows like $\log T$, the best local rate is no faster than 1 divided by the square root of $\log T$, even after T transactions. Exact impossibility is the limiting case; weak learning begins much earlier. The Technical Appendix gives the complete two-point derivation.

Strong reinforcement supplies a sharp primitive case. In a reinforced market, past reward enters through a positive feedback function $g(a + N_i(t))$. If

$$\sum_{m=0}^{\infty} \frac{1}{g(a+m)} < \infty. \quad (7)$$

then an exponential-clock construction yields a finite random time after which one alternative receives every subsequent reward (Oliveira, 2009; Gottfried and Grosskinsky, 2024). The remaining probability outside the winner is summable, so B_{∞} is finite and Theorem 1 applies. For polynomial feedback with exponent ρ , this is the superlinear regime $\rho > 1$. Linear preferential attachment may generate power laws without satisfying the exact impossibility condition. The Technical Appendix gives both the exponential-clock proof and an independent common-support proof.

A separate problem appears when initial position is latent. Let λ_i summarize inherited visibility, distribution access, sponsorship, or reputation. Allocation becomes

$$p_{it}(\beta, \lambda) = \frac{\exp(x_i^{\top} \beta + \lambda_i + s_{it})}{\sum_{j=1}^n \exp(x_j^{\top} \beta + \lambda_j + s_{jt})}. \quad (8)$$

For any admissible vector d , transform the parameters as

$$\beta^{(d)} = \beta + d, \quad \lambda_i^{(d)} = \lambda_i - x_i^{\top} d. \quad (9)$$

The composite index $x_i^{\top} \beta + \lambda_i$ is unchanged, so every observable allocation probability is unchanged. Contribution and position are then exactly observationally equivalent. The single-history theorem and this attribution gauge are distinct: the first is a learning impossibility even with observed position; the second is point non-identification when position is latent.

4. Shadow Futures and the Self-Reinforcing Market

A shadow future is a positive-probability allocation history generated by the same verified inputs, structural parameters, and initial state, but a different sequence of allocation shocks, under which an agent receives a different terminal reward. Shadow futures do not show that

work lacked a causal effect. They are the missing experimental repetitions required to estimate that effect.

Some modern labor and entrepreneurial markets combine real productive effort, repeated entry, and highly skewed rewards. The participant does not merely buy a lottery ticket. The participant builds the ticket. A creator produces a video, a contractor submits a proposal, a developer releases an app, and a founder starts a company. Each action can improve the probability of success. The allocation channel can then convert an early result into future exposure. A recommendation generates views that train the next recommendation. A first customer creates the record needed to acquire the second. A first citation creates the visibility that produces later citations.

The winner may have contributed more. The realized history cannot, by itself, establish how much more. Risk-bearing may justify a prospective premium and society may rationally offer large prizes to induce experimentation (Kerr, Nanda, and Rhodes-Kropf, 2014). It does not follow that the realized jackpot is a sufficient statistic for the winner's causal contribution. Incentive compatibility and retrospective attribution are different questions.

The information absorption barrier begins when the record becomes increasingly self-confirming. The winner's history grows longer at the same time that the marginal evidentiary value of the history shrinks. Volume is not replication.

This motivates a distinction between economic and epistemic monopoly. Economic monopoly concerns durable control over allocation conditions such as price, access, compatibility, or a route to market. Epistemic monopoly concerns control over the production of the comparison paths needed to explain allocation. The phrase has prior uses in knowledge governance (Meghani, 2021); here it is defined narrowly, relative to a stated contribution estimand and

information set. The two can coincide, but neither implies the other. Thousands of nominal sellers may share one ranking and one distribution history. A regulated monopoly could preserve learning through randomized pilots, independent evaluation, and open experimentation.

Competition therefore has an attribution role in addition to its familiar price and discovery roles (Hayek, 2002). Independent distributors, marketplaces, journals, funders, procurement channels, and evaluators are potential replications. They allow similar productive inputs to meet different audiences, rankings, and early shocks. Interoperability, multihoming, portability, and randomized exposure are identification institutions because they keep alternate histories open.

The Music Lab experiments give a concrete analogue. The same songs were placed into parallel cultural markets, and stronger social influence increased inequality and cross-market unpredictability while quality remained only partially determinative (Salganik, Dodds, and Watts, 2006). The experiment does not prove the theorem for labor or entrepreneurship. It shows what a single market deletes: the terminal outcomes the same object would have occupied in other reinforced worlds.

5. Competition Policy and Merit-Sensitive Taxation

The theorem separates measurement from intervention. Better time sheets, audits, credentials, reputation scores, and cryptographic proofs can reduce fraud or hidden action. They do not create the allocation history that would have occurred under another early sequence. Verification proves execution. Attribution requires comparison.

Competition policy can therefore be understood partly as evidence policy. A merger or platform design may reduce the number or independence of routes through which alternatives receive exposure and accumulate records. Two channels that look duplicative in a static cost calculation may be informational complements. Combining them can save fixed costs while eliminating a

control group. The theorem does not make every merger harmful, and the welfare effect depends on decision errors and the cost of preserving independent channels. It identifies a question ordinary concentration measures can miss: how many genuinely independent histories remain available?

The same information limit reaches taxation. Let O denote the observable record, $Y_i(O)$ the realized reward, and $C_i(\theta, O)$ the contribution assigned under structural economy θ . Define positional rent residually as

$$\mathcal{R}_i(\theta, O) = Y_i(O) - C_i(\theta, O). \quad (10)$$

A merit-separating tax would satisfy

$$\tau_i(O) = \mathcal{R}_i(\theta, O) \quad (11)$$

Theorem 2 - No universal merit-separating tax. If two structural economies generate the same observable law but assign different contribution to the same realized reward on a set of positive probability, no tax rule measurable from O can be merit-separating in both.

Proof. The same observable record requires the same tax. The two observationally equivalent economies require different residual rent taxes. One number cannot equal both.

The theorem does not say that all high income is rent, that risk is fictitious, or that taxation is impossible. It says exact desert-sensitive taxation is not merely administratively difficult on this information set. The target itself may be unrecoverable. Policy must then rely on identified sets, observable compounding mechanisms, or institutions that create new variation.

This changes the role of redistribution without uniquely determining a tax rate. A reinforcement-neutral levy can target observable superlinear compounding rather than a disputed moral decomposition. An unconditional social dividend is attribution-invariant because it does not rebuild the same impossible merit ranking on the transfer side. Interoperability, randomized

exposure, public options, structural separation, and independent procurement can go further by preserving the comparisons from which contribution may be learned.

6. Conclusion

A market can allocate reward without being able to explain its allocation. Productive inputs may be perfectly observed and may affect every reward probability directly. Once reward changes future access, however, the market can consume the comparisons needed to estimate how much those inputs contributed.

The comparison budget formalizes the boundary. When total comparison is finite, distinct contribution parameters generate equivalent complete-history laws. No estimator, test, or shrinking confidence procedure can recover a nonconstant contribution quantity from one realized path. When comparison grows slowly, attribution remains weak long before literal absorption. Strong reinforcement is one sharp way to exhaust the budget; latent position can make contribution and position exactly observationally equivalent.

Akerlof's lemons problem turns on hidden quality. The shadow-futures problem turns on missing history. Competition matters because it can generate the alternate paths that one self-reinforcing market deletes. Tax policy faces the same constraint when it attempts to separate earned contribution from positional rent.

Markets produce rewards. They also produce, or fail to produce, the evidence by which those rewards are later explained. A self-reinforcing market can keep paying after it has stopped learning.

References

- Akerlof, G. A. (1970). The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3), 488-500.
- Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394), 116-131.
- David, P. A. (1985). Clio and the economics of QWERTY. *American Economic Review*, 75(2), 332-337.
- Gabrielyan, S. S. (2011). Absolute continuity and singularity of two probability measures on a filtered space. *Journal of Theoretical Probability*, 24(3), 595-614.
- Gottfried, T., and Grosskinsky, S. (2024). Asymptotics of generalized Pólya urns with non-linear feedback. *Electronic Journal of Probability*, 29, 1-56.
- Hayek, F. A. (2002). Competition as a discovery procedure. *The Quarterly Journal of Austrian Economics*, 5(3), 9-23.
- Jacod, J., and Shiryaev, A. N. (2003). *Limit Theorems for Stochastic Processes* (2nd ed.). Springer.
- Kerr, W. R., Nanda, R., and Rhodes-Kropf, M. (2014). Entrepreneurship as experimentation. *Journal of Economic Perspectives*, 28(3), 25-48.
- Le Cam, L., and Yang, G. L. (2000). *Asymptotics in Statistics: Some Basic Concepts* (2nd ed.). Springer.
- Le Goff, L. C., and Soulier, P. (2017). Parameter estimation of a two-colored urn model class. *The International Journal of Biostatistics*, 13(1), Article 20160029.
- Meghani, Z. (2021). Regulations matter: Epistemic monopoly, domination, patents, and the public interest. *Philosophy & Technology*, 34(4), 1449-1474.

- Merton, R. K. (1968). The Matthew effect in science. *Science*, 159(3810), 56-63.
- Oliveira, R. I. (2009). The onset of dominance in balls-in-bins processes with feedback. *Random Structures & Algorithms*, 34(4), 454-477.
- Pemantle, R. (2007). A survey of random processes with reinforcement. *Probability Surveys*, 4, 1-79.
- Piketty, T., Saez, E., and Stantcheva, S. (2014). Optimal taxation of top labor incomes: A tale of three elasticities. *American Economic Journal: Economic Policy*, 6(1), 230-271.
- Rothschild, C., and Scheuer, F. (2016). Optimal taxation with rent-seeking. *The Review of Economic Studies*, 83(3), 1225-1262.
- Salganik, M. J., Dodds, P. S., and Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science*, 311(5762), 854-856.
- van de Rijt, A., Kang, S. M., Restivo, M., and Patil, A. (2014). Field experiments of success-breeds-success dynamics. *Proceedings of the National Academy of Sciences*, 111(19), 6934-6939.

Technical Appendix

This appendix supplies the complete mathematical details for the information bound, the finite comparison-budget theorem, the finite-horizon lower bound, and the strong-reinforcement result. It also records the short proofs of the auxiliary replication, contribution-position, and tax results stated in the main article.

Canonical experiment and notation

All probability laws are defined on the same raw filtered measurable space $(\Omega, \mathcal{F}^\infty, (\mathcal{F}_t)_{t \geq 0})$, where $\mathcal{F}^\infty = \sigma(\cup_{t \geq 0} \mathcal{F}_t)$. The sigma-fields are not completed separately under different parameter values, and the initial law on $\mathcal{F}_0$ is common across parameters.

The verified-input profiles x_{it} and state indices s_{it} are common predictable functions of the canonical past. On a fixed realized history, their values do not change when β changes. All parameter-dependent randomness in the next observed allocation enters through the conditional recipient kernel $p_t(\beta)$. If additional observed covariates or state innovations have β -dependent conditional laws, their conditional Hellinger increments must be included in the statistical experiment.

Appendix A. Information Bounds for Conditional Logit

A.1 Fisher information

Let

$$\mu_t(\beta) = \sum_i p_{it}(\beta) x_{it}.$$

For any fixed vector a ,

$$\text{trVar}_{p_t(\beta)}(x_{J,t}) = \mathbb{E}_{p_t(\beta)} \|x_{J,t} - \mu_t(\beta)\|^2 \leq \mathbb{E}_{p_t(\beta)} \|x_{J,t} - a\|^2.$$

Choose $a = x_{w_t t}$, where w_t maximizes $p_{it}(\beta)$. Then

$$\text{tr} \mathcal{I}_t(\beta) \leq \sum_{j \neq w_t} p_{jt}(\beta) \|x_{jt} - x_{w_t t}\|^2 \leq D_X^2 \varepsilon_t(\beta),$$

which proves Proposition 1's information bound.

A.2 KL divergence

Write $\delta = \beta' - \beta$ and

$$A_t(u) = \log \sum_i \exp(x_{it}^\top u + s_{it}).$$

The conditional KL divergence is the Bregman divergence

$$D_{\text{KL}}(p_t(\beta) \| p_t(\beta')) = A_t(\beta') - A_t(\beta) - \nabla A_t(\beta)^\top \delta.$$

Taylor's theorem gives

$$D_{\text{KL}}(p_t(\beta) \| p_t(\beta')) = \int_0^1 (1-r) \delta^\top \text{Var}_{p_t(\beta+r\delta)}(x_{J,t}) \delta \, dr. \quad (\text{A1})$$

Let w_t maximize $p_{it}(\beta)$. Along the exponential tilt,

$$\frac{1 - p_{w_t t}(\beta + r\delta)}{p_{w_t t}(\beta + r\delta)} \leq \exp(D_X \|\delta\|) \frac{\varepsilon_t(\beta)}{1 - \varepsilon_t(\beta)}.$$

Because $1 - \varepsilon_t(\beta) = \max_i p_{it}(\beta) \geq 1/n$,

$$1 - \max_i p_{it}(\beta + r\delta) \leq n \exp(D_X \text{diam}(\Theta)) \varepsilon_t(\beta).$$

Apply the information bound at the intermediate parameter in (A1), bound the quadratic form by its trace, and integrate $1 - r$ over $[0,1]$. This gives the KL inequality used in equation (6) of the main text.

Finally,

$$\log \frac{p_{J,t}(\beta')}{p_{J,t}(\beta)} = \delta^\top x_{J,t} - [A_t(\beta') - A_t(\beta)].$$

The bracketed term is constant conditional on $\mathcal{F}_t$. Its conditional variance is therefore

$$\delta^\top \mathcal{I}_t(\beta) \delta \leq \|\delta\|^2 D_X^2 \varepsilon_t(\beta),$$

which proves the auxiliary likelihood-ratio variance bound.

Appendix B. Proof Details for Theorem 1

Fix $\beta, \beta' \in \Theta$. Write $Q = \mathbb{P}_\beta$ and $P = \mathbb{P}_{\beta'}$, and let Q_T and P_T denote their restrictions to $\mathcal{F}_T$. By local equivalence, $Q_T \sim P_T$ for every finite T . Define the finite-horizon likelihood ratio

$$Z_T = \frac{dQ_T}{dP_T} = \prod_{t=0}^{T-1} \frac{p_{J_{t+1},t}(\beta)}{p_{J_{t+1},t}(\beta')} \quad (\text{B1})$$

Let

$$R_{t+1} = \frac{Z_{t+1}}{Z_t}.$$

Conditional on $\mathcal{F}_t$,

$$\mathbb{E}_P[\sqrt{R_{t+1}} \mid \mathcal{F}_t] = \sum_{i=1}^n \sqrt{p_{it}(\beta)p_{it}(\beta')} \quad (\text{B2})$$

Hence

$$2(1 - \mathbb{E}_P[\sqrt{R_{t+1}} \mid \mathcal{F}_t]) = \sum_{i=1}^n \left(\sqrt{p_{it}(\beta)} - \sqrt{p_{it}(\beta')} \right)^2 = h_t^2(\beta, \beta') \quad (\text{B3})$$

In this discrete-time, locally equivalent filtered experiment, the order-one-half Hellinger criterion states that

$$Q \ll P \quad \Leftrightarrow \quad \sum_{t=0}^{\infty} h_t^2(\beta, \beta') < \infty \quad Q\text{-almost surely} \quad (\text{B4})$$

This is the discrete-time specialization of the filtered-space absolute-continuity criterion in Jacod and Shiryaev (2003, Chapter IV) and Gabrielyan (2011). Assumption 1 and the finite-budget condition (5) in Theorem 1 give

$$\sum_{t=0}^{\infty} h_t^2(\beta, \beta') \leq K_{\beta, \beta'} B_{\infty}(\beta) < \infty \quad \mathbb{P}_{\beta}\text{-almost surely.}$$

Applying (B4) with $Q = \mathbb{P}_{\beta}$ and $P = \mathbb{P}_{\beta'}$ yields $\mathbb{P}_{\beta} \ll \mathbb{P}_{\beta'}$. Reversing β and β' , and using $B_{\infty}(\beta') < \infty$ $\mathbb{P}_{\beta'}$ -almost surely, gives $\mathbb{P}_{\beta'} \ll \mathbb{P}_{\beta}$. Thus

$$\mathbb{P}_{\beta} \sim \mathbb{P}_{\beta'} \quad \text{on } \mathcal{F}_{\infty} \quad (\text{B5})$$

For the estimator conclusion, let $\hat{\psi}_T$ be $\mathcal{F}_T$ -measurable and suppose $\hat{\psi}_T \rightarrow \psi(\beta)$ in $\mathbb{P}_{\beta}$ -probability. For $\epsilon > 0$, define

$$A_T(\epsilon) = \{d(\hat{\psi}_T, \psi(\beta)) > \epsilon\}.$$

Then $\mathbb{P}_{\beta}(A_T(\epsilon)) \rightarrow 0$. Since $\mathbb{P}_{\beta'} \ll \mathbb{P}_{\beta}$, absolute continuity implies $\mathbb{P}_{\beta'}(A_T(\epsilon)) \rightarrow 0$. The same estimator therefore converges to $\psi(\beta)$ under β' and cannot also be consistent for a different value $\psi(\beta')$.

For testing, let the $\mathcal{F}_T$ -measurable variable $\phi_T = 1$ denote rejection of β in favor of β' . If $\mathbb{P}_{\beta}(\phi_T = 1) \rightarrow 0$, absolute continuity implies $\mathbb{P}_{\beta'}(\phi_T = 1) \rightarrow 0$. Hence $\mathbb{P}_{\beta'}(\phi_T = 0) \rightarrow 1$, so the type-II error tends to one.

For confidence sets, choose β, β' with

$$\Delta = d(\psi(\beta), \psi(\beta')) > 0.$$

Suppose the $\mathcal{F}_T$ -measurable random sets C_T have coverage tending to one at every parameter and $\text{diam}_d(C_T) \rightarrow 0$ in probability at every parameter. Under β , define

$$E_T = \left\{ \psi(\beta) \in C_T, \text{diam}_d(C_T) < \frac{\Delta}{2} \right\}.$$

Then $\mathbb{P}_\beta(E_T) \rightarrow 1$. Equivalence implies $\mathbb{P}_{\beta'}(E_T) \rightarrow 1$. But on E_T , the set C_T cannot contain $\psi(\beta')$, contradicting coverage under β' . This proves the confidence-set conclusion.

Appendix C. Finite-Horizon Lower Bound

The chain rule for conditional KL divergence gives

$$D_{\text{KL}}(\mathbb{P}_\beta^T \parallel \mathbb{P}_{\beta'}^T) = \mathbb{E}_\beta \sum_{t=0}^{T-1} D_{\text{KL}}(p_t(\beta) \parallel p_t(\beta')).$$

The finite-horizon KL bound (6) follows from Proposition 1 and the chain rule. Le Cam's two-point inequality gives

$$\max_{\vartheta \in \{\beta, \beta'\}} \mathbb{P}_\vartheta(|\hat{\psi}_T - \psi(\vartheta)| \geq \Delta_\psi/2) \geq \frac{1 - \|\mathbb{P}_\beta^T - \mathbb{P}_{\beta'}^T\|_{\text{TV}}}{2}.$$

Pinsker's inequality gives

$$\|\mathbb{P}_\beta^T - \mathbb{P}_{\beta'}^T\|_{\text{TV}} \leq \sqrt{\frac{1}{2} D_{\text{KL}}(\mathbb{P}_\beta^T \parallel \mathbb{P}_{\beta'}^T)},$$

which yields the two-point lower bound stated in Section 3 of the main article.

For Corollary 1, choose

$$\beta_T = \beta_0 + \frac{hv}{\sqrt{b_T}}$$

for sufficiently small fixed h . Differentiability gives

$$|\psi(\beta_T) - \psi(\beta_0)| = \frac{|h \nabla \psi(\beta_0)^\top v|}{\sqrt{b_T}} + o(b_T^{-1/2}).$$

The finite-horizon KL bound keeps the KL divergence bounded by a constant proportional to h^2 . Choosing h small enough keeps total variation bounded away from one. The two-point inequality then yields a local risk lower bound of order $b_T^{-1/2}$.

Appendix D. Strong Reinforcement

Consider a finite set of agents with fixed verified-input profiles x_i . Let

$$R_{i,t+1} = \mathbf{1}\{J_{t+1} = i\}, \quad N_i(t) = \sum_{r=1}^t R_{i,r}, \quad N_i(0) = 0. \quad (D1)$$

Let $a > 0$ and let g be positive and nondecreasing. Conditional allocation is

$$p_i(\beta) = \frac{\exp(x_i^\top \beta) g(a + N_i(t))}{\sum_{j=1}^n \exp(x_j^\top \beta) g(a + N_j(t))}. \quad (D2)$$

Assume strong reinforcement:

$$\sum_{m=0}^{\infty} \frac{1}{g(a + m)} < \infty. \quad (D3)$$

Write

$$q_i = \exp(x_i^\top \beta) > 0.$$

For each agent i and occupancy level $m \geq 0$, let $E_{i,m}$ be independent exponential random variables with rate $q_i g(a + m)$, and define

$$T_i = \sum_{m=0}^{\infty} E_{i,m}.$$

Condition (D3) gives

$$\mathbb{E}[T_i] = \sum_{m=0}^{\infty} \frac{1}{q_i g(a + m)} < \infty,$$

so $T_i < \infty$ almost surely. The variables T_i are independent and continuously distributed, so there is almost surely a unique minimizer I^* .

The exponential-clock construction has the same embedded recipient sequence as equation (D2).

Agent I^* experiences infinitely many clock rings by time T_{I^*} . Every other agent has only finitely many rings before that time. Hence all but finitely many discrete rewards go to I^* .

Let n_j denote the final count of loser j and let the winner's post-absorption count grow as m .

Then

$$\varepsilon_t(\beta) \leq \frac{\sum_{j \neq I^*} q_j g(a + n_j)}{q_{I^*} g(a + m)}.$$

Summing over m and using condition (D3) gives $B_\infty(\beta) < \infty$ almost surely.

Direct common-support proof for strong reinforcement. Let $\mathcal{M}$ be the countable set of eventually constant recipient histories. The absorption argument above gives $\mathbb{P}_\beta(\mathcal{M}) = 1$ for every finite β . Fix $\omega \in \mathcal{M}$. Suppose that after a finite prefix ending at date T , the history assigns every later reward to agent i , and let $n_j = N_j(T)$.

The finite prefix has strictly positive probability under every finite β . Conditional on that prefix, the probability that agent i receives every later reward is

$$\prod_{m=0}^{\infty} \frac{q_i g(a + n_i + m)}{q_i g(a + n_i + m) + \sum_{j \neq i} q_j g(a + n_j)}.$$

Let

$$C_{\beta, \omega} = \sum_{j \neq i} q_j g(a + n_j).$$

The tail probability becomes

$$\prod_{m=0}^{\infty} \left[1 + \frac{C_{\beta, \omega}}{q_i g(a + n_i + m)} \right]^{-1}.$$

Condition (D3) implies

$$\sum_{m=0}^{\infty} \frac{c_{\beta,\omega}}{q_i g(a + n_i + m)} < \infty.$$

The infinite-product criterion therefore makes the tail probability strictly positive. Hence $\mathbb{P}_{\beta}(\{\omega\}) > 0$ for every $\omega \in \mathcal{M}$ and every finite β . Since every $\mathbb{P}_{\beta}$ assigns probability one to the same countable set and positive probability to every atom in that set, the strong-reinforcement laws have the same positive atomic support and are mutually absolutely continuous. This independently proves the strong-reinforcement result and its estimator and testing consequences stated in the main article.

Appendix E. Auxiliary Results

E.1 Independent replication

Suppose M independent markets share contribution parameter β and each is observed for a fixed horizon L . If $\beta \mapsto P^{\beta}(J_1, \dots, J_L)$ is injective and Θ is compact, the uniform law of large numbers makes the normalized log likelihood converge to an expectation uniquely maximized at the truth by Gibbs' inequality. The maximum-likelihood estimator is therefore consistent as $M \rightarrow \infty$.

E.2 Contribution-position invariance

In the latent-position model, substitute $\beta + d$ for β and $\lambda_i - x_i^T d$ for λ_i . Then $x_i^T(\beta + d) + \lambda_i - x_i^T d = x_i^T \beta + \lambda_i$. Every conditional allocation probability and every observable history law is unchanged, so the data identify only the composite index.

E.3 Merit-separating taxation

If two structural economies generate the same observable law but assign different contribution to the same reward, any record-measurable tax must choose the same liability in both. Exact rent

taxation instead requires liabilities $Y_i - C_i$, which differ. One rule cannot satisfy both requirements, proving the result in Section 5.