

Screening and the Price of a Stationary Type Distribution

Dalin Sheng*

May 2026

Abstract

A stationary mechanism uses a protocol to screen agents and form replacement types. We characterize implementable allocation and distribution pairs and output costs of cohort composition. With three trials, implementation is equivalent to a monotone effective allocation and three moment constraints. Finitely many coordinates suffice exactly when kernel differences are state independent and have finite span; its dimension is the flow rank. At fixed mean, the second moment measures variance. Incentive compatibility requires a minimum; additional variance improves sorting until saturation. The third moment controls its placement. Targets preserving the two-moment value form an interval; departures incur a location tax. Under locally stable IC constraints and independent moment gradients, this tax is locally quadratic. Its coefficient equals the second-moment shadow price per unit of relocation capacity times an IC amplification factor, which exceeds one exactly when IC blocks a first-order change preserving the first two moments.

Keywords: Stationary screening; mechanism design; cross-cohort type formation; stationary type distributions; flow rank; location tax.

JEL classification: C61; D02; D82.

*Southwestern University of Finance and Economics, Chengdu 611130, China. Email: shengdl@swufe.edu.cn. This work was supported by the National Natural Science Foundation of China (grant number 72503198).

1 Introduction

Mechanism design typically treats the distribution of private types as a primitive. In many institutions that operate across cohorts, the rule applied today also shapes the population that arrives tomorrow. A work-and-training protocol determines how an employee’s skill is used today and what skills her replacements acquire; an educational track affects a student’s payoff and the skill distribution of the next cohort. The same policy is then an allocation rule for the current population and a transition rule for the future population. Implementation is therefore defined over an allocation–distribution pair (q, μ) : one policy must screen the current population and reproduce the next one.

This paper characterizes which target pairs can be sustained by a stationary mechanism when one report-contingent protocol governs both the slope of current utility and the distribution of next-period types. For sustainable targets, it also solves the institution’s output-maximizing protocol assignment. Incentive compatibility orders mean protocol intensity across reports, while stationarity of the type distribution restricts the aggregate moments of assigned protocol intensity. Feasibility and optimal design therefore depend on how both restrictions apply to the same policy.

Professional training makes this joint role concrete. In a surgical residency, the operative role assigned to a trainee governs both the current use of her skill and the experience transmitted through the training process. Accreditation also requires each graduating resident to complete a specified profile of cases across defined categories.¹ The German dual system similarly combines production with formation: firms appoint qualified in-company trainers, while occupation-specific training regulations organize the competencies assessed in final examinations.²

The model has a unit mass of short-lived agents in every period. An agent privately observes an inherited type, reports it, and receives a material allocation, a protocol, and a transfer. The institution takes the material-allocation rule q as fixed and must preserve a target type distribution μ . Protocol intensity $t \in \mathcal{T} \subset [0, 1]$ sets the slope of the agent’s current payoff and the success probability in each Bernoulli formation trial that produces the next cohort. The protocol set is any closed subset of $[0, 1]$ with at least two elements. The full protocol interval $[0, 1]$ contains every technologically available success probability; a finite protocol set restricts the institution to finitely many protocol intensities.

In the baseline technology, a protocol generates three independent Bernoulli formation trials. If the type distributions conditional on zero through three successes are affinely independent, preserving μ is equivalent to fixing the first three moments of assigned protocol intensity. Current incentives depend on the conditional mean intensity. The envelope characterization for one-dimensional screening makes truthful reporting equivalent to monotonicity of the effective allocation $q(z)[\alpha + \beta \mathbb{E}(T \mid z)]$. Population stationarity fixes unconditional moments of protocol intensity, and incentive compatibility (IC) orders its conditional mean across reports.

The conflict already appears in a two-moment example. Let types be uniform on $[0, 1]$, set $\alpha = \beta = 1$,

¹Current ACGME requirements set a minimum of 850 major cases, require residents to document their operative experience in the Case Log System, and specify separate category minimums, including alimentary tract, abdominal, and vascular surgery (ACGME, 2026, 2024).

²Pedagogical aptitude is usually demonstrated under the Ausbildereignungsverordnung through its examination, recognized equivalent qualifications, or an exemption granted by the competent body. Final certificates are issued by competent bodies, usually chambers, under occupation-specific training regulations. The 2026 BIBB register lists 324 recognized occupations (Federal Republic of Germany, 2009; Federal Institute for Vocational Education and Training, 2024, 2026).

take $q(z) = 1/(1 + \frac{4}{3}z)$, and require protocol mean and second moment $(m, k) = (1/2, 1/4)$. The stationary identities imply

$$\mathbb{E}T = \frac{1}{2}, \quad \mathbb{E}T^2 = \frac{1}{4} \implies T = \frac{1}{2} \text{ almost surely.}$$

The resulting effective allocation is strictly decreasing, so it violates incentive compatibility. Among incentive-compatible schedules with mean $1/2$, the unique minimum-dispersion schedule is $r(z) = (1 + 12z)/14$, and its second moment is $61/196 > 1/4$. The gap $12/196$ is the minimum increase in the second moment required to offset the decline in q . The conflict is already present in the first two moment constraints, so changing the target third moment cannot restore feasibility.

The moments enter the full design problem in sequence. At a fixed mean, the second protocol moment measures *protocol variance*. Incentive compatibility requires a minimum amount of this variance to offset the shape of the mandated allocation. The remainder supports productive sorting across types until the saturation threshold; additional protocol variance is output neutral. The third protocol moment then determines how that variance is allocated across types and cohorts formed after different numbers of successes. In the three-trial technology it equals the steady-state share formed after all three successes. Holding the first two moments fixed, changing it reallocates cohort shares in the proportions $(-1, 3, -3, 1)$. The associated output loss is the *location tax* imposed by the residual cohort-composition target.

Theorem 1 derives this design program from the mechanism. An allocation–distribution pair is implementable exactly when its effective allocation is monotone and its protocol assignment satisfies the three moment constraints. The joint dual places the shadow prices on these constraints and the multiplier on the IC order restriction in one pointwise protocol comparison; the design problems below therefore inherit their shadow prices from the full mechanism problem. This finite reduction is intrinsic to the transition technology. More generally, finitely many aggregate protocol coordinates determine the next population distribution exactly when differences between protocol kernels do not depend on inherited type and their span is finite-dimensional. The minimal dimension is the *flow rank*, and minimal coordinates are unique up to an invertible affine transformation. Requiring finitely many coordinates to determine the next distribution only for ordered policies gives the same condition. With N independent formation trials, the first N power moments are minimal whenever the protocol set contains at least $N + 1$ intensities.

The first two moments determine how protocol variance is used. For any feasible mean and closed protocol set, output is concave in the second moment, increases strictly below an endogenous saturation threshold when that region is nonempty, and is flat above it. At and below saturation, every optimum exhausts the second-moment constraint. In that region an interval protocol set yields a unique deterministic schedule, while a finite set implements conditional means by mixing adjacent protocols. With the full protocol interval and a declining material-allocation rule, protocol variance divides into the minimum required by IC, an output-increasing range, and an output-neutral range. Along $q_\delta(z) = 1/(1 + \delta z)$, a steeper decline raises the IC variance requirement and lowers saturation, shrinking the range available for productive sorting.

The third-moment constraint selects among policies with the prescribed first two moments. The third moments that preserve the two-moment value form the image of its optimal face, hence a compact zero-loss interval. This interval is a singleton when $\mathcal{T} = [0, 1]$ and the second moment lies strictly between its minimum and saturation. Feasible target values outside it incur the location tax.

The same stationarity and IC multipliers price the tax. At a two-moment optimum strictly below

saturation, an exact loss identity charges schedule displacement and conditional protocol variance at the shadow price of the second-moment constraint and movement away from the active IC face at the IC multiplier. Along any sequence with quadratic loss, conditional randomization changes the third moment only at second order, so the leading movement lies on the face selected jointly by these prices. When the active IC constraints are locally unchanged, the tax is locally quadratic, and its coefficient factors as

$$\kappa^{\text{IC}} = \kappa^0 \Omega^{\text{IC}}, \quad \Omega^{\text{IC}} \geq 1.$$

Here κ^0 prices third-moment movement under the first two moment constraints without IC, while Ω^{IC} is the ratio of that relocation capacity to the capacity remaining on the active IC face. The amplification factor exceeds one exactly when incentive compatibility removes a first-order direction that changes the third moment while preserving the first two. In an exact benchmark with a declining allocation rule, about 8.25% of the relocation capacity available without IC remains, and the local tax coefficient is multiplied by about 12.12. The shadow price measures the marginal value of protocol variance, while the amplification factor measures how much IC reduces the capacity to reallocate it.

The results use progressively stronger assumptions. The representation theorem covers compact protocol sets with jointly Feller kernels, and the two-moment design theorem covers every closed scalar protocol set. The variance decomposition assumes the full protocol interval and declining allocation rules. The local tax formula also requires local stability of the active IC constraints and independent moment gradients.

Related literature. The closest conceptual connection is dynamic mechanism design with endogenous types. [Pavan \(2017\)](#) develops a framework in which allocations inside a continuing relationship can shape a participant’s future private information, as under learning by doing, experimentation, or habit formation. [Athey and Segal \(2013\)](#) construct an efficient dynamic mechanism, [Pavan et al. \(2014\)](#) give a Myersonian characterization of dynamic incentives, and [Bergemann and Välimäki \(2019\)](#) synthesize work on persistent private information and intertemporal contracting. Here type formation occurs across cohorts. Each agent reports once and exits, while her assigned protocol changes the type distribution of her replacements. Strategic incentives are therefore one-period, and the protocol must also reproduce the target population distribution. These two requirements generate moment constraints, a saturation threshold, and the location tax. Replacement also connects the mechanism problem to intergenerational welfare and mobility ([Phelan, 2006](#); [Farhi and Werning, 2007](#)), stationary distributions in dynamic economies ([Hopenhayn and Prescott, 1992](#)), and firm-sponsored skill formation ([Acemoglu and Pischke, 1998](#)).

A second connection is stationary dynamic market design, which [Che \(2026\)](#) organizes around reductions from stochastic systems to stationary programs. Within this approach, [Baccara et al. \(2020\)](#) characterize dynamic matching under the trade-off between market thickness and waiting costs. [Che and Tercieux \(2026\)](#) jointly design queue entry, priority, and information and establish the optimality of first-come-first-served service in a rich domain. [Che and Choi \(2025\)](#) study auctions with queue- and inventory-dependent admission. [Madsen and Shmaya \(2025\)](#) analyze collective upkeep under participation and reporting constraints and derive an optimum with at most two membership tiers. These stationary programs control queues, inventories, admission, or upkeep. The stationary program in this paper controls a target distribution of future private types produced by the same report-contingent protocol that screens

the current population. Flow rank establishes when aggregate moments of protocol use are sufficient statistics for that distribution, and the location tax prices the remaining distributional shape.

The analysis also draws on scalar screening and envelope methods (Mussa and Rosen, 1978; Myerson, 1981; Milgrom and Segal, 2002), implementation duality (Nöldeke and Samuelson, 2018), dual approaches to feasible distributions and posterior moments (Kolotilin, 2018; Dworczak and Martini, 2019; Dworczak and Kolotilin, 2024), and occupation measures for constrained Markov control (Hernández-Lerma and Lasserre, 1996; Altman, 1999). The first theorem places the multipliers on stationarity and IC in one dual problem at the protocol level, while the flow-rank converse identifies the transition technologies that admit an exact finite representation. The design results use isotonic regression (Barlow et al., 1972), moment systems (Karlin and Studden, 1966), and the extreme-point geometry of monotone functions ordered by majorization (Kleiner et al., 2021). Economically, stationarity of the type distribution determines the feasible protocol moments, and incentive compatibility restricts how they can be allocated across types. Together they produce the saturation threshold, the variance decomposition, and the IC amplification factor in the location tax.

Section 2 presents the model, implementation with a stationary type distribution, and canonical flow coordinates. Section 3 solves the two-moment relaxation for any closed protocol set. Section 4 establishes the variance decomposition and its comparative statics. Section 5 adds the third-moment constraint and develops the location tax. Section 6 concludes.

2 Model and Stationary Screening

2.1 Protocol set, payoffs, and formation

Time is discrete. In each period a unit mass of short-lived agents privately observes an inherited type $z \in \mathcal{Z} = [\underline{z}, \bar{z}]$. The target type distribution μ , which the mechanism must preserve, is atomless and has full support. Each agent participates once, reports her type, receives a material allocation, a protocol, and a transfer, and exits.

The institution fixes a positive continuous material-allocation rule $q : \mathcal{Z} \rightarrow (0, \bar{q}]$. A protocol has intensity

$$t \in \mathcal{T}, \quad \emptyset \neq \mathcal{T} \subset [0, 1] \text{ closed}, \quad a := \min \mathcal{T} < b := \max \mathcal{T}. \quad (1)$$

The full protocol interval is $\mathcal{T} = [0, 1]$. An interval protocol set $\mathcal{T} = [a, b]$ has no gaps over its attainable range. A finite $\mathcal{T}$ restricts the institution to finitely many intensities.

If report r receives protocol t , allocation $q(r)$, and signed payment $p(r)$, a true type z obtains

$$zq(r)(\alpha + \beta t) - p(r), \quad \beta > 0, \quad \alpha + \beta a > 0. \quad (2)$$

Terms independent of the true type can be absorbed into the unrestricted transfer, and the outside option is normalized to zero. Because (2) is affine in t , a lottery over protocols affects current incentives only through its conditional mean. A more training-intensive protocol widens the set of trainees who benefit from an agent's knowledge, making current output more sensitive to her type. This channel also appears in knowledge hierarchies, where the reach of individual knowledge is a design variable (Garicano, 2000).

Protocol intensity is also the success probability in each of three independent Bernoulli formation

trials. Let R_j be the distribution of next-period types conditional on exactly $j \in \{0, 1, 2, 3\}$ successes. The protocol-specific replacement distribution is

$$\widehat{Q}_t = (1-t)^3 R_0 + 3t(1-t)^2 R_1 + 3t^2(1-t) R_2 + t^3 R_3. \quad (3)$$

In the training reading, each replacement is formed through three stages, for instance successive rotations or apprenticeship years. Thus $[0, 1]$ in (1) contains every feasible success probability. The three formation trials are a fixed feature of the baseline environment.

Fix target moments (m, k, m_3) . The shares of the next cohort formed after zero through three successes are

$$(b_0, b_1, b_2, b_3) = (1 - 3m + 3k - m_3, 3(m - 2k + m_3), 3(k - m_3), m_3). \quad (4)$$

Assume these coefficients are probabilities and define

$$\mu = \sum_{j=0}^3 b_j R_j. \quad (5)$$

Equivalently, the target moments are

$$\mathbb{E}T = m, \quad \mathbb{E}T^2 = k, \quad \mathbb{E}T^3 = m_3. \quad (6)$$

These moments also determine the distribution of formation successes. Let J be the number of successful trials for a member of the next cohort. Then $\mathbb{E}J = 3m$, $\mathbb{E}[J(J-1)] = 6k$, and $\mathbb{E}[J(J-1)(J-2)] = 6m_3$. Thus m , k , and m_3 are the success rates for a randomly selected trial, ordered pair of trials, and ordered triple of trials. In particular, $b_3 = m_3$ is the steady-state share formed after three successes, the mandated share of fully trained members in the training interpretation.

Assume that $R_1 - R_0, R_2 - R_0, R_3 - R_0$ are linearly independent signed measures, equivalently that $R_0, \dots, R_3$ are affinely independent. Let $S_0(dz' | z)$ be a common Markov kernel satisfying $\mu S_0 = \mu$, and fix $\zeta \in [0, 1)$. The transition under protocol t is

$$Q_t(dz' | z) = \zeta S_0(dz' | z) + (1 - \zeta) \widehat{Q}_t(dz'). \quad (7)$$

The persistent component may depend on inherited type and is common to every protocol. Protocol choice affects population formation only through the component generated by the three independent trials.

2.2 Stationary direct mechanisms

A stationary direct mechanism consists of a Borel kernel $\sigma(dt | r)$ on $\mathcal{T}$ and a Borel payment $p(r)$. Define the conditional mean intensity and the incentive-relevant effective allocation by

$$r_\sigma(z) := \int_{\mathcal{T}} t \sigma(dt | z), \quad (8)$$

$$g_\sigma(z) := q(z)[\alpha + \beta r_\sigma(z)]. \quad (9)$$

The mechanism *implements q while preserving μ* if truthful reporting is globally incentive compatible and individually rational and the truthful transition leaves μ stationary.

The allocation rule q and population distribution μ are mandated, as when accreditation fixes the distribution of cases a residency must deliver or chamber examinations fix the profile an occupation must reproduce.

Let $\mathcal{M}_\uparrow$ be the cone of essentially bounded functions with a nondecreasing representative. The dual cone associated with this monotonicity restriction is

$$\mathcal{H}_\mu := \left\{ \xi \in L^1(\mu) : \mathbb{E}\xi = 0, \quad X_\xi(x) := \int_{[\underline{z}, x]} \xi(z) \mu(dz) \geq 0 \text{ for every } x \right\}. \quad (10)$$

For a Carathéodory integrand $\ell(z, t)$ satisfying

$$\int_{\mathcal{Z}} \sup_{t \in \mathcal{T}} |\ell(z, t)| \mu(dz) < \infty,$$

consider its minimum over protocol policies that can be combined with transfers to implement q while preserving μ . Thus ℓ is an element of $L^1(\mu; C(\mathcal{T}))$. Denote the value by $C_{m,k,m_3}(\ell)$.

Theorem 1 (Implementation, stationarity, and duality). *Under the preceding assumptions, the following statements hold.*

- (i) *A pointwise protocol policy can be combined with transfers to implement q while preserving μ if and only if its effective allocation is nondecreasing at every report and its moments satisfy the three equalities below. Equivalently, modulo μ -null modifications, a policy is feasible if and only if*

$$g_\sigma \in \mathcal{M}_\uparrow, \quad \mathbb{E}T = m, \quad \mathbb{E}T^2 = k, \quad \mathbb{E}T^3 = m_3. \quad (11)$$

Every almost-everywhere witness has a Borel null-set modification that defines such a pointwise direct mechanism. For any pointwise version, truthful payments are

$$V(z) = V_0 + \int_{\underline{z}}^z g_\sigma(x) dx, \quad (12)$$

$$p(z) = zg_\sigma(z) - V(z), \quad (13)$$

and individual rationality is equivalent to $V_0 \geq 0$.

- (ii) *If (11) is feasible, then*

$$\begin{aligned} C_{m,k,m_3}(\ell) = & \sup_{\substack{\eta, \lambda, \theta \in \mathbb{R} \\ \xi \in \mathcal{H}_\mu}} \left\{ -\eta m - \lambda k - \theta m_3 \right. \\ & \left. + \int \left[\alpha q(z) \xi(z) + \min_{t \in \mathcal{T}} \left\{ \ell(z, t) + [\eta + \beta q(z) \xi(z)] t \right. \right. \right. \\ & \quad \left. \left. \left. + \lambda t^2 + \theta t^3 \right\} \right] \mu(dz) \right\}. \end{aligned} \quad (14)$$

There is no duality gap.

(iii) For output, take $\ell(z, t) = -w(z)t$, where $w(z) := zq(z)$. The dual variables η, λ, θ are the shadow prices on the first three moment constraints. The variable ξ is the multiplier on the IC order restriction. Whenever a dual maximizer $(\eta^*, \lambda^*, \theta^*, \xi^*)$ exists, every primal optimizer σ^* satisfies, for μ -almost every z ,

$$\sigma^* \left(\arg \min_{s \in \mathcal{T}} \{ \theta^* s^3 + \lambda^* s^2 + [\eta^* - w(z) + \beta q(z) \xi^*(z)] s \} \mid z \right) = 1. \quad (15)$$

The theorem derives the moment program from the transition technology. Integrating (7) under any policy gives

$$\begin{aligned} \mu P_\sigma &= \zeta \mu + (1 - \zeta) \{ [1 - 3\mathbb{E}T + 3\mathbb{E}T^2 - \mathbb{E}T^3] R_0 \\ &\quad + 3[\mathbb{E}T - 2\mathbb{E}T^2 + \mathbb{E}T^3] R_1 \\ &\quad + 3[\mathbb{E}T^2 - \mathbb{E}T^3] R_2 + \mathbb{E}T^3 R_3 \}. \end{aligned} \quad (16)$$

Affine independence makes equality with (5) equivalent to the three moment constraints. The same conditional distribution of T determines g_σ , so stationarity of μ and incentive compatibility cannot be assigned to separate policies.

2.3 Minimal flow rank and canonical coordinates

The reduction rests on a scalar conditional mean summarizing current payoff responsiveness and finitely many unconditional protocol coordinates summarizing population flow. On the payoff side, let τ, ρ be continuous protocol indices, with τ nonconstant. Equality of $\int \tau d\pi$ implies equality of $\int \rho d\pi$ for every pair of protocol lotteries if and only if

$$\rho(t) = \alpha_0 + \beta_0 \tau(t) \quad \text{for every } t \in \mathcal{T}. \quad (17)$$

Thus an affine payoff link is necessary for a single conditional protocol index to carry all current incentive effects.

For population flow, let $\{\tilde{Q}_t : t \in \mathcal{T}\}$ be a *jointly Feller* family: for every $\varphi \in C(\mathcal{Z})$, the map $(z, t) \mapsto \int \varphi(z') \tilde{Q}_t(dz' \mid z)$ is continuous. For a continuous coordinate map $x : \mathcal{T} \rightarrow \mathbb{R}^d$ and a Borel policy σ , define

$$\bar{x}_\sigma := \int_{\mathcal{Z}} \int_{\mathcal{T}} x(t) \sigma(dt \mid z) \mu(dz), \quad (18)$$

$$\mu \tilde{P}_\sigma := \int_{\mathcal{Z}} \int_{\mathcal{T}} \tilde{Q}_t(\cdot \mid z) \sigma(dt \mid z) \mu(dz). \quad (19)$$

The map x is *flow sufficient* if $\bar{x}_\sigma = \bar{x}_{\sigma'}$ implies $\mu \tilde{P}_\sigma = \mu \tilde{P}_{\sigma'}$ for all Borel policies. It is *order-restricted flow sufficient* if the same implication is required only when both conditional first coordinates $\int x_1(t) \sigma(dt \mid z)$ and $\int x_1(t) \sigma'(dt \mid z)$ are nondecreasing. A finite-dimensional flow-sufficient coordinate is *minimal* if no flow-sufficient coordinate has smaller affine dimension $\dim \text{aff } x(\mathcal{T})$.

Fix $t_0 \in \mathcal{T}$. If every difference $\tilde{Q}_t(\cdot \mid z) - \tilde{Q}_{t_0}(\cdot \mid z)$ is independent of z , denote it by δ_t and define the

flow rank

$$\text{rank}_F(\tilde{Q}) := \dim \text{span}\{\delta_t : t \in \mathcal{T}\}. \quad (20)$$

The dimension in (20) may be infinite. If some difference depends on z , set $\text{rank}_F(\tilde{Q}) = +\infty$. The definition does not depend on the reference protocol.

Theorem 2 (Minimal flow-rank representation). *Let $\mathcal{T}$ be compact, let μ be atomless with full support on the compact interval $\mathcal{Z}$, and let $\{\tilde{Q}_t : t \in \mathcal{T}\}$ be jointly Feller. Suppose $x : \mathcal{T} \rightarrow \mathbb{R}^d$ is continuous, $\text{aff } x(\mathcal{T}) = \mathbb{R}^d$, and x_1 is nonconstant.*

(i) *The following statements are equivalent:*

(a) *x is flow sufficient.*

(b) *x is order-restricted flow sufficient.*

(c) *For every reference protocol t_0 , there are unique signed measures $\eta_1, \dots, \eta_d$. They are independent of inherited state, have total mass zero, and satisfy*

$$\tilde{Q}_t(\text{d}z' \mid z) = \tilde{Q}_{t_0}(\text{d}z' \mid z) + \sum_{i=1}^d [x_i(t) - x_i(t_0)] \eta_i(\text{d}z') \quad (21)$$

for every $t \in \mathcal{T}$ and $z \in \mathcal{Z}$.

(ii) *Whenever these statements hold, $\text{rank}_F(\tilde{Q}) \leq d$. If the flow rank is a finite number r , then r is the smallest affine dimension of any continuous flow-sufficient protocol coordinate, where redundant affine coordinates are discarded. For $r \geq 1$, any two minimal coordinate maps $a, b : \mathcal{T} \rightarrow \mathbb{R}^r$ satisfy*

$$b(t) = Aa(t) + c \quad \text{for every } t \in \mathcal{T} \quad (22)$$

for some nonsingular matrix A and vector c .

(iii) *Let $a : \mathcal{T} \rightarrow \mathbb{R}^r$ be minimal and let $\bar{a} \in \text{co } a(\mathcal{T})$. The centered kernel*

$$\tilde{K}_{\bar{a}} := \int_{\mathcal{T}} \tilde{Q}_t \pi(\text{d}t), \quad \int_{\mathcal{T}} a(t) \pi(\text{d}t) = \bar{a}, \quad (23)$$

is independent of the representing lottery π . If $\mu \tilde{K}_{\bar{a}} = \mu$, then

$$\mu \tilde{P}_{\sigma} = \mu \iff \bar{a}_{\sigma} = \bar{a}. \quad (24)$$

Flow rank makes the dimension of a moment reduction intrinsic. State-independent kernel differences are necessary and sufficient for finitely many aggregate statistics of protocol use to determine the next population distribution, and requiring the implication only among ordered policies does not weaken the condition. The number of independent aggregate restrictions needed for stationarity equals the affine dimension in which protocols move the transition kernel, and minimal coordinates are unique up to an invertible affine change of variables. When a protocol change affects formation differently across inherited types, stationarity also depends on how protocol use is allocated across reports, beyond its unconditional

distribution.

Proof idea. Integrating the kernel-difference representation proves sufficiency. For necessity, zero-mean state-dependent swaps perturb a strictly ordered baseline policy between arbitrary protocols. Order-restricted flow sufficiency then makes every difference between two protocol kernels orthogonal to all zero-mean state functions, hence independent of inherited state. Finite-dimensional linear algebra gives minimality and affine uniqueness.

Corollary 1 (Independent formation trials). *Suppose $N \geq 1$ and a protocol generates N independent Bernoulli formation trials,*

$$\widehat{Q}_t^N = \sum_{j=0}^N \binom{N}{j} t^j (1-t)^{N-j} R_j, \quad Q_t^N = \zeta S_0 + (1-\zeta) \widehat{Q}_t^N, \quad (25)$$

where $\zeta \in [0, 1)$, the common kernel S_0 does not depend on the protocol, and $R_0, \dots, R_N$ are affinely independent. If $\mathcal{T}$ contains at least $N + 1$ distinct intensities, then

$$\text{rank}_F(Q^N) = N, \quad x_N(t) := (t, t^2, \dots, t^N) \quad (26)$$

is a minimal flow-sufficient coordinate system. Every other minimal system is an invertible affine transformation of x_N . Hence, for any target moment vector $\bar{m} = (m_1, \dots, m_N)$ whose centered kernel preserves μ ,

$$\mu P_\sigma = \mu \iff \mathbb{E}_\sigma T^j = m_j, \quad j = 1, \dots, N. \quad (27)$$

In the baseline with three trials, whenever the protocol set contains at least four distinct intensities (in particular under the full protocol interval), Corollary 1 selects (t, t^2, t^3) as the unique minimal protocol-only flow coordinates up to affine relabeling. A finite protocol set with fewer than $N + 1$ points has the lower rank of the restricted moment curve, so redundant moment constraints drop out automatically.

All four multipliers in (14) act inside one protocol comparison, and when IC binds its multiplier remains a separate term in (15) alongside the shadow prices on the moment constraints. The analysis below first evaluates the two-moment relaxation at third-moment multiplier $\theta = 0$ to identify how protocol variance is used. Section 5 then imposes the actual target m_3 .

3 Protocol Variance and Saturation

The full target distribution fixes (m, k, m_3) . Conditioning on the first two moments (m, k) , while retaining exact incentive compatibility for the current population distribution μ , isolates how those moments govern protocol assignment. Adding the third-moment target completes the full stationary design problem. Formally, define

$$\Sigma_2(m, k) := \{\sigma : g_\sigma \in \mathcal{M}_\uparrow, \mathbb{E}T = m, \mathbb{E}T^2 = k\}, \quad (28)$$

$$\Sigma_3(m, k, m_3) := \{\sigma \in \Sigma_2(m, k) : \mathbb{E}T^3 = m_3\}. \quad (29)$$

Theorem 1 identifies $\Sigma_3(m, k, m_3)$ as the protocol policies that implement q while preserving μ . The set $\Sigma_2(m, k)$ defines the *two-moment relaxation*: it fixes m and k while leaving m_3 unrestricted, or equivalently sets $\theta = 0$ in the joint dual. We suppress (m, k) from Σ_2 when there is no ambiguity. In cohort terms, the relaxation fixes the aggregate success rates for a randomly selected trial and ordered pair of trials while leaving the ordered-triple success rate unrestricted.

Protocol policies allocate the second moment between schedule variation across types and randomization within reports. The program's objective is current output

$$\mathbb{E}\{Zq(Z)[\alpha + \beta T]\} = \alpha \mathbb{E}w + \beta \mathbb{E}[wT]. \quad (30)$$

Only the second term varies across policies. Every policy that also satisfies the third-moment constraint is feasible for this program, so its value is an upper bound on the full value. Section 5 imposes m_3 and characterizes exactly when that bound is attained and, otherwise, the resulting loss.

3.1 Moment geometry with a closed protocol set

For $r \in [a, b]$, let

$$\mathcal{P}_{\mathcal{T}}(r) := \left\{ \pi \in \Delta(\mathcal{T}) : \int t \pi(dt) = r \right\} \quad (31)$$

and define

$$\underline{K}_{\mathcal{T}}(r) := \min_{\pi \in \mathcal{P}_{\mathcal{T}}(r)} \int t^2 \pi(dt), \quad (32)$$

$$\overline{K}_{\mathcal{T}}(r) := \max_{\pi \in \mathcal{P}_{\mathcal{T}}(r)} \int t^2 \pi(dt). \quad (33)$$

Write

$$t_-(r) := \max(\mathcal{T} \cap [a, r]), \quad t_+(r) := \min(\mathcal{T} \cap [r, b]). \quad (34)$$

Lemma 1 (Second-moment bounds with a closed protocol set). *For every $r \in [a, b]$,*

$$\underline{K}_{\mathcal{T}}(r) = \begin{cases} r^2, & r \in \mathcal{T}, \\ r^2 + [r - t_-(r)][t_+(r) - r], & r \notin \mathcal{T}, \end{cases} \quad \overline{K}_{\mathcal{T}}(r) = (a + b)r - ab. \quad (35)$$

The lower lottery is uniquely degenerate at $r \in \mathcal{T}$ and otherwise uniquely mixes the two bracketing protocols. The upper lottery uniquely mixes a and b . The lower envelope is continuous and convex; it is strictly convex on the whole interval if and only if $\mathcal{T} = [a, b]$.

For any Borel mean schedule $r(z)$, an aggregate second moment k can be generated if and only if

$$\mathbb{E}\underline{K}_{\mathcal{T}}(r(Z)) \leq k \leq \mathbb{E}\overline{K}_{\mathcal{T}}(r(Z)). \quad (36)$$

Whenever this condition holds, atomlessness permits a policy using at most two protocols at every type.

The last statement is an aggregate Carathéodory claim; its pointwise analogue fails. On $\{0, 1/2, 1\}$, conditional mean $1/2$ and conditional second moment $3/8$ require all three protocols. Atomlessness

restores binary support by assigning lower- and upper-envelope lotteries to different type sets while preserving the prescribed conditional means.

3.2 Feasibility, output-increasing variance, and saturation

Define the incentive-compatible schedule set

$$\mathcal{D}_{q,\mathcal{T}} := \{r \in L^2(\mu) : a \leq r \leq b, q(\alpha + \beta r) \text{ has a nondecreasing representative}\}, \quad (37)$$

$$C_{q,\mathcal{T}}(m) := \{r \in \mathcal{D}_{q,\mathcal{T}} : \mathbb{E}r = m\}. \quad (38)$$

For a schedule r , its minimum required second moment is

$$\mathcal{K}_{\mathcal{T}}(r) := \mathbb{E}\underline{K}_{\mathcal{T}}(r(Z)). \quad (39)$$

Assume $C_{q,\mathcal{T}}(m) \neq \emptyset$, and put

$$L_{q,\mathcal{T}}(m) := \min_{r \in C_{q,\mathcal{T}}(m)} \mathcal{K}_{\mathcal{T}}(r), \quad (40)$$

$$B_{q,\mathcal{T}}(m, k) := \max_{\substack{r \in C_{q,\mathcal{T}}(m) \\ \mathcal{K}_{\mathcal{T}}(r) \leq k}} \mathbb{E}[wr]. \quad (41)$$

Let

$$\mathcal{M}_{q,\mathcal{T}}(m) := \arg \max_{r \in C_{q,\mathcal{T}}(m)} \mathbb{E}[wr], \quad (42)$$

$$k_{q,\mathcal{T}}^{\text{sat}}(m) := \min_{r \in \mathcal{M}_{q,\mathcal{T}}(m)} \mathcal{K}_{\mathcal{T}}(r). \quad (43)$$

The saturation threshold is the smallest protocol second moment that permits an output maximizer subject only to the mean constraint.

Theorem 3 (Two-moment design with a closed protocol set). *Suppose $C_{q,\mathcal{T}}(m) \neq \emptyset$.*

(i) *The two-moment relaxation is feasible if and only if*

$$L_{q,\mathcal{T}}(m) \leq k \leq \overline{K}_{\mathcal{T}}(m). \quad (44)$$

On this interval the policy problem is exactly (41), and current output is $\alpha \mathbb{E}w + \beta B_{q,\mathcal{T}}(m, k)$.

(ii) *The map $k \mapsto B_{q,\mathcal{T}}(m, k)$ is continuous, concave, and nondecreasing. If $L_{q,\mathcal{T}}(m) < k_{q,\mathcal{T}}^{\text{sat}}(m)$, its strictly increasing and constant regions are, respectively,*

$$\begin{aligned} & [L_{q,\mathcal{T}}(m), k_{q,\mathcal{T}}^{\text{sat}}(m)], \\ & [k_{q,\mathcal{T}}^{\text{sat}}(m), \overline{K}_{\mathcal{T}}(m)]. \end{aligned}$$

If the thresholds coincide, the whole feasible interval is flat.

(iii) At and below saturation, the second-moment constraint binds for every optimal schedule:

$$\mathcal{K}_{\mathcal{T}}(r) = k. \quad (45)$$

Every optimal policy then uses the unique conditional lower-envelope lottery. At every feasible k , some optimal policy uses at most two protocols at each type.

(iv) If $\mathcal{T} = [a, b]$, then $\underline{K}_{\mathcal{T}}(r) = r^2$. The minimum-dispersion schedule is unique, and the set of mean-constrained output maximizers has a unique minimum-norm member. At and below saturation the optimal schedule r_k is unique and the optimal policy is deterministic:

$$T = r_k(Z) \quad \text{almost surely.} \quad (46)$$

For $L_{q,[a,b]}(m) < k < k_{q,[a,b]}^{\text{sat}}(m)$, there is $\lambda > 0$ such that the first equality below holds. If, in addition, $m \in \text{ri}\{\mathbb{E}r : r \in \mathcal{D}_{q,[a,b]}\}$, there is an $\eta \in \mathbb{R}$ such that the second equality holds:

$$r_k = \text{proj}_{\mathcal{C}_{q,[a,b]}(m)} \left(\frac{w}{2\lambda} \right) \quad (47)$$

$$= \text{proj}_{\mathcal{D}_{q,[a,b]}} \left(\frac{w - \eta}{2\lambda} \right), \quad \mathbb{E}r_k = m. \quad (48)$$

The multiplier λ is the shadow price on the second-moment constraint. Whenever the joint two-moment dual is attained, the same value can be used for its second-moment multiplier.

(v) If $\mathcal{T} = \{t_1 < \dots < t_J\}$, then $\underline{K}_{\mathcal{T}}$ is the piecewise-affine interpolation of t_j^2 . At and below saturation every optimal policy uses the adjacent protocols bracketing its conditional mean. The optimal schedule and the minimum-dispersion member of the mean-constrained optimal set need not be unique.

Interval and finite protocol sets share the same value structure. On an interval, strict convexity yields a deterministic optimal assignment and a Hilbert-space representation. On a finite protocol set, the adjacent-protocol lottery implements each conditional mean. Below saturation the second-moment constraint binds, and additional protocol variance raises output through finer sorting. Above saturation, the second moment already permits an output maximizer subject only to the mean constraint, and output depends on a policy only through its conditional means, so conditional randomization absorbs the excess without changing output.

Suppose $\mathcal{T} = [a, b]$, $L_{q,\mathcal{T}}(m) < k < k_{q,\mathcal{T}}^{\text{sat}}(m)$, and the joint two-moment dual is attained. Let (η, λ, ξ) be an optimal multiplier vector with $\lambda > 0$. Specializing (15) to $\theta = 0$ and combining it with determinism gives, for μ -almost every z ,

$$r_k(z) = \Pi_{[a,b]} \left(\frac{w(z) - \eta - \beta q(z)\xi(z)}{2\lambda} \right). \quad (49)$$

The schedule is therefore shaped jointly by the allocation rule's output weight, the shadow prices on the first two moment constraints, and the IC multiplier. Section 4 shows how a declining allocation rule divides protocol variance into the minimum required by IC, the range over which output rises, and an output-neutral range. Section 5 uses this identity to locate the binding blocks of the optimal schedule.

4 Declining Allocations and the Use of Protocol Variance

On the full protocol interval $\mathcal{T} = [0, 1]$, the two-moment relaxation has a deterministic optimal policy through saturation, so the second moment is the squared norm of its conditional-mean schedule. In this section and the next, abbreviate

$$\mathcal{D}_q := \mathcal{D}_{q,[0,1]}, \quad L_q(m) := L_{q,[0,1]}(m), \quad k_q^{\text{sat}}(m) := k_{q,[0,1]}^{\text{sat}}(m).$$

Since every schedule with mean m has $m^2 \leq \mathbb{E}r^2 \leq m$, the target k selects a point in a total variance range of $m - m^2$. Decompose that range as

$$m - m^2 = \underbrace{[L_q(m) - m^2]}_{C_q(m)} + \underbrace{[k_q^{\text{sat}}(m) - L_q(m)]}_{P_q(m)} + \underbrace{[m - k_q^{\text{sat}}(m)]}_{N_q(m)}. \quad (50)$$

The three terms are the minimum variance required by IC, the range over which output rises, and the output-neutral range. Feasibility requires at least $C_q(m)$. By Theorem 3, the next $P_q(m)$ units raise output strictly, and only units beyond saturation enter $N_q(m)$.

For declining allocation rules, the three components have closed-form expressions. Suppose throughout this section that $\alpha > 0$, q is nonincreasing, and put

$$h(z) := \frac{1}{q(z)}, \quad H := \mathbb{E}h, \quad S := \alpha + \beta m, \quad \bar{g} := \frac{S}{H}. \quad (51)$$

The global IC-feasible interval is

$$A := \alpha q(\underline{z}), \quad B := (\alpha + \beta)q(\bar{z}), \quad (52)$$

and define the distribution obtained by reweighting μ by $1/q$:

$$\nu(dz) := \frac{h(z)}{H} \mu(dz). \quad (53)$$

Proposition 1 (Variance decomposition under declining allocations). *Suppose q is nonincreasing.*

(i) *The IC-feasible interval is nonempty if and only if $A \leq B$. Given this condition, mean m is feasible if and only if $\bar{g} \in [A, B]$. The unique minimum-dispersion schedule and its second moment are*

$$r^L(z) = \frac{\bar{g}h(z) - \alpha}{\beta} = \frac{Sh(z)/H - \alpha}{\beta}, \quad (54)$$

$$L_q(m) = m^2 + \frac{S^2}{\beta^2} \left(\frac{\mathbb{E}h^2}{H^2} - 1 \right). \quad (55)$$

(ii) *Suppose $A < B$, $\bar{g} \in (A, B)$, and $w(z) = zq(z)$ is strictly increasing. There is a unique cutoff $z^* \in (\underline{z}, \bar{z})$ satisfying*

$$\nu((z^*, \bar{z}]) = \frac{\bar{g} - A}{B - A}. \quad (56)$$

The output maximizer subject only to the mean constraint is

$$g^*(z) = A + (B - A)\mathbf{1}\{z > z^*\}, \quad (57)$$

$$\hat{r}_q(z) = \frac{h(z)g^*(z) - \alpha}{\beta}, \quad (58)$$

and the saturation threshold is

$$k_q^{\text{sat}}(m) = \frac{1}{\beta^2} \left[A^2 \mathbb{E} h^2 + (B^2 - A^2) \mathbb{E} \{h^2 \mathbf{1}(Z > z^*)\} - 2\alpha S + \alpha^2 \right]. \quad (59)$$

(iii) Under the conditions in part (ii), the output-increasing range has the decomposition

$$P_q(m) = \frac{H}{\beta^2} \{ \mathbb{E}_\nu [h(g^* - \bar{g})^2] + 2\bar{g} \text{Cov}_\nu(h, g^*) \} > 0. \quad (60)$$

If, in addition, q is strictly decreasing, then all three components in (50) are strictly positive. In the corresponding feasible boundary cases:

$$q \text{ constant} \implies (C_q, P_q, N_q) = (0, m(1 - m), 0); \quad (61)$$

$$A = B \text{ and hence } \bar{g} = A \implies P_q(m) = 0, \quad L_q(m) = k_q^{\text{sat}}(m). \quad (62)$$

The shape of the allocation rule affects feasibility and productivity through different schedules. The minimum-dispersion schedule in (54) keeps the effective allocation constant and uses the minimum possible protocol variance to offset the decline in q . The saturation schedule assigns the upper effective-allocation bound to the highest output weights under the reweighted distribution ν . The covariance term in (60) is nonnegative because a declining allocation rule makes both h and the optimally sorted g^* nondecreasing in type.

At the boundary of the IC-feasible interval

$$\frac{q(\underline{z})}{q(\bar{z})} = \frac{\alpha + \beta}{\alpha}, \quad (63)$$

the feasible effective allocation is forced to be constant. The minimum variance required by IC then reaches the saturation threshold: additional protocol variance can only be placed in conditional randomization and leaves current output unchanged.

The reciprocal family below gives strict comparative statics along a one-parameter path of targets.

Corollary 2 (A steeper decline raises the IC variance requirement). *Let Z be uniform on $[0, 1]$, set $\alpha = \beta = 1$, $m = 1/2$, and*

$$q_\delta(z) := \frac{1}{1 + \delta z}, \quad \delta \in [0, 1]. \quad (64)$$

Put $s = \sqrt{1 + \delta}$. The minimum and saturation schedules are

$$r_\delta^L(z) = \frac{1 - \delta + 3\delta z}{2 + \delta}, \quad (65)$$

$$\widehat{r}_\delta(z) = \begin{cases} \delta z, & z < z_\delta^*, \\ \frac{1 - \delta + 2\delta z}{1 + \delta}, & z > z_\delta^*, \end{cases} \quad z_\delta^* = \frac{1}{1 + s}. \quad (66)$$

Their squared norms are

$$L_\delta = \frac{\delta^2 + \delta + 1}{(\delta + 2)^2}, \quad (67)$$

$$k_\delta^{\text{sat}} = \frac{s^3 - s^2 - s + 4}{3s(s + 1)}. \quad (68)$$

As δ rises, L_δ increases strictly, k_δ^{sat} decreases strictly, and the output-increasing range decreases strictly. The three components move continuously from

$$(C_0, P_0, N_0) = \left(0, \frac{1}{4}, 0\right) \quad \text{to} \quad (C_1, P_1, N_1) = \left(\frac{1}{12}, 0, \frac{1}{6}\right). \quad (69)$$

A larger δ raises the minimum variance compatible with truth telling, lowers the point at which productive sorting saturates, and expands the output-neutral range. A steeper decline tilts the minimum-dispersion schedule more sharply to hold the effective allocation constant, and it narrows the IC-feasible interval, so less spread remains available for sorting.

5 The Third Moment and the Location Tax

The baseline model contains three formation trials and the three moment constraints implied by stationarity of the type distribution in (6). Sections 3–4 solved the two-moment relaxation. The target m_3 is the steady-state share of members formed after all three trials succeed. Imposing it completes the stationary design problem. Holding the first two moments fixed, an increase ε in m_3 changes the cohort shares in (4) by

$$\varepsilon(-1, 3, -3, 1), \quad (70)$$

raising the shares formed after one and three successes at the expense of those formed after zero and two. The third-moment constraint therefore governs the incidence of three successes: for a fixed amount of protocol variance, it determines how that variance is allocated across cohorts and along the ordered protocol set.

5.1 The zero-loss set

Recall the two-moment feasible set Σ_2 from (28). Throughout this section, fix a feasible pair of first and second moments (m, k) , so that $\Sigma_2(m, k) \neq \emptyset$; this set is compact and convex. For $\sigma \in \Sigma_2$, define

$$Y(\sigma) := \mathbb{E}\{Zq(Z)[\alpha + \beta T]\}, \quad M_3(\sigma) := \mathbb{E}T^3. \quad (71)$$

Put

$$Y_2 := \max_{\sigma \in \Sigma_2} Y(\sigma), \quad \Sigma_2^* := \arg \max_{\sigma \in \Sigma_2} Y(\sigma), \quad (72)$$

$$I_3 := M_3(\Sigma_2), \quad I_3^* := M_3(\Sigma_2^*). \quad (73)$$

For $x \in I_3$, the value after imposing the third-moment constraint is

$$Y_3(x) := \max_{\sigma \in \Sigma_3(m, k, x)} Y(\sigma). \quad (74)$$

If $m_3 \in I_3$, then $Y_3(m_3)$ is the value of the full problem; if $m_3 \notin I_3$, that problem is infeasible. Values at other $x \in I_3$ give the value function for a counterfactual third-moment target evaluated at the current type distribution and output weights.

Proposition 2 (Third-moment zero-loss set). *The sets I_3 and I_3^* are nonempty compact intervals, and Y_3 is concave and upper semicontinuous. Moreover,*

$$Y_3(x) = Y_2 \iff x \in I_3^*. \quad (75)$$

Suppose every two-moment optimizer has the same conditional-mean schedule r^ . Define*

$$L_{\mathcal{T}, r^*}(k) := \min\{\mathbb{E}T^3 : \mathbb{E}[T | Z] = r^*(Z), \mathbb{E}T^2 = k\}, \quad (76)$$

$$U_{\mathcal{T}, r^*}(k) := \max\{\mathbb{E}T^3 : \mathbb{E}[T | Z] = r^*(Z), \mathbb{E}T^2 = k\}. \quad (77)$$

Then

$$I_3^* = [L_{\mathcal{T}, r^*}(k), U_{\mathcal{T}, r^*}(k)]. \quad (78)$$

If the two-moment optimal policy σ^ is unique, the interval is the singleton $\{\mathbb{E}_{\sigma^*}T^3\}$; if that policy is deterministic, this value equals $\mathbb{E}(r^*)^3$.*

Equation (78) reads the zero-loss set off the two-moment solution: it is the interval of third moments that the optimal schedule can span by relocating its conditional variance across types. When $\mathcal{T} = [0, 1]$ and the second moment lies strictly between $L_q(m)$ and $k_q^{\text{sat}}(m)$, Theorem 3 makes the two-moment policy unique and deterministic, so

$$I_3^* = \{m_3^\dagger\}, \quad m_3^\dagger := \mathbb{E}(r^*)^3. \quad (79)$$

For an interval protocol set, the optimal policy remains unique through saturation. Above saturation, output-neutral conditional variance can make the set of optimal policies, and hence I_3^* , non-singleton.

With a finite protocol set, nonuniqueness can occur already at the threshold.

5.2 Shadow prices and the local tax when IC binds

For a closed convex set $C \subset L^2(\mu)$, we use the outward-normal convention

$$N_C(x) := \{v : \langle v, y - x \rangle \leq 0 \text{ for every } y \in C\}.$$

The tangent cone $T_C(x)$ is the closure of the radial cone $\{s(y - x) : y \in C, s \geq 0\}$.

Define the location tax and the third-moment penalized value by

$$\Delta_3(x) := Y_2 - Y_3(x), \quad (80)$$

$$\Psi_3(\theta) := \max_{\sigma \in \Sigma_2} \{Y(\sigma) - \theta M_3(\sigma)\}. \quad (81)$$

In the interpretation based on the number of successes, $\Delta_3(x)$ is the output cost of meeting the residual cohort-composition requirement x while the success rates for a randomly selected trial and ordered pair of trials remain fixed. For a feasible full target, the full value is

$$Y_3(m_3) = Y_2 - \Delta_3(m_3). \quad (82)$$

Theorem 4 (The location tax). *The following statements hold.*

(i) For every $x \in I_3$,

$$Y_3(x) = \inf_{\theta \in \mathbb{R}} \{\Psi_3(\theta) + \theta x\}. \quad (83)$$

For every $x \in \text{ri } I_3$, the infimum is attained. Whenever a minimizer exists, every minimizing θ is positive when $x < \inf I_3^*$ and negative when $x > \sup I_3^*$. At every $x \in I_3^*$, $\theta = 0$ is a minimizer; at an endpoint of this interval, additional nonzero minimizing shadow prices may coexist.

(ii) Suppose $\mathcal{T} = [0, 1]$ and $L_q(m) < k < k_q^{\text{sat}}(m)$, and let r^* be the two-moment optimum. Suppose there exist $\eta \in \mathbb{R}$, $\lambda > 0$, and $n \in N_{\mathcal{D}_q}(r^*)$ satisfying

$$w - \eta \mathbf{1} - 2\lambda r^* = n. \quad (84)$$

For any $x \in I_3$ and $\sigma \in \Sigma_3(m, k, x)$, put $d = \mathbb{E}[T \mid Z] - r^*$ and $V = \mathbb{E} \text{Var}(T \mid Z)$. Its output loss satisfies the exact identity

$$\frac{Y_2 - Y(\sigma)}{\beta} = \lambda(\|d\|_2^2 + V) - \langle n, d \rangle. \quad (85)$$

If the normal takes the IC form $n = \beta q \xi$, the last term is $-\langle \xi, g_\sigma - g_{r^*} \rangle$. Both terms are nonnegative. Writing $a_* := 3 \|(r^*)^2\|_2$, every feasible displacement ε also satisfies

$$\Delta_3(m_3^\dagger + \varepsilon) \geq \beta \lambda \left[\frac{\sqrt{a_*^2 + 16|\varepsilon|} - a_*}{8} \right]^2. \quad (86)$$

Hence $m_3^\dagger$ is the only zero-loss third moment under these conditions.

(iii) Under the conditions of part (ii), let

$$\mathcal{K}^* := T_{\mathcal{D}_q}(r^*) \cap \mathbf{1}^\perp \cap (r^*)^\perp \cap n^\perp \quad (87)$$

and, for $s \in \{-1, +1\}$, define

$$\chi_s := \inf\{\|v\|_2^2 : v \in \mathcal{K}^*, \ 3 \langle (r^*)^2, v \rangle = s\}. \quad (88)$$

The convention is $\chi_s = +\infty$ when the constraint set is empty. Then, along feasible displacements of sign s ,

$$\liminf_{\varepsilon \rightarrow 0, \ s \varepsilon > 0} \frac{\Delta_3(m_3^\dagger + \varepsilon)}{\varepsilon^2} \geq \beta \lambda \chi_s. \quad (89)$$

(iv) Maintain the conditions of part (ii). Suppose the IC critical cone

$$\mathcal{V} := T_{\mathcal{D}_q}(r^*) \cap n^\perp \quad (90)$$

is a stable closed linear space: it has a dense bounded core whose finite-dimensional subsets permit two-sided local perturbations within the same active IC constraints, and the three moment gradients are independent on that cone. Put

$$\mathcal{W} := \mathcal{V} \cap \mathbf{1}^\perp \cap (r^*)^\perp, \quad \psi := P_{\mathcal{W}}[(r^*)^2]. \quad (91)$$

If $\psi \neq 0$, then

$$\Delta_3(m_3^\dagger + \varepsilon) = \frac{\beta \lambda}{9 \|\psi\|_2^2} \varepsilon^2 + o(\varepsilon^2). \quad (92)$$

A finite strictly complementary pattern of bound blocks, pooling blocks, and strictly increasing blocks of the optimal effective allocation satisfies the stated local active-set stability condition. Independence of the three moment gradients on that cone gives the remaining nondegeneracy condition.

Proof idea. Concave biconjugacy identifies the shadow price on the third-moment constraint, and the KKT decomposition at an interior two-moment optimum below saturation gives the exact loss identity. Rescaling nearly optimal policies produces a critical direction and the local lower bound. For the matching upper bound, local stability of the active IC constraints provides a finite-dimensional feasible core; the implicit-function theorem then constructs moment-preserving perturbations along the minimum-energy direction.

Corollary 3 (The IC amplification factor). *Maintain the assumptions of Theorem 4(iv), suppose $0 < r^* < 1$ almost everywhere and $\psi \neq 0$, and let $\mathcal{W}^{\text{IC}} := \mathcal{W}$. Define the moment tangent without the local IC restriction by*

$$\mathcal{W}^0 := \mathbf{1}^\perp \cap (r^*)^\perp \quad (93)$$

and, for any closed linear subspace $\mathcal{U}$, its local third-moment capacity

$$S_3(\mathcal{U}) := \sup_{\substack{v \in \mathcal{U} \\ \|v\|_2 \leq 1}} [3 \langle (r^*)^2, v \rangle]^2 = 9 \|P_{\mathcal{U}}[(r^*)^2]\|_2^2. \quad (94)$$

Put

$$S_3^0 := S_3(\mathcal{W}^0), \quad S_3^{\text{IC}} := S_3(\mathcal{W}^{\text{IC}}), \quad \Omega^{\text{IC}} := \frac{S_3^0}{S_3^{\text{IC}}}. \quad (95)$$

The local tax coefficient factors as

$$\underbrace{\kappa^{\text{IC}}}_{\beta\lambda/S_3^{\text{IC}}} = \underbrace{\kappa^0}_{\beta\lambda/S_3^0} \underbrace{\Omega^{\text{IC}}}_{\geq 1}. \quad (96)$$

More precisely, let

$$\mathcal{B}^{\text{IC}} := \mathcal{W}^0 \cap (\mathcal{W}^{\text{IC}})^\perp. \quad (97)$$

Then

$$\Omega^{\text{IC}} = 1 + \frac{9 \|P_{\mathcal{B}^{\text{IC}}}[(r^*)^2]\|_2^2}{S_3^{\text{IC}}}, \quad (98)$$

and the following conditions are equivalent:

$$\Omega^{\text{IC}} > 1 \iff P_{\mathcal{B}^{\text{IC}}}[(r^*)^2] \neq 0 \iff P_{\mathcal{W}^0}[(r^*)^2] \notin \mathcal{W}^{\text{IC}}. \quad (99)$$

If the IC order restriction is locally slack, $\mathcal{W}^{\text{IC}} = \mathcal{W}^0$ and $\Omega^{\text{IC}} = 1$.

Up to the output normalization β , Theorem 4(i) identifies θ as the shadow price on the third-moment constraint in (14). The shadow prices on the first two moments and the IC multiplier remain in the same protocol comparison.

The loss identity separates stationarity and IC. The multiplier λ prices schedule displacement and conditional variance; the IC normal prices departure from the active face. The conditional Jensen gap in the third moment is at most $3V$. Hence a sequence with quadratic loss has $V = O(\varepsilon^2)$, and its first-order third-moment movement lies in $\mathcal{K}^*$.

The capacity in (94) is the squared first-order change in the third moment per unit of schedule movement that preserves the first two moments. Dividing $\beta\lambda$ by this capacity gives the local location tax. Corollary 3 holds the two-moment schedule and λ fixed: κ^0 uses the tangent without the local IC restriction, and the amplification factor records the capacity lost to the active IC face. The factor exceeds one exactly when that face rules out a third-moment direction; a positive IC multiplier alone is insufficient. Active IC constraints restrict how protocol intensity can be reallocated across types. When they remove a third-moment direction, a given change in cohort composition requires more schedule movement and therefore costs more output. If the face changes with the sign of the displacement, the two one-sided coefficients in (89) can differ.

5.3 A binding declining allocation rule and two boundary cases

Return to $Z \sim U[0, 1]$, $\alpha = \beta = 1$, and the full protocol interval. For the reciprocal allocation rule $q_\delta(z) = 1/(1 + \delta z)$, the unconstrained projection target in effective-allocation coordinates is

$$y_{\eta,\lambda}(z) = \frac{1}{1 + \delta z} + \frac{z - \eta(1 + \delta z)}{2\lambda(1 + \delta z)^2}. \quad (100)$$

It is monotone or single peaked. We solve the regime in which the lower bound of the IC-feasible interval and a tail order constraint both bind strictly. There the two-moment optimum has three blocks:

$$g^*(z) = \begin{cases} A, & z \leq \ell, \\ y_{\eta,\lambda}(z), & \ell < z < c, \\ p, & z \geq c, \end{cases} \quad 0 < \ell < c < 1, \quad A < p < B. \quad (101)$$

The five unknowns $(\ell, c, p, \eta, \lambda)$ are pinned down by the lower contact, the pool contact, the zero total pool residual, and the first two moment constraints. The strictly positive cumulative pool residual is the cumulative IC multiplier X_ξ . Appendix D gives the equations and the closed Gram-matrix formula for $\|\psi\|_2^2$.

Appendix D constructs an exact rational benchmark with $\delta = 4/5$ and

$$\begin{aligned} (\ell, c, p) &\simeq (0.142857, 0.384615, 1.088436), \\ (\lambda, \eta) &\simeq (0.268420, 0.066852). \end{aligned}$$

Its schedule lies strictly inside the protocol bounds almost everywhere. Let (m, k) be the induced first two moments; numerically, $m \simeq 0.500534$ and $k \simeq 0.329878$. Here $L_q(m) \simeq 0.311803 < k < k_q^{\text{sat}}(m) \simeq 0.347747$, so k lies strictly above the minimum feasible second moment and below saturation. The appendix verifies algebraically that a strictly binding order constraint creates the tail pool. The upper protocol bound remains slack.

At this pair, the residuals without and with the local IC restriction satisfy

$$\|P_{\mathcal{W}^0}[(r^*)^2]\|_2^2 \simeq 0.005095966, \quad \|P_{\mathcal{W}^{\text{IC}}}[(r^*)^2]\|_2^2 \simeq 0.000420469. \quad (102)$$

Consequently,

$$\Omega^{\text{IC}} \simeq 12.1197, \quad \kappa^0 \simeq 5.85257, \quad \kappa^{\text{IC}} \simeq 70.9314. \quad (103)$$

The active IC face leaves only 8.25% of the unrestricted third-moment capacity and multiplies the local tax by about 12.12. The amplification factor exceeds one by the equivalent conditions in (99), as the appendix verifies directly. The free block makes the first two projected moment gradients independent, and $\|\psi\|_2 > 0$ makes the third independent; hence all stability and nondegeneracy conditions in Theorem 4(iv) hold at this binding pair.

In two boundary cases, this coefficient simplifies. When the IC constraint is locally slack, the IC normal vanishes; away from protocol clipping the moment tangent space is unrestricted, and at clipping it reduces to the box face. When the IC-feasible interval collapses, the output-increasing range vanishes and

the second-moment shadow price can be chosen zero; the relevant object is then the zero-loss interval of Proposition 2.

Example 1 (A zero-loss interval when the IC-feasible interval collapses). Let Z be uniform, $q(z) = 2/(1+z)$, $\alpha = \beta = 1$, and $\mathcal{T} = \{0, 1/3, 2/3, 1\}$. IC forces $g(z) = 2$ and $r(z) = z$. Fix $(m, k) = (1/2, 7/18)$. Moving the required conditional variance from $[0, 1/3]$ to $[2/3, 1]$ leaves the first two moments and output unchanged and shifts the third moment from $53/162$ to $55/162$. Hence

$$[53/162, 55/162] \subset I_3^*. \quad (104)$$

Example 2 (An exact location tax with a finite protocol set). Let Z be uniform, $q \equiv 1$, $\alpha = \beta = 1$, use the same four-point protocol set, and again fix $(m, k) = (1/2, 7/18)$. The two-moment optimum assigns equal quartile blocks to the four protocols and selects $m_3^\dagger = 1/3$. For every feasible $x \in [17/54, 19/54]$,

$$\Delta_3(x) = \frac{81}{4} \left(x - \frac{1}{3} \right)^2. \quad (105)$$

The finite-protocol example and the binding continuum case measure the same restriction through different local geometries: the residual capacity to move the third moment while holding the first two fixed.

6 Conclusion

This paper characterizes stationary screening when one report-contingent protocol screens the current population and forms the next one. Stationarity of the type distribution and incentive compatibility then restrict the same assignment. The first theorem derives their joint program and dual, while the flow-rank theorem identifies exactly when population flow admits a finite set of protocol-only coordinates and makes the minimal number of coordinates intrinsic.

With independent Bernoulli formation trials, those coordinates are protocol moments. At fixed mean, the second moment measures protocol variance. Declining allocation rules require a minimum amount of that variance for IC; the remainder supports productive sorting until saturation and becomes output neutral thereafter. The third moment allocates the fixed protocol variance across types and cohorts formed after different numbers of successes. Its zero-loss values are the third moments generated by two-moment optimal policies. Outside that set, the location tax equals the output cost of the mandated residual composition. Its local coefficient combines the shadow price on the second-moment constraint with the relocation capacity left by IC. The IC amplification factor exceeds one exactly when the active IC face rules out a relocation direction that satisfies the moment constraints.

Short-lived agents and unrestricted transfers isolate the interaction between IC and population flow. With long-lived agents, limited liability, or budget balance, continuation and rent constraints would enter the same policy problem. Local stability of the active IC constraints yields the two-sided location-tax formula; one-sided critical cones govern changes in the active face, and the zero-loss set may expand above saturation. The resulting shadow prices on the moment constraints are the marginal values relevant when the institution chooses its target or the speed of convergence toward it.

A Proofs for Stationary Screening and Flow Rank

Proof of Theorem 1. We divide the proof into representation, population flow, and duality.

Representation and payments. Fix a pointwise protocol policy. If $V(z)$ is truthful utility, the two IC inequalities for $z > r$ give

$$V(r) + (z - r)g_\sigma(r) \leq V(z) \leq V(r) + (z - r)g_\sigma(z). \quad (106)$$

They imply that g_σ is nondecreasing. Darboux sums for a bounded monotone function then yield

$$V(z) - V(r) = \int_r^z g_\sigma(x) dx.$$

This proves (12)–(13). Conversely, monotonicity and the same integral bounds prove every global IC inequality. Since $g_\sigma > 0$, truthful utility is minimized at the lowest type, so IR is equivalent to $V_0 \geq 0$.

The convex analysis below identifies policies only μ -almost everywhere. Let g^0 be a nondecreasing representative of their effective allocation. On a dense full-measure set,

$$q(z)(\alpha + \beta a) \leq g^0(z) \leq q(z)(\alpha + \beta b). \quad (107)$$

Take the right-continuous version of g^0 , with its left limit at the upper endpoint. Atomlessness leaves its equivalence class unchanged; full support, density of the full-measure set, and continuity of q extend (107) to every report. On the exceptional null set, mix a and b to produce conditional mean $[g^0/q - \alpha]/\beta$. The mixing weight is Borel. This modification does not alter any population integral and gives a pointwise direct mechanism.

Population flow. Integrating (7) under the truthful occupation measure, using $\mu S_0 = \mu$, gives (16). Subtracting (5) yields

$$\begin{aligned} \mu P_\sigma - \mu &= (1 - \zeta) \{ 3(\mathbb{E}T - m)(R_1 - R_0) \\ &\quad + 3(\mathbb{E}T^2 - k)(R_0 - 2R_1 + R_2) \\ &\quad + (\mathbb{E}T^3 - m_3)(-R_0 + 3R_1 - 3R_2 + R_3) \}. \end{aligned} \quad (108)$$

Affine independence of R_0, R_1, R_2, R_3 is equivalent to linear independence of the three directions in braces. Since $\zeta < 1$, stationarity of μ is therefore equivalent to the three moment constraints in (6). Together with the reporting argument this proves part (i).

Compact policy space and the order polar. A policy is a Young measure

$$\gamma(dz, dt) = \mu(dz)\sigma(dt | z) \quad (109)$$

with fixed first marginal μ . Equip these measures with the stable topology generated by

$$\gamma \mapsto \int f(z, t) \gamma(\mathrm{d}z, \mathrm{d}t), \quad f \in L^1(\mu; C(\mathcal{T})). \quad (110)$$

Their set is compact and convex. One direct proof identifies a Young measure with a positive norm-one functional on $L^1(\mu; C(\mathcal{T}))$ that agrees with integration under μ on functions of z alone. Positivity and the marginal identities are weak-star closed; Banach–Alaoglu gives compactness, and the restricted weak-star topology is exactly (110).

The cone $\mathcal{M}_\uparrow$ is weak-star closed in $L^\infty(\mu)$. Indeed, atomlessness and full support make the cdf of μ a continuous strictly increasing map onto $[0, 1]$. After this change of variable, distributional monotonicity is the family of weak-star closed inequalities

$$\int_0^1 s(u) \varphi'(u) \mathrm{d}u \leq 0 \quad \text{for every nonnegative } \varphi \in C_c^1(0, 1).$$

Its polar is $\mathcal{H}_\mu$. To see this, testing positive and negative constants forces $\mathbb{E}\xi = 0$, while testing upper interval indicators gives $X_\xi(x) \geq 0$. Conversely, for a right-continuous nondecreasing s , Stieltjes integration gives

$$\int \xi s \mathrm{d}\mu = - \int_{(\underline{z}, \bar{z}]} X_\xi(x) \mathrm{d}s(x) \leq 0. \quad (111)$$

Lagrangian minimax. Let Γ denote the compact convex set of Young policies with first marginal μ . For $y = (\eta, \lambda, \theta, \xi) \in \mathbb{R}^3 \times \mathcal{H}_\mu$, define

$$\begin{aligned} \mathcal{L}(\gamma; y) := & \int \ell \mathrm{d}\gamma + \eta(\mathbb{E}_\gamma T - m) + \lambda(\mathbb{E}_\gamma T^2 - k) + \theta(\mathbb{E}_\gamma T^3 - m_3) \\ & + \int \xi(z) q(z) [\alpha + \beta \mathbb{E}_\gamma(T \mid Z = z)] \mu(\mathrm{d}z). \end{aligned} \quad (112)$$

If γ satisfies the moment equalities and its effective allocation is monotone, the supremum of $\mathcal{L}(\gamma; y)$ over y equals $\int \ell \mathrm{d}\gamma$: the polar term is nonpositive and $\xi = 0$ is available. If a moment equality fails, an unrestricted scalar multiplier can be scaled to make the supremum infinite. If the effective allocation is not monotone, weak-star closedness and the bipolar theorem supply a $\xi \in \mathcal{H}_\mu$ with strictly positive pairing, and scaling it makes the supremum infinite. Hence

$$C_{m,k,m_3}(\ell) = \min_{\gamma \in \Gamma} \sup_y \mathcal{L}(\gamma; y).$$

For fixed y , the Lagrangian is continuous affine in γ under the stable topology; for fixed γ , it is continuous affine in y under the product of the Euclidean and L^1 topologies. Sion's minimax theorem therefore gives

$$C_{m,k,m_3}(\ell) = \sup_y \min_{\gamma \in \Gamma} \mathcal{L}(\gamma; y).$$

The integrand is Carathéodory and dominated by an integrable envelope, so a measurable minimizing selection permits pointwise minimization over the compact protocol set. Expanding (112) gives

$$- \eta m - \lambda k - \theta m_3 + \int \left[\alpha q \xi + \min_{t \in \mathcal{T}} \{ \ell(z, t) + [\eta + \beta q \xi] t + \lambda t^2 + \theta t^3 \} \right] d\mu.$$

This proves (14) without a duality gap. Substituting $\ell = -wt$ identifies the shadow prices on the three moment constraints and the IC multiplier. If a maximizing multiplier vector and a primal optimizer exist, equality in the integrated pointwise inequality makes the nonnegative pointwise optimality gap vanish almost everywhere. Hence, for μ -almost every report, the optimizer's conditional lottery assigns probability one to the pointwise argmin in (15). This proves parts (ii) and (iii). $\square$

A.1 Proof of the minimal flow-rank representation

We first verify the statement about scalar payoffs preceding Theorem 2. Let t^- and t^+ minimize and maximize the nonconstant continuous index τ . Choose α_0, β_0 so that $\rho(t^\pm) = \alpha_0 + \beta_0 \tau(t^\pm)$. For any t , the degenerate lottery at t and the unique lottery on $\{t^-, t^+\}$ with the same mean $\tau(t)$ must have the same mean responsiveness. Hence $\rho(t) = \alpha_0 + \beta_0 \tau(t)$. Conversely, this affine form makes mean responsiveness depend only on mean protocol intensity.

Proof of Theorem 2. We prove the representation, minimality, and claims about the moment conditions for invariance in turn.

From order-restricted sufficiency to state-independent kernel differences. Flow sufficiency implies order-restricted flow sufficiency, and (21) implies flow sufficiency after integration. It remains to derive the representation from the order-restricted property.

Fix $t, u \in \mathcal{T}$. Choose $t^+, t^- \in \mathcal{T}$ with $x_1(t^+) > x_1(t^-)$, and let π^0 put strictly positive mass on every distinct protocol in $\{t, u, t^+, t^-\}$. Write

$$s(z) := \frac{z - \underline{z}}{\bar{z} - \underline{z}},$$

and let δ_v^D denote Dirac mass at protocol v . For sufficiently small $\gamma > 0$,

$$\sigma^b(dv | z) := \pi^0(dv) + \gamma[s(z) - 1/2][\delta_{t^+}^D - \delta_{t^-}^D](dv) \quad (113)$$

is a probability kernel whose conditional first coordinate is strictly increasing.

Its slope in the first coordinate is

$$L_b := \frac{\gamma[x_1(t^+) - x_1(t^-)]}{\bar{z} - \underline{z}} > 0.$$

Take any $f \in C^1(\mathcal{Z})$ with $\mathbb{E}f = 0$. For sufficiently small $\varepsilon > 0$, both kernels

$$\sigma^\pm(dv | z) := \sigma^b(dv | z) \pm \varepsilon f(z)[\delta_t^D - \delta_u^D](dv) \quad (114)$$

remain nonnegative, and their conditional first coordinates remain increasing: in addition to preserving nonnegativity, choose ε so that

$$\varepsilon |x_1(t) - x_1(u)| \|f'\|_\infty < L_b,$$

with the restriction vacuous when the coefficient is zero. Their aggregate x -coordinates agree because $\mathbb{E}f = 0$. Order-restricted sufficiency therefore gives, for every $\varphi \in C(\mathcal{Z})$,

$$\int f(z) \{ \tilde{Q}_t \varphi(z) - \tilde{Q}_u \varphi(z) \} \mu(dz) = 0. \quad (115)$$

The expression in braces is continuous by the joint Feller property. The C^1 zero-mean functions are dense in the zero-mean subspace of $L^2(\mu)$, so this expression is constant μ -almost everywhere. Full support makes it constant everywhere. Since continuous functions separate finite Borel measures on the compact state space,

$$\tilde{Q}_t(\cdot | z) - \tilde{Q}_u(\cdot | z)$$

is independent of z .

Fix t_0 , abbreviate this state-independent difference by $\delta_t := \tilde{Q}_t - \tilde{Q}_{t_0}$, and choose $t_1, \dots, t_d$ so that

$$x(t_1) - x(t_0), \dots, x(t_d) - x(t_0)$$

form a basis of $\mathbb{R}^d$. There are unique signed measures $\eta_1, \dots, \eta_d$ solving

$$\delta_{t_r} = \sum_{i=1}^d [x_i(t_r) - x_i(t_0)] \eta_i, \quad r = 1, \dots, d. \quad (116)$$

For an arbitrary t , choose coefficients c_r satisfying

$$x(t) - x(t_0) = \sum_{r=1}^d c_r [x(t_r) - x(t_0)]$$

and form the finite signed protocol measure

$$\nu := \delta_t^D - \delta_{t_0}^D - \sum_{r=1}^d c_r (\delta_{t_r}^D - \delta_{t_0}^D). \quad (117)$$

It has total mass zero and $\int x d\nu = 0$. If $\nu = 0$, integrating this signed protocol identity against $v \mapsto \delta_v$ gives $\int \delta_v d\nu(v) = 0$. Otherwise, its positive and negative parts have the same mass and, after normalization, are two protocol lotteries with the same mean x . The corresponding constant policies lie in the order-restricted domain. Sufficiency again gives $\int \delta_v d\nu(v) = 0$. In either case,

$$\delta_t = \sum_{r=1}^d c_r \delta_{t_r}.$$

Together with (116), this is (21). Nonsingularity in (116) and zero total mass of every δ_{t_r} imply $\eta_i(\mathcal{Z}) = 0$.

The same basis argument proves uniqueness. This completes part (i).

Minimal dimension and uniqueness. Every representation (21) places the flow space inside $\text{span}\{\eta_1, \dots, \eta_d\}$, so its rank is at most d . For an arbitrary sufficient coordinate map, first restrict its codomain to the affine hull of its image and apply an affine reparametrization; this removes redundant coordinates without changing any coordinate means or the sufficiency property. If the resulting dimension is positive, one coordinate can be chosen nonconstant, so part (i) gives the same rank lower bound. If a zero-dimensional coordinate is flow sufficient, fix protocols t, u and compare the policy using t on a Borel type set A and u elsewhere with the policy using u everywhere. Flow sufficiency gives

$$\int_A [\tilde{Q}_t(\cdot | z) - \tilde{Q}_u(\cdot | z)] \mu(dz) = 0 \quad \text{for every } A.$$

After pairing with continuous test functions, joint Feller continuity and full support imply $\tilde{Q}_t(\cdot | z) = \tilde{Q}_u(\cdot | z)$ for every z . Thus the flow rank is zero.

Conversely, suppose the flow rank is the finite number r . When $r = 0$, all protocol kernels coincide and a constant coordinate is sufficient. For $r \geq 1$, choose a basis $\eta_1, \dots, \eta_r$ of the flow space and write uniquely

$$\delta_t = \sum_{i=1}^r a_i(t) \eta_i, \quad a(t_0) = 0. \quad (118)$$

The coordinate map a is continuous. To see this, linear independence of the measures permits continuous test functions $\varphi_1, \dots, \varphi_r$ for which the matrix $(\int \varphi_j d\eta_i)_{j,i}$ is nonsingular. Pairing (118) with these functions expresses $a(t)$ as an invertible linear transformation of the continuous vector

$$(\tilde{Q}_t \varphi_j(z_0) - \tilde{Q}_{t_0} \varphi_j(z_0))_{j=1}^r$$

at any fixed state z_0 . Because the δ_t 's span the flow space, $a(\mathcal{T})$ spans $\mathbb{R}^r$. Representation (118) makes a flow sufficient, so the minimum coordinate dimension is exactly r .

Let a and b be two minimal systems. After subtracting their values at t_0 , each representation uses r measures spanning the same r -dimensional flow space, hence a basis of that space. A change of basis therefore gives

$$b(t) - b(t_0) = A[a(t) - a(t_0)]$$

for a nonsingular A , which is (22). This proves part (ii).

Moment conditions for stationarity. For a minimal system, the representing measures $\eta_1, \dots, \eta_r$ are linearly independent. Integrating its kernel-difference representation under any lottery with mean coordinate $\bar{a}$ gives

$$\tilde{K}_{\bar{a}} = \tilde{Q}_{t_0} + \sum_{i=1}^r [\bar{a}_i - a_i(t_0)] \eta_i, \quad (119)$$

which is independent of the representing lottery. Under a policy σ ,

$$\mu \tilde{P}_\sigma - \mu \tilde{K}_{\bar{a}} = \sum_{i=1}^r (\bar{a}_{\sigma,i} - \bar{a}_i) \eta_i. \quad (120)$$

If the centered kernel preserves μ , linear independence proves (24). This proves part (iii). $\square$

Proof of Corollary 1. The Bernstein polynomials form a basis of the polynomials of degree at most N . Hence there are signed measures $A_1, \dots, A_N$ such that

$$\widehat{Q}_t^N = R_0 + \sum_{\ell=1}^N t^\ell A_\ell. \quad (121)$$

Affine independence of $R_0, \dots, R_N$, together with the invertible Bernstein-to-power change of basis, makes $A_1, \dots, A_N$ linearly independent. Each has total mass zero. The common persistent component cancels across protocols, so

$$Q_t^N - Q_{t_0}^N = (1 - \zeta) \sum_{\ell=1}^N (t^\ell - t_0^\ell) A_\ell.$$

For any $N + 1$ distinct protocol intensities, the corresponding power vectors are affinely independent by the Vandermonde determinant. Thus the flow space has dimension N , and $(t, \dots, t^N)$ is flow sufficient and minimal. The uniqueness and claims about the moments required for invariance follow from Theorem 2(ii)–(iii).

The same Vandermonde argument shows that a protocol set of $J \leq N + 1$ distinct protocols has flow rank $J - 1$; for $J \geq N + 1$, the rank is N . This proves the final statement following the corollary in the main text. $\square$

B Proofs for the Two-Moment Relaxation

Proof of Lemma 1. Let T have a law in $\mathcal{P}_{\mathcal{T}}(r)$. If $r \in \mathcal{T}$, strict convexity gives $\mathbb{E}T^2 \geq r^2$, with equality only at δ_r .

If $r \notin \mathcal{T}$, put $\ell = t_-(r) < r < u = t_+(r)$. The gap contains no feasible protocol, so

$$t^2 - (\ell + u)t + \ell u = (t - \ell)(t - u) \geq 0 \quad \text{for every } t \in \mathcal{T}. \quad (122)$$

Taking expectations yields $\mathbb{E}T^2 \geq (\ell + u)r - \ell u$. The unique lottery on $\{\ell, u\}$ with mean r attains the bound, which equals the lower formula in (35). Similarly,

$$t^2 - (a + b)t + ab = (t - a)(t - b) \leq 0 \quad \text{for } t \in [a, b] \quad (123)$$

gives the upper formula and its unique endpoint lottery.

Mixing the two extremal lotteries while preserving their common mean generates every conditional second moment between the envelopes. Convexity of $\underline{K}_{\mathcal{T}}$ follows by mixing minimizing laws at two means. It is affine on every gap of $\mathcal{T}$. At a feasible protocol, the excess $[r - t_-(r)][t_+(r) - r]$ tends to

zero, proving continuity. Thus the lower envelope is strictly convex everywhere exactly when the protocol set has no gap, namely when $\mathcal{T} = [a, b]$. The bracketing maps are monotone and Borel, so both extremal lotteries admit Borel selections.

Now fix a Borel schedule $r(z)$. Conditional necessity and affinity of the upper envelope give

$$\mathbb{E}\underline{K}_{\mathcal{T}}(r(Z)) \leq \mathbb{E}T^2 \leq \mathbb{E}\bar{K}_{\mathcal{T}}(r(Z)) = \bar{K}_{\mathcal{T}}(\mathbb{E}r(Z)).$$

For sufficiency, let π_z^L and π_z^U be the selected lower- and upper-envelope lotteries and define

$$d(z) := \bar{K}_{\mathcal{T}}(r(z)) - \underline{K}_{\mathcal{T}}(r(z)).$$

The measure $d\mu$ is atomless. For every

$$c \in [0, \mathbb{E}d]$$

there is a Borel type set A_c with $\int_{A_c} d \, d\mu = c$. Use π_z^U on A_c and π_z^L elsewhere. Conditional means are unchanged, every conditional support is binary, and the aggregate second moment rises from the lower endpoint by exactly c . This proves (36). $\square$

Lemma 2 (Schedule compactness). *The sets $\mathcal{D}_{q,\mathcal{T}}$ and $C_{q,\mathcal{T}}(m)$ are closed and convex in $L^2(\mu)$, and $C_{q,\mathcal{T}}(m)$ is weakly compact. The functional $\mathcal{K}_{\mathcal{T}}$ is convex, norm-continuous, and weakly lower semicontinuous.*

Proof. The cone of L^2 classes with a nondecreasing representative is closed and convex. The affine map $r \mapsto q(\alpha + \beta r)$ is continuous, so its inverse image, intersected with the box and mean hyperplane, gives the first claim. The box makes the mean set bounded; a closed bounded convex subset of the Hilbert space L^2 is weakly compact.

Every affine piece of $\underline{K}_{\mathcal{T}}$ has slope in $[2a, 2b]$; hence it is globally Lipschitz. Its integral is norm-continuous and convex, and every norm-lower-semicontinuous convex functional is weakly lower semicontinuous. $\square$

Proof of Theorem 3. Take any policy and put $r(z) = \mathbb{E}[T \mid Z = z]$. The reporting part of Theorem 1 makes IC equivalent to $r \in C_{q,\mathcal{T}}(m)$. The two-moment relaxation imposes its first two moment constraints. Lemma 1 gives

$$\mathcal{K}_{\mathcal{T}}(r) \leq k \leq \bar{K}_{\mathcal{T}}(m). \quad (124)$$

Minimizing the left side proves necessity of (44). Conversely, apply the construction in Lemma 1 to a minimizer defining $L_{q,\mathcal{T}}(m)$. This proves feasibility throughout the displayed interval.

Since

$$\mathbb{E}[w(Z)T] = \mathbb{E}[w(Z)r(Z)], \quad (125)$$

the same construction sends every feasible schedule in (41) to an exact-moment policy without changing output. Conversely, every policy gives such a schedule. This proves part (i), including attainment by Lemma 2.

For concavity, take schedules r_i feasible at k_i . Convexity gives

$$\mathcal{K}_{\mathcal{T}}(\theta r_1 + (1 - \theta)r_2) \leq \theta k_1 + (1 - \theta)k_2. \quad (126)$$

Output is affine under the same mixture. Set inclusion makes the value nondecreasing. A finite concave function is continuous in the interior. Weak compactness and lower semicontinuity give one-sided continuity at the lower endpoint; mixing an upper-endpoint optimizer with a minimum-dispersion schedule gives continuity at the upper endpoint.

The mean-constrained maximum is first attainable at k^{sat} and remains attainable thereafter, so the value is flat above saturation. If it were flat between $k_1 < k_2 \leq k^{\text{sat}}$, concavity and monotonicity would propagate that flat segment through saturation. The mean-constrained maximum would then be attainable at $k_1 < k^{\text{sat}}$, contradicting the definition of the threshold. This proves part (ii).

If an optimizer at or below saturation had $\mathcal{K}_{\mathcal{T}}(r) < k$, the same output would be feasible at a lower second-moment level, contradicting strict increase. When the output-increasing interval is empty, saturation equals the smallest feasible second moment, so equality follows from the definition. Conditional second moments are bounded below by $\underline{K}_{\mathcal{T}}(r(z))$; equality of their integrals forces pointwise equality almost everywhere. Uniqueness of the conditional lower lottery then proves the support claim. The same construction yields a binary optimal policy at every other second-moment level. This proves part (iii).

Suppose now $\mathcal{T} = [a, b]$. Then $\mathcal{K}_{\mathcal{T}}(r) = \mathbb{E}r^2$. Strict convexity of the squared norm gives a unique minimum-dispersion schedule and a unique minimum-norm member of the mean-constrained optimal set. If two distinct schedules were optimal at or below saturation, their midpoint would have the same output and

$$\mathbb{E} \left(\frac{r_1 + r_2}{2} \right)^2 = \frac{\mathbb{E}r_1^2 + \mathbb{E}r_2^2}{2} - \frac{1}{4}\mathbb{E}(r_1 - r_2)^2 < k, \quad (127)$$

contradicting the binding second-moment constraint and strict increase. Conditional Jensen, together with $\mathbb{E}T^2 = \mathbb{E}r_k^2$, forces zero conditional variance, so the policy is deterministic.

At a second moment strictly between the minimum and saturation, the unique minimum-dispersion schedule is a strict feasible point for the quadratic inequality. Separation gives

$$B(m, k) = \min_{\lambda \geq 0} \left\{ \lambda k + \max_{r \in C_{q, [a, b]}(m)} \mathbb{E}[wr - \lambda r^2] \right\}. \quad (128)$$

A zero multiplier would make the optimizer a mean-constrained maximizer below saturation, so $\lambda > 0$. Completing the square gives (47). Under the relative-interior condition stated in part (iv), the finite-dimensional mean-map qualification gives the normal-cone sum rule and scalar multiplier η , and yields (48). Whenever the joint two-moment dual is attained, comparing (128) with the $\theta = 0$ restriction of (14) permits the same value to be chosen for its second-moment multiplier. This proves part (iv).

Finally, when $\mathcal{T}$ is finite, the formula in (35) is the adjacent linear interpolation of squared protocol points. Parts (i)–(iii) give adjacent binary support at and below saturation. The affine pieces of the lower envelope on each gap of the protocol set permit nonunique schedules. The example below exhibits nonuniqueness of both the optimal schedule and saturation minimizer. This proves part (v). $\square$

Example 3 (Nonuniqueness with a finite protocol set). Let Z be uniform on $[1, 2]$, set $q(z) = 1/z$, $\alpha = 0$,

$\beta = 1$, and $\mathcal{T} = \{1/4, 3/4, 1\}$. Put $m = 9/20$, and define

$$r_1(z) = \frac{3z}{10}, \quad (129)$$

$$r_2(z) = z \left[\frac{3}{10} + \frac{1}{100} \left(z - \frac{14}{9} \right) \right]. \quad (130)$$

Both schedules take values strictly between $1/4$ and $3/4$, and

$$q(z)r_1(z) = \frac{3}{10}, \quad q(z)r_2(z) = \frac{3}{10} + \frac{1}{100} \left(z - \frac{14}{9} \right), \quad (131)$$

so both satisfy IC. Since $\mathbb{E}Z = 3/2$, $\mathbb{E}Z^2 = 7/3$, and $(14/9)\mathbb{E}Z = 7/3$, both means equal $9/20$. Moreover $w(z) = zq(z) = 1$, so every feasible schedule with this mean has the same output.

On $[1/4, 3/4]$, the lower envelope is $\underline{K}_{\mathcal{T}}(r) = r - 3/16$, and convexity makes this affine function a global supporting lower bound. Every feasible schedule therefore requires at least

$$m - \frac{3}{16} = \frac{21}{80} \quad (132)$$

units of the second moment. Both r_1 and r_2 attain the bound. Thus $L_{q,\mathcal{T}}(m) = k_{q,\mathcal{T}}^{\text{sat}}(m) = 21/80$, and the two distinct schedules are both minimum-dispersion members of the mean-constrained optimal set.

C Proofs for Declining Allocations

Proof of Proposition 1. Write $g = q(\alpha + \beta r)$. Because q is nonincreasing, every nondecreasing feasible g satisfies

$$g(z) \geq \sup_{t \leq z} \alpha q(t) = A, \quad g(z) \leq \inf_{t \geq z} (\alpha + \beta) q(t) = B. \quad (133)$$

Conversely, any nondecreasing g valued in $[A, B]$ respects the pointwise bounds. The mean restriction becomes

$$\mathbb{E}[hg] = S. \quad (134)$$

It is feasible exactly when $AH \leq S \leq BH$, and every intermediate value is attained by a constant effective allocation.

The constant $g \equiv \bar{g} = S/H$ generates (54). To prove that it minimizes protocol variance, take another feasible g . The schedule difference

$$r - r^L = \frac{h(g - \bar{g})}{\beta} \quad (135)$$

has mean zero and changes sign at most once, from negative to positive. Both r and r^L are nondecreasing. Choose a constant γ between the left and right values of $r + r^L$ at the crossing. Then

$$\mathbb{E}[r^2 - (r^L)^2] = \mathbb{E}[(r - r^L)(r + r^L - \gamma)] \geq 0. \quad (136)$$

Strict convexity of the squared norm gives uniqueness. Since

$$r^L - m = \frac{S}{\beta} \left(\frac{h}{H} - 1 \right),$$

its variance gives (55).

For output maximization subject only to the mean constraint, observe

$$\mathbb{E}[Zg] = H\mathbb{E}_\nu[wg], \quad \mathbb{E}_\nu g = \bar{g}. \quad (137)$$

The bathtub principle assigns B to the highest values of the strictly increasing output weight w and A to the rest. Atomlessness and full support give the unique cutoff in (56). The assignment is nondecreasing, so it is IC. Substituting $\hat{r}_q = (hg^* - \alpha)/\beta$, using $\mathbb{E}[hg^*] = S$, gives (59).

The difference between the saturation threshold and the smallest feasible second moment is

$$\begin{aligned} P_q(m) &= \frac{1}{\beta^2} \mathbb{E}\{h^2[(g^*)^2 - \bar{g}^2]\} \\ &= \frac{H}{\beta^2} \left\{ \mathbb{E}_\nu[h(g^* - \bar{g})^2] + 2\bar{g}\mathbb{E}_\nu[h(g^* - \bar{g})] \right\}. \end{aligned} \quad (138)$$

Since $\mathbb{E}_\nu(g^* - \bar{g}) = 0$, the last expectation is $\text{Cov}_\nu(h, g^*)$. Both functions are nondecreasing, so the covariance is nonnegative. The first term is strictly positive when $A < B$ and $\bar{g} \in (A, B)$.

If q is strictly decreasing, h is nonconstant, and (55) gives $C_q(m) > 0$. On a positive-mass interval below the cutoff, $0 < \hat{r}_q < 1$, except possibly at the endpoint. Hence

$$N_q(m) = m - k_q^{\text{sat}}(m) = \mathbb{E}[\hat{r}_q(1 - \hat{r}_q)] > 0.$$

For constant q , the minimum schedule is $r \equiv m$. The mean-constrained optimizer is the upper-tail indicator of mass m , giving (61). If $A = B$, every feasible effective allocation is that constant, so the minimum and output-maximizing schedules coincide. This proves (62). $\square$

Proof of Corollary 2. For (64),

$$h(z) = 1 + \delta z, \quad H = 1 + \frac{\delta}{2}, \quad S = \frac{3}{2}, \quad \bar{g} = \frac{3}{2 + \delta}, \quad A = 1, \quad B = \frac{2}{1 + \delta}. \quad (139)$$

Substitution gives (65). For $\delta < 1$, the cutoff equation is

$$\int_{z_\delta^*}^1 (1 + \delta z) dz = \frac{1 + \delta}{2}. \quad (140)$$

It reduces to $\delta(z_\delta^*)^2 + 2z_\delta^* - 1 = 0$, whose positive root is $1/(1 + \sqrt{1 + \delta})$. The two values A, B give (66). At $\delta = 1$, the IC-feasible interval collapses and both schedules reduce to $r(z) = z$.

Direct integration gives, with $s = \sqrt{1 + \delta}$,

$$\int_0^1 (r_\delta^L)^2 dz = \frac{\delta^2 + \delta + 1}{(\delta + 2)^2},$$

$$\int_0^{z_\delta^*} (\delta z)^2 dz + \int_{z_\delta^*}^1 \left(\frac{1 - \delta + 2\delta z}{1 + \delta} \right)^2 dz = \frac{s^3 - s^2 - s + 4}{3s(s + 1)}. \quad (141)$$

The derivative of the first expression is

$$\frac{3\delta}{(\delta + 2)^3} > 0 \quad \text{for } \delta > 0.$$

Viewing the second as a function of $s \in [1, \sqrt{2}]$, its derivative is

$$\frac{s^4 + 2s^3 - 8s - 4}{3s^2(s + 1)^2} < 0. \quad (142)$$

The numerator is increasing on this interval and equals $-4\sqrt{2}$ at its upper endpoint, so it is negative throughout. Evaluating at $\delta = 0, 1$ gives (69). $\square$

D Proofs for the Third Moment and the Location Tax

D.1 The zero-loss set

The reduction of stationarity of μ to three moment constraints is proved with Theorem 1 in Appendix A. The argument below starts from the two-moment policy set.

Proof of Proposition 2. The maps Y and M_3 are continuous affine functionals of the compact convex policy set Σ_2 . The image I_3 is therefore a nonempty compact interval. The optimal face Σ_2^* is also nonempty, compact, and convex, so its image I_3^* has the same properties.

Mixing policies proves concavity of Y_3 , and compactness gives attainment and upper semicontinuity. If $x \in I_3^*$, a two-moment optimizer remains feasible under the third-moment constraint, so $Y_3(x) = Y_2$. Conversely, equality and attainment produce a policy in $\Sigma_3(m, k, x)$ that also belongs to Σ_2^* , hence $x \in I_3^*$. This proves (75).

If every two-moment optimizer has mean schedule r^* , output depends only on that schedule. Every policy with conditional mean r^* and second moment k is therefore two-moment optimal, and every two-moment optimizer belongs to this conditional policy set. Its affine third-moment image is exactly the interval between (76) and (77). Policy uniqueness collapses the image to $\{\mathbb{E}_{\sigma^*} T^3\}$; determinism further identifies this number with $\mathbb{E}(r^*)^3$. $\square$

D.2 The general shadow price

Proof of Theorem 4(i). Partitioning Σ_2 by its third moment gives, for every $\theta \in \mathbb{R}$,

$$\Psi_3(\theta) = \max_{x \in I_3} \{Y_3(x) - \theta x\}. \quad (143)$$

The function Y_3 is proper, concave, and upper semicontinuous on the compact interval I_3 . Concave biconjugacy applied to (143) gives (83). A finite concave function has a supergradient at every relative-interior point; equality holds at exactly those θ that are supergradients. Concavity and (75) imply that every supergradient is positive strictly to the left of the maximizing interval and negative strictly to its right, while zero is a supergradient on that interval. $\square$

D.3 Exact loss and localization

Proof of Theorem 4(ii). Let $r = \mathbb{E}[T \mid Z]$ and $d = r - r^*$. The first-moment constraint gives $\langle \mathbf{1}, d \rangle = 0$. Conditional variance decomposition and binding of the second-moment constraint at r^* give

$$k = \|r\|_2^2 + V = \|r^*\|_2^2, \quad -2\langle r^*, d \rangle = \|d\|_2^2 + V. \quad (144)$$

Substitute $w = \eta \mathbf{1} + 2\lambda r^* + n$ from (84):

$$\begin{aligned} \frac{Y_2 - Y(\sigma)}{\beta} &= \langle w, r^* - r \rangle \\ &= -2\lambda \langle r^*, d \rangle - \langle n, d \rangle = \lambda(\|d\|_2^2 + V) - \langle n, d \rangle. \end{aligned} \quad (145)$$

The normal inequality gives $\langle n, d \rangle \leq 0$. If $n = \beta q \xi$, then $\beta q d = g_\sigma - g_{r^*}$, yielding the expression involving the IC multiplier. $\square$

We record two bounds used in the local proof. Conditional Jensen gives the lower bound. The upper bound uses $\sup_{t \in [0,1]} (t^3)'' = 6$:

$$0 \leq \mathbb{E}[T^3 \mid Z] - r^3 \leq 3 \text{Var}(T \mid Z). \quad (146)$$

Moreover,

$$r^3 - (r^*)^3 = 3(r^*)^2 d + 3r^* d^2 + d^3. \quad (147)$$

Since $r^*, r \in [0, 1]$,

$$\varepsilon = 3 \langle (r^*)^2, d \rangle + R, \quad |R| \leq 4 \|d\|_2^2 + 3V. \quad (148)$$

Put $a_* := 3 \|(r^*)^2\|_2$ and $e := \|d\|_2^2 + V$. By Cauchy–Schwarz,

$$|\varepsilon| \leq a_* \sqrt{e} + 4e.$$

Solving this quadratic inequality for $\sqrt{e}$ and applying (85) proves (86). In particular, the zero-loss set is a singleton under the conditions of Theorem 4(ii); the analysis of the critical cone below gives the exact local rate.

D.4 Lower bound from the critical cone and upper bound from active-set stability

Proof of Theorem 4(iii). Fix a sign s , take any feasible $\varepsilon_j = s t_j \rightarrow 0$, and choose optimal policies. Pass to a subsequence along which the tax divided by t_j^2 converges to the liminf of the original sequence. If that

$\liminf$ is $+\infty$, the claimed bound holds. Otherwise these quotients are bounded, and (85) implies

$$\|d_j\|_2 = O(t_j), \quad V_j = O(t_j^2), \quad -\langle n, d_j \rangle = O(t_j^2). \quad (149)$$

The normalized directions $v_j = d_j/t_j$ have a weakly convergent subsequence, say to v .

Because each $r^* + d_j \in \mathcal{D}_q$, the normalized directions lie in the radial cone. Its closure $T_{\mathcal{D}_q}(r^*)$ is closed and convex, hence weakly closed, so the weak limit belongs to that tangent cone. The first-moment constraint gives $\langle \mathbf{1}, v \rangle = 0$. Dividing (144) by t_j gives $\langle r^*, v \rangle = 0$, and the last estimate in (149) gives $\langle n, v \rangle = 0$. Finally, (148) divided by t_j gives

$$3 \langle (r^*)^2, v \rangle = s.$$

Thus v is feasible in (88). Weak lower semicontinuity and the loss identity yield

$$\liminf_j \frac{\Delta_3(m_3^\dagger + \varepsilon_j)}{\beta t_j^2} \geq \lambda \liminf_j \|v_j\|_2^2 \geq \lambda \|v\|_2^2 \geq \lambda \chi_s.$$

□

Part (iv) uses the following stability condition. A closed linear face $\mathcal{V}$ is *stable* if it contains a dense linear core $\mathcal{S} \subset L^\infty(\mu)$ such that:

- (a) for every finite-dimensional $E \subset \mathcal{S}$, all sufficiently small $u \in E$ satisfy $r^* + u \in \mathcal{D}_q$ and $\langle n, u \rangle = 0$;
- (b) the map

$$u \mapsto (\langle \mathbf{1}, u \rangle, 2 \langle r^*, u \rangle, 3 \langle (r^*)^2, u \rangle) \quad (150)$$

maps $\mathcal{V}$ onto $\mathbb{R}^3$.

Proof of Theorem 4(iv). Because $\mathcal{V}$ is linear, $\mathcal{K}^* = \mathcal{W}$. For $v \in \mathcal{W}$,

$$\langle (r^*)^2, v \rangle = \langle \psi, v \rangle.$$

Hence Cauchy–Schwarz shows that the minimum norm under $3 \langle (r^*)^2, v \rangle = s$ is attained at

$$v_s = \frac{s\psi}{3 \|\psi\|_2^2}, \quad \|v_s\|_2^2 = \frac{1}{9 \|\psi\|_2^2}. \quad (151)$$

Thus $\chi_s = 1/(9 \|\psi\|_2^2)$ for both signs.

The derivative image of the dense core is dense in $\mathbb{R}^3$ by surjectivity on $\mathcal{V}$. Since that image is a linear subspace of a finite-dimensional space, it equals $\mathbb{R}^3$. First suppose v_s belongs to the stable core. Enlarge its span, if necessary, by three core directions on which the moment derivative is invertible; the resulting finite-dimensional subspace E remains stable. Define

$$F(u) := \begin{pmatrix} \mathbb{E}(r^* + u) - m \\ \mathbb{E}(r^* + u)^2 - k \\ \mathbb{E}(r^* + u)^3 - m_3^\dagger \end{pmatrix}. \quad (152)$$

Its derivative at zero is (150). Choose a three-dimensional complement on which that derivative is invertible. The implicit-function theorem gives

$$u_s(t) = tv_s + o(t), \quad F(u_s(t)) = (0, 0, st)^\top.$$

Stability keeps $r^* + u_s(t)$ under the same active IC constraints. The deterministic policy generated by this schedule is feasible and, by (85), loses

$$\beta\lambda \|u_s(t)\|_2^2 = \frac{\beta\lambda}{9\|\psi\|_2^2} t^2 + o(t^2).$$

Part (iii) gives the matching lower bound.

For a general $v_s \in \mathcal{V}$, approximate it by core directions. Surjectivity and finite dimensionality of the moment codomain permit a vanishing correction with three fixed core directions so that every approximation has the exact derivative moments $(0, 0, s)$. Given $\rho > 0$, choose such an approximation whose squared norm is at most $\|v_s\|_2^2 + \rho$, and apply the preceding finite-dimensional construction. It gives

$$\limsup_{t \downarrow 0} \frac{\Delta_3(m_3^\dagger + st)}{t^2} \leq \beta\lambda(\|v_s\|_2^2 + \rho).$$

Letting $\rho \downarrow 0$ and combining this inequality with part (iii) proves (92). $\square$

Proof of Corollary 3. For any closed linear subspace $\mathcal{U}$, the Riesz representation theorem and Cauchy–Schwarz give

$$\sup_{\substack{v \in \mathcal{U} \\ \|v\|_2 \leq 1}} [3 \langle (r^*)^2, v \rangle]^2 = 9 \|P_{\mathcal{U}}[(r^*)^2]\|_2^2.$$

In particular, $S_3^{\text{IC}} = 9\|\psi\|_2^2$, so Theorem 4(iv) implies

$$\kappa^{\text{IC}} = \frac{\beta\lambda}{S_3^{\text{IC}}}.$$

The almost-everywhere protocol-interiority assumption makes the closed two-sided tangent generated by pointwise protocol feasibility equal to $L^2(\mu)$. After imposing the derivatives of the first two moment constraints, the moment tangent without IC is therefore $\mathcal{W}^0$.

Because the active IC constraints add restrictions, $\mathcal{W}^{\text{IC}} \subseteq \mathcal{W}^0$. Both spaces are closed, and hence

$$\mathcal{W}^0 = \mathcal{W}^{\text{IC}} \oplus \mathcal{B}^{\text{IC}}.$$

For $f = (r^*)^2$, orthogonal projection and Pythagoras give

$$\begin{aligned} P_{\mathcal{W}^0} f &= P_{\mathcal{W}^{\text{IC}}} f + P_{\mathcal{B}^{\text{IC}}} f, \\ S_3^0 &= S_3^{\text{IC}} + 9\|P_{\mathcal{B}^{\text{IC}}} f\|_2^2. \end{aligned}$$

Dividing by the positive S_3^{IC} proves (96) and (98). The Pythagorean remainder is strictly positive exactly

when $P_{\mathcal{B}^{\text{IC}}} f \neq 0$, equivalently when $P_{\mathcal{W}^0} f \notin \mathcal{W}^{\text{IC}}$. This proves (99). If the IC restriction is locally slack, its critical cone is the full moment tangent without IC, so the two spaces and capacities coincide. $\square$

D.5 Strict active patterns under declining allocation rules

For declining q , put $h = 1/q$ and $g = q(\alpha + \beta r)$. The feasible effective allocations are nondecreasing functions in the global IC-feasible interval $[A, B]$. Suppose g^* has finitely many blocks of the following kinds:

- (i) bound blocks on which $g^* = A$ or $g^* = B$, with a strictly positive bound multiplier;
- (ii) interior pooling blocks on which g^* is constant, with a cumulative IC multiplier strictly positive in the block interior and zero at its endpoints;
- (iii) free blocks on which g^* is continuously differentiable, separated from the bounds, and has derivative bounded away from zero.

Positive jumps remain separated; at zero-gap junctions perturbations use matching traces.

The IC critical cone in effective-allocation coordinates is then

$$\mathcal{V}_g = \{u : u = 0 \text{ on bound blocks and } u \text{ is constant on each pooling block}\}. \quad (153)$$

To verify this, a tangent direction must have the sign permitted by the active bound; strict bound complementarity and zero normal pairing force it to vanish on bound blocks. On a pooling block the tangent is nondecreasing. If X is the strictly positive cumulative IC multiplier, Stieltjes integration gives

$$\langle \xi, u \rangle = - \int X \, du.$$

Zero pairing forces $du = 0$, hence constancy. Conversely, piecewise smooth members of (153) form a dense core. Their derivatives are dominated by the strict slope on free blocks, they are constant on pools and zero on bounds, and small perturbations preserve every positive jump. Every finite-dimensional subset is therefore stable. Transforming back gives

$$\mathcal{V} = \{hu/\beta : u \in \mathcal{V}_g\}, \quad (154)$$

which gives a stable core for the IC critical cone. If the three moment gradients are independent on this face, the derivative map is onto and the full stability condition in part (iv) follows.

D.6 Equations for the binding pool under a reciprocal allocation rule

Return to $Z \sim U[0, 1]$, $\alpha = \beta = 1$, and $\mathcal{T} = [0, 1]$. For $q_\delta = 1/(1 + \delta z)$, write

$$h(z) = 1 + \delta z, \quad A = 1, \quad B = \frac{2}{1 + \delta}, \quad S = 1 + m.$$

When $L_q(m) < k < k_q^{\text{sat}}(m)$, completing the square in the g -Lagrangian shows that the weighted projection target is (100). Equivalently,

$$y_{\eta,\lambda}(z) = \frac{A_0}{h(z)} - \frac{B_0}{h(z)^2}, \quad A_0 = 1 - \frac{\eta}{2\lambda} + \frac{1}{2\lambda\delta}, \quad B_0 = \frac{1}{2\lambda\delta}. \quad (155)$$

Its derivative has the sign of $2B_0 - A_0h(z)$, so it is monotone or single peaked.

In the binding three-block regime (101), the five unknowns solve

$$\begin{aligned} y_{\eta,\lambda}(\ell) &= A, \\ y_{\eta,\lambda}(c) &= p, \\ \int_c^1 h(z)^2 [y_{\eta,\lambda}(z) - p] dz &= 0, \\ \int_0^1 h(z) g^*(z) dz &= S, \\ \int_0^1 [h(z) g^*(z) - 1]^2 dz &= k. \end{aligned} \quad (156)$$

The regime is strict when

$$y < A \text{ below } \ell, \quad A < y < p, \quad y' > 0 \text{ on } (\ell, c),$$

and

$$\int_c^x h(z)^2 [y_{\eta,\lambda}(z) - p] dz > 0 \quad (c < x < 1). \quad (157)$$

These inequalities verify the weighted projection variational inequality: the lower residual pairs negatively with every feasible upward perturbation, the free residual is zero, and on the tail pool Stieltjes integration turns the pairing into minus the integral of the positive cumulative residual against dg .

To compute the coefficient in closed form, define

$$(P_{\mathcal{V}}f)(z) = \begin{cases} 0, & z \leq \ell, \\ f(z), & \ell < z < c, \\ h(z) \frac{\int_c^1 h(x)f(x) dx}{\int_c^1 h(x)^2 dx}, & z \geq c. \end{cases} \quad (158)$$

This is the orthogonal projection on the critical schedule subspace: perturbations on the pooled tail are proportional to h . Let

$$\phi_j = P_{\mathcal{V}}[(r^*)^j], \quad G_{ij} = \langle \phi_i, \phi_j \rangle, \quad i, j = 0, 1, 2.$$

Then

$$\|\psi\|_2^2 = G_{22} - \begin{pmatrix} G_{02} & G_{12} \end{pmatrix} \begin{pmatrix} G_{00} & G_{01} \\ G_{01} & G_{11} \end{pmatrix}^{-1} \begin{pmatrix} G_{02} \\ G_{12} \end{pmatrix}. \quad (159)$$

For the exact benchmark, set

$$\begin{aligned} \delta &= \frac{4}{5}, & \ell &= \frac{1}{7}, & c &= \frac{5}{13}, & p &= \frac{771147}{708491}, \\ \lambda &= \frac{17712275}{65987064}, & \eta &= \frac{551420}{8248383}. \end{aligned} \quad (160)$$

Then direct substitution gives

$$y_{\eta,\lambda}(z) = \frac{1521(29848z + 10195)}{708491(4z + 5)^2}, \quad (161)$$

$$y'_{\eta,\lambda}(z) = -\frac{48672(3731z - 2115)}{708491(4z + 5)^3}. \quad (162)$$

The derivative is strictly positive through $c = 5/13 < 2115/3731$. Since $y(1/7) = 1$ and $y(5/13) = p$, the lower and free-block inequalities hold exactly. The pool residual and its cumulative are

$$h(z)^2[y_{\eta,\lambda}(z) - p] = -\frac{24336(13z - 5)(39z - 31)}{17712275}, \quad (163)$$

$$\int_{5/13}^x h(z)^2[y_{\eta,\lambda}(z) - p] dz = \frac{24336(1 - x)(13x - 5)^2}{17712275}. \quad (164)$$

The latter is strictly positive for $5/13 < x < 1$ and returns to zero at one. Also

$$1 < p = \frac{771147}{708491} < \frac{10}{9} = B.$$

Thus every strict condition is verified algebraically.

Define $m = \int r^*$ and $k = \int (r^*)^2$ for this profile. The profile differs from both the minimum-dispersion schedule with constant effective allocation and the bang-bang mean-constrained maximizer. The positive-multiplier KKT system is sufficient for the convex two-moment relaxation, so $L_q(m) < k < k_q^{\text{sat}}(m)$. Finally, the residual in (159) is strictly positive. Otherwise $(r^*)^2$ would be affine in $1, r^*$ against every perturbation supported on the free block, forcing the strictly increasing continuous r^* to take at most two values there. Numerical evaluation of this strictly positive residual gives

$$\|P_{\mathcal{W}^{\text{IC}}}[(r^*)^2]\|_2^2 \simeq 0.0004204694053. \quad (165)$$

For comparison, let $m_j = \int (r^*)^j$ and project only off the first two moment gradients. The residual without the local IC restriction is

$$\psi^0 = (r^*)^2 - a - br^*, \quad \|\psi^0\|_2^2 = m_4 - \begin{pmatrix} m_2 & m_3 \end{pmatrix} \begin{pmatrix} 1 & m_1 \\ m_1 & m_2 \end{pmatrix}^{-1} \begin{pmatrix} m_2 \\ m_3 \end{pmatrix}. \quad (166)$$

On the lower block $[0, \ell]$, every direction that preserves the active IC constraints vanishes, whereas $r^*(z) = \delta z$. The monic quadratic $(r^*)^2 - a - br^*$ cannot vanish on this entire interval. Therefore $\psi^0 \notin \mathcal{W}^{\text{IC}}$,

which verifies the strict condition in (99) algebraically. Numerical integration gives

$$(m_1, m_2, m_3, m_4) \simeq (0.5005341583, 0.3298779943, 0.2405755730, 0.1856826924),$$

$$\|\psi^0\|_2^2 \simeq 0.005095966168, \quad \Omega^{\text{IC}} \simeq 12.1197074. \quad (167)$$

Combining these values with $\lambda = 17712275/65987064$ yields the two coefficients in (103).

D.7 Verification of the two boundary examples

Verification of Example 1. The lower effective-allocation bound at the bottom and upper bound at the top both equal two. Monotonicity therefore forces $g(z) = 2$, hence $r(z) = z$.

The adjacent-protocol lottery with conditional mean z has

$$\mathbb{E}T^2 = \frac{19}{54}, \quad \mathbb{E}T^3 = \frac{5}{18}. \quad (168)$$

On $[0, 1/3]$, replacing it by the endpoint lottery with success probability z raises the second moment by $1/27$ and the third by $4/81$. Making the same replacement on $[2/3, 1]$ raises the second moment by $1/27$ and the third by $5/81$. The resulting policies have third moments $53/162$ and $55/162$, respectively. Both have $k = 7/18$ and the same forced output. Their mixtures give (104). $\square$

Verification of Example 2. Every marginal share vector satisfying the first two moments is indexed by $x = \pi_4 \in [1/6, 1/3]$:

$$(\pi_1, \pi_2, \pi_3, \pi_4) = \left(\frac{1}{2} - x, 3x - \frac{1}{2}, 1 - 3x, x \right). \quad (169)$$

Sorting these shares increasingly across uniform types is IC and output optimal for the given marginal distribution. Direct integration gives

$$\mathbb{E}[ZT] = \frac{17}{48} - \left(x - \frac{1}{4} \right)^2, \quad (170)$$

$$\mathbb{E}T^3 = \frac{5}{18} + \frac{2x}{9}. \quad (171)$$

Thus $x = 1/4$ is the unique two-moment optimum and selects $m_3^\dagger = 1/3$. Eliminating x gives (105) on the feasible interval. $\square$

References

- ACEMOGLU, D. AND J.-S. PISCHKE (1998): “Why Do Firms Train? Theory and Evidence,” *Quarterly Journal of Economics*, 113, 79–119.
- ACGME (2024): “Defined Category Minimums and Credit Role for General Surgery,” Accreditation Council for Graduate Medical Education.
- (2026): “ACGME Program Requirements for Graduate Medical Education in General Surgery Including FAQs,” Accreditation Council for Graduate Medical Education.
- ALTMAN, E. (1999): *Constrained Markov Decision Processes*, Boca Raton, FL: Chapman & Hall/CRC.
- ATHEY, S. AND I. SEGAL (2013): “An Efficient Dynamic Mechanism,” *Econometrica*, 81, 2463–2485.
- BACCARA, M., S. LEE, AND L. YARIV (2020): “Optimal Dynamic Matching,” *Theoretical Economics*, 15, 1221–1278.
- BARLOW, R. E., D. J. BARTHOLOMEW, J. M. BREMNER, AND H. D. BRUNK (1972): *Statistical Inference under Order Restrictions: The Theory and Application of Isotonic Regression*, New York: Wiley.
- BERGEMANN, D. AND J. VÄLIMÄKI (2019): “Dynamic Mechanism Design: An Introduction,” *Journal of Economic Literature*, 57, 235–274.
- CHE, Y.-K. (2026): “Dynamic Market Design,” Forthcoming in the Econometric Society monograph for the 2025 World Congress, vol. 2, chap. 3.
- CHE, Y.-K. AND A. B. CHOI (2025): “Optimal Auction Design for Dynamic Stochastic Environments: Myerson Meets Naor,” Working paper.
- CHE, Y.-K. AND O. TERCIEUX (2026): “Optimal Queue Design,” *Journal of Political Economy*, forthcoming.
- DWORCZAK, P. AND A. KOLOTILIN (2024): “The Persuasion Duality,” *Theoretical Economics*, 19, 1701–1755.
- DWORCZAK, P. AND G. MARTINI (2019): “The Simple Economics of Optimal Persuasion,” *Journal of Political Economy*, 127, 1993–2048.
- FARHI, E. AND I. WERNING (2007): “Inequality and Social Discounting,” *Journal of Political Economy*, 115, 365–402.
- FEDERAL INSTITUTE FOR VOCATIONAL EDUCATION AND TRAINING (2024): “Examination and Certification in Germany,” Federal Institute for Vocational Education and Training.
- (2026): *Verzeichnis der anerkannten Ausbildungsberufe 2026*, Bonn: Federal Institute for Vocational Education and Training.
- FEDERAL REPUBLIC OF GERMANY (2009): “Ordinance on Trainer Aptitude (Ausbilder-Eignungsverordnung),” Federal law.
- GARICANO, L. (2000): “Hierarchies and the Organization of Knowledge in Production,” *Journal of Political Economy*, 108, 874–904.
- HERNÁNDEZ-LERMA, O. AND J. B. LASSERRE (1996): *Discrete-Time Markov Control Processes: Basic Optimality Criteria*, New York: Springer.
- HOPENHAYN, H. A. AND E. C. PRESCOTT (1992): “Stochastic Monotonicity and Stationary Distributions for Dynamic Economies,” *Econometrica*, 60, 1387–1406.
- KARLIN, S. AND W. J. STUDDEN (1966): *Tchebycheff Systems: With Applications in Analysis and Statistics*, New York: Interscience Publishers.

- KLEINER, A., B. MOLDOVANU, AND P. STRACK (2021): “Extreme Points and Majorization: Economic Applications,” *Econometrica*, 89, 1557–1593.
- KOLOTLIN, A. (2018): “Optimal Information Disclosure: A Linear Programming Approach,” *Theoretical Economics*, 13, 607–635.
- MADSEN, E. AND E. SHMAYA (2025): “Collective Upkeep,” Working paper.
- MILGROM, P. AND I. SEGAL (2002): “Envelope Theorems for Arbitrary Choice Sets,” *Econometrica*, 70, 583–601.
- MUSSA, M. AND S. ROSEN (1978): “Monopoly and Product Quality,” *Journal of Economic Theory*, 18, 301–317.
- MYERSON, R. B. (1981): “Optimal Auction Design,” *Mathematics of Operations Research*, 6, 58–73.
- NÖLDEKE, G. AND L. SAMUELSON (2018): “The Implementation Duality,” *Econometrica*, 86, 1283–1324.
- PAVAN, A. (2017): “Dynamic Mechanism Design: Robustness and Endogenous Types,” in *Advances in Economics and Econometrics: Eleventh World Congress*, ed. by B. E. Honoré, A. Pakes, M. Piazzesi, and L. Samuelson, Cambridge: Cambridge University Press, vol. 1, 1–62.
- PAVAN, A., I. SEGAL, AND J. TOIKKA (2014): “Dynamic Mechanism Design: A Myersonian Approach,” *Econometrica*, 82, 601–653.
- PHELAN, C. (2006): “Opportunity and Social Mobility,” *Review of Economic Studies*, 73, 487–504.