

Screening and the Cost of a Prescribed Quantity Schedule

Dalin Sheng^{a,*}

^aSouthwestern University of Finance and Economics

555 Liutai Avenue, Wenjiang District, Chengdu, Sichuan 611130, China

*Corresponding author: shengdl@swufe.edu.cn

September 2026

Abstract

An institution assigns task actions to privately informed agents subject to a prescribed quantity schedule and the requirement of reproducing a stationary type distribution. Actions determine current output and shape successor types, so stationarity constrains the aggregate action distribution while incentive compatibility governs its allocation across private types. Incentives order the effective allocation: the prescribed quantity scaled by the action multiplier. Under a nonincreasing schedule, an action marginal is implementable if and only if it dominates in the convex order the distribution that holds this effective allocation constant. A stationary target is implementable if and only if some marginal above this bound reproduces it; we characterize when finitely many unconditional action statistics suffice for this test. For interval actions, the bound implies a sharp variance floor determined by the relative dispersion of prescribed quantities. Variance above the floor enables sorting until output attains its mean-constrained maximum, beyond which within-type lotteries absorb residual variance costlessly. Strict incentive compatibility preserves target moments at interior points and approaches the weak output maximum arbitrarily closely. Additional moment restrictions can lower output and preclude strict incentives at boundary targets. These losses stem from the quantity commitment. For nonnegative types and interval actions, freely chosen quantities allow monotone sorting to preserve any technologically feasible action marginal.

Keywords: mechanism design; screening; endogenous type distributions; stationary distributions; sufficient statistics; convex order.

JEL classification: C61; D02; D82.

1 Introduction

An institution staffs a fixed set of positions and has committed in advance to the quantity of work assigned to each reported productivity level. Collective agreements, civil-service grade structures, and statutory caseload or staffing ratios commonly fix workloads on a slower cycle than individual task assignments. The institution retains discretion over the character of the work—such as the

level of autonomy, task complexity, or supervisory oversight—and the associated compensation. These assignments play two roles: they determine current output and shape the pool of successors who will occupy the positions in the next period. An assignment rule designed to reproduce the desired composition of that pool need not induce truthful reporting of productivity. We characterize the mechanisms that satisfy both individual incentives and aggregate reproduction, and quantify the output loss imposed by the quantity commitment.

The tension reflects a structural conflict. The aggregate reproduction requirement constrains the marginal distribution of the action, whereas incentive compatibility orders the action scaled by a quantity the institution cannot adjust. When the prescribed quantity schedule declines in reported productivity, incentive compatibility forces dispersion into any action distribution capable of reproducing the population. This paper characterizes that dispersion and measures its output cost.

The institution takes the quantity schedule and the target type distribution as given. Each incumbent privately observes productivity, reports it, receives the prescribed quantity along with an action lottery, and makes a signed payment to the institution. After production, the incumbent departs, and a successor’s type is drawn from an action-dependent transition kernel. Incumbents are short-lived, so each faces a static reporting problem. Stationarity links successive cohorts by requiring the assignment rule to reproduce the target type distribution among incoming occupants.

A two-type example illustrates the tension. Low and high types each occupy half the positions, under a prescribed schedule with quantities $q_L = 2$ and $q_H = 1$. The action is $t \in \{0, 1\}$ and scales output by $1 + t$, where action 1 represents a complex, demanding assignment and action 0 a routine one. Let r_i denote the probability that type i is assigned action 1. With quasilinear utility and transfers, incentive compatibility requires the effective allocation—quantity multiplied by the expected output multiplier—to be nondecreasing in type. Because the low type already handles twice the volume, truthful reporting requires

$$2(1 + r_L) \leq 1 + r_H,$$

which admits only $r_L = 0$ and $r_H = 1$: exactly half the positions must receive action 1. Suppose that reproducing the target population instead requires assigning action 1 to one quarter of the positions. While such an aggregate share is technologically feasible, it cannot be implemented: truthful reporting cannot be sustained when higher actions are decoupled from reports. Thus, the quantity commitment, rather than the transition technology, places the stationary target out of reach.

Absent demographic constraints, the institution would raise actions as far as truthful reporting permits: for nonnegative types, higher actions weakly increase output and entail no direct resource cost. Stationarity constrains the aggregate action distribution, rationing an otherwise costless input. The design problem is how to allocate that input across private reports under the prescribed quantity schedule.

This boundary example isolates the sharpest form of the tension: the ratio of quantities matches

the maximal ratio of output multipliers, forcing the incentive constraint to bind and leaving a unique feasible assignment. A milder decline in the quantity schedule leaves an interval of feasible mean actions and a rich family of implementable action distributions. Under such a schedule, incentive compatibility selects among action distributions through dispersion alone, as offsetting the decline in quantity requires assigning lower actions to lower types and higher actions to higher types.

The same logic governs the general model. An agent reporting $\hat{z}$ receives quantity $q(\hat{z})$, and the action scales output by $\alpha + \beta t$. Let $r(z) = \mathbb{E}[T \mid Z = z]$ denote the conditional mean action under truthful reporting. With quasilinear preferences and unrestricted transfers, incentive compatibility is equivalent to monotonicity of the effective allocation

$$g(z) = q(z)[\alpha + \beta r(z)].$$

When q declines, higher types receive smaller quantities, so they must receive higher mean actions to keep g nondecreasing. Consequently, sustaining truthful reporting requires dispersing actions across types.

Our first main result gives an exact characterization. Fix a nonincreasing schedule q and a feasible mean action m , and let

$$r_m^0(z) = \frac{\alpha + \beta m}{\beta q(z) \mathbb{E}[1/q(Z)]} - \frac{\alpha}{\beta}, \quad \pi_m^0 = \mathcal{L}_\mu(r_m^0(Z))$$

denote the conditional-mean schedule that holds the effective allocation constant and the distribution it induces under the prior μ . An action marginal π with mean m is implementable under weak incentive compatibility if and only if $\pi_m^0 \preceq_{\text{cx}} \pi$ (Theorem 1). Implementable action distributions are precisely those at least as dispersed as π_m^0 in the convex order. Incentives bound dispersion from below; any upper bound arises solely from the action domain. The proof exploits a single-crossing property: any incentive-compatible conditional-mean schedule crosses r_m^0 at most once from below, which establishes convex-order dominance among conditional expectations. Conditional Jensen's inequality then extends the bound to realized actions. For sufficiency, Strassen's theorem (Strassen, 1965) constructs a lottery assignment with marginal π and conditional mean r_m^0 . This characterization requires no topological structure on the action set beyond closedness; when the action set is finite, it reduces to a finite system of linear inequalities.

Combining this convex-order bound with the transition technology characterizes stationary implementation. When the action shifts successor types through a component common to all current types, the action marginal alone determines the successor distribution. The institution can then reproduce the population if and only if some marginal that replicates the target dominates π_m^0 in the convex order (Corollary 1). This criterion restricts the transition family only through weak continuity.

The dimension of aggregate information required by the institution is governed by the transition technology. For jointly Feller transitions, finitely many unconditional action statistics determine the successor distribution if and only if action-induced kernel differences are independent of the current

type and span a finite-dimensional space (Theorem 2). The dimension of that space is the minimal number of statistics needed. Polynomial transitions with linearly independent coefficient measures form the canonical class, as stationarity then fixes the corresponding action moments. In particular, a quadratic transition law reduces the design problem to a two-moment allocation problem. When kernel differences depend on the current type, no finite-dimensional statistic of the action alone suffices, and the successor distribution must be tracked through joint moments of type and action.

For actions in $[0, 1]$ and a nonincreasing schedule q , the convex-order bound implies a sharp variance floor. At a feasible mean m , incentive compatibility requires an action variance of at least

$$\left(m + \frac{\alpha}{\beta}\right)^2 \text{CV}^2\left(\frac{1}{q}\right),$$

attained by the schedule that holds the effective allocation constant (Corollary 2). Two economic forces determine this floor. The squared coefficient of variation of $1/q$ measures the relative dispersion of the prescribed quantity schedule. The scale factor $(m + \alpha/\beta)^2$ decreases in β : when the action exerts a stronger effect on output, less dispersion is needed to offset a given decline in q . The floor vanishes if and only if the quantity schedule is constant almost everywhere. A specified action marginal must still satisfy the full convex-order condition; its variance alone does not determine implementability.

Above the floor, variance is partitioned into distinct economic regimes: it initially enhances sorting and raises output, but eventually becomes redundant and is absorbed costlessly by within-type lotteries. Specifically, additional variance permits greater separation across types: assigning higher actions to higher types increases output until it attains the mean-constrained optimum. Beyond this threshold, further variance is absorbed by within-type randomization without affecting output. We characterize the value function for general quantity schedules and closed action sets, and compute both thresholds in closed form when quantities decline (Theorems 3 and 4). These comparative statics vary the transition technology while holding the population fixed, so that stationarity calls for a different second moment of the same population.

Strict incentive compatibility can be attained at an arbitrarily small output cost. At any interior feasible mean and any feasible second moment strictly above the floor, we construct strictly incentive-compatible mechanisms that match both moments exactly. Mixing such a mechanism with a weak optimum makes truthful reporting uniquely optimal and approaches the weak output maximum (Proposition 4). With interval actions, these slackness conditions are also necessary, so strict incentives preclude only boundary means and the minimum second moment. They remain available at the maximal second moment, where the output supremum need not be attained.

Richer transition technologies impose higher-order moment restrictions, and an additional restriction operates through two distinct channels. First, it can preclude implementation entirely: on a four-action set, a mean and second moment may be jointly implementable while the third moment required by stationarity is not (Example 2). Second, when implementation remains feasible, the additional restriction can reduce output. For interval actions and a two-moment optimum in

the productive range, Theorem 5 provides an exact three-channel decomposition of this output loss using dual multipliers: the reallocation of conditional means, the variance lost to within-type randomization, and the shadow cost of binding incentive and boundary constraints. Strict incentives survive whenever the additional target lies in the relative interior of the attainable moment set (Corollary 3). At its boundary, however, a target can force pooling even when it is interior to the unrestricted moment polytope.

While the benchmark objective is current gross output with internal transfers, the value and loss results in Sections 5–7 also apply to revenue after replacing $zq(z)$ with the virtual-value weight $\phi(z)q(z)$. The upper-tail characterization then requires ϕq to be strictly increasing.

These efficiency losses stem from the quantity commitment rather than from stationarity itself. An institution free to set quantities, facing nonnegative types and interval actions, would set every quantity at the common upper bound and sort actions monotonically. This unconstrained policy preserves the entire marginal distribution of actions and hence any technologically feasible stationary target under the common transition law (Proposition 2). A prescribed schedule eliminates that flexibility. What remains is the dispersion the schedule forces into the action distribution and the output the institution forgoes in producing it.

Related literature. The closest formulation is that of Bergemann et al. (2026b), where an exogenous inventory of qualities constrains screening and feasibility is expressed through majorization. Our model separates the constrained object from the ordered one: the distributional restriction falls on the physical action T , while incentives order the effective allocation $g = q(\alpha + \beta r)$, and the prescribed schedule links the two through $r = (g/q - \alpha)/\beta$. The two restrictions also run in opposite directions. An inventory is scarcity, an endowment the designer cannot exceed, and it bounds the allocation from above in majorization. Here the restriction is an equality target produced by a fixed point across cohorts, and incentives bound the action marginal from below in the convex order, with any ceiling supplied by the action set rather than by supply. A feasibility test built for a scarcity ceiling does not answer which populations a prescribed schedule puts out of reach. Utility keeps the standard screening form $zg - p$, so the entire departure lies in the schedule: it decides which action assignments deliver a monotone g , and that is what generates the convex-order bound. Stationarity then selects the marginals above the bound that also reproduce the target population.

The characterization builds on recent advances in mechanism design and monotone optimization under distributional constraints (Kleiner et al., 2021; Nikzad, 2022; Yang and Zentefis, 2024; Kleiner et al., 2026). The main departure is the constructive derivation of the lower-bound distribution π_m^0 directly from the prescribed quantity schedule, which makes the feasibility test explicit. Strassen’s theorem, in the kernel formulation of Leskelä and Vihola (2017), establishes the existence of the corresponding lottery implementation.

Aggregate restrictions serve as an incentive instrument in the literature on linked decisions and quotas (Jackson and Sonnenschein, 2007; Frankel, 2014, 2016; Escobar and Toikka, 2013; Ball and Kattwinkel, 2026; Bergemann et al., 2026a). In our environment, the economic role of the

constraint is reversed: the aggregate restriction is not an instrument designed to discipline reports, but an exogenous demographic requirement dictated by the transition technology, and the prescribed quantity schedule constrains how the designer can meet it.

By modeling how current task assignments train successors, the paper connects to the literature on task allocation and expertise accumulation (Baccara et al., 2023) and on apprenticeship contracts (Fudenberg and Rayo, 2019). Dynamic mechanism design typically studies long-lived agents whose allocations affect their own future private information (Pavan et al., 2014; Makris and Pavan, 2021). Here, incumbents are short-lived and their actions affect the types of incoming agents, so each incumbent faces a static screening problem while stationarity accounts for the entire intertemporal link. Baccara et al. (2020) and Leshno (2022) study related stationary allocation problems.

The incentive reduction uses standard screening and envelope arguments (Mussa and Rosen, 1978; Myerson, 1981; Milgrom and Segal, 2002). The optimization uses moment geometry (Karlin and Studden, 1966), isotonic regression (Barlow et al., 1972), and ironing (Toikka, 2011). The dual formulation is related to Nöldeke and Samuelson (2018); Dworczak and Kolotilin (2024).

Section 2 sets out the model and the incentive constraint. Sections 3 and 4 characterize implementable action marginals and the transition laws for which finitely many action statistics suffice. Section 5 derives optimal output under two moment restrictions. Sections 6 and 7 take up strict incentives and additional moment restrictions. Proofs are in the appendices.

2 Model and Incentive Compatibility

We consider a dynamic framework that couples a static screening problem within each period to an intertemporal law of motion across successive cohorts. Incumbents face a one-shot private reporting problem, and the task assignments selected today shape the productivity distribution of tomorrow's cohort. Stationarity imposes an aggregate constraint on the distribution of actions, while the prescribed quantity schedule governs how actions can be assigned to types. The efficiency losses characterized below stem entirely from this quantity schedule. We begin with polynomial transition laws, under which stationarity reduces to a finite system of moment restrictions on the action distribution. Section 4 characterizes the general class of transition laws that admit such a finite-dimensional reduction.

2.1 Environment and preferences

Time is discrete. A unit mass of positions is occupied each period by short-lived incumbents. Their types have distribution μ on the compact interval $\mathcal{Z} = [\underline{z}, \bar{z}]$. Each incumbent privately observes type z and reports $\hat{z}$. The institution must reproduce μ among the next cohort of occupants.

Report $\hat{z}$ determines the incumbent's prescribed quantity $q(\hat{z})$, action lottery, and signed payment to the institution. The quantity schedule $q : \mathcal{Z} \rightarrow (0, \bar{q}]$ is continuous and committed prior to the choice of lotteries and payments. Each report receives exactly the workload prescribed by the

schedule. The action set $\mathcal{T} \subset [0, 1]$ is a nonempty closed set with endpoints

$$a := \min \mathcal{T} < b := \max \mathcal{T}.$$

An assignment rule $\sigma(dt \mid \hat{z})$ specifies a Borel probability measure on $\mathcal{T}$, and $p(\hat{z})$ is a signed payment from the incumbent to the institution. We use lowercase t for a realized action and uppercase T for the induced random variable.

For true type z , report $\hat{z}$, and action t , gross output is $zq(\hat{z})(\alpha + \beta t)$. The incumbent's utility is

$$zq(\hat{z})(\alpha + \beta t) - p(\hat{z}), \quad \beta > 0, \quad \alpha + \beta a > 0. \quad (1)$$

Payoff terms independent of the true type are absorbed into the transfer. The outside option is normalized to zero.

The conditional mean action determines both expected utility and expected output:

$$r_\sigma(z) := \int_{\mathcal{T}} t \sigma(dt \mid z).$$

Define the effective allocation and the output weight by

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

A type z reporting $\hat{z}$ obtains expected utility $zg_\sigma(\hat{z}) - p(\hat{z})$. Thus, g_σ plays the role of the allocation schedule in the associated one-dimensional screening problem. For any conditional-mean schedule r , we write $g_r := q(\alpha + \beta r)$.

The institution commits to the mechanism before incumbents observe their types. The incumbent observes z , reports $\hat{z}$, receives the prescribed quantity and an action drawn from $\sigma(\cdot \mid \hat{z})$, and makes payment $p(\hat{z})$. After production, the incumbent departs. The successor's type is drawn from the transition kernel specified below.

The action represents the qualitative character of the assignment: the degree of autonomy, the novelty of assigned tasks, or the intensity of supervision. It operates through two distinct channels. First, it scales current output through the multiplier $\alpha + \beta t$. Second, it develops human capital for the incoming cohort: more demanding assignments provide richer training opportunities that prepare successors. Because incumbents depart after one period, they capture none of this human capital return; their reporting problem is strictly static, and the training channel operates purely through the stationarity constraint.

2.2 Polynomial transitions and stationarity

The action-dependent component of the next-period type distribution is

$$\widehat{Q}_t = R_0 + \sum_{\ell=1}^N t^\ell A_\ell, \quad (2)$$

where $R_0 := \widehat{Q}_0$. The signed measures $A_1, \dots, A_N$ have total mass zero and are linearly independent, and $\widehat{Q}_t$ is a probability measure for every $t \in [0, 1]$.

The action-induced component is common across current types. Type persistence is captured by an action-independent Markov kernel S_0 :

$$Q_t(dz' | z) = \zeta S_0(dz' | z) + (1 - \zeta) \widehat{Q}_t(dz'), \quad \zeta \in [0, 1], \quad \mu S_0 = \mu. \quad (3)$$

Under truthful reporting, an assignment rule σ induces the next-period distribution

$$\mu P_\sigma := \int_{\mathcal{Z}} \int_{\mathcal{T}} Q_t(\cdot | z) \sigma(dt | z) \mu(dz). \quad (4)$$

Stationarity requires $\mu P_\sigma = \mu$.

A target distribution reproducible by the transition technology has a unique representation

$$\mu = R_0 + \sum_{\ell=1}^N m_\ell A_\ell, \quad (m_1, \dots, m_N) \in \mathcal{M}_N(\mathcal{T}), \quad (5)$$

where $\mathcal{M}_N(\mathcal{T})$ denotes the moment space generated by probability measures on $\mathcal{T}$ (Karlin and Studden, 1966). Integrating (3) and using $\mu S_0 = \mu$ yields

$$\mu P_\sigma = \zeta \mu + (1 - \zeta) \left[R_0 + \sum_{\ell=1}^N \mathbb{E}[T^\ell] A_\ell \right].$$

Linear independence of the coefficient measures ensures that stationarity is equivalent to the moment equalities

$$\mathbb{E}[T^\ell] = m_\ell, \quad \ell = 1, \dots, N. \quad (6)$$

This corresponds to the polynomial case of Theorem 2(iii). These equalities constitute a demographic reproduction requirement rather than a discretionary performance target. While a higher mean action would raise current output, it would also alter the composition of the successor pool, violating stationarity. Condition (6) thus fixes the aggregate action resources available to the designer. The subsequent sections examine the efficiency cost of distributing this fixed aggregate across types under the prescribed quantity schedule.

Assumption 1 (Benchmark environment). *(i) The type space $\mathcal{Z}$ is compact; μ is atomless with full support; q is positive and continuous; and $\mathcal{T}$ is a nonempty closed set with distinct endpoints.*

- (ii) Incumbents are short-lived, preferences are given by (1), transfers are unrestricted, and the action enters current payoff through $\alpha + \beta t$.
- (iii) The transition law satisfies (2) and (3); the target has the unique representation (5).
- (iv) Stationarity requires $\mu P_\sigma = \mu$ after one period, with μP_σ defined in (4). The quantity schedule q and target μ are fixed.

Two features of Assumption 1 warrant economic emphasis. Short-lived incumbents in (ii) isolate the static screening dimension: because the action trains a successor rather than returning to the incumbent, the effective allocation $q(\alpha + \beta r)$ fully determines agent incentives, and the envelope formula contains no continuation value. This specification naturally captures rotating professional appointments—such as judicial clerkships, medical residencies, and junior civil-service postings—where incumbents perform fixed-term assignments that prepare the next generation. Second, part (iv) imposes exact period-by-period reproduction, matching institutions governed by periodic accreditation standards or cohort quotas. The action moments implied by stationarity must therefore be met for each cohort.

Unless stated otherwise, expectations are taken over $Z \sim \mu$ and the conditional law of T . We write $\langle f, g \rangle = \int f g \, d\mu$ whenever the product is integrable, and use $\|f\|_2$ for the $L^2(\mu)$ norm. In finite-dimensional spaces, brackets denote the Euclidean inner product. We write co for convex hull, ri for relative interior, δ_t for the Dirac measure at t , and $\mathcal{L}_\mu(X)$ for the law of X under μ .

A concrete microfoundation for this preparation channel generates polynomial transitions: the quantity specifies a base volume of tasks, and the action governs the success probability at each training stage.

Example 1 (Independent preparation stages). A successor passes through N preparation stages. Conditional on action t , each stage succeeds independently with probability t . Let R_j denote the next-period type distribution after exactly j successes. Then

$$\widehat{Q}_t^{(N)} = \sum_{j=0}^N \binom{N}{j} t^j (1-t)^{N-j} R_j.$$

This family belongs to class (2); its coefficient measures are linearly independent if and only if $R_0, \dots, R_N$ are affinely independent.

Let J denote the number of successful stages and b_j the cohort share with $J = j$. Then

$$b_j = \binom{N}{j} \mathbb{E}[T^j (1-T)^{N-j}].$$

For $N = 3$, write $m = m_1$ and $k = m_2$. The cohort shares are

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

The corresponding factorial moments are $\mathbb{E}[J] = 3m$, $\mathbb{E}[J(J-1)] = 6k$, and $\mathbb{E}[J(J-1)(J-2)] = 6m_3$.

Thus, m , k , and m_3 represent the success rates of a randomly selected stage, ordered pair of stages, and ordered triple of stages; $m_3 = b_3$ is the fully trained share.

The moment-space condition $(m_1, \dots, m_N) \in \mathcal{M}_N(\mathcal{T})$ is strictly stronger than nonnegativity of the cohort shares. For example, $(b_0, b_1, b_2, b_3) = (0, 1, 0, 0)$ implies $m = 1/3$ and $k = 0 < m^2$, which cannot be generated by any probability distribution on T . If R_{j+1} first-order stochastically dominates R_j for each j , then $\widehat{Q}_{t'}^{(N)}$ first-order stochastically dominates $\widehat{Q}_t^{(N)}$ whenever $t' > t$.

The stage model provides a concrete interpretation: the ℓ th moment is the probability that an ordered ℓ -tuple of stages all succeed, so stationarity prescribes target training rates. This moment interpretation extends to any transition satisfying (2) and (3).

2.3 Quantity commitment and objective

The institution selects (σ, p) to implement the prescribed schedule q while preserving μ . The quantity schedule is fixed in advance: each report commands a predetermined workload. Only the action lottery and the transfer schedule remain within the designer's control, and the mechanism (q, σ, p) must induce truthful reporting.

Schedules fixed this way typically decline in reported productivity because institutions use task volume to equalize baseline effort across grades. Routine, high-volume responsibilities are concentrated in lower tiers, whereas higher tiers receive greater autonomy and discretion over task execution rather than larger physical caseloads.

A declining schedule compensates for lower reported productivity with a larger volume of base tasks. On $[0, 1]$, for instance, $q_\delta(z) = (1 + \delta z)^{-1}$ with $0 < \delta < 1$ assigns higher reports fewer tasks while ensuring that the output weight $zq_\delta(z) = z/(1 + \delta z)$ remains strictly increasing. Assigning more productive actions to higher types remains efficient under such a schedule; what the schedule eliminates is the ability to reward higher types through volume.

For nonnegative types, quantities represent productive opportunities: gross utility and gross output both equal $zq(\widehat{z})(\alpha + \beta t)$, and transfers divide this output between the incumbent and the institution.

Because preserving μ is an explicit constraint of the design problem, a stationary assignment rule reproduces μ each period. Its discounted surplus is $(1 - \rho)^{-1}$ times current surplus for any discount factor $\rho \in (0, 1)$. Ranking stationary rules by current output is therefore strategically equivalent to ranking them by discounted surplus, justifying the objective below.

Assumption 2 (Current-period design criterion). *For fixed q and μ , the institution maximizes current output. Transfers are internal to the incumbent-institution pair. Randomization over actions has no direct resource or risk cost.*

The results draw on these primitives at different levels of generality. The characterization of implementable action distributions requires q to be nonincreasing. The two-moment value theorem holds for any positive continuous q and any closed action set. Closed-form thresholds require interval actions and a nonincreasing q to yield an explicit minimum second moment, and a strictly increasing

output weight $w(z) = zq(z)$ to determine the second moment that achieves the mean-constrained output maximum. The transition assumptions vary analogously: Assumption 1(iii) establishes the polynomial benchmark; Section 3.2 requires only weak continuity of a common action-induced kernel; Section 4 characterizes when a finite-dimensional reduction exists under joint continuity; and Sections 5–7 analyze quadratic and higher-order moment problems.

2.4 Incentive compatibility

Let $L_{\uparrow}^{\infty}(\mu)$ denote the space of essentially bounded functions with a nondecreasing representative. A direct mechanism *implements q while preserving μ* if truthful reporting is globally incentive compatible and individually rational, and the next-period distribution is μ . With unrestricted transfers, an assignment rule admits incentive-compatible payments if and only if its effective allocation g_{σ} is nondecreasing. Payments then satisfy the standard envelope formula.

Stationarity is evaluated under truthful reporting. Weak incentive compatibility makes truthfulness a best response; Section 6 requires it to be the unique best response.

For a nondecreasing effective allocation g , transfers satisfy $V(z) = V_0 + \int_{\underline{z}}^z g(x) dx$ and $p(z) = zq(z) - V(z)$. Since $g > 0$, individual rationality is equivalent to $V_0 \geq 0$. Online Appendix D establishes duality under moment and monotonicity constraints.

2.5 Feasible mean actions and freely chosen quantities

Define

$$\Lambda_{\mathcal{T}} := \frac{\alpha + \beta b}{\alpha + \beta a},$$

which represents the ratio of the maximal to minimal output multiplier. On the unit interval, $\Lambda_{[0,1]} = 1 + \beta/\alpha$.

Proposition 1 (Range condition for incentive compatibility). *An incentive-compatible conditional-mean schedule exists if and only if*

$$\sup_z \frac{\max_{\tilde{z} \leq z} q(\tilde{z})}{q(z)} \leq \Lambda_{\mathcal{T}}. \quad (7)$$

If q is nonincreasing, the left side is $q(\underline{z})/q(\bar{z})$. In that case define

$$g_L := (\alpha + \beta a)q(\underline{z}), \quad g_U := (\alpha + \beta b)q(\bar{z}).$$

The range condition is $g_L \leq g_U$. Every mean action m for which

$$\frac{\alpha + \beta m}{\mathbb{E}[1/q(Z)]} \in [g_L, g_U]$$

is attainable by an incentive-compatible schedule.

Proposition 1 characterizes when a nondecreasing effective allocation can be fitted between the pointwise bounds $(\alpha + \beta a)q$ and $(\alpha + \beta b)q$ (cf. Yang and Zentefis, 2024, Remark 1). Economically, it compares the range of the output multiplier with the variation in the quantity schedule: the action multiplier must be flexible enough to absorb the maximal proportional decline in q . Feasibility of the mean is a precondition, after which the design problem must accommodate the higher-order moment restrictions required by stationarity.

The economic cost of the quantity commitment is underscored by removing it.

Proposition 2 (A freely chosen quantity schedule). *Maintain Assumption 1 except that q is not fixed. Suppose $\underline{z} \geq 0$, $\mathcal{T} = [0, 1]$, and Assumption 2 holds. The institution chooses q and the assignment rule subject to $0 < q \leq \bar{q}$, while preserving the target μ .*

- (i) *Every feasible assignment rule is weakly dominated in output by one with $q \equiv \bar{q}$ and a deterministic nondecreasing action assignment that has the same marginal distribution of T . Hence $q \equiv \bar{q}$ is optimal and preserves every aggregate moment of the action distribution.*
- (ii) *In the auxiliary problem with only $\mathbb{E}[T] = m$, $m \in [0, 1]$, imposed, the least second moment compatible with incentive compatibility is m^2 , attained by $T \equiv m$. The range condition in (7) then holds with strict slack.*

Under the conditions of Proposition 2, setting $q \equiv \bar{q}$ and sorting actions comonotonically with type satisfies incentive compatibility while preserving any given action marginal. Thus, screening adds no restriction on technologically feasible stationary targets. All dispersion requirements and output losses derived below arise from giving up that freedom: under a prescribed schedule q , sorting must operate through the monotonicity of $q(\alpha + \beta r)$. We now characterize the action distributions that can be implemented under this restriction.

3 Implementable Action Distributions

Under a nonincreasing quantity schedule, incentive compatibility restricts the action distribution in a single way: it imposes a lower bound in the convex order. This section derives the bound, shows that nothing else is required, and combines it with the transition law to obtain a criterion for stationary implementation.

3.1 A convex-order lower bound

Maintain the screening primitives in Assumption 1(i)–(ii), and suppose q is nonincreasing. In particular, $\mathcal{T} \subset [0, 1]$ is closed, its endpoints satisfy $a < b$, and $\alpha + \beta a > 0$.

Let $\mathcal{P}(\mathcal{T})$ be the Borel probability measures on $\mathcal{T}$. A marginal $\pi \in \mathcal{P}(\mathcal{T})$ is *implementable* if a Borel assignment rule with marginal π has a nondecreasing effective allocation $g_\sigma = q(\alpha + \beta r_\sigma)$. Such an assignment can be paired with transfers to satisfy weak incentive compatibility. For probability

measures π_0, π_1 on $[a, b]$, write $\pi_0 \preceq_{\text{cx}} \pi_1$ if they have the same mean and

$$\int \phi(t) \pi_0(dt) \leq \int \phi(t) \pi_1(dt) \quad \text{for every continuous convex } \phi : [a, b] \rightarrow \mathbb{R}.$$

Convex order ranks dispersion at a common mean.

For $m \in [a, b]$, define

$$H := \mathbb{E}[1/q(Z)], \quad \bar{g}_m := \frac{\alpha + \beta m}{H},$$

$$r_m^0(z) := \frac{\bar{g}_m/q(z) - \alpha}{\beta}, \quad \pi_m^0 := \mathcal{L}_\mu(r_m^0(Z)).$$

The schedule r_m^0 has mean m and holds the effective allocation constant at $\bar{g}_m$. Set

$$g_L := (\alpha + \beta a)q(\underline{z}), \quad g_U := (\alpha + \beta b)q(\bar{z}).$$

Theorem 1 (Implementable action marginals). *Under the screening primitives stated above, suppose q is nonincreasing.*

(i) *A mean action $m \in [a, b]$ is attainable by an incentive-compatible assignment if and only if*

$$g_L \leq \bar{g}_m \leq g_U. \quad (8)$$

Equivalently, $r_m^0(z) \in [a, b]$ for every z .

(ii) *Suppose (8) holds. For every $\pi \in \mathcal{P}(\mathcal{T})$ with mean m , the following are equivalent:*

- (a) *π is implementable;*
- (b) *$\pi_m^0 \preceq_{\text{cx}} \pi$;*
- (c) *for every $s \in [a, b]$,*

$$\mathbb{E}[(s - r_m^0(Z))_+] \leq \int_{\mathcal{T}} (s - t)_+ \pi(dt), \quad x_+ := \max\{x, 0\}.$$

Every implementable marginal has an implementation satisfying $\mathbb{E}[T \mid Z] = r_m^0(Z)$ and $g_\sigma \equiv \bar{g}_m$.

(iii) *If $\mathcal{T} = \{t_1 < \dots < t_J\}$, then, for probabilities $p_j \geq 0$ with $\sum_j p_j = 1$ and $\sum_j p_j t_j = m$, the marginal $\sum_j p_j \delta_{t_j}$ is implementable if and only if*

$$\sum_{i=1}^J p_i (t_j - t_i)_+ \geq \mathbb{E}[(t_j - r_m^0(Z))_+], \quad j = 2, \dots, J-1. \quad (9)$$

Condition (8) also applies in this part. For $J = 2$, there are no interior inequalities.

The schedule that holds the effective allocation constant represents the least dispersed conditional-mean schedule consistent with incentive compatibility. For any incentive-compatible schedule r , the

difference $r - r_m^0 = (g_r - \bar{g}_m)/(\beta q)$ has the same sign as the nondecreasing function $g_r - \bar{g}_m$, so it changes sign at most once, from negative to positive. Because two nondecreasing schedules with a common mean that cross once from below are ranked in the convex order, this single-crossing property establishes the lower bound on conditional means. Conditional Jensen's inequality then extends the bound to realized actions. Conversely, sufficiency follows from Strassen's theorem (Strassen, 1965): in the kernel formulation of Leskelä and Vihola (2017, Theorem 1.1), any marginal π that dominates π_m^0 in the convex order can be implemented by an assignment rule whose conditional mean is exactly r_m^0 .

On an interval $\mathcal{T} = [a, b]$, the deterministic assignment $T = r_m^0(Z)$ is feasible. Its induced distribution π_m^0 is the minimal implementable marginal in the convex order and minimizes the expectation of every continuous convex function of the action at mean m . If $\mathcal{T}$ exhibits discrete gaps, r_m^0 remains an implementable conditional-mean schedule, although its law may assign mass outside $\mathcal{T}$. Part (ii) characterizes all implementable marginals supported on $\mathcal{T}$ in both cases.

An assignment that holds g_σ constant leaves every agent indifferent across reports. Alternative implementations of the same marginal may raise output through further sorting. The economic force behind the lower bound is the decline in the prescribed quantity schedule. If q were nondecreasing, the constant assignment $r(z) \equiv m$ would be incentive compatible. Its degenerate law is weakly dominated by π_m^0 in the convex order, with strict dominance when q is nonconstant.

3.2 Stationary targets and the marginal test

Let the transition law take the form

$$Q_t(\cdot | z) = \zeta S_0(\cdot | z) + (1 - \zeta)\hat{Q}_t(\cdot), \quad 0 \leq \zeta < 1, \quad \mu S_0 = \mu, \quad (10)$$

where $\{\hat{Q}_t : t \in \mathcal{T}\}$ is any weakly continuous family of probability measures. For a feasible mean m , define

$$\begin{aligned} \mathcal{A}_\mu(m) &:= \left\{ \pi \in \mathcal{P}(\mathcal{T}) : \int t\pi(dt) = m, \quad \int \hat{Q}_t \pi(dt) = \mu \right\}, \\ \mathcal{U}_q(m) &:= \left\{ \pi \in \mathcal{P}(\mathcal{T}) : \pi_m^0 \preceq_{\text{cx}} \pi \right\}. \end{aligned} \quad (11)$$

Equality between measures in (11) means equality on every Borel set of successor types.

Corollary 1 (Stationary implementation through action marginals). *Maintain the hypotheses of Theorem 1 and (10). Fix a mean m satisfying (8). There exists a weakly incentive-compatible assignment of mean m that preserves μ if and only if*

$$\mathcal{A}_\mu(m) \cap \mathcal{U}_q(m) \neq \emptyset. \quad (12)$$

For polynomial transitions with independent coefficient measures and target $\mu = R_0 + \sum_{\ell=1}^N m_\ell A_\ell$,

this is equivalent to the existence of $\pi \in \mathcal{P}(\mathcal{T})$ such that

$$\pi_{m_1}^0 \preceq_{\text{cx}} \pi, \quad \int t^\ell \pi(dt) = m_\ell \quad (\ell = 1, \dots, N), \quad (13)$$

provided the mean m_1 satisfies (8).

When the target leaves the mean action undetermined, this criterion applies at each mean satisfying (8). With a finite action set, stationary implementation reduces to a standard linear programming feasibility problem: find probabilities $p_j \geq 0$ that sum to one, match the target moments, and satisfy (9). The transition technology determines what stationarity requires of the action distribution, while the convex-order bound delineates what incentives permit under the prescribed schedule.

Example 2 (Infeasibility under a third-moment restriction). Let $Z \sim U[0, 1]$, $q(z) = 1/(1 + z/2)$, $\alpha = \beta = 1$, and $\mathcal{T} = \{0, 3/10, 7/10, 1\}$. At $m = 1/2$,

$$r_m^0(z) = \frac{1}{5} + \frac{3}{5}z.$$

Consider the marginal

$$\pi^\circ = \frac{1}{100}\delta_0 + \frac{211}{400}\delta_{3/10} + \frac{161}{400}\delta_{7/10} + \frac{3}{50}\delta_1.$$

Its first three moments are

$$(m_1, m_2, m_3) = \left(\frac{1}{2}, \frac{3047}{10000}, \frac{2123}{10000} \right).$$

The first two moments are implementable on this four-action set. Lotteries between adjacent actions with conditional mean r_m^0 induce the marginal

$$\pi^- := \frac{1}{36}\delta_0 + \frac{17}{36}\delta_{3/10} + \frac{17}{36}\delta_{7/10} + \frac{1}{36}\delta_1, \quad \int t^2 \pi^-(dt) = \frac{181}{600}.$$

Mix this assignment with the endpoint lottery having the same mean at each type. A weight of $13/850$ on the endpoint lottery gives second moment $3047/10000$.

The marginal π° violates the convex-order inequality at $s = 3/10$:

$$\int (3/10 - t)_+ \pi^\circ(dt) = \frac{3}{1000} < \frac{1}{120} = \mathbb{E}[(3/10 - r_m^0(Z))_+].$$

On four distinct actions, normalization and the first three moments uniquely determine the probabilities. The third-moment target therefore forces π° and rules out incentive-compatible implementation.

In the context of the training technology in Example 1, the example shows that a finer multi-stage training process restricts the designer rather than providing additional instruments. With N preparation stages, stationarity requires the action marginal to satisfy N simultaneous moment restrictions while maintaining convex-order dominance over π_m^0 . Consequently, each additional

stage adds a moment constraint without expanding the designer's degrees of freedom. Two stages leave this target attainable; three stages render it impossible under the same quantity schedule and target population. Appendix B verifies the computations and constructs a cubic transition law that imposes these three moments under stationarity.

4 Sufficient Statistics for Stationarity

Corollary 1 characterizes stationary implementation via the action marginal alone, exploiting the property that transition law (10) allows the marginal distribution of actions to fully determine the successor distribution. Whether finitely many unconditional action statistics determine every successor distribution is a property of the transition technology. Such statistics suffice if and only if action-induced kernel differences are independent of current types and span a finite-dimensional space. Its dimension is the minimal number of statistics needed to identify the successor law. Equation (6), which supplies the moment restrictions analyzed from Section 5 onward, emerges as the canonical polynomial instance of this representation.

Let $\mathcal{Z}$ be a compact interval with probability measure μ , and let $\mathcal{T}$ be compact. A transition family is *jointly Feller* if $(z, t) \mapsto \int f(z')Q_t(dz' | z)$ is continuous for every $f \in C(\mathcal{Z})$. For a continuous $x : \mathcal{T} \rightarrow \mathbb{R}^d$, write $\bar{x}_\sigma := \mathbb{E}_\sigma[x(T)]$, with μP_σ defined by (4). The statistic is *sufficient* if $\bar{x}_\sigma = \bar{x}_{\sigma'}$ implies $\mu P_\sigma = \mu P_{\sigma'}$ for all Borel assignment rules. It is sufficient *among monotone assignments* if this implication holds when both conditional means of $x_1(T)$ are nondecreasing. A sufficient statistic is *minimal* if its image has the smallest affine dimension.

Fix $t_0 \in \mathcal{T}$. If every difference $Q_t(\cdot | z) - Q_{t_0}(\cdot | z)$ is independent of z , denote it by Δ_t and define

$$\mathcal{D}_Q := \text{span}\{\Delta_t : t \in \mathcal{T}\}, \quad d_Q := \dim \mathcal{D}_Q.$$

If any kernel difference depends on current type, set $d_Q = +\infty$. Thus infinite dimension covers both type-dependent effects and a family of type-independent differences with infinite-dimensional span. The definition is independent of the reference action.

Theorem 2 (Finite-dimensional sufficient statistics). *Let $\mathcal{T}$ be compact, let μ be atomless with full support on the compact interval $\mathcal{Z}$, and let $\{Q_t : t \in \mathcal{T}\}$ be jointly Feller.*

(i) *Suppose that $x : \mathcal{T} \rightarrow \mathbb{R}^d$ is continuous, $\text{aff } x(\mathcal{T}) = \mathbb{R}^d$, and x_1 is nonconstant. The following statements are equivalent.*

- (a) *x is sufficient for the next-period distribution.*
- (b) *x is sufficient among monotone assignments.*
- (c) *For every reference action t_0 , there are unique signed measures $D_1, \dots, D_d$, independent of current type and with total mass zero, such that*

$$Q_t(dz' | z) = Q_{t_0}(dz' | z) + \sum_{i=1}^d [x_i(t) - x_i(t_0)] D_i(dz') \quad (14)$$

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

(ii) A finite-dimensional continuous sufficient statistic exists if and only if $d_Q < \infty$. In that case the minimum affine dimension of a continuous sufficient statistic is d_Q . If $d_Q \geq 1$, any two minimal sufficient statistics $x, y : \mathcal{T} \rightarrow \mathbb{R}^{d_Q}$ satisfy

$$y(t) = Mx(t) + c \quad \text{for every } t \in \mathcal{T} \quad (15)$$

for some nonsingular matrix M and vector c .

(iii) Suppose that $d_Q < \infty$. Let $x : \mathcal{T} \rightarrow \mathbb{R}^{d_Q}$ be minimal and let $\bar{x} \in \text{co } x(\mathcal{T})$. For any lottery π on $\mathcal{T}$ satisfying $\int_{\mathcal{T}} x(t)\pi(dt) = \bar{x}$, the kernel

$$Q_{\bar{x}}(\cdot | z) := \int_{\mathcal{T}} Q_t(\cdot | z)\pi(dt)$$

does not depend on the choice of π . If $\mu Q_{\bar{x}} = \mu$, then

$$\mu P_{\sigma} = \mu \iff \bar{x}_{\sigma} = \bar{x}. \quad (16)$$

Theorem 2 shows that unconditional statistics of the action suffice if and only if the action shifts successor types through a component that is invariant to incumbent types, and the family of such shifts has finite rank. The proof of necessity constructs perturbations that reallocate actions across types while holding the aggregate statistic and monotonicity fixed: if all such assignments must induce the identical successor distribution, kernel differences cannot depend on the current type. The affine dimension d_Q counts the precise number of independent demographic targets, explaining why a quadratic transition law yields a two-moment screening problem. Appendix C proves the theorem, and Online Appendix A develops the polynomial case alongside a smooth transition family of infinite rank.

Example 3 (Type-dependent transition effects). Let $\mathcal{Z} = [0, 1]$, let μ be uniform, and let $\mathcal{T} = \{0, 1\}$. Define

$$Q_0(\cdot | z) = \delta_0, \quad Q_1(\cdot | z) = (1 - z)\delta_0 + z\delta_1,$$

where δ_y is the Dirac measure at y . The action-induced difference is

$$Q_1(\cdot | z) - Q_0(\cdot | z) = z(\delta_1 - \delta_0),$$

which varies directly with current type z . Consider

$$\sigma^A(\{1\} | z) = z, \quad \sigma^B(\{1\} | z) = \mathbf{1}\{z \geq 1/2\}.$$

Both conditional mean actions are nondecreasing, and both assignment rules induce the identical Bernoulli(1/2) marginal distribution of T . They nevertheless produce distinct successor distributions:

the probability of successor state 1 is

$$\int_0^1 z^2 dz = \frac{1}{3} \quad \text{under } \sigma^A, \quad \int_{1/2}^1 z dz = \frac{3}{8} \quad \text{under } \sigma^B.$$

The action marginal therefore fails to determine the successor distribution, even among monotone assignments. In this example, the joint moment $\mathbb{E}[ZT]$ constitutes a minimal sufficient statistic.

Two structural considerations govern the interface between transition sufficiency and incentive compatibility. First, the monotonicity in Theorem 2(i) applies to the conditional mean of $x_1(T)$, whereas incentive compatibility orders the effective allocation $q(\alpha + \beta r)$. Online Appendix B verifies the representation theorem directly on the incentive-compatible domain when the range condition holds with strict slack and the mean action is affine in the statistic; incorporating the mean action into the statistic can increase the minimal dimension by one. Second, the dimension required to identify every successor distribution can exceed the number of moment equalities needed to pin down a specific boundary target, as characterized in Online Appendix C.

When action-induced kernel differences depend on incumbent types, the action marginal alone no longer determines the successor distribution. For a representation $Q_t(\cdot | z) = Q^0(\cdot | z) + \sum_j \varphi_j(z, t) A_j(\cdot)$, the relevant statistics are the joint moments $\mathbb{E}[\varphi_j(Z, T)]$. Stationarity then constrains how actions are matched to types. Online Appendix D develops the corresponding duality theory and characterizes optimal lottery support when the dual optimum is attained.

5 Optimal Output under Two Moment Restrictions

Under a quadratic transition law, stationarity fixes the first two action moments, $\mathbb{E}[T] = m$ and $\mathbb{E}[T^2] = k$. We characterize the feasible moment pairs and the assignments that maximize output. For transitions of higher degree, the two-moment problem is a relaxation, so its value bounds what the institution can achieve. The bound is not generally attained: Section 7 shows that where the two-moment optimum is unique and deterministic, exactly one additional-moment target preserves the two-moment value, and every other feasible target costs output. Throughout this and the next two sections the objective is current output, as specified in Assumption 2.

Incentives and output both run through the conditional-mean schedule. The second moment does not run through that schedule alone: it splits into the dispersion of the schedule and the randomization carried out within types, and only the first component affects output.

$$\mathbb{E}[T^2] = \mathbb{E}[r_\sigma(Z)^2] + \mathbb{E}[\text{Var}(T | Z)], \quad Y(\sigma) = \alpha \mathbb{E}[w] + \beta \mathbb{E}[w r_\sigma].$$

At a fixed mean, a change in the second moment is a change in the variance of the action, and randomization within types can deliver that change while leaving the conditional-mean schedule and output alone. Define the two-moment implementation set

$$\Sigma_2(m, k) := \{\sigma : g_\sigma \in L_1^\infty(\mu), \mathbb{E}[T] = m, \mathbb{E}[T^2] = k\}.$$

In the preparation-stage reading, m and k are the success rates of one stage and of an ordered pair of stages.

Comparative statics. We compare second-moment restrictions while holding q , μ , and m fixed. For a fixed quadratic transition law with independent coefficients, the target $\mu = R_0 + mA_1 + kA_2$ determines both moments uniquely. To vary the required k at the same population, consider the technology family

$$R_0(k) := \mu - mA_1 - kA_2, \quad \widehat{Q}_t^{[k]} := R_0(k) + tA_1 + t^2A_2.$$

Keep the persistence kernel fixed and restrict k to values for which $\widehat{Q}_t^{[k]}$ is a probability measure for every $t \in [0, 1]$. Stationarity then requires the corresponding second moment at mean m . This comparison measures the output cost of the moment restriction at a fixed type distribution. Varying the stationary population under a fixed technology would also change the type weights in the objective.

Revenue. If μ has density f , distribution function F , and virtual value $\phi(z) = z - [1 - F(z)]/f(z)$, the envelope formula gives revenue $\alpha\mathbb{E}[w^R] + \beta\mathbb{E}[w^RT] - V_0$, where $w^R = \phi q$. Every result below applies with w^R in place of w ; in particular the upper-tail rule of Theorem 4(ii) then requires ϕq to be strictly increasing.

5.1 Second-moment bounds for a closed action set

For a conditional mean $r \in [a, b]$, the action set determines the minimum and maximum conditional second moments. Let

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

and define

$$\begin{aligned} \underline{K}_{\mathcal{T}}(r) &:= \min_{\pi \in \mathcal{P}_{\mathcal{T}}(r)} \int t^2 \pi(dt), \\ \overline{K}_{\mathcal{T}}(r) &:= \max_{\pi \in \mathcal{P}_{\mathcal{T}}(r)} \int t^2 \pi(dt). \end{aligned}$$

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

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

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

$$\underline{K}_{\mathcal{T}}(r) = \begin{cases} r^2, & r \in \mathcal{T}, \\ r^2 + [r - t_-(r)][t_+(r) - r], & r \notin \mathcal{T}, \end{cases} \quad \overline{K}_{\mathcal{T}}(r) = (a + b)r - ab. \quad (17)$$

The minimizing lottery is degenerate when $r \in \mathcal{T}$ and otherwise mixes the two actions bracketing r . The maximizing lottery mixes a and b . The lower envelope is continuous and convex. It is strictly convex on $[a, b]$ if and only if $\mathcal{T} = [a, b]$.

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

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

Whenever these inequalities hold, there is an assignment rule that uses at most two actions at each type.

Atomlessness permits two-action support at each type. Assigning the lower- and upper-envelope lotteries to different sets of types matches the aggregate second moment while preserving every conditional mean. A second-moment restriction at each type can require more support: on $\{0, 1/2, 1\}$, mean $1/2$ and second moment $3/8$ require all three actions.

5.2 Feasibility, value, and optimal assignments

Lemma 1 reduces the problem to the choice of a conditional-mean schedule. Define

$$\begin{aligned} \mathcal{D}_{q,\mathcal{T}} &:= \{r \in L^2(\mu) : a \leq r \leq b, q(\alpha + \beta r) \text{ has a nondecreasing representative}\}, \\ \mathcal{C}_{q,\mathcal{T}}(m) &:= \{r \in \mathcal{D}_{q,\mathcal{T}} : \mathbb{E}[r] = m\}. \end{aligned}$$

For a schedule r , let

$$\mathcal{K}_{\mathcal{T}}(r) := \mathbb{E}[\underline{K}_{\mathcal{T}}(r(Z))]$$

be the smallest aggregate second moment needed to implement its conditional means. Assume $\mathcal{C}_{q,\mathcal{T}}(m) \neq \emptyset$, and define

$$\begin{aligned} L_{q,\mathcal{T}}(m) &:= \min_{r \in \mathcal{C}_{q,\mathcal{T}}(m)} \mathcal{K}_{\mathcal{T}}(r), \\ B_{q,\mathcal{T}}(m, k) &:= \max_{\substack{r \in \mathcal{C}_{q,\mathcal{T}}(m) \\ \mathcal{K}_{\mathcal{T}}(r) \leq k}} \mathbb{E}[wr]. \end{aligned} \tag{18}$$

The threshold $L_{q,\mathcal{T}}(m)$ is the minimum second moment consistent with mean m and incentive compatibility. The value $B_{q,\mathcal{T}}(m, k)$ measures the largest feasible $\mathbb{E}[wr]$; its contribution to output is $\beta B_{q,\mathcal{T}}(m, k)$. At a feasible target k , conditional randomization can raise the second moment of any schedule with $\mathcal{K}_{\mathcal{T}}(r) \leq k$ to exactly k .

Let

$$\begin{aligned} \mathcal{M}_{q,\mathcal{T}}(m) &:= \arg \max_{r \in \mathcal{C}_{q,\mathcal{T}}(m)} \mathbb{E}[wr], \\ k_{q,\mathcal{T}}^{\text{mc}}(m) &:= \min_{r \in \mathcal{M}_{q,\mathcal{T}}(m)} \mathcal{K}_{\mathcal{T}}(r). \end{aligned}$$

The threshold $k_{q,\mathcal{T}}^{\text{mc}}(m)$ is the least second moment required to attain the mean-constrained output maximum, with incentive compatibility and the quantity schedule retained. Thus

$$L_{q,\mathcal{T}}(m) \leq k_{q,\mathcal{T}}^{\text{mc}}(m) \leq \bar{K}_{\mathcal{T}}(m).$$

Theorem 3 (Feasibility and optimal output). *Suppose $\mathcal{C}_{q,\mathcal{T}}(m) \neq \emptyset$.*

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

$$L_{q,\mathcal{T}}(m) \leq k \leq \bar{K}_{\mathcal{T}}(m). \quad (19)$$

For every feasible k , the maximum in (18) is attained, and current output is $\alpha\mathbb{E}[w] + \beta B_{q,\mathcal{T}}(m, k)$. An optimal assignment rule can be chosen to use at most two actions at each type.

(ii) *On the interval in (19), the function $k \mapsto B_{q,\mathcal{T}}(m, k)$ is continuous, concave, and nondecreasing. If $L_{q,\mathcal{T}}(m) < k_{q,\mathcal{T}}^{\text{mc}}(m)$, it is strictly increasing on*

$$[L_{q,\mathcal{T}}(m), k_{q,\mathcal{T}}^{\text{mc}}(m)]$$

and constant on

$$[k_{q,\mathcal{T}}^{\text{mc}}(m), \bar{K}_{\mathcal{T}}(m)].$$

If the two thresholds coincide, the value is constant on the entire feasible interval.

Raising k relaxes the second-moment restriction on conditional-mean schedules until an output-maximizing schedule in $\mathcal{M}_{q,\mathcal{T}}(m)$ first becomes feasible, which occurs at $k_{q,\mathcal{T}}^{\text{mc}}(m)$. Beyond this threshold, the institution maintains the output-maximizing conditional-mean schedule and satisfies the higher second-moment target entirely through within-type randomization; further dispersion yields no additional output.

Proposition 3 (Optimal assignment structure). *Suppose $\mathcal{C}_{q,\mathcal{T}}(m) \neq \emptyset$ and k satisfies (19).*

(i) *If $k \leq k_{q,\mathcal{T}}^{\text{mc}}(m)$, the second-moment constraint binds at every optimal schedule:*

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

Every optimal assignment rule therefore uses the conditional lower-envelope lottery at almost every type.

(ii) *Suppose $\mathcal{T} = [a, b]$. Then $\underline{K}_{\mathcal{T}}(r) = r^2$. The minimizer defining $L_{q,[a,b]}(m)$ is unique, and the mean-constrained output maximizers have a unique member with minimum L^2 norm. For $k \leq k_{q,[a,b]}^{\text{mc}}(m)$, the optimal schedule r_k is unique and every optimal assignment is deterministic:*

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

If $L_{q,[a,b]}(m) < k < k_{q,[a,b]}^{\text{mc}}(m)$, there is $\lambda > 0$ such that

$$r_k = \Pi_{\mathcal{C}_{q,[a,b]}(m)} \left(\frac{w}{2\lambda} \right), \quad (20)$$

where $\Pi_{\mathcal{C}}$ denotes the $L^2(\mu)$ projection onto a closed convex set $\mathcal{C}$. If, in addition, $m \in \text{ri}\{\mathbb{E}[r] : r \in \mathcal{D}_{q,[a,b]}\}$, there is $\eta \in \mathbb{R}$ such that

$$r_k = \Pi_{\mathcal{D}_{q,[a,b]}} \left(\frac{w - \eta}{2\lambda} \right), \quad \mathbb{E}[r_k] = m. \quad (21)$$

The multiplier λ is a supergradient of $k \mapsto B_{q,[a,b]}(m, k)$. At points of differentiability, it is the marginal value of the second moment. If the joint two-moment dual is attained, λ can be chosen as the multiplier on the second-moment restriction.

(iii) Suppose $\mathcal{T} = \{t_1 < \dots < t_J\}$. Then $\underline{K}_{\mathcal{T}}$ is the piecewise-affine interpolation of t_j^2 . For $k \leq k_{q,\mathcal{T}}^{\text{mc}}(m)$, every optimal assignment uses the two adjacent actions bracketing its conditional mean. The optimal schedule need not be unique. The set of mean-constrained output maximizers that minimize the required second moment may also contain more than one schedule.

For a quadratic transition law with independent coefficients, the target $\mu = R_0 + mA_1 + kA_2$ is implementable exactly when $\mathcal{C}_{q,\mathcal{T}}(m) \neq \emptyset$ and (19) holds. This combines the moment equalities implied by stationarity with Theorem 3(i).

The difference between interval and finite action sets comes from the lower second-moment envelope. On an interval it is strictly convex, which gives a unique deterministic optimum up to $k_{q,[a,b]}^{\text{mc}}(m)$. With finite actions it is affine between adjacent actions, so distinct conditional-mean schedules can attain the same value. Assumption 2 treats randomization as costless; direct lottery costs would change the output comparison while leaving moment feasibility unchanged.

5.3 Explicit thresholds for nonincreasing quantities

For the rest of this section, let q be nonincreasing, $\mathcal{T} = [0, 1]$, and $\alpha > 0$. These conditions give an explicit minimum second moment. If $w = zq$ is strictly increasing, we can also calculate the second moment required for the mean-constrained output maximum. Write

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

At a fixed mean m , differences in second moments are also differences in variances. Since an action in $[0, 1]$ has second moment between m^2 and m , the total variance range has length $m(1 - m)$. The two thresholds divide this range as follows:

$$m - m^2 = \underbrace{[L_q(m) - m^2]}_{\mathcal{C}_q(m)} + \underbrace{[k_q^{\text{mc}}(m) - L_q(m)]}_{\mathcal{P}_q(m)} + \underbrace{[m - k_q^{\text{mc}}(m)]}_{\mathcal{N}_q(m)}. \quad (22)$$

The first term $C_q(m)$ represents the compliance floor: the minimal action variance required to sustain incentive compatibility. The second term $P_q(m)$ constitutes the productive range, where additional variance enables sorting higher actions to higher types and strictly increases output. The third term $N_q(m)$ forms the neutral range, where the output-maximizing sorting is already achieved and additional variance is absorbed costlessly by within-type lotteries.

Let

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

For an effective allocation $g = q(\alpha + \beta r)$, the mean restriction $\mathbb{E}[r] = m$ is equivalent to

$$\mathbb{E}[hg] = S.$$

Because q is nonincreasing, h is nondecreasing. On the full action interval, the feasible nondecreasing effective allocations take values in

$$g_L := \alpha q(\underline{z}), \quad g_U := (\alpha + \beta)q(\bar{z}).$$

Define

$$\bar{g}_m := \frac{S}{H}, \quad \nu(dz) := \frac{h(z)}{H} \mu(dz).$$

Then $\mathbb{E}_\nu[g] = \bar{g}_m$. Here h is the reciprocal quantity schedule, H is its mean, and S is the total effective allocation the mean restriction requires. Interpreting $\mathbb{E}[hg] = S$ as a resource budget illuminates this change of measure: increasing the effective allocation g for a type with a smaller committed quantity requires a larger adjustment in the action, thereby consuming more of the aggregate mean action m . The probability measure ν reweights the population measure μ by h/H , giving more weight to types with smaller prescribed quantities. The mean constraint, and hence the cutoff below, is evaluated under ν . A constant effective allocation satisfies the mean restriction and is feasible exactly when $\bar{g}_m \in [g_L, g_U]$. For the remainder of this section, maintain the range condition in Proposition 1.

Theorem 4 (Second-moment thresholds for a nonincreasing quantity schedule). *Let $\mathcal{T} = [0, 1]$, $\alpha > 0$, and suppose that q is nonincreasing.*

- (i) *Suppose $g_L \leq g_U$ and $\bar{g}_m \in [g_L, g_U]$. The unique incentive-compatible schedule with mean m that minimizes the second moment is*

$$r_m^0(z) = \frac{\bar{g}_m h(z) - \alpha}{\beta} = \frac{Sh(z)/H - \alpha}{\beta},$$

and the minimum is

$$L_q(m) = m^2 + \frac{S^2}{\beta^2} \left(\frac{\mathbb{E}[h^2]}{H^2} - 1 \right). \quad (23)$$

- (ii) *Suppose $g_L < g_U$, $\bar{g}_m \in (g_L, g_U)$, and $w(z) = zq(z)$ is strictly increasing. There is a unique*

cutoff $z^* \in (z, \bar{z})$ satisfying

$$\nu((z^*, \bar{z}]) = \frac{\bar{g}_m - g_L}{g_U - g_L}. \quad (24)$$

Up to null sets, the effective allocation that maximizes output subject to the mean and incentive constraints is

$$g^*(z) = g_L + (g_U - g_L)\mathbf{1}\{z > z^*\}.$$

The corresponding conditional-mean schedule is

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

and the least second moment supporting this output maximum is

$$k_q^{\text{mc}}(m) = \frac{1}{\beta^2} \left[g_L^2 \mathbb{E}[h^2] + (g_U^2 - g_L^2) \mathbb{E}[h^2 \mathbf{1}\{Z > z^*\}] - 2\alpha S + \alpha^2 \right]. \quad (25)$$

(iii) Under the conditions in part (ii), the length of the interval over which output increases is

$$P_q(m) = \frac{H}{\beta^2} \{ \mathbb{E}_\nu[h(g^* - \bar{g}_m)^2] + 2\bar{g}_m \text{Cov}_\nu(h, g^*) \} > 0. \quad (26)$$

If q is strictly decreasing, then all three terms in (22) are strictly positive. The two boundary cases are

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

and

$$g_L = g_U \implies P_q(m) = 0, \quad L_q(m) = k_q^{\text{mc}}(m).$$

The minimum-variance schedule holds the effective allocation, and hence the payment, constant across reports, securing truthful reporting by making every type indifferent. At an interior feasible mean, this variance is also the infimum under strict incentive compatibility, as Proposition 4 shows.

Part (ii) characterizes an upper-tail assignment rule: output is maximized by allocating the maximal feasible effective allocation to types with the highest output weights. Because w is strictly increasing, these types form an upper interval, and the resulting two-valued allocation is automatically monotone. The cutoff z^* is determined under ν rather than μ due to the budgetary weighting $\mathbb{E}[hg] = S$.

Both terms in (26) are nonnegative. The covariance term is nonnegative because h and g^* are both nondecreasing, and the first term is positive in the interior case. The length $P_q(m)$ therefore quantifies the variance capacity that the designer can productively channel into sorting.

The optimal effective allocation in part (ii) is constant on nondegenerate type intervals, so every assignment attaining the maximum leaves some types indifferent across reports. Along the entire

output plateau, strictly incentive-compatible mechanisms approach the weak maximum arbitrarily closely without attaining it; Proposition 4 formalizes this limiting behavior.

5.4 Variance required by incentive compatibility

Equation (23) yields a simple expression for the first term in (22).

Corollary 2 (Minimum variance under incentive compatibility). *Under the hypotheses of Theorem 4(i),*

$$C_q(m) = L_q(m) - m^2 = \left(m + \frac{\alpha}{\beta}\right)^2 \text{CV}^2\left(\frac{1}{q}\right),$$

where

$$\text{CV}^2(1/q) = \frac{\text{Var}_\mu(1/q)}{(\mathbb{E}_\mu[1/q])^2}.$$

The minimum variance is zero if and only if q is constant μ -almost everywhere.

Each component in the compliance floor reflects a distinct economic force. The squared coefficient of variation of $1/q$ measures the relative dispersion of prescribed quantities: a schedule with greater cross-grade variation forces the effective allocation to be reconstructed from a wider spread of actions. The scale factor $(m + \alpha/\beta)^2$ reflects the effectiveness of the action multiplier: it decreases in β , meaning that when the action exerts a stronger effect on output, less action dispersion is required to offset a given decline in q . This factor also increases quadratically in m , so the compliance cost is not a fixed overhead. Incentive compatibility ties the proportional spread of $\alpha + \beta r$ to the spread of $1/q$, so a higher aggregate mean translates that fixed proportion into a larger absolute variance. Consequently, an institution requiring a higher average training level faces quadratically escalating variance costs under the same prescribed schedule.

Equally revealing is what the floor excludes. It depends solely on the marginal distribution of $1/q$ under μ : two institutions whose quantity schedules share the same distribution face the identical variance floor, regardless of how quantities are mapped to types. Monotonicity of q determines whether the bound applies; the distribution of q determines its magnitude. For an action set with gaps, $\underline{K}_{\mathcal{T}}(r)$ exceeds r^2 precisely when $r \notin \mathcal{T}$: such conditional means require randomization between the adjacent available actions.

Parts (ii) and (iii) of Theorem 4 require $w(z) = zq(z)$ to be strictly increasing. For differentiable q and nonnegative types, $q(z) + zq'(z) > 0$ suffices, which for $z > 0$ reads $zq'(z)/q(z) > -1$: quantities decline less than proportionately with type, so higher types remain the more valuable place to put the action. Without that ordering, assigning the upper bound to the largest output weights can violate monotonicity in type, and the cutoff rule fails even though the minimum-variance formula in part (i) still applies.

Illustration. Take $Z \sim U[0, 1]$, $\alpha = \beta = 1$, $m = 1/2$, and $q(z) = (1 + z/2)^{-1}$. Theorem 4 gives $L_q(m) = 0.28$ and $k_q^{\text{mc}}(m) \simeq 0.3808$. Second moments below 0.28 are not implementable at this mean; output rises as k moves from 0.28 to 0.3808 and is flat from there to $1/2$, where within-type

lotteries supply the additional variance. Read along the second-moment axis, the two thresholds are the interior boundaries of the compliance, productive, and neutral ranges in (22). The two boundary cases bracket this one. With $q \equiv 1$ the thresholds are $1/4$ and $1/2$, so the entire variance range is productive; with $q(z) = (1+z)^{-1}$ both equal $1/3$, incentive compatibility pins down the effective allocation, and no variance is left to allocate. All three are fixed-population comparisons of the kind defined at the start of this section.

6 Strict Incentive Compatibility

Weak incentive compatibility makes truthful reporting optimal but allows other best responses. Under the minimum-variance schedule, the effective allocation is constant, leaving every reporting strategy as a best response. Because stationarity is evaluated under truthful reporting, alternative best responses can change the realized action marginal and hence the successor distribution through (3). The output-maximizing allocation in Theorem 4(ii) pools types on intervals and exhibits the same indifference property. Requiring strict incentive compatibility ensures that truthful reporting is the unique best response, thereby robustly securing the stationary reproduction of the population. This robustness is available under the slackness conditions below at an arbitrarily small output cost: the strict supremum need not be attained, but the shortfall can be made vanishingly small while matching both prescribed moments exactly. A mechanism is *strictly incentive compatible* if

$$zg_\sigma(z) - p(z) > zg_\sigma(\hat{z}) - p(\hat{z}) \quad \text{for all } z \neq \hat{z}.$$

With the envelope payments, this is equivalent to g_σ being strictly increasing. Write $\Sigma_2^{\text{str}}(m, k)$ for the assignments in $\Sigma_2(m, k)$ that admit pointwise strict implementation. Throughout this section, q is nonincreasing and $\mathcal{T}$ is a general closed action set, as in Theorem 1. Thus

$$g_L = (\alpha + \beta a)q(\underline{z}), \quad g_U = (\alpha + \beta b)q(\bar{z}), \quad \bar{g}_m = \frac{\alpha + \beta m}{\mathbb{E}[1/q(Z)]}.$$

Proposition 4 (Strict implementation and the two-moment value). *Maintain the screening primitives of Theorem 1, with q nonincreasing.*

(i) *Suppose*

$$g_L < \bar{g}_m < g_U, \quad L_{q, \mathcal{T}}(m) < k \leq \bar{K}_{\mathcal{T}}(m). \quad (27)$$

Then $\Sigma_2^{\text{str}}(m, k) \neq \emptyset$. A member can be chosen with a continuous effective allocation $\hat{g}$ satisfying $\hat{g}(z') - \hat{g}(z) \geq c(z' - z)$ for all $z' > z$, for some $c > 0$.

(ii) *Under (27), fix such an assignment $\hat{\sigma}$. Let σ be any pointwise weakly implementable version of an assignment in $\Sigma_2(m, k)$. For every $0 < \theta \leq 1$, the conditional lottery*

$$\sigma_\theta(dt | z) = (1 - \theta)\sigma(dt | z) + \theta\hat{\sigma}(dt | z) \quad (28)$$

belongs to $\Sigma_2^{\text{str}}(m, k)$. It preserves both moments exactly and converges to σ uniformly in total variation as $\theta \downarrow 0$. With the same lower-type utility $V_0 \geq 0$, the envelope payments converge uniformly as well. In particular,

$$\sup_{\sigma \in \Sigma_2^{\text{str}}(m, k)} Y(\sigma) = \max_{\sigma \in \Sigma_2(m, k)} Y(\sigma), \quad Y(\sigma) := \mathbb{E}[Zq(Z)(\alpha + \beta T)]. \quad (29)$$

The supremum on the left need not be attained.

(iii) If $\mathcal{T} = [a, b]$, condition (27) is also necessary for $\Sigma_2^{\text{str}}(m, k) \neq \emptyset$. At a weakly feasible mean,

$$L_{q, [a, b]}(m) = \mathbb{E}[(r_m^0(Z))^2] = m^2 + \frac{(\alpha + \beta m)^2}{\beta^2} \text{CV}^2(1/q). \quad (30)$$

Thus, on an interval, strict implementation excludes the minimum second moment but does not exclude the maximum second moment $\bar{K}_{\mathcal{T}}(m) = (a + b)m - ab$.

The construction perturbs the effective allocation by adding a strictly positive slope. Centering this perturbation preserves its budget-weighted mean under ν , while the strict inequalities in (27) ensure interior feasibility for the resulting conditional means. Lower-envelope lotteries implement these conditional means at a second moment strictly below k , and mixing with endpoint lotteries restores the aggregate second moment to exactly k . Finally, mixing this strictly incentive-compatible mechanism with any weak optimum preserves both moments exactly while converging to the weak output maximum as the perturbation weight vanishes. Under quadratic transitions, every such mixture exactly preserves the target moments, ensuring that the approximation never destabilizes the stationary population.

The two excluded cases are exactly the two configurations that force a constant effective allocation: a boundary of the feasible mean range, and, with interval actions, the minimum second moment. With gaps in the action set, the strict second-moment inequality is sufficient but not necessary. With constant q and two actions, for instance, the mean already fixes the second moment, and strictly increasing conditional probabilities still implement an interior mean strictly.

Corollary 3 extends the construction to further moments, and Example 4 shows why a boundary target can rule out strict implementation. Both sides of the trade-off are linear in θ : the lower bound on an agent's gain from truthful reporting is proportional to the weight placed on the strict assignment, and so is the output forgone relative to the weak maximum. The institution therefore buys uniqueness of the truthful best response at a constant rate.

7 The Cost of an Additional Moment Restriction

An additional moment restriction can alter the set of feasible assignments and strictly reduce output below the two-moment maximum. This section characterizes the targets that preserve the two-moment value, decomposes the output loss when the two-moment optimum is deterministic and

its second-moment constraint has a positive multiplier, and establishes when strict incentives can be sustained.

Under the three-stage preparation technology of Example 1, stationarity prescribes more than the success rate m of a single stage and the rate k of an ordered pair. It also prescribes $m_3 = b_3$, the share of successors who complete every stage. An institution that reproduces its population must deliver a fully trained share as well as an average level of preparation, and this third requirement can bind where the first two do not. In general the additional restriction takes the form

$$\mathbb{E}[\varphi(T)] = x,$$

where $\varphi : \mathcal{T} \rightarrow \mathbb{R}$ is continuous and is not in $\text{span}\{1, t, t^2\}$ on $\mathcal{T}$. The restriction arises from a transition law with one additional coordinate,

$$\hat{Q}_t = R_0 + tA_1 + t^2A_2 + \varphi(t)A_\varphi,$$

where A_1, A_2, A_φ are linearly independent signed measures of total mass zero.

A fixed technology and target determine x . To compare restrictions at the same q, μ, m, k , hold the coefficient measures fixed and set $R_0(x) = \mu - mA_1 - kA_2 - xA_\varphi$. For each x such that $R_0(x) + tA_1 + t^2A_2 + \varphi(t)A_\varphi$ is a probability measure at every $t \in \mathcal{T}$, stationarity requires $\mathbb{E}[\varphi(T)] = x$ among assignments that match m and k . The allocation value below is defined at every IC-feasible moment value; the reading of x as a stationarity requirement applies on the subrange that yields probability measures.

Fix a feasible pair (m, k) , write $\Sigma_2 := \Sigma_2(m, k)$, and define

$$\begin{aligned} Y(\sigma) &:= \mathbb{E}[Zq(Z)(\alpha + \beta T)], & M_\varphi(\sigma) &:= \mathbb{E}[\varphi(T)], \\ Y_2 &:= \max_{\sigma \in \Sigma_2} Y(\sigma), & \Sigma_2^* &:= \arg \max_{\sigma \in \Sigma_2} Y(\sigma), \\ I_\varphi &:= M_\varphi(\Sigma_2), & I_\varphi^* &:= M_\varphi(\Sigma_2^*). \end{aligned}$$

For $x \in I_\varphi$, let

$$\Sigma_\varphi(x) := \{\sigma \in \Sigma_2 : M_\varphi(\sigma) = x\}, \quad Y_\varphi(x) := \max_{\sigma \in \Sigma_\varphi(x)} Y(\sigma),$$

and call $Y_2 - Y_\varphi(x)$ the output loss at x . Thus I_φ contains all attainable additional-moment values and I_φ^* those attained by a two-moment optimizer.

7.1 Attainable moment values

Proposition 5 (Moment values attained at the two-moment optimum). *The sets I_φ and I_φ^* are nonempty compact intervals, and Y_φ is concave and upper semicontinuous. For every $x \in I_\varphi$,*

$$Y_\varphi(x) = Y_2 \iff x \in I_\varphi^*. \tag{31}$$

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

$$\begin{aligned} L_{\varphi, \mathcal{T}, r^*}(k) &:= \min\{\mathbb{E}[\varphi(T)] : \mathbb{E}[T \mid Z] = r^*(Z), \mathbb{E}[T^2] = k\}, \\ U_{\varphi, \mathcal{T}, r^*}(k) &:= \max\{\mathbb{E}[\varphi(T)] : \mathbb{E}[T \mid Z] = r^*(Z), \mathbb{E}[T^2] = k\}. \end{aligned}$$

Then

$$I_\varphi^* = [L_{\varphi, \mathcal{T}, r^*}(k), U_{\varphi, \mathcal{T}, r^*}(k)]. \quad (32)$$

If the two-moment optimal assignment rule σ^* is unique, this interval is the singleton $\{\mathbb{E}_{\sigma^*}[\varphi(T)]\}$. If this unique assignment is deterministic, write its moment value as

$$x_\varphi^\dagger := \mathbb{E}[\varphi(r^*(Z))].$$

Within-type lotteries can change M_φ while preserving conditional means and output, so I_φ^* can contain several values. If $\mathcal{T} = [0, 1]$ and $L_q(m) < k < k_q^{\text{mc}}(m)$, the two-moment optimum is unique and deterministic. In that case $I_\varphi^* = \{x_\varphi^\dagger\}$ for every continuous φ .

Corollary 3 (Strict implementation of additional interior moments). *Maintain the hypotheses and inequalities of Proposition 4(i). For every $x \in \text{ri } I_\varphi$, a strictly incentive-compatible assignment exists in $\Sigma_\varphi(x)$, and*

$$\sup_{\substack{\sigma \in \Sigma_\varphi(x) \\ g_\sigma \text{ strictly increasing}}} Y(\sigma) = Y_\varphi(x).$$

More generally, the same conclusion holds for a finite vector of continuous action statistics when its target lies in the relative interior of their joint image over $\Sigma_2(m, k)$.

Relative interior is taken in the moment set attainable under IC and the first two moment restrictions. At such a target, one can construct a strict implementation and mix it with a weak optimum without changing any moments. A boundary target can instead force pooling, as the next example shows.

Example 4 (A boundary of strict implementation). Let $Z \sim U[0, 1]$, $q(z) = 1/(1 + z/2)$, $\alpha = \beta = 1$, and $\mathcal{T} = \{0, 3/10, 7/10, 1\}$. Fix

$$m = \frac{1}{2}, \quad k = \frac{33}{100}.$$

Here $\bar{g}_m = 6/5$ lies strictly between 1 and $4/3$, and $L_{q, \mathcal{T}}(m) = 181/600 < 33/100 < 1/2$, so the two-moment constraints admit strict truthfulness. For $\varphi(t) = t^3$, the weakly attainable third moments are exactly

$$I_3 = \left[\frac{277}{1200}, \frac{311}{1200} \right]. \quad (33)$$

The third moments attainable under strict incentive compatibility form the *open* interval with the same endpoints. At the upper endpoint the unique action marginal is

$$\pi^\dagger = \frac{1}{36}\delta_0 + \frac{289}{504}\delta_{3/10} + \frac{17}{72}\delta_{7/10} + \frac{41}{252}\delta_1. \quad (34)$$

All four probabilities are positive, so the three-moment vector is interior to the unrestricted moment polytope. The convex-order inequality of Theorem 1(iii) nonetheless binds at $s = 3/10$ for $\pi^\dagger$, forcing the effective allocation to be constant almost everywhere. Strict incentive compatibility is therefore impossible. The lower endpoint has the same obstruction at $s = 7/10$. Appendix F verifies the example and constructs a cubic transition law under which either endpoint is a stationary target admitting weak but not strict implementation.

An interval-action version appears in the proof. Prescribing the first six moments of $\pi^\dagger$ on $[0, 1]$ forces the same finite support and excludes strict implementation. Its first two moments alone satisfy Proposition 4(i).

7.2 Shadow prices and exact output loss

Define the penalized two-moment value

$$\Psi_\varphi(\theta) := \max_{\sigma \in \Sigma_2} \{Y(\sigma) - \theta M_\varphi(\sigma)\}.$$

For a closed convex set $\mathcal{C} \subset L^2(\mu)$, let

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

denote its normal cone at x .

Theorem 5 (Shadow price and output loss). *(i) For every $x \in I_\varphi$,*

$$Y_\varphi(x) = \inf_{\theta \in \mathbb{R}} \{\Psi_\varphi(\theta) + \theta x\}. \quad (35)$$

For every $x \in \text{ri } I_\varphi$, the infimum is attained. Whenever a minimizer exists, every minimizing θ is positive when $x < \inf I_\varphi^$ and negative when $x > \sup I_\varphi^*$. At every $x \in I_\varphi^*$, $\theta = 0$ is a minimizer. At an endpoint of this interval, nonzero minimizing multipliers may also exist.*

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

$$w - \eta \mathbf{1} - 2\lambda r^* = n. \quad (36)$$

The relative-interior condition in Proposition 3(ii) is sufficient for such a KKT triple to exist. For any $x \in I_\varphi$ and $\sigma \in \Sigma_\varphi(x)$, let

$$d := \mathbb{E}[T \mid Z] - r^*, \quad \mathbf{V}_\sigma := \mathbb{E}[\text{Var}(T \mid Z)].$$

Then

$$\frac{Y_2 - Y(\sigma)}{\beta} = \lambda(\|d\|_2^2 + \mathbf{V}_\sigma) - \langle n, d \rangle. \quad (37)$$

If $n = \beta q \xi$ for an IC multiplier ξ , then $-\langle n, d \rangle = -\langle \xi, g_\sigma - g_{r^*} \rangle$. Each term on the right side of (37) is nonnegative. Equality holds only when $T = r^*(Z)$ almost surely. Hence $x_\varphi^\dagger$ is the unique moment value with no output loss.

An optimal θ represents a supergradient of Y_φ ; at points of differentiability, it measures the marginal output effect of increasing the moment target x . Part (ii) gives an exact accounting of the output loss relative to the deterministic two-moment optimum, decomposing it into three distinct economic channels. In output units divided by β , $\lambda \|d\|_2^2$ represents the cost of distorting conditional means away from r^* ; λV_σ measures the loss from variance absorbed by within-type randomization; and $-\langle n, d \rangle$ reflects the shadow cost of binding allocation constraints.

The first and third terms depend exclusively on the departure of conditional means d , quantifying the output forgone by sorting types suboptimally relative to r^* . The second term depends only on within-type randomization V_σ . Within-type randomization affects no conditional mean and generates no direct output; however, it absorbs variance that could otherwise be deployed to separate conditional means across types.

The multiplier λ prices the second-moment budget in the first two terms. It is the shadow price identified in Proposition 3(ii). Concavity of the two-moment value implies that this price is nonincreasing in k . Within the productive range, a given amount of within-type variance or squared displacement of conditional means therefore carries a weakly higher charge when the second moment is closer to $L_q(m)$.

The normal cone $N_{\mathcal{D}_q}(r^*)$ encapsulates both the action bounds and the monotonicity of the effective allocation $q(\alpha + \beta r)$. When $n = \beta q \xi$, the final term isolates the shadow price of the incentive compatibility constraint. This decomposition holds for any continuous function φ , which enters only by selecting the feasible assignments that satisfy target x .

Proof idea. Concave duality gives the shadow-price formula. The loss identity follows from the two-moment optimality condition: w equals a constant plus $2\lambda r^*$ and the normal vector n . Since both assignments have the same second moment, their difference satisfies $-2\langle r^*, d \rangle = \|d\|_2^2 + V_\sigma$. Substitution gives the three nonnegative loss terms.

8 Conclusion

Reproducing a population is an aggregate requirement; screening is an allocative requirement governing how actions are matched to types. This paper examines how the two interact when an institution has committed in advance to a prescribed quantity schedule. Under a nonincreasing schedule, incentive compatibility restricts the action distribution in a single fundamental way: through a convex-order lower bound generated by the conditional-mean schedule that holds the effective allocation constant. A stationary target distribution is implementable if and only if some action marginal dominating this bound reproduces it, and the transition technology dictates whether this condition reduces to finitely many aggregate action statistics.

This convex-order bound directly translates into an economic cost. On an action interval, it implies a sharp variance floor that increases with the relative dispersion of committed quantities and decreases with the effectiveness of the action multiplier. Variance above this floor is initially productive, enabling sorting until output reaches its mean-constrained maximum; beyond this point, within-type lotteries absorb any residual variance required by stationarity without output loss. At an interior feasible mean and a feasible second moment above the floor, strict incentives approach the weak output maximum. Further moment restrictions can lower output and compel pooling at boundary targets. In the productive range, the loss relative to the deterministic two-moment optimum separates into three terms under the stated multiplier conditions. The value and loss results also apply to revenue with ϕq replacing zq ; the upper-tail rule requires ϕq to be strictly increasing.

This cost is driven by the quantity commitment rather than by the stationarity requirement itself. With nonnegative types and interval actions, an institution free to choose quantities would set them at the common upper bound and sort actions comonotonically, preserving any technologically feasible action marginal. Under a prescribed schedule, sorting must operate through the effective allocation, and the commitment binds on two margins: it imposes a variance floor that can rule out incentive-compatible implementation of technologically feasible targets, and it limits the extent to which required dispersion can be converted into sorting, pricing that loss in forgone output.

Three modeling features define the scope of this analysis and suggest natural directions for future inquiry. First, randomization is costless in our benchmark (Assumption 2); direct administrative costs of lotteries would change the designer’s objective while leaving moment feasibility unchanged, including where lotteries supply residual variance. Second, short-lived incumbents (Assumption 1(ii)) keep the reporting problem static, ensuring that incentives order the effective allocation alone. Extending the framework to long-lived agents would introduce continuation values, coupling current reports to future private types. Third, exact period-by-period reproduction (Assumption 1(iv)) fixes the aggregate action requirements for each cohort. The analysis therefore concerns assignments that reproduce the target in every period.

More broadly, the methodology developed here applies whenever an aggregate distributional constraint falls on one variable, incentive compatibility orders another, and an inflexible schedule links the two.

A Proofs for Screening Feasibility

Proof of Proposition 1. Incentive compatibility requires a nondecreasing effective allocation g between $(\alpha + \beta a)q$ and $(\alpha + \beta b)q$. Thus, for $\tilde{z} \leq z$,

$$(\alpha + \beta a)q(\tilde{z}) \leq g(\tilde{z}) \leq g(z) \leq (\alpha + \beta b)q(z),$$

which gives (7). Conversely, under that condition,

$$g_0(z) := (\alpha + \beta a) \max_{\tilde{z} \leq z} q(\tilde{z})$$

is continuous, nondecreasing, and lies between the required bounds. Its conditional mean $(g_0/q - \alpha)/\beta$ lies in $[a, b]$ and is implemented by an endpoint lottery.

If q is nonincreasing, the running maximum is $q(\underline{z})$. The range condition becomes $g_L \leq g_U$, and every constant $g \in [g_L, g_U]$ is feasible. The mean equality $\mathbb{E}[hg] = S$ selects $g = S/H = \bar{g}_m$. $\square$

Proof of Proposition 2. Let π be the action marginal of any feasible assignment. Atomlessness and full support of μ imply that $T^\uparrow(z) = F_\pi^{-1}(F_\mu(z))$ is nondecreasing and has marginal π , so it preserves all transition moments. Since $Z \geq 0$, $q \leq \bar{q}$, and $\alpha + \beta T > 0$, the rearrangement inequality gives

$$\mathbb{E}[Zq(Z)(\alpha + \beta T)] \leq \bar{q}\{\alpha\mathbb{E}[Z] + \beta\mathbb{E}[ZT]\} \leq \bar{q}\{\alpha\mathbb{E}[Z] + \beta\mathbb{E}[ZT^\uparrow]\}.$$

The last value is attained under $q \equiv \bar{q}$, whose effective allocation is nondecreasing. This proves (i). For (ii), Jensen gives $\mathbb{E}[T^2] \geq m^2$, attained by $T \equiv m$ on $[0, 1]$. The constant quantity schedule has range ratio $1 < (\alpha + \beta)/\alpha$, since $\alpha > 0$, so the range condition holds strictly. $\square$

B Proofs for Implementable Action Distributions

Proof of Theorem 1. Mean feasibility. Any nondecreasing effective allocation with an action in $[a, b]$ lies between g_L and g_U . To see the lower bound, use $g(z) \geq (\alpha + \beta a)q(z)$, monotonicity of g , continuity of q , and full support of μ , and let types approach $\underline{z}$. The upper bound follows by approaching $\bar{z}$ in $g(z) \leq (\alpha + \beta b)q(z)$. Since $\mathbb{E}[(1/q)g] = \alpha + \beta m = H\bar{g}_m$, their weighted mean $\bar{g}_m$ must lie between these bounds.

Conversely, (8) implies $r_m^0 \in [a, b]$ pointwise. This follows by evaluating its smallest value at $\underline{z}$ and its largest value at $\bar{z}$; $\bar{g}_m > 0$ and $1/q$ is nondecreasing. For each type, randomize between a and b with probability $(r_m^0(z) - a)/(b - a)$ on b . The resulting mean is r_m^0 , so the aggregate mean is m and $g \equiv \bar{g}_m$. Constant payments then implement the assignment. This proves part (i).

The convex-order lower bound. Fix any incentive-compatible assignment of mean m , and write $r = \mathbb{E}[T \mid Z]$, $r_0 = r_m^0$, and $g = q(\alpha + \beta r)$. Choose a nondecreasing representative of g . Both r and r_0 are nondecreasing. Indeed, $\alpha + \beta r = g/q$ is the product of two positive nondecreasing functions,

and $\alpha + \beta r_0 = \bar{g}_m/q$. Also,

$$r - r_0 = \frac{g - \bar{g}_m}{\beta q}, \quad \mathbb{E}[r - r_0] = 0.$$

The difference can change sign only from negative to positive.

To establish convex order, let F be the distribution function of μ and define, for $0 < u < 1$,

$$R(u) := r(F^{-1}(u)), \quad R_0(u) := r_0(F^{-1}(u)).$$

Atomlessness and full support make F continuous and strictly increasing on the type interval. The difference $R - R_0$ has integral zero and the same sign pattern as $r - r_0$. Consequently

$$\int_0^u [R(v) - R_0(v)] dv \leq 0 \quad \text{for all } u \in [0, 1], \quad (\text{B.1})$$

with equality at $u = 1$. Before the sign change, the integrand is nonpositive. After the sign change, the integral equals minus the remaining nonnegative tail. For any bounded nondecreasing function X on $(0, 1)$,

$$\int_0^1 (s - X(v))_+ dv = \max_{0 \leq u \leq 1} \left\{ su - \int_0^u X(v) dv \right\}.$$

The maximum includes precisely the initial segment where $X(v) < s$, with any part of a flat segment at s immaterial. Applying this identity to (B.1) gives

$$\mathbb{E}[(s - r_0(Z))_+] \leq \mathbb{E}[(s - r(Z))_+] \quad \text{for every } s \in [a, b].$$

Every continuous convex function on $[a, b]$ is a uniform limit of convex piecewise-affine functions. Each approximating function is affine plus a finite nonnegative combination of $(s - t)_+$. Equality of means therefore implies $\mathcal{L}(r_0(Z)) \preceq_{\text{cx}} \mathcal{L}(r(Z))$. Conditional Jensen's inequality then yields

$$\pi_m^0 \preceq_{\text{cx}} \mathcal{L}(r(Z)) \preceq_{\text{cx}} \mathcal{L}(T).$$

This proves necessity in part (ii), as well as the equivalence of its convex-order condition and the expectation inequalities.

Sufficiency of the convex-order bound. Suppose $\pi_m^0 \preceq_{\text{cx}} \pi$. By the martingale-coupling theorem of Strassen (1965), as stated in Leskelä and Vihola (2017, Theorem 1.1), there is a probability measure γ on $[a, b] \times \mathcal{T}$ with marginals π_m^0 and π for which $\mathbb{E}_\gamma[T \mid X] = X$ almost surely. Disintegrate γ with respect to its first marginal to obtain a Borel probability kernel $K(dt \mid x)$. On the π_m^0 -null exceptional set where its mean is not x , replace K by the endpoint lottery with mean x . This replacement is measurable and leaves γ unchanged.

Now set

$$\sigma(dt \mid z) := K(dt \mid r_m^0(z)).$$

It is a Borel assignment rule, its marginal is π , and its conditional mean is $r_m^0(z)$ at every type. Its effective allocation is $\bar{g}_m$. For example, the constant payment $p(z) = z\bar{g}_m$ gives truthful utility $(z - \underline{z})\bar{g}_m \geq 0$, independent of the report. The mechanism is pointwise weakly incentive compatible and individually rational. This proves sufficiency and the final claim in part (ii).

The finite-action test. Define

$$D(s) := \sum_{i=1}^J p_i(s - t_i)_+ - \mathbb{E}[(s - r_m^0(Z))_+].$$

On each interval $[t_j, t_{j+1}]$, the first term is affine and the second term is convex, so D is concave. It is nonnegative on that interval whenever it is nonnegative at both endpoints. Moreover $D(a) = D(b) = 0$ because both distributions are supported on $[a, b]$ and have mean m . Thus the expectation inequalities hold for every s exactly when (9) holds at the interior action levels. This proves part (iii). $\square$

Proof of Corollary 1. For any assignment with marginal π ,

$$\mu P_\sigma = \zeta \mu + (1 - \zeta) \int \hat{Q}_t \pi(dt).$$

Since $\zeta < 1$, stationarity is equivalent to the measure equality in $\mathcal{A}_\mu(m)$. Theorem 1(ii) characterizes the implementable marginals as $\mathcal{U}_q(m)$, proving (12). For sufficiency, implement any marginal in the intersection with a constant effective allocation. The aggregate transition depends only on that marginal. For the polynomial specialization, integrating the polynomial kernel and using independence of $A_1, \dots, A_N$ makes the measure equality equivalent to the displayed moment equalities. This gives (13). $\square$

Verification of Example 2. Here $H = 5/4$, $\bar{g}_m = 6/5$, $g_L = 1$, and $g_U = 4/3$, so the mean is feasible and $r_m^0 = 1/5 + 3z/5$. Its law is uniform on $[1/5, 4/5]$. On $[1/5, 3/10]$, the adjacent-action lottery puts probability $(3/10 - r)/(3/10)$ on zero. Integrating with density $5/3$ yields marginal mass $1/36$ on zero. Symmetry gives the same mass on one and mass $17/36$ on each of $3/10$ and $7/10$. These lotteries have conditional mean r_m^0 and constant effective allocation. The endpoint lottery with that same conditional mean has marginal masses $(1/2, 1/2)$ on $(0, 1)$, and hence second moment $1/2$. The mixture weight solves

$$(1 - \theta) \frac{181}{600} + \theta \frac{1}{2} = \frac{3047}{10000}, \quad \theta = \frac{13}{850} \in (0, 1).$$

It preserves the conditional mean pointwise and thus establishes two-moment feasibility on the same action set.

Direct summation gives the first three moments of π° . Its expectation of $(3/10 - T)_+$ is

$(3/10)(1/100) = 3/1000$. For the benchmark with constant effective allocation, the expectation is

$$\frac{5}{3} \int_{1/5}^{3/10} (3/10 - r) dr = \frac{1}{120}.$$

The four-by-four Vandermonde matrix with columns $(1, t_j, t_j^2, t_j^3)^\top$ is nonsingular. Hence the first three moments and normalization identify π° uniquely on $\mathcal{T}$. The violated convex-order inequality rules out an incentive-compatible assignment matching all three moments.

To exhibit a compatible stationary environment, let

$$\begin{aligned} f_1(z') &= 2z' - 1, \\ f_2(z') &= 6(z')^2 - 6z' + 1, \\ f_3(z') &= 20(z')^3 - 30(z')^2 + 12z' - 1. \end{aligned}$$

These linearly independent polynomials have zero integrals on $[0, 1]$ and absolute value at most one there. Let μ remain uniform, use the displayed target moments, and define, for $t \in [0, 1]$,

$$\widehat{Q}_t(dz') = \left[1 + \frac{1}{4} \sum_{\ell=1}^3 (t^\ell - m_\ell) f_\ell(z') \right] dz'.$$

The density integrates to one and is at least $1/4$, because $|t^\ell - m_\ell| \leq 1$. Its coefficient measures are independent and its dependence on (z', t) is continuous. Set $\zeta = 0$. Stationarity of the uniform target then holds exactly when $\mathbb{E}[T^\ell] = m_\ell$, $\ell = 1, 2, 3$. The marginal π° reproduces that target without incentive constraints, but no truthful assignment does. The two-moment relaxation remains feasible by the construction above. $\square$

C Proof of the Transition Representation

Proof of Theorem 2. (i) Representation. Integration proves (c) $\Rightarrow$ (a), and (a) $\Rightarrow$ (b) is immediate. Assume (b). Fix $t, u \in \mathcal{T}$, choose t^+, t^- with $x_1(t^+) > x_1(t^-)$, and let π^0 put positive mass on every distinct action in $\{t, u, t^+, t^-\}$. With $s(z) = (z - \underline{z})/(\bar{z} - \underline{z})$, set

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

For small $\gamma > 0$, this is a probability kernel whose conditional mean of x_1 has strictly positive constant slope. For any $f \in C^1(\mathcal{Z})$ with $\mathbb{E}[f] = 0$, the perturbations

$$\sigma^\pm(dv | z) = \sigma^b(dv | z) \pm \varepsilon f(z)(\delta_t - \delta_u)(dv)$$

remain nonnegative with nondecreasing conditional means of x_1 for sufficiently small $\varepsilon > 0$. Their aggregate statistics coincide. Writing $Q_t\phi(z) = \int \phi(z')Q_t(dz' | z)$, sufficiency gives

$$\int f(z)[Q_t\phi(z) - Q_u\phi(z)]\mu(dz) = 0 \quad (\phi \in C(\mathcal{Z})).$$

Zero-mean C^1 functions are dense in the zero-mean subspace of $L^2(\mu)$. Hence the bracketed continuous function is constant μ -almost everywhere, and full support makes it constant everywhere. Continuous test functions separate measures, so $Q_t - Q_u$ is independent of current type.

Fix t_0 . Choose $t_1, \dots, t_d$ such that the vectors $x(t_j) - x(t_0)$ form a basis, and define the unique signed measures D_i by

$$\Delta_{t_j} = \sum_{i=1}^d [x_i(t_j) - x_i(t_0)]D_i, \quad j = 1, \dots, d. \quad (\text{C.1})$$

They have zero total mass. For any t , write $x(t) - x(t_0) = \sum_j c_j[x(t_j) - x(t_0)]$. The signed measure

$$\omega = \delta_t - \delta_{t_0} - \sum_j c_j(\delta_{t_j} - \delta_{t_0})$$

has zero mass and zero x -moment. If $\omega \neq 0$, its normalized positive and negative parts define constant assignment rules with the same statistic. Thus $\int \Delta_v \omega(dv) = 0$, also true when $\omega = 0$. Combining this equality with (C.1) gives (14), including uniqueness.

(ii) *Minimal dimension.* Removing redundant affine coordinates and applying (i) gives $d \geq d_Q$ for every sufficient statistic of positive affine dimension d . If a zero-dimensional statistic suffices, compare the constant action u with an assignment using t on any Borel set A and u elsewhere. Then $\int_A (Q_t - Q_u) d\mu = 0$ for every A ; continuity and full support imply $Q_t = Q_u$. Thus $d_Q = 0$, which also plainly suffices.

For $0 < d_Q < \infty$, choose a basis $D_1, \dots, D_{d_Q}$ of $\mathcal{D}_Q$ and define

$$\Delta_t = \sum_{i=1}^{d_Q} x_i(t)D_i, \quad x(t_0) = 0. \quad (\text{C.2})$$

Choose continuous test functions ϕ_j for which $(\int \phi_j dD_i)_{j,i}$ is nonsingular. Testing (C.2) against these functions shows that x is continuous by the Feller property. Its image spans $\mathbb{R}^{d_Q}$, and integration makes it sufficient, proving minimality. Any two minimal statistics represent bases of $\mathcal{D}_Q$; the change of basis gives (15).

(iii) *Stationarity.* For a minimal statistic, the D_i are linearly independent. Integrating the representation under any lottery of mean statistic $\bar{x}$ gives

$$Q_{\bar{x}} = Q_{t_0} + \sum_i [\bar{x}_i - x_i(t_0)]D_i, \quad \mu P_{\sigma} - \mu Q_{\bar{x}} = \sum_i (\bar{x}_{\sigma,i} - \bar{x}_i)D_i.$$

The first identity makes $Q_{\bar{x}}$ well defined; the second and linear independence give (16) when $\mu Q_{\bar{x}} = \mu$. $\square$

D Proofs for Optimal Output under Two Moment Restrictions

Proof of Lemma 1. If $r \notin \mathcal{T}$, let $\ell = t_-(r) < r < u = t_+(r)$. Since $(t - \ell)(t - u) \geq 0$ on $\mathcal{T}$, every lottery of mean r satisfies

$$\mathbb{E}[T^2] \geq (\ell + u)r - \ell u = r^2 + (r - \ell)(u - r).$$

Equality requires the unique mean- r lottery on $\{\ell, u\}$. For $r \in \mathcal{T}$, Jensen gives the bound r^2 , attained only by δ_r . Similarly, $(t - a)(t - b) \leq 0$ gives $\mathbb{E}[T^2] \leq (a + b)r - ab$, with equality only for the mean- r endpoint lottery. Mixing these extremal lotteries attains every intermediate moment. This proves (17). The lower envelope linearly interpolates r^2 across gaps of $\mathcal{T}$; it is continuous, convex, and Lipschitz, and is strictly convex on $[a, b]$ exactly when $\mathcal{T} = [a, b]$.

For a Borel conditional-mean schedule r , integration gives

$$\mathbb{E}[K_{\mathcal{T}}(r(Z))] \leq \mathbb{E}[T^2] \leq \overline{K}_{\mathcal{T}}(\mathbb{E}[r(Z)]).$$

The extremal lotteries are Borel in r . Let $d(z) = \overline{K}_{\mathcal{T}}(r(z)) - \underline{K}_{\mathcal{T}}(r(z))$. The measure $\vartheta(A) = \int_A d \, d\mu$ is atomless, so its range is $[0, \mathbb{E}[d]]$ by Lyapunov's theorem (Aliprantis and Border, 2006, Theorem 13.33). Choosing a set A of the required ϑ -mass and using the upper lottery on A and the lower lottery elsewhere attains the target second moment with at most two actions at each type, preserving every conditional mean. $\square$

Proof of Theorem 3(i). The set $\mathcal{C}_{q,\mathcal{T}}(m)$ is bounded, convex, and norm closed in $L^2(\mu)$. For closedness, an L^2 -convergent sequence of effective allocations has a subsequence converging almost everywhere; its monotone limit extends by one-sided limits from the full-measure, dense set of convergence. The pointwise bounds and mean equality are also closed. Reflexivity therefore makes $\mathcal{C}_{q,\mathcal{T}}(m)$ weakly compact. By Lemma 1, $\mathcal{K}_{\mathcal{T}}(r) = \mathbb{E}[\underline{K}_{\mathcal{T}}(r)]$ is convex and norm continuous, hence weakly lower semicontinuous. Its minimum is attained, and every nonempty sublevel set is weakly compact. The linear objective $\mathbb{E}[wr]$ attains its maximum on each such set.

The lemma also shows that a conditional-mean schedule $r \in \mathcal{C}_{q,\mathcal{T}}(m)$ can implement exactly the second moments $\mathcal{K}_{\mathcal{T}}(r) \leq k \leq \overline{K}_{\mathcal{T}}(m)$. Minimizing the lower bound proves (19). Applying the lemma to an optimal r implements the exact moment k with at most two actions at each type and preserves $\mathbb{E}[wT] = \mathbb{E}[wr]$. This proves (18) and the output formula. $\square$

Proof of Theorem 3(ii). Write $L = L_{q,\mathcal{T}}(m)$, $K = k_{q,\mathcal{T}}^{\text{mc}}(m)$, $U = \overline{K}_{\mathcal{T}}(m)$, and $B(k) = B_{q,\mathcal{T}}(m, k)$. Mixing feasible mean schedules, using convexity of $\mathcal{K}_{\mathcal{T}}$ and linearity of output, proves concavity of B . Expanding sublevel sets give monotonicity.

If $L < U$, concavity gives continuity on (L, U) . For $k_n \downarrow L$, weak compactness gives a weak limit of a subsequence of optimal schedules. Lower semicontinuity makes the limit feasible at L , and continuity of output gives $B(k_n) \rightarrow B(L)$. At U , mix an optimizer r_U with a minimizer r_L of $\mathcal{K}_{\mathcal{T}}$. For $\theta < 1$,

$$\mathcal{K}_{\mathcal{T}}(\theta r_U + (1 - \theta)r_L) \leq \theta U + (1 - \theta)L < U,$$

while output tends to $B(U)$ as $\theta \uparrow 1$. This proves left continuity. The case $L = U$ is immediate.

A mean-constrained maximizer with required moment K is feasible for every $k \geq K$, so B is constant on $[K, U]$. If $B(k_1) = B(k_2)$ for $L \leq k_1 < k_2 \leq K$, concavity and monotonicity give $B(k_1) = B(K)$. An optimizer at $k_1 < K$ would then attain the mean-constrained maximum with required moment below K , a contradiction. Hence B is strictly increasing on $[L, K]$ when $L < K$, and is constant throughout when $L = K$. $\square$

Proof of Proposition 3. (i) Binding and lower-envelope lotteries. Write $L = L_{q,\mathcal{T}}(m)$, $K = k_{q,\mathcal{T}}^{\text{mc}}(m)$, and $B(k) = B_{q,\mathcal{T}}(m, k)$. Any optimizer at $k \leq K$ has $\mathcal{K}_{\mathcal{T}}(r) = k$: a smaller required moment contradicts the definition of L at $k = L$, strict increase of the value at $L < k < K$, or the definition of K at $k = K$. The nonnegative conditional gap $\mathbb{E}[T^2 \mid Z] - \underline{K}_{\mathcal{T}}(r(Z))$ therefore vanishes almost everywhere. The lemma's equality conditions give the lower-envelope lottery at almost every type.

(ii) Interval actions. Here $\mathcal{K}_{\mathcal{T}}(r) = \|r\|_2^2$. Strict convexity gives a unique minimum-norm feasible schedule and a unique minimum-norm mean-constrained maximizer. For any $k \leq K$, two distinct optimal schedules would have an optimal midpoint with norm squared below k , contradicting (i). Conditional Jensen then makes every optimal assignment deterministic.

At $L < k < K$, the minimum-norm schedule is strictly feasible. Slater's condition gives an attained multiplier $\lambda \geq 0$ for which r_k maximizes $\mathbb{E}[wr - \lambda r^2]$ over $\mathcal{C}_{q,[a,b]}(m)$. Since $B(k) < B(K)$, $\lambda > 0$. Completing the square gives (20). Define

$$\Phi_{\lambda}(s) = \sup_{\substack{r \in \mathcal{D}_{q,[a,b]} \\ \mathbb{E}[r] = s}} \mathbb{E}[wr - \lambda r^2].$$

This finite concave function has a supergradient η at the relative-interior mean m . For any $r \in \mathcal{D}_{q,[a,b]}$,

$$\mathbb{E}[(w - \eta)r - \lambda r^2] \leq \Phi_{\lambda}(m) - \eta m = \mathbb{E}[(w - \eta)r_k - \lambda r_k^2].$$

Completing the square proves (21) and its normal-cone form.

If the joint dual is attained at $(\varpi_1^*, \varpi_2^*, \xi^*)$, its second-moment component is a supergradient of B ; $B(K) > B(k)$ implies $\varpi_2^* > 0$. Complementary slackness gives $\langle \xi^*, g_{r_k} \rangle = 0$, and the support condition makes r_k maximize $[w - \varpi_1^* - \beta q \xi^*]r - \varpi_2^* r^2$ pointwise. Integration, cancellation of the mean terms, and $\langle \xi^*, g_r \rangle \leq 0$ give

$$\mathbb{E}[wr_k - \varpi_2^* r_k^2] \geq \mathbb{E}[wr - \varpi_2^* r^2] \quad (r \in \mathcal{C}_{q,[a,b]}(m)).$$

Thus (20) holds with $\lambda = \varpi_2^*$.

(iii) *Finite actions.* The lower envelope is affine between adjacent actions, so (i) gives the claimed adjacent-action lotteries. Nonuniqueness occurs for $Z \sim U[1, 2]$, $q(z) = 1/z$, $\alpha = 0$, $\beta = 1$, $\mathcal{T} = \{1/4, 3/4, 1\}$, and $m = 9/20$. The distinct schedules

$$r_1(z) = \frac{3z}{10}, \quad r_2(z) = z \left[\frac{3}{10} + \frac{z - 14/9}{100} \right]$$

are incentive compatible, have mean m , and lie in $[1/4, 3/4]$, where $\underline{K}_{\mathcal{T}}(r) = r - 3/16$. Both attain the Jensen bound $\mathcal{K}_{\mathcal{T}}(r) \geq \underline{K}_{\mathcal{T}}(m) = 21/80$. Since $w \equiv 1$, every mean- m schedule also maximizes output. Thus $L = K = 21/80$, with both schedules optimal. $\square$

Proof of Theorem 4. Take a nondecreasing representative of the effective allocation

$$g = q(\alpha + \beta r).$$

By Proposition 1, the feasible effective allocations are the nondecreasing functions with values in $[g_L, g_U]$. The mean-action restriction is

$$\mathbb{E}[hg] = S.$$

Part (i): the minimum second moment. The schedule r_m^0 has mean m , is feasible under the hypotheses, and makes $g \equiv \bar{g}_m = S/H$. By Theorem 1(ii), its distribution is below that of every incentive-compatible action in convex order. Applying the convex test t^2 shows that the deterministic assignment $T = r_m^0(Z)$ minimizes the second moment. The set of feasible mean schedules is convex and the squared L^2 -norm is strictly convex, so the minimizing conditional-mean schedule is unique. Equality in conditional Jensen for t^2 also requires the minimizing assignment to be deterministic. Finally,

$$r_m^0 - m = \frac{S}{\beta} \left(\frac{h}{H} - 1 \right),$$

which gives (23).

Part (ii): the mean-constrained output maximum. Let

$$\nu(\mathrm{d}z) = \frac{h(z)}{H} \mu(\mathrm{d}z).$$

The mean constraint becomes $\mathbb{E}_{\nu}[g] = \bar{g}_m$, and

$$\mathbb{E}[wr] = \frac{H}{\beta} \mathbb{E}_{\nu}[wg] - \frac{\alpha}{\beta} \mathbb{E}[w].$$

It is therefore enough to maximize $\mathbb{E}_{\nu}[wg]$ subject to

$$g_L \leq g \leq g_U, \quad \mathbb{E}_{\nu}[g] = \bar{g}_m.$$

First drop monotonicity. Maximizing the linear objective assigns g_U to the largest values of w and g_L to the remaining types. Since w is strictly increasing, the maximizing set is an upper interval, unique up to a ν -null set. The resulting allocation is itself nondecreasing and therefore satisfies incentive compatibility.

The required upper-tail probability is

$$\frac{\bar{g}_m - g_L}{g_U - g_L}.$$

The measure ν is atomless and has full support, so its upper-tail mass is continuous and strictly decreasing in the cutoff. This gives the unique z^* in (24) and the effective allocation g^* stated in the theorem. The corresponding conditional-mean schedule is $\hat{r}_q = (hg^* - \alpha)/\beta$. It is unique up to null sets. Since the action set is the full interval, the deterministic assignment $T = \hat{r}_q(Z)$ has the smallest second moment among assignments with this conditional mean. Computing $\mathbb{E}[\hat{r}_q^2]$ gives (25).

Part (iii): the distance between the thresholds. The two schedules have the same mean. Hence

$$\begin{aligned}\beta^2 \mathbf{P}_q(m) &= \beta^2 \mathbb{E}[(\hat{r}_q)^2 - (r_m^0)^2] \\ &= \mathbb{E}[h^2((g^*)^2 - \bar{g}_m^2)],\end{aligned}$$

where the term linear in $g^* - \bar{g}_m$ vanishes because $\mathbb{E}[h(g^* - \bar{g}_m)] = 0$. Expanding around $\bar{g}_m$ and using $\mathbb{E}_\nu[g^*] = \bar{g}_m$ yields

$$\mathbf{P}_q(m) = \frac{H}{\beta^2} \left\{ \mathbb{E}_\nu[h(g^* - \bar{g}_m)^2] + 2\bar{g}_m \text{Cov}_\nu(h, g^*) \right\}.$$

The first term is strictly positive because g^* takes both boundary values on sets of positive ν -measure. The covariance is nonnegative because h and g^* are nondecreasing. This proves (26) and its strict positivity.

Suppose now that q is strictly decreasing. Then h is nonconstant, so (23) gives $\mathbf{C}_q(m) > 0$. On the nonempty interval $(\underline{z}, z^*)$, the optimal effective allocation is g_L . Strict decline and $g_L < g_U$ imply

$$0 < \hat{r}_q(z) < 1$$

except possibly at an endpoint. Full support therefore gives

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

Together with the strict positivity of $\mathbf{P}_q(m)$, all three terms in (22) are positive.

If q is constant, then h is constant and the minimum schedule is $r_m^0 \equiv m$. The mean-constrained optimum assigns actions 0 and 1 by an upper-tail cutoff of mass m , so $k_q^{\text{mc}}(m) = m$. Hence

$$(\mathbf{C}_q(m), \mathbf{P}_q(m), \mathbf{N}_q(m)) = (0, m(1 - m), 0).$$

If $g_L = g_U$, every feasible effective allocation is constant. The second-moment minimizer and the mean-constrained output maximizer then coincide, so $P_q(m) = 0$ and $L_q(m) = k_q^{\text{mc}}(m)$. $\square$

E Proof of Strict Implementation

Proof of Proposition 4. Strict truthfulness and effective allocations. Adding the strict reporting inequalities for any two different types gives $(z' - z)(g(z') - g(z)) > 0$, so strict truthfulness requires g to be strictly increasing. Conversely, the envelope payment gives, when $z > \hat{z}$,

$$[zg(z) - p(z)] - [zg(\hat{z}) - p(\hat{z})] = \int_{\hat{z}}^z [g(x) - g(\hat{z})] dx > 0.$$

For $z < \hat{z}$, the corresponding expression is $\int_z^{\hat{z}} [g(\hat{z}) - g(x)] dx > 0$. The envelope integral is with respect to Lebesgue measure, so the argument applies to all type distributions allowed by the assumptions. Positivity of g and $V_0 \geq 0$ give individual rationality. Choosing $V_0 > 0$ makes participation strictly preferable to the outside option.

Part (i): a strictly truthful assignment with exact moments. Write $h = 1/q$, $H = \mathbb{E}[h]$, and $\nu(dz) = h(z)\mu(dz)/H$. First observe that

$$L_{q,\mathcal{T}}(m) = \mathbb{E}[\underline{K}_{\mathcal{T}}(r_m^0(Z))]. \quad (\text{E.1})$$

Indeed, the conditional lower-envelope lottery at r_m^0 is weakly implementable and attains the right side. Conversely, $\underline{K}_{\mathcal{T}}$ is continuous and convex, and $\underline{K}_{\mathcal{T}}(t) = t^2$ for $t \in \mathcal{T}$. Applying Theorem 1(ii) to this convex function gives the right side as a lower bound on the second moment of every weak implementation. The conditional envelope characterization in Lemma 1 then proves (E.1).

Let $\ell_z = \bar{z} - \underline{z} > 0$ and define

$$s(z) = \frac{z - \underline{z}}{\ell_z}, \quad v(z) = s(z) - \mathbb{E}_{\nu}[s(Z)].$$

Then $|v| \leq 1$, $\mathbb{E}[hv] = 0$, and $v(z') - v(z) = (z' - z)/\ell_z$. For $\delta > 0$, set

$$g_{\delta}(z) = \bar{g}_m + \delta v(z), \quad r_{\delta}(z) = \frac{h(z)g_{\delta}(z) - \alpha}{\beta}.$$

If $\delta < \min\{\bar{g}_m - g_L, g_U - \bar{g}_m\}$, then g_{δ} lies strictly between the global effective-allocation bounds, $r_{\delta}(z) \in (a, b)$ for every z , and $\mathbb{E}[r_{\delta}] = m$. The schedule g_{δ} is continuous and has slope δ/ℓ_z .

Let $\sigma_{\delta}^{-}(dt | z)$ be the measurable conditional lower-envelope lottery with mean $r_{\delta}(z)$, whose existence follows from Lemma 1. Its aggregate second moment is

$$k_{\delta} = \mathbb{E}[\underline{K}_{\mathcal{T}}(r_{\delta}(Z))] \longrightarrow L_{q,\mathcal{T}}(m) \quad \text{as } \delta \downarrow 0,$$

since r_{δ} converges uniformly to r_m^0 and $\underline{K}_{\mathcal{T}}$ is continuous on a compact interval. Choose δ small

enough that $k_\delta < k$. At the same conditional mean, the endpoint lottery is

$$e_{r_\delta(z)} = \frac{b - r_\delta(z)}{b - a} \delta_a + \frac{r_\delta(z) - a}{b - a} \delta_b.$$

Its aggregate second moment is $\bar{K}_\mathcal{T}(m) = (a + b)m - ab$. With

$$\tau = \frac{k - k_\delta}{\bar{K}_\mathcal{T}(m) - k_\delta} \in (0, 1], \quad \hat{\sigma} = (1 - \tau)\sigma_\delta^- + \tau e_{r_\delta},$$

the first two moments are exactly (m, k) , and the effective allocation is still g_δ . Envelope transfers make this mechanism strictly truthful, with $c = \delta/\ell_z$. The case $k = \bar{K}_\mathcal{T}(m)$ simply gives $\tau = 1$.

Part (ii): approximation within the same constraints. Choose a pointwise weakly implementable version of σ , modifying it on a μ -null set if necessary as in Theorem OA.1(i). The conditional mixture (28) preserves both moments by linearity. Its effective allocation is

$$g_\theta = (1 - \theta)g_\sigma + \theta g_\delta, \quad g_\theta(z') - g_\theta(z) \geq \theta c(z' - z) \quad (z' > z).$$

Hence the envelope payments satisfy the quantitative strict-reporting bound

$$[zg_\theta(z) - p_\theta(z)] - [zg_\theta(\hat{z}) - p_\theta(\hat{z})] \geq \frac{\theta c}{2}(z - \hat{z})^2.$$

The lottery is drawn conditional on the report, so incentive comparisons are made before its realization.

Define total variation distance by $d_{\text{TV}}(P, Q) = \sup_A |P(A) - Q(A)|$. Then

$$\sup_z d_{\text{TV}}(\sigma_\theta(\cdot | z), \sigma(\cdot | z)) \leq \theta.$$

The effective allocations are bounded, so $g_\theta \rightarrow g_\sigma$ uniformly. Since the envelope operator is linear at a common V_0 , $p_\theta = (1 - \theta)p_\sigma + \theta p_{\hat{\sigma}}$, and payments also converge uniformly. Output satisfies

$$Y(\sigma_\theta) = (1 - \theta)Y(\sigma) + \theta Y(\hat{\sigma}).$$

Apply the construction to a weak optimum to obtain (29). Set inclusion gives the reverse inequality. This equates the strict supremum with the weak maximum.

Part (iii): necessity for interval actions. For any pointwise weak implementation, monotonicity and the action bounds imply $g_L \leq g(z) \leq g_U$, while $\mathbb{E}_\nu[g] = \bar{g}_m$. If g is strictly increasing, full support of ν makes this expectation strictly between the two global bounds. For interval actions, $\underline{K}_\mathcal{T}(r) = r^2$, so (E.1) gives (30). The feasible mean-schedule set is convex, and its squared $L^2(\mu)$ norm is strictly convex; thus r_m^0 is its unique norm minimizer. If $\mathbb{E}[T^2] = L_{q,[a,b]}(m)$, conditional Jensen forces $T = r_m^0(Z)$ almost surely and $g = \bar{g}_m$ almost everywhere. A strictly increasing function cannot be constant almost everywhere under a full-support measure. Therefore strict implementation requires

$k > L_{q,[a,b]}(m)$. The upper bound on k is necessary for any action in $[a, b]$. Sufficiency was proved in part (i). $\square$

F Proofs for the Additional Moment Restriction

F.1 Attainable moment values

Proof of Proposition 5. Write $\gamma(dz, dt) = \mu(dz)\sigma(dt | z)$ for the Young measure induced by an assignment rule. The set of Young measures with first marginal μ is compact and convex in the stable topology. The two moment equalities are closed in the stable topology. The IC restriction is also closed because, by (OA.8) in Online Appendix D, it is the intersection over $\xi \in \mathcal{H}_\mu$ of the closed half-spaces $\langle \xi, g_\sigma \rangle \leq 0$. Hence Σ_2 is compact and convex.

The maps Y and M_φ are continuous and affine on this set: their integrands belong to $L^1(\mu; C(\mathcal{T}))$. The image $I_\varphi = M_\varphi(\Sigma_2)$ is therefore a nonempty compact interval. The optimizer set Σ_2^* is a nonempty compact convex face of Σ_2 , so its image I_φ^* is also a compact interval.

Convex combinations of assignment rules with different moment values show that Y_φ is concave. To prove upper semicontinuity, let $x_n \rightarrow x$, and pass to a subsequence along which $Y_\varphi(x_n)$ converges to its upper limit. Choose $\sigma_n \in \Sigma_\varphi(x_n)$ attaining $Y_\varphi(x_n)$. Compactness gives a stably convergent subsequence. Its limit belongs to $\Sigma_\varphi(x)$, and continuity of Y yields

$$\limsup_n Y_\varphi(x_n) \leq Y_\varphi(x).$$

Finally, imposing $M_\varphi = x$ leaves the two-moment value unchanged if and only if some two-moment optimizer has moment value x . This proves (31).

Now suppose every two-moment optimizer has conditional-mean schedule r^* . An assignment rule with this conditional mean and aggregate second moment k satisfies the two moment constraints, has the same effective allocation as an optimizer, and therefore belongs to Σ_2^* . Conversely, every element of Σ_2^* has conditional mean r^* by hypothesis.

This set of assignments is compact and convex. The conditional-mean restriction is the intersection, over $f \in L^\infty(\mu)$, of equalities that are closed in the stable topology:

$$\int_{\mathcal{Z} \times \mathcal{T}} f(z)[t - r^*(z)] \gamma(dz, dt) = 0.$$

The aggregate second-moment equality is also closed in this topology. The image under M_φ is therefore the compact interval in (32). If the optimal assignment rule is unique, the image is a singleton. If that assignment is deterministic, its value is $\mathbb{E}[\varphi(r^*)]$. $\square$

Proof of Corollary 3. We prove the result for a finite vector of continuous action statistics. The scalar result follows as a special case. Let M be the associated affine moment map and $C = M(\Sigma_2(m, k))$. This is a nonempty compact convex subset of a finite-dimensional space. By Proposition 4(i), choose a strictly implementable $\hat{\sigma} \in \Sigma_2(m, k)$ and write $\hat{x} = M(\hat{\sigma})$. Fix $x \in \text{ri } C$. If $x = \hat{x}$, the

desired strictly truthful assignment already exists. Otherwise, relative interior implies that, for all sufficiently small $\epsilon > 0$,

$$y = \frac{x - \epsilon \hat{x}}{1 - \epsilon} \in C.$$

Choose a weakly implementable σ_y with $M(\sigma_y) = y$. The conditional mixture $\tau = (1 - \epsilon)\sigma_y + \epsilon\hat{\sigma}$ is strictly implementable and satisfies $M(\tau) = x$, with the first two moments unchanged. This argument also covers a singleton image, whose relative interior is itself. Finally, mix τ with a weak output optimizer at the exact target x , using a mixing weight that tends to zero. Linearity preserves every target moment, strict monotonicity gives strict truthfulness, and output tends to the weak optimum, exactly as in Proposition 4(ii). $\square$

Verification of Example 4. The schedule with constant effective allocation is $r_m^0(z) = 1/5 + 3z/5$. At the same first two moments, let $x = \mathbb{E}[T^3]$. Inverting the four-action Vandermonde system gives the only possible probabilities:

$$\begin{aligned} p_0 &= \frac{53}{42} - \frac{100x}{21}, & p_{3/10} &= \frac{250x}{21} - \frac{211}{84}, \\ p_{7/10} &= \frac{93}{28} - \frac{250x}{21}, & p_1 &= \frac{100x}{21} - \frac{15}{14}. \end{aligned} \tag{F.1}$$

The convex-order slacks at the only two interior action levels are

$$\sum_t p_t(3/10 - t)_+ - \frac{1}{120} = \frac{311 - 1200x}{840}, \quad \sum_t p_t(7/10 - t)_+ - \frac{5}{24} = \frac{1200x - 277}{840}.$$

All four probabilities are positive throughout the closed interval (33). Theorem 1(iii) therefore gives exactly that weakly attainable interval. At its upper endpoint, (F.1) yields (34); the lower-endpoint probabilities are $(41/252, 17/72, 289/504, 1/36)$. The four-action moment map is an affine isomorphism from the probability simplex to a full-dimensional three-moment polytope. Positive probabilities therefore place the target in the interior of the unrestricted moment set.

For a direct weak implementation of the upper endpoint, use the conditional adjacent-action lottery at r_m^0 , whose marginal is $(1/36, 17/36, 17/36, 1/36)$. An alternative lottery uses $\{0, 3/10\}$ below the conditional mean $3/10$ and $\{3/10, 1\}$ above it, retaining the mean at every type. Its probabilities on the four actions are $(1/36, 85/126, 0, 25/84)$. Mixing these two conditional rules with equal weights gives $\pi^\dagger$, with a constant effective allocation. The lower-envelope second moment is $181/600$, as in Example 2. Thus $33/100 > 181/600$, and the strict two-moment conclusion applies.

It remains to exclude strict truthfulness at the endpoints. Let s be either $3/10$ or $7/10$, and let $z_s \in (0, 1)$ satisfy $r_m^0(z_s) = s$. Write $h = 1/q$, $\bar{g}_m = 6/5$, and $d = r - r_m^0 = h(g - \bar{g}_m)$. For a weak implementation, $\mathbb{E}[h(g - \bar{g}_m)] = 0$ and g is nondecreasing. Its h -weighted average on $[0, z_s]$ is no larger than its overall average $\bar{g}_m$. It is strictly smaller unless g is constant almost everywhere: equality of the lower and upper weighted averages, together with monotonicity, forces equality of g

across almost every ordered pair of types. Consequently

$$\int_0^{z_s} d(z) dz \leq 0,$$

with strict inequality whenever the effective allocation is not constant almost everywhere. Conditional Jensen and $(s - r)_+ \geq (s - r)\mathbf{1}\{z \leq z_s\}$ imply

$$\begin{aligned} \mathbb{E}[(s - T)_+] &\geq \mathbb{E}[(s - r)_+] \\ &\geq \int_0^{z_s} (s - r(z)) dz = \mathbb{E}[(s - r_m^0)_+] - \int_0^{z_s} d(z) dz. \end{aligned}$$

Thus a binding convex-order inequality at either value of s forces g to be constant almost everywhere. Full support then rules out pointwise strict incentive compatibility. At interior x , Corollary 3 supplies strict implementation. This proves that the strict attainable interval is exactly the open interval in the example.

To connect the obstruction to stationarity, write $m_\ell = \int t^\ell \pi^\dagger(dt)$ and, for $N = 3$, define

$$\hat{Q}_t(dz') = \left[1 + \frac{1}{2N} \sum_{\ell=1}^N (t^\ell - m_\ell) \sin(2\pi \ell z') \right] dz'. \quad (\text{F.2})$$

The density integrates to one, is at least $1/2$, and is jointly continuous. Its coefficient measures have zero mass and are linearly independent by orthogonality of the sine functions. With $\zeta = 0$, the uniform target is stationary exactly when all prescribed moments are matched. The upper endpoint is therefore a stationary target that is weakly but not strictly implementable. The same construction applies to the lower endpoint using its moments.

Finally, take $\mathcal{T} = [0, 1]$ and prescribe the first six moments of $\pi^\dagger$:

$$(m_1, \dots, m_6) = \left(\frac{1}{2}, \frac{33}{100}, \frac{311}{1200}, \frac{6721}{30000}, \frac{8151}{40000}, \frac{143171}{750000} \right).$$

The polynomial

$$P(t) = t(1-t)(t-3/10)^2(t-7/10)^2$$

is nonnegative on $[0, 1]$, has degree six, and vanishes exactly on the four actions used above. Matching the first six moments gives $\mathbb{E}[P(T)] = \int P(t)\pi^\dagger(dt) = 0$, so the action must lie on those four points almost surely. The first three moments then identify $\pi^\dagger$ uniquely, and the preceding obstruction applies. The weak implementation remains feasible, and the first two moments alone admit strict implementation because $L_{q,[0,1]}(1/2) = 7/25 < 33/100$. Using $N = 6$ in (F.2) gives a valid polynomial stationary technology for this interval example as well. Here $\mathbb{E}[P(T)] = 0$ places the six-moment vector on the boundary of the unrestricted moment set. The finite-action three-moment target, by contrast, is interior to its unrestricted moment polytope. $\square$

F.2 Shadow price and exact loss

Proof of Theorem 5(i). Partitioning Σ_2 by the value of M_φ gives

$$\Psi_\varphi(\theta) = \max_{y \in I_\varphi} \{Y_\varphi(y) - \theta y\}.$$

Because Y_φ is proper, concave, and upper semicontinuous on the compact interval I_φ , the concave biconjugacy theorem gives (35).

A scalar θ minimizes the right side of (35) at x if and only if

$$Y_\varphi(y) \leq Y_\varphi(x) + \theta(y - x) \quad \text{for every } y \in I_\varphi;$$

that is, if and only if θ is a supergradient of Y_φ at x . The superdifferential is nonempty at every relative-interior point, which proves attainment there.

Suppose $x < \inf I_\varphi^*$, and take $y \in I_\varphi^*$. Then $y > x$ and $Y_\varphi(y) = Y_2 > Y_\varphi(x)$. The supergradient inequality therefore implies $\theta > 0$. The argument to the right of I_φ^* is symmetric and gives $\theta < 0$. If $x \in I_\varphi^*$, then $Y_\varphi(y) \leq Y_\varphi(x)$ for all y , so $\theta = 0$ is a supergradient. At an endpoint of the maximizing interval, the one-sided superdifferential may contain additional nonzero multipliers. $\square$

Proof of Theorem 5(ii). Let $\sigma \in \Sigma_\varphi(x)$, and write

$$r := \mathbb{E}[T \mid Z], \quad d := r - r^*, \quad \mathbf{V}_\sigma := \mathbb{E}[\text{Var}(T \mid Z)].$$

The first-moment constraint gives $\langle \mathbf{1}, d \rangle = 0$. Since the benchmark is deterministic and both assignments have second moment k ,

$$-2 \langle r^*, d \rangle = \|d\|_2^2 + \mathbf{V}_\sigma. \tag{F.3}$$

The KKT identity (36) gives

$$\frac{Y_2 - Y(\sigma)}{\beta} = -\langle w, d \rangle = -2\lambda \langle r^*, d \rangle - \langle n, d \rangle.$$

Substituting (F.3) gives (37).

Because $r \in \mathcal{D}_q$ and $n \in N_{\mathcal{D}_q}(r^*)$, the normal-cone inequality gives $\langle n, d \rangle \leq 0$. Thus all three terms on the right side of the loss identity are nonnegative. If $n = \beta q \xi$, then $\beta q d = g_\sigma - g_{r^*}$, which yields the stated IC form. Since $\lambda > 0$, equality requires $d = 0$ and $\mathbf{V}_\sigma = 0$. Hence $T = r^*(Z)$ almost surely, and the only moment value attained without loss is $x_\varphi^\dagger = \mathbb{E}[\varphi(r^*)]$. $\square$

Declaration of competing interests

The author declares no competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No empirical or experimental data were used.

Acknowledgments

I thank participants in the SWUFE Brownbag Seminar for useful discussions. This work was supported by the National Natural Science Foundation of China (grant number 72503198).

References

- Charalambos D. Aliprantis and Kim C. Border. *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer, 3 edition, 2006. doi: 10.1007/3-540-29587-9.
- Mariagiovanna Baccara, SangMok Lee, and Leeat Yariv. Optimal dynamic matching. *Theoretical Economics*, 15(3):1221–1278, 2020. doi: 10.3982/TE3740.
- Mariagiovanna Baccara, SangMok Lee, and Leeat Yariv. Task allocation and on-the-job training. *Journal of Economic Theory*, 207:105587, 2023. doi: 10.1016/j.jet.2022.105587.
- Ian Ball and Deniz Kattwinkel. Quota mechanisms: Finite-sample optimality and robustness. *Journal of Political Economy*, 2026. doi: 10.1086/742716. In press.
- R. E. Barlow, D. J. Bartholomew, J. M. Bremner, and H. D. Brunk. *Statistical Inference under Order Restrictions: The Theory and Application of Isotonic Regression*. John Wiley & Sons, 1972.
- Dirk Bergemann, Tan Gan, and Yingkai Li. Managing persuasion robustly: The optimality of quota rules. *Journal of Economic Theory*, 236:106222, 2026a. doi: 10.1016/j.jet.2026.106222.
- Dirk Bergemann, Tibor Heumann, and Stephen Morris. Information design and mechanism design: An integrated framework. Discussion Paper 2494, Cowles Foundation for Research in Economics, Yale University, 2026b. URL <https://arxiv.org/abs/2601.17267>.
- Piotr Dworczak and Anton Kolotilin. The persuasion duality. *Theoretical Economics*, 19(4): 1701–1755, 2024. doi: 10.3982/TE5900.
- Juan F. Escobar and Juuso Toikka. Efficiency in games with Markovian private information. *Econometrica*, 81(5):1887–1934, 2013. doi: 10.3982/ECTA9557.
- Alex Frankel. Delegating multiple decisions. *American Economic Journal: Microeconomics*, 8(4): 16–53, 2016. doi: 10.1257/mic.20130270.
- Alexander Frankel. Aligned delegation. *American Economic Review*, 104(1):66–83, 2014. doi: 10.1257/aer.104.1.66.

- Drew Fudenberg and Luis Rayo. Training and effort dynamics in apprenticeship. *American Economic Review*, 109(11):3780–3812, 2019. doi: 10.1257/aer.20171939.
- Matthew O. Jackson and Hugo F. Sonnenschein. Overcoming incentive constraints by linking decisions. *Econometrica*, 75(1):241–257, 2007. doi: 10.1111/j.1468-0262.2007.00737.x.
- Samuel Karlin and William J. Studden. *Tchebycheff Systems: With Applications in Analysis and Statistics*. Interscience Publishers, 1966.
- Andreas Kleiner, Benny Moldovanu, and Philipp Strack. Extreme points and majorization: Economic applications. *Econometrica*, 89(4):1557–1593, 2021. doi: 10.3982/ECTA18312.
- Andreas Kleiner, Benny Moldovanu, and Philipp Strack. Extreme points and majorization. Discussion Paper 2492, Cowles Foundation for Research in Economics, Yale University, 2026. URL <https://cowles.yale.edu/research/cfdp-2492-extreme-points-and-majorization>.
- Jacob D. Leshno. Dynamic matching in overloaded waiting lists. *American Economic Review*, 112(12):3876–3910, 2022. doi: 10.1257/aer.20201111.
- Lasse Leskelä and Matti Vihola. Conditional convex orders and measurable martingale couplings. *Bernoulli*, 23(4A):2784–2807, 2017. doi: 10.3150/16-BEJ827.
- Miltiadis Makris and Alessandro Pavan. Taxation under learning by doing. *Journal of Political Economy*, 129(6):1878–1944, 2021. doi: 10.1086/713745.
- Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002. doi: 10.1111/1468-0262.00296.
- Michael Mussa and Sherwin Rosen. Monopoly and product quality. *Journal of Economic Theory*, 18(2):301–317, 1978. doi: 10.1016/0022-0531(78)90085-6.
- Roger B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981. doi: 10.1287/moor.6.1.58.
- Afshin Nikzad. Constrained majorization: Applications in mechanism design. Working Paper 4030091, SSRN, 2022. Revised October 2023.
- Georg Nöldeke and Larry Samuelson. The implementation duality. *Econometrica*, 86(4):1283–1324, 2018. doi: 10.3982/ECTA13307.
- Alessandro Pavan, Ilya Segal, and Juuso Toikka. Dynamic mechanism design: A Myersonian approach. *Econometrica*, 82(2):601–653, 2014. doi: 10.3982/ECTA10269.
- Volker Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36(2):423–439, 1965. doi: 10.1214/aoms/1177700153.

Juuso Toikka. Ironing without control. *Journal of Economic Theory*, 146(6):2510–2526, 2011. doi: 10.1016/j.jet.2011.06.003.

Kai Hao Yang and Alexander K. Zentefis. Monotone function intervals: Theory and applications. *American Economic Review*, 114(8):2239–2270, 2024. doi: 10.1257/aer.20230330.