

The Cabinet as Yardstick

Comparison, Contest, and Composition

Torun Dewan

Preliminary working note. Comments welcome.

Abstract

A leader judges her ministers not against a fixed standard but against each other. We set signal-jamming career concerns inside a cabinet: each minister exerts hidden effort to lift a noisy signal of his competence, the signals share a common government fortune, and the leader's instrument is coarse – a scarce promotion, not a wage, and in its sharpest form a *succession*, the elevation of one minister to a single great office. Two things follow. Her promotion ranks ministers by the contrast of their performances, which differences the common fortune out, so a minister's career concerns are governed by the idiosyncratic noise that blurs the comparison and not by the shared fortune she filters away: the cabinet is a *yardstick*, the sharper the more its members' fortunes co-move. And because promotion is scarce, effort is a contest – highest for the minister whose rank is contestable, slack for the shoo-in and the no-hoper. Anticipating both, the leader composes her cabinet for comparison. She correlates her ministers' fortunes, which sharpens the yardstick and serves effort and selection alike; but in choosing how level to make the field she must trade them off, for the clustering that keeps the contest alive blunts her choice of whom to raise, while a graded field that elevates the ablest slackens the contest. The optimal cabinet is partly clustered – a graded contest – and it reconciles the single agent's recruitment of extremes with the team's tournament of equals as two limits of one objective. The same comparison that names the abler so cleanly is, by construction, silent after a government-wide reverse, which moves no contrast and so gives the leader no one to dismiss. The model carries the attribution logic of the coalition into the executive, and stands apart from single-agent career concerns (Ashworth, 2005) and continuous-wage comparisons (Meyer and Vickers, 1997) by being at once coarse, scarce, and political.

1 Introduction

A minister works to look able. In the classic account (Holmström, 1999) he does so to lift the market's estimate of his competence, and his reward – a wage – rises smoothly with how able he seems. Politics is coarser. A leader does not pay her ministers by perceived quality; she promotes one and not another, retains or sacks, hands a great office to a few, and at the sharpest anoints a single successor. Ashworth (2005) took the single politician's hidden effort to

such an up-or-down judgement, and a companion to this note (Dewan, 2026a) studies the same coarse career concern for one minister over time. But a leader rarely judges a minister alone. She runs a team, sees them side by side, and the question she asks of a poor month is not only “is he able?” but “is he the one to let go?”. That comparison is the subject here.

Three features make the cabinet its own object. First, ministers share a *common fortune* – the government’s standing rises and falls on all of them together – on top of their own competence and luck. Second, the leader’s controls are *coarse and scarce*: a promotion is a seat there is not enough of, so to elevate one is to pass over another. Third, the leader is herself strategic about who is in the room. Each feature, we show, changes the career concern, and together they make a model distinct from the single agent and from the continuous-wage theory of performance comparison (Meyer and Vickers, 1997).

None of these features is, on its own, ours. That a comparison nets out a common shock is the logic of relative performance evaluation (Holmström, 1982), of the rank-order tournament that filters common uncertainty (Lazear and Rosen, 1981; Green and Stokey, 1983), and of yardstick competition among regulated firms (Shleifer, 1985), whose term we borrow; that performance comparisons carry dynamic career concerns is the subject of Meyer and Vickers (1997). What is new is the instrument. These are theories of the continuous wage – a principal who tunes a smooth schedule of rewards – and a leader has no such dial. She rewards with seats, not salaries: a scarce promotion, given to one minister only by being withheld from another. That coarseness is no simplification but the politics of the thing, and it is what turns the comparison into a pivotal contest and makes the composition of the cabinet, rather than the slope of a pay schedule, the leader’s instrument. Our contribution is to carry the yardstick into that coarse and scarce reward, and to show what a leader does with it.

The coarseness invites a second comparison, now to the theory of contests. There are models of the indivisible prize – the all-pay auction and the Tullock contest (Baye et al., 1996), and in politics the competition for office that breeds a *mediocracy* (Mattozzi and Merlo, 2015) – and these share our scarce, all-or-nothing reward. But they award it on output: the prize goes to the highest bid or the greatest effort. Ours turns on a *posterior*, on who is judged abler once effort is netted away, so that selection and incentive are settled by an inference and not by a race. The cell we occupy is accordingly narrow, a reward at once coarse and posterior-based, and its other occupants are the career-concerns models of electoral accountability – Ashworth (2005), lately Ashworth et al. (2017) and Kartik et al. (2025) – which inherit Fearon’s question of selecting good types against sanctioning poor performance (Fearon, 1999). In every one a single office-holder is kept or replaced against a fixed bar. What none has, and what the cabinet supplies, is a *field*: a scarce place contested by several, a common fortune the comparison filters away, and a bench the leader composes. Scarcity makes the judgement relative, comparison makes it clean, and composition makes it hers.

The argument is short. Because the signals share a common fortune, the leader who wants

to rank her ministers must net it out – and she does so by reading each against the others. Ranking by the contrast of performances differences the common shock away, exactly as the difference of partners’ polls does in a companion theory of coalitions (Dewan, 2026b). So the cabinet becomes a yardstick whose resolution is the idiosyncratic noise, not the shared fortune; a minister’s incentive to exert effort is sharpest where the comparison is sharp, which is where ministerial fortunes co-move. Because the promotion is scarce, the incentive is a tournament: effort tracks the chance that effort changes one’s rank, highest for the minister on the cusp. And because the leader foresees all this, she assembles a cabinet correlated enough to keep the yardstick clean and comparable enough to keep ranks contestable – recruitment and motivation becoming one act of composition. A tension emerges: the promotion that drives the contest is also the one that fills high office, and the two part company, for the very levelling that keeps the contest alive blunts her choice of whom to raise. The cabinet she settles on is a *graded contest*, and at its limits it recovers both parents – the lone agent’s elevation of the clearly ablest and the pure tournament of equals. The result rhymes with the executive-structure and cabinet work this builds on (Dewan and Myatt, 2007), and imports, into the executive, the attribution apparatus developed for parties and coalitions.

2 A cabinet of ministers

A leader (“she”) heads a cabinet of n ministers $i = 1, \dots, n$ (“he”). Minister i enters with a publicly known *reputation* μ_i and an uncertain *competence*

$$\theta_i \sim N(\mu_i, \sigma^2), \quad i = 1, \dots, n, \quad (1)$$

independent across ministers. Neither the leader nor the minister learns the realised θ_i : as in the single-agent career-concerns tradition (Holmström, 1999; Ashworth, 2005) the two share a common prior, and a minister works to look able without knowing how able he is. The reputations μ_i may differ, and that heterogeneity – the standing of the names a leader gathers around her table – is what she will later choose (Section 6); until then it is fixed and common knowledge.

Timing. (i) The reputations $\{\mu_i\}$ are given. (ii) Each minister takes a hidden effort $a_i \geq 0$ at cost $C(a_i)$, with C strictly increasing and convex, $C'(0) = 0$ and $C'(a) \rightarrow \infty$. (iii) Performances realise. (iv) The leader observes the whole vector $y = (y_1, \dots, y_n)$, forms her posteriors, and fills $k < n$ scarce senior places with the k ministers she judges ablest; a promoted minister enjoys $b > 0$, and effort is sunk whether or not he rises.

Minister i ’s performance is the signal

$$y_i = \theta_i + a_i + u + \eta_i, \quad u \sim N(0, \sigma_u^2), \quad \eta_i \sim N(0, \sigma_\eta^2) \text{ i.i.d.}, \quad (2)$$

where u is a shock common to the whole cabinet – the government’s shared fortune – and η_i is the minister’s own luck, independent of u and of competence.¹ The correlation of any two ministers’ performance shocks is

$$\rho = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_\eta^2} \in (0, 1). \quad (3)$$

Strategies and equilibrium. A pure strategy for minister i is an effort $a_i \geq 0$; a strategy for the leader is a promotion set $P(y) \subseteq \{1, \dots, n\}$ with $|P(y)| = k$. We study pure-strategy perfect Bayesian equilibria in which the leader promotes the k ministers of highest posterior mean competence and each minister best-responds, with beliefs consistent: the leader anticipates the equilibrium efforts a_j^* and nets them from the signals she reads. Because effort is anticipated it is absorbed one-for-one and deceives no one – the signal-jamming property – so the leader ranks the *effort-netted* signals $z_i \equiv y_i - a_i^*$. Histories record performances, never actions.

3 A single minister

Before the cabinet, the lone minister – the case to which the rest of the paper is the answer. Set $n = 1$, and let the leader judge one minister against a fixed standard θ^\dagger : the calibre his office demands, or the worth of the replacement she would otherwise turn to. She promotes him when his performance clears the bar. Holding a single signal $y = \theta + a + u + \eta$, her posterior competence is

$$\mathbb{E}[\theta | y] = \mu + \lambda_0 (y - a^* - \mu), \quad \lambda_0 = \frac{\sigma^2}{\sigma^2 + \sigma_u^2 + \sigma_\eta^2}, \quad (4)$$

and she promotes iff it exceeds θ^\dagger . Effort is anticipated, hence signal-jammed: the minister lifts his expected verdict only at the rate λ_0 , and his effort solves

$$\frac{b\sqrt{\lambda_0}}{\sigma} \phi\left(\frac{\theta^\dagger - \mu}{\sqrt{\lambda_0}\sigma}\right) = C'(a^*), \quad (5)$$

where ϕ is the standard normal density. The right-hand side is his marginal cost of effort; the left is what a unit of effort buys – the prize b times the chance the extra effort tips the verdict, which is the density of his posterior at the bar. That density is largest when his reputation sits at the standard, $\mu = \theta^\dagger$, and falls away as μ departs the bar in either direction: it is the apparently average minister, neither safe nor hopeless, who works hardest, for only with him is promotion in the balance. The model and this pivotal logic are due to Ashworth (2005) – a single politician’s hidden effort, judged up or down against a fixed standard, with a unit of effort worth the prize times the probability mass it shifts across the bar. He follows the incumbent

¹That the common shock u cancels exactly from every comparison (Lemma 1) relies on its loading on each signal being identical. With heterogeneous loadings, $y_i = \theta_i + a_i + \xi_i u + \eta_i$, the differencing nets out only the shared component and a residual $(\xi_i - \xi_j)u$ survives; the yardstick is then sharp in proportion to how evenly the government’s fortune is borne. Equal loadings are the clean benchmark.

already past that standard, however – the reelected, whose expected ability clears it – and on that range reports the action as simply decreasing in reputation. The two-sided shape, with effort peaking at the bar so that it is the middling minister who tries hardest, is present in his density but is not a result he draws; we state it because it is the single-agent shadow of the contest to come.² Either way a minister’s incentive here turns on his standing relative to a fixed number, and on no one else.

Two things change when the leader judges a team, and the rest of the paper turns on them. The standard ceases to be a fixed external number and becomes the field: a minister is promoted not when he clears θ^\dagger but when he outscores his rivals for a scarce seat, so the bar is endogenous and moves with the company he keeps (Section 5). And the lone verdict (4) is contaminated by the whole variance $\sigma^2 + \sigma_u^2 + \sigma_\eta^2$ – the government’s shared fortune u muddies it along with the minister’s own luck – whereas reading the team against itself strips the common σ_u^2 away, leaving the sharper weight $\lambda > \lambda_0$ (Section 4). Absolute judgement against a bar becomes relative judgement against a field: that is the step from the single minister to the cabinet.

4 The yardstick

The leader cannot separate a minister’s competence from the government’s luck out of his signal alone; the common u sits in every y_i . But she holds the whole vector, and the vector lets her compare. Joint normality makes the posterior linear, and the comparison explicit.

Lemma 1 (Posterior competence and the yardstick). *Write $z_i = y_i - a_i^*$ and $\lambda = \sigma^2 / (\sigma^2 + \sigma_\eta^2) \in (0, 1)$. The leader’s posterior mean competence is*

$$\mathbb{E}[\theta_i | y] = \mu_i + \lambda \left[(z_i - \mu_i) - \beta \sum_{j=1}^n (z_j - \mu_j) \right], \quad \beta = \frac{\sigma_u^2}{\sigma^2 + \sigma_\eta^2 + n \sigma_u^2}. \quad (6)$$

The bracketed common-sum term is identical across ministers, so the leader’s ranking by posterior competence coincides with the ranking by the score

$$s_i = (1 - \lambda) \mu_i + \lambda z_i, \quad i = 1, \dots, n. \quad (7)$$

For any pair, $s_i - s_j = (1 - \lambda)(\mu_i - \mu_j) + \lambda[(\theta_i - \theta_j) + (\eta_i - \eta_j)]$, free of the common fortune u for every realisation and every n . The noise that blurs a comparison is thus $\lambda(\eta_i - \eta_j)$, with variance $2\lambda^2\sigma_\eta^2$, governed by the idiosyncratic σ_η and never by the common σ_u .

The common shock enters every posterior with the same weight and so cancels from every comparison; what orders the ministers is the part of performance that is theirs, weighed against their colleagues. The leader judges by comparison not because she elects to but because comparison

²This two-sided reading, that effort peaks at the bar so that the apparently average type works hardest, was first stated in Dewan and Myatt (2012), a now defunct note from which this paper descends.

is what the signal structure leaves her: holding the vector, she filters the government’s fortune out whatever its size.

Proposition 1 (The cabinet is a yardstick). *Reading the team, the leader nets the common fortune from her verdict automatically: by Lemma 1 promotion turns on the score (7), whose comparative noise is set by σ_η alone. A minister’s career concern is therefore immune to the shared volatility σ_u – however turbulent the government’s fortunes, they do not cloud who is judged abler. The contrast with isolated judgement is sharpest when fortunes co-move: were the leader forced to read each minister against a fixed bar, as in the single-agent coarse career concern of Ashworth (2005), the common shock would contaminate her reading with variance σ_u^2 , the more so as $\rho \rightarrow 1$; the cabinet removes exactly this term. As $\rho \rightarrow 0$ there is no common fortune to net out, comparison buys nothing, and the team reading collapses to the lone-agent one. Where the lone-agent account has shared volatility dull career concerns by adding noise, the cabinet is insulated from it – and the gain from comparison rises with the very co-movement that hurts the isolated judge.*

This is the attribution result of the coalition, relocated, and at bottom the relative-performance sufficient-statistic logic of Holmström (1982): what is new is not the differencing-out of the common shock, which is standard, but that it operates within a reward that is coarse, scarce, and posterior-based. There the difference of partners’ polls names a defector by cancelling the national tide; here the difference of ministers’ performances names the abler by cancelling the government’s fortune. The same ρ governs the value of the comparison in both, and the same moral holds: what a principal can see is set by how much of her agents’ fortunes is shared.

5 The contest

Because the places are scarce, the career concern is a tournament. A unit of effort lifts a minister’s score s_i by λ ; what it buys is the chance that the lift carries him across the k -th place in the field.

Proposition 2 (Contestable ranks try hardest). *In equilibrium each minister’s effort satisfies the first-order condition*

$$b \lambda p_i = C'(a_i^*), \tag{8}$$

where p_i is the density of the event that minister i ’s score s_i equals the marginal (k -th highest) score in the field – the probability that a small change in his performance changes his promotion. Effort is greatest for the minister whose standing is contestable, near the cutoff, and slack for the shoo-in and the no-hoper, for whom no effort moves the verdict. The lone-agent finding that the average minister tries hardest becomes, in the team, that the marginal minister does – where the margin is set by the field. Effort also scales with the yardstick’s resolution λ : the sharper the comparison (Proposition 1), the harder the contestable minister works.

Derivation. Suppose minister i takes effort a_i while the leader nets the conjectured a_i^* . By (7) his score is $s_i = \mu_i + \lambda(a_i - a_i^*) + \lambda u + w_i$, where $w_i = \lambda(\theta_i - \mu_i) + \lambda\eta_i$ collects his idiosyncratic noise. Since $\theta_i + \eta_i \sim N(\mu_i, \sigma^2/\lambda)$ we have $w_i \sim N(0, v)$ with

$$v = \lambda^2(\sigma^2 + \sigma_\eta^2) = \lambda\sigma^2, \quad (9)$$

so at equilibrium the *de-fortune* score $\tilde{s}_i \equiv s_i - \lambda u$ is $N(\mu_i, v)$, centred at reputation with common spread. The common term λu is shared by all scores and cancels from the ranking, as in Lemma 1, so the contest runs on the \tilde{s}_i .

Minister i wins a place iff \tilde{s}_i exceeds the k -th highest rival score $c_i = \text{top}_k\{\tilde{s}_j\}_{j \neq i}$. Writing $D_i = c_i - w_i$ with density ϕ_i , his promotion probability is $\Pi_i(a_i) = \Pr(\mu_i + \lambda(a_i - a_i^*) > D_i) = \Phi_i(\mu_i + \lambda(a_i - a_i^*))$, so $\Pi'_i(a_i^*) = \lambda\phi_i(\mu_i)$. The minister maximises $b\Pi_i(a_i) - C(a_i)$; with $p_i \equiv \phi_i(\mu_i)$ this is the first-order condition (8). The factor λ is the signal-jamming damping: effort lifts the raw signal one-for-one, but the leader's score only by λ , because she shrinks performance toward the prior. Convexity of C makes effort increasing in the pivotal density p_i .

For the leading case of a single succession (n ministers, $k = 1$) with two contenders, $c_i = \tilde{s}_j$ and $D_i = \mu_j + w_j - w_i \sim N(\mu_j, 2v)$, so the pivotal density is closed-form:

$$p_i = \frac{1}{2\sqrt{\pi v}} \exp\left(-\frac{(\mu_i - \mu_j)^2}{4v}\right), \quad v = \lambda\sigma^2. \quad (10)$$

It depends on the *reputation gap* $|\mu_i - \mu_j|$ alone, is greatest at $\mu_i = \mu_j$, and falls away as the gap widens – the shoo-in and the no-hoper. Because (10) is symmetric in the gap, evenly-matched ministers face equal pivotal densities and so, by (8), exert equal effort: a contest of equals. This two-contender characterisation is exact. For a general field (n, k) the pivotal density p_i is that of an order statistic of the rivals' scores; it remains single-peaked in μ_i when that density is log-concave, but we do not characterise the multi-contender contest in full and state the result for the leading two-contender succession. Finally, since $\lambda p_i = \lambda/(2\sqrt{\pi\lambda\sigma^2}) e^{-\cdot} \propto \sqrt{\lambda}$, the marginal incentive rises with the yardstick's resolution λ : a cleaner comparison makes the contestable minister work harder, tying Proposition 2 to Proposition 1. \square

Two forces thus set effort, and both run through comparison: the *resolution* of the yardstick (Proposition 1), carried by λ , and the *contestability* of a minister's rank within the field, carried by p_i . The leader, we now let choose the field, controls both.

6 Composition for comparison

Foreseeing Propositions 1 and 2, the leader's recruitment is no longer a list of separate hires but the design of a field.

Proposition 3 (The effort-maximising field). *The cabinet that maximises induced effort is*

clustered and correlated: the leader gathers ministers of comparable perceived competence, so that ranks are contestable and pivotal densities high, and ministers whose fortunes co-move, so that the yardstick is clean. This is the contest motive acting alone: were induced effort her only concern she would gather exactly the comparable in-between, assembled into a contest, and shun the clear winners and losers. It is to be set against the selection motive, which wants the field spread so that the ablest is surely raised. Both motives are present already for the single agent – whose effort, too, peaks at the bar – so the tension is between effort and selection, not between one agent and a team; the next subsection resolves it.

Derivation. Let the leader value the effort her cabinet supplies and choose the field to maximise aggregate induced effort $\sum_i a_i^*$, with $a_i^* = (C')^{-1}(b\lambda p_i)$ increasing in p_i by (8). Two levers follow.

Clustering. By (10) – and, in general, by the single-peakedness of p_i in the reputation gap – pivotal densities are higher the closer ministers’ reputations lie. Bringing the field together raises every contender’s p_i toward its peak at $\mu_i = \mu_j$, and with it every a_i^* ; dispersing reputations into clear winners and losers drives pivotal densities, and effort, to zero. The effort-maximising field is therefore clustered around a common standing – a bench of comparable rivals – rather than a spread of sure things and lost causes.

Correlating. Hold the total performance-shock variance fixed, $\sigma_u^2 + \sigma_\eta^2 = \tau^2$, and let the leader choose how much of it is shared by assigning ministers to common or to separate fortunes. Then $\sigma_\eta^2 = (1 - \rho)\tau^2$ and

$$\lambda = \frac{\sigma^2}{\sigma^2 + (1 - \rho)\tau^2} \quad \text{is increasing in } \rho. \quad (11)$$

Since the marginal incentive $\lambda p_i \propto \sqrt{\lambda}$ rises with λ , the leader prefers fortunes as correlated as possible: co-movement converts idiosyncratic noise – the only kind that blurs the comparison – into common noise she filters out, sharpening the yardstick and lifting effort. This is the precise content of “correlate to keep the yardstick clean,” and it is why a leader who judges by comparison wants her ministers exposed to a shared fate. \square

A leader who judges in isolation seeks types she can be sure of; a leader who judges by comparison seeks types she can compare. But effort is not all she wants from a cabinet, and the clustering that maximises it is not, in general, what she chooses.

The leader’s objective: effort against selection

A scarce promotion does two things at once. It is the prize that drives the contest, and it is the decision that fills a senior office, where competence tells most. We make the leader’s design problem explicit for the leading two-contender succession ($k = 1$). Her pool fixes the field’s mean standing m ; what she composes is its *spread* $\Delta = \mu_1 - \mu_2 \in [0, \bar{\Delta}]$, where $\bar{\Delta}$ is the dispersion the available names allow. (A general field is composed the same way, by the contestability and

selection forces below, but without a single scalar knob.) She values the effort her field induces and the calibre of whom she promotes, the latter through an increasing value g :

$$V(\Delta) = \sum_i a_i^*(\Delta) + \gamma g(\mathbb{E}[\theta_P]), \quad \gamma > 0, \quad (12)$$

with g the identity in the linear benchmark and *concave* when an office's worth has diminishing returns to its holder's calibre. The curvature of g is a property of the office: a singular, high-stakes post whose value keeps rising in calibre is near-linear, an ordinary portfolio for which a competent-enough minister suffices is concave. The two terms pull against each other, and the spread Δ trades them off.

Lemma 2 (The selection premium rises with the gap). *In the succession ($k = 1$) between two contenders with reputations $\mu_1 \geq \mu_2$, mean $m = \frac{1}{2}(\mu_1 + \mu_2)$ and gap $\Delta = \mu_1 - \mu_2$, the expected competence promoted is*

$$\mathbb{E}[\theta_P] = m + \underbrace{\Delta \left(\Phi(\delta) - \frac{1}{2} \right) + \sqrt{2v} \phi(\delta)}_{SP(\Delta)}, \quad \delta = \frac{\Delta}{\sqrt{2v}}, \quad v = \lambda \sigma^2, \quad (13)$$

and the selection premium is increasing in the gap, $SP'(\Delta) = \Phi(\delta) - \frac{1}{2} \geq 0$, from $SP(0) = \sqrt{v/\pi}$ (a level field still promotes the abler, on the strength of performance) to $SP(\Delta) \rightarrow \Delta/2$ (a graded field simply elevates the star).

The two-contender closed form is the tractable face of a general fact: for any n and k the expected competence promoted equals the expected sum of the k highest scores in the field, proved as (15) in Appendix A; selection is governed by the order statistics of the de-fortune scores whatever the size of the cabinet.

So the spread is the knob of the trade-off. Effort wants it closed: by (10) the pivotal density, and with it $\sum_i a_i^*$, is greatest at $\Delta = 0$ and decays as the field grades, its slope vanishing as the field spreads. Selection wants it open: by Lemma 2 the competence promoted climbs with Δ , and its marginal value $SP'(\Delta) = \Phi(\delta) - \frac{1}{2}$ rises toward $\frac{1}{2}$. Correlation, by contrast, is no trade-off at all: raising ρ (hence λ and $v = \lambda \sigma^2$) lifts both effort, via $\lambda p_i \propto \sqrt{\lambda}$, and selection, via $SP(0) = \sqrt{v/\pi}$. A clean yardstick is wanted on every count; only the levelness of the field is contested.

Proposition 4 (The optimal field: a graded contest, and its all-or-nothing limit). *The leader sets the field's mean as high as her pool allows and chooses the spread $\Delta^* \in [0, \bar{\Delta}]$.*

- (i) (Diminishing returns: a graded contest.) *When the value of office-holder calibre is concave – an ordinary portfolio, for which a competent-enough minister suffices – g' tapers as calibre rises, so the selection gain $\gamma g'(\mathbb{E}[\theta_P]) (\Phi(\delta) - \frac{1}{2})$ is bounded. Once the selection motive is active – γ above a threshold γ_0 , below which the contest is kept level ($\Delta^* = 0$) –*

and office-value satiates enough that grading the whole field does not pay, it balances the effort loss at an interior $\Delta^* \in (0, \bar{\Delta})$ satisfying

$$2 \underbrace{\frac{\partial}{\partial \Delta} (C')^{-1}(b\lambda p(\Delta))}_{\text{effort lost as the field grades } (<0)} + \gamma g'(\mathbb{E}[\theta_P]) (\Phi(\delta^*) - \frac{1}{2}) = 0. \quad (14)$$

The optimal cabinet is partly clustered – levelled enough to keep the contest alive, graded enough to raise able ministers – and Δ^* widens as selection matters more (larger γ) and narrows as the prize sharpens the contest (larger b).

- (ii) (The singular office: bang-bang.) When calibre does not satiate – a singular, high-stakes post whose value rises without tapering, g near the identity – the effort slope vanishes as the field grades while the selection slope rises toward $\frac{1}{2}$, so the interior stationary point of (14) is a trough. The optimum is then a corner – a level tournament of equals when the prize dominates, the maximally-graded elevation of a lone heir when calibre dominates – switching discontinuously at a threshold γ^* where $V(0) = V(\bar{\Delta})$.

A compact choice set guarantees a maximiser in either case.

Either way the team is reconciled with its parents. The single-agent recruitment rule – hire the clearly able or the clearly mediocre, shun the middle – is the selection motive acting alone; the tournament of equals is the effort motive acting alone; and the cabinet a leader actually builds is a level tournament, a graded elevation, or – when an office’s worth tapers in its holder’s calibre – the graded contest between them. What was a puzzle in isolation – why elevate the in-between at all? – is the price of keeping the contest live, paid up to the point where the forgone calibre of office is no longer worth the effort it buys. The curvature of office-value thus maps to the shape of the field: the ordinary run of portfolios, where ability has diminishing returns, are surrounded by a graded contest, while the most consequential, singular offices – where calibre never satiates – are filled by an all-or-nothing field, a flat slate of equals or a lone clear heir. What a leader builds around an office is read off how that office values calibre.

The yardstick read from below

We have read the cabinet from above, where the comparison elevates the abler; the same comparison governs dismissal, and there it falls silent in a particular way. Decompose what the leader’s signals let her know.

Lemma 3 (What the cabinet identifies). Write $\bar{z} = \frac{1}{n} \sum_j z_j$, with $\bar{\theta}, \bar{\mu}, \bar{\eta}$ the corresponding averages. (i) Relative competence is clean. For any pair, $z_i - z_j = (\theta_i - \theta_j) + (\eta_i - \eta_j)$ carries no common fortune, so the posterior over the relative competence $\theta_i - \theta_j$ is free of u for every realisation and every n , its only noise the idiosyncratic η . A purely common disturbance – u

moving with the η_i held fixed – shifts no contrast and leaves every relative posterior unchanged.

(ii) Collective competence is confounded. *The average $\bar{z} = \bar{\theta} + u + \bar{\eta}$ places collective competence $\bar{\theta}$ and the common fortune u in a single term; from one cross-section the leader identifies only their sum. Since $\text{Var}(\bar{\theta} - \bar{\mu}) = \sigma^2/n$ and $\text{Var}(\bar{\eta}) = \sigma_\eta^2/n$ both vanish as the cabinet grows, $\bar{z} - \bar{\mu} \rightarrow u$: a large cabinet recovers the fortune, and with it the absolute level, while a small one cannot tell an able government from a lucky one. At every n the collective is identified less precisely than the relative, the shortfall being the common-fortune variance σ_u^2 .*

The cabinet is a sharp instrument for one question and a blunt one for the other. It tells the leader who among her ministers is abler with a clarity the common fortune cannot cloud; it tells her whether her government is able or merely fortunate only as the bench grows large enough to average the fortune away.

Corollary 1 (Silence on a common reverse). *A disturbance that is purely common moves no contrast. The comparative apparatus therefore yields neither a reason to dismiss – it has filtered the common level away – nor a basis for choosing whom – no contrast has moved. The yardstick names a minister to sack only in so far as his misfortune was his own.*

So the asymmetry is exact. Promotion rides on contrasts the common fortune leaves clean, and co-movement sharpens it; dismissal for cause, after a government-wide reverse, has no contrast to ride on at all. The instrument that names the abler with such clarity cannot, when the whole government stumbles, name anyone, and that silence is the price of the very differencing that keeps the promotion clean.

7 Concluding remarks

A cabinet is a comparison. Once a leader judges ministers against one another, the common fortune they share drops out of her verdict and the idiosyncratic noise that separates them rules it; the scarce promotion makes effort a contest won at the margin; and the leader, foreseeing both, builds a correlated team so that to recruit is already to motivate. But the same scarce promotion that drives the contest is the act that fills high office, and there the two motives part: the clustering that keeps the contest live blunts her choice of whom to raise. The cabinet she settles on is a graded contest – partly levelled, partly ranked – and in its two limits it recovers both parents, the single agent’s elevation of the clearly ablest and the pure tournament of equals. Its companion in time turns the lesson over. Where this contest makes the middling minister its centre – the clearly ablest a shoo-in who exerts least, the leader’s lever the comparable field she gathers around him – the dynamic problem of which talent to spend, and when, makes the exceptional minister the one to withhold, husbanded in reserve as the threat that keeps a lesser incumbent honest (Dewan, 2026a). Mediocrity is where the contest lives; talent is what the schedule conserves. The lone agent of Ashworth (2005) and the smooth wage of Holmström

(1999) cannot show any of this, and the continuous-wage yardstick of Meyer and Vickers (1997) cannot show its coarse, scarce, political form. What it does show is that the apparatus a companion paper builds to explain when a coalition can see a defector (Dewan, 2026b) explains, unchanged, when a leader can see an able minister: the same differencing of a common shock, read once for blame in the coalition and here for competence in the cabinet, so that the executive becomes one more theatre of attribution.

This is a small step, and a first-shot: one cabinet, judged at a single date, with the bench taken as given before the leader composes it. What it offers is a way to see a promotion not as a prize but as a verdict, and the cabinet a leader gathers as the instrument that makes the verdict hers.

A Proofs

On the equilibrium path each minister exerts the conjectured effort, so the leader's effort-netted signal is $z_i = y_i - a_i^* = \theta_i + u + \eta_i$ and $(\theta_1, \dots, \theta_n, z_1, \dots, z_n)$ is jointly normal.

Proof of Lemma 1. Write $A = \sigma^2 + \sigma_\eta^2$ and let $J = \mathbf{1}\mathbf{1}^\top$ be the $n \times n$ matrix of ones. From $z_j = \theta_j + u + \eta_j$,

$$\mathbb{E}[z_j] = \mu_j, \quad \text{Cov}(\theta_i, z_j) = \sigma^2 \delta_{ij}, \quad \text{Cov}(z_j, z_l) = A \delta_{jl} + \sigma_u^2,$$

so the covariance matrix of z is $\Sigma = AI + \sigma_u^2 J$ and the cross-covariance of θ_i with z is $\sigma^2 e_i$. By the Sherman–Morrison formula, $\Sigma^{-1} = A^{-1}(I - \frac{\sigma_u^2}{A+n\sigma_u^2} J)$, so the Gaussian conditional mean is

$$\mathbb{E}[\theta_i | z] = \mu_i + \sigma^2 e_i^\top \Sigma^{-1}(z - \mu) = \mu_i + \frac{\sigma^2}{A} \left[(z_i - \mu_i) - \frac{\sigma_u^2}{A + n\sigma_u^2} \sum_j (z_j - \mu_j) \right],$$

which is (6) with $\lambda = \sigma^2/A$ and $\beta = \sigma_u^2/(A + n\sigma_u^2)$. The bracketed sum is common to all i , so ranking by $\mathbb{E}[\theta_i | z]$ ranks by $s_i = (1 - \lambda)\mu_i + \lambda z_i$. Finally $z_i - z_j = (\theta_i - \theta_j) + (\eta_i - \eta_j)$ is free of u , so $s_i - s_j = (1 - \lambda)(\mu_i - \mu_j) + \lambda[(\theta_i - \theta_j) + (\eta_i - \eta_j)]$ with comparison noise $\lambda(\eta_i - \eta_j)$ of variance $2\lambda^2\sigma_\eta^2$. \square

Proof of Proposition 1. By Lemma 1 the ranking turns on $\{s_i\}$, whose pairwise differences are free of u and carry noise of variance $2\lambda^2\sigma_\eta^2$, a function of σ_η alone. Reading a minister against a fixed bar would instead expose the estimate to the whole shock $u + \eta_i$, of variance $\sigma_u^2 + \sigma_\eta^2$; the comparison removes the term σ_u^2 , whose share $\rho = \sigma_u^2/(\sigma_u^2 + \sigma_\eta^2)$ rises to one as fortunes co-move and falls to zero as they separate, where the team reading and the lone-agent reading coincide. \square

Proof of Proposition 2. Fix the conjectured profile and let i deviate to a_i . By (7) and $z_i = \theta_i + (a_i - a_i^*) + u + \eta_i$ his score is $s_i = \mu_i + \lambda(a_i - a_i^*) + \lambda u + w_i$, where $w_i = \lambda[(\theta_i - \mu_i) + \eta_i] \sim N(0, v)$

and $v = \lambda^2(\sigma^2 + \sigma_\eta^2) = \lambda\sigma^2$. The common λu cancels from every comparison, so i is promoted iff his de-fortune score $\tilde{s}_i = \mu_i + \lambda(a_i - a_i^*) + w_i$ exceeds the k -th highest rival score c_i . Writing $D_i = c_i - w_i$ with density ϕ_i , the promotion probability is $\Pi_i(a_i) = \Phi_i(\mu_i + \lambda(a_i - a_i^*))$, so $\Pi_i'(a_i^*) = \lambda\phi_i(\mu_i) = \lambda p_i$. Maximising $b\Pi_i(a_i) - C(a_i)$ yields $b\lambda p_i = C'(a_i^*)$, equation (8); convexity of C makes this a maximum with a_i^* increasing in p_i . For $k = 1$ with two in contention, $c_i = \tilde{s}_j$ and $D_i = \mu_j + w_j - w_i \sim N(\mu_j, 2v)$, so $p_i = \phi_i(\mu_i) = \frac{1}{2\sqrt{\pi v}} \exp(-(\mu_i - \mu_j)^2/4v)$ as in (10), single-peaked in the gap and maximal at $\mu_i = \mu_j$; and $\lambda p_i \propto \sqrt{\lambda}$ rises with the yardstick's resolution. \square

Proof of Proposition 3. By (8), $a_i^* = (C')^{-1}(b\lambda p_i)$ is increasing in p_i . *Clustering:* p_i is single-peaked in the reputation gap, maximal at coincident reputations (explicitly (10) for the leading case), so $\sum_i a_i^*$ rises as reputations are brought together and falls to zero as they disperse into sure winners and losers. *Correlating:* holding $\sigma_u^2 + \sigma_\eta^2 = \tau^2$ fixed and writing $\sigma_\eta^2 = (1 - \rho)\tau^2$, equation (11) gives $\lambda = \sigma^2/(\sigma^2 + (1 - \rho)\tau^2)$ increasing in ρ ; since the marginal incentive $b\lambda p_i$ increases in λ (indeed $\lambda p_i \propto \sqrt{\lambda}$), aggregate effort rises with ρ . The effort-maximising field is therefore clustered and maximally correlated. \square

Proof of Lemma 2. Write each de-fortune score as $\tilde{s}_i = \mu_i + w_i$, $w_i \sim N(0, v)$ independent across i , $v = \lambda\sigma^2$, and recall that P collects the k highest scores (the common λu does not affect the ranking). For any n and k ,

$$\mathbb{E}\left[\sum_{i \in P} \theta_i\right] = \mathbb{E}[\text{sum of the } k \text{ largest of } \tilde{s}_1, \dots, \tilde{s}_n]. \quad (15)$$

Indeed $\mathbf{1}\{i \in P\}$ is a function of the scores, so by the tower property and (6), $\mathbb{E}[\sum_{i \in P} \theta_i] = \mathbb{E}[\sum_{i \in P} \mathbb{E}[\theta_i | z]] = \mathbb{E}[\sum_{i \in P} (s_i - T)]$, where $T = \lambda\beta \sum_j (z_j - \mu_j)$ is common to every minister. Since $|P| = k$ we have $\sum_{i \in P} T = kT$ and $\mathbb{E}[T] = 0$; and $s_i = \tilde{s}_i + \lambda u$ with $\sum_{i \in P} \lambda u = k\lambda u$ and $\mathbb{E}[u] = 0$. Both common terms vanish in expectation, leaving (15). For $k = 1$ and two contenders this reads $\mathbb{E}[\theta_P] = \mathbb{E}[\max(\tilde{s}_1, \tilde{s}_2)]$ with $\tilde{s}_i \sim N(\mu_i, v)$ independent. With $\Delta = \mu_1 - \mu_2$ and $\delta = \Delta/\sqrt{2v}$, the expected maximum of two independent normals is

$$\mathbb{E}[\max(\tilde{s}_1, \tilde{s}_2)] = \mu_1\Phi(\delta) + \mu_2\Phi(-\delta) + \sqrt{2v}\phi(\delta) = m + \Delta(\Phi(\delta) - \frac{1}{2}) + \sqrt{2v}\phi(\delta),$$

substituting $\mu_1 = m + \Delta/2$, $\mu_2 = m - \Delta/2$; this is (13). Using $\phi'(\delta) = -\delta\phi(\delta)$, $SP'(\Delta) = (\Phi(\delta) - \frac{1}{2}) + \delta\phi(\delta) - \delta\phi(\delta) = \Phi(\delta) - \frac{1}{2} \geq 0$, and $SP(0) = \sqrt{2v}\phi(0) = \sqrt{v/\pi}$, while $SP(\Delta) \rightarrow \Delta/2$ as $\Delta \rightarrow \infty$. \square

Proof of Proposition 4. For $k = 1$ and two contenders both face the pivotal density $p(\Delta) = \frac{1}{2\sqrt{\pi v}} \exp(-\Delta^2/4v)$ of (10), so aggregate effort is $E(\Delta) \equiv 2(C')^{-1}(b\lambda p(\Delta))$ and, by Lemma 2, $\mathbb{E}[\theta_P] = m + SP(\Delta)$. Thus

$$V(\Delta) = E(\Delta) + \gamma g(m + SP(\Delta)), \quad V'(\Delta) = E'(\Delta) + \gamma g'(m + SP(\Delta)) SP'(\Delta).$$

Since $\partial V/\partial m = \gamma g'(\cdot) > 0$ the leader raises m to the frontier her pool allows. For the spread, record four facts from the closed forms. (a) $p'(\Delta) = -(\Delta/2v)p(\Delta)$, so $p'(0) = 0$, $p'(\Delta) < 0$ for $\Delta > 0$, and $p'(\Delta) \rightarrow 0$ as $\Delta \rightarrow \infty$; with $(C')^{-1}$ increasing this gives $E'(0) = 0$, $E'(\Delta) < 0$ on $(0, \bar{\Delta}]$, and $E'(\Delta) \rightarrow 0$ as the field grades. (b) $SP'(\Delta) = \Phi(\delta) - \frac{1}{2}$ with $SP'(0) = 0$ and $SP'(\Delta) \uparrow \frac{1}{2}$; hence $V'(0) = 0$ – a level field is always stationary. (c) The curvature there is $V''(0) = E''(0) + \gamma g'(m + SP(0)) SP''(0)$, with $E''(0) = 2(C')^{-1'}(b\lambda p(0)) b\lambda p''(0) < 0$ (since $p''(0) = -p(0)/2v < 0$) and $SP''(0) = \frac{1}{2\sqrt{\pi v}} > 0$, so $V''(0) > 0$ exactly when $\gamma > \gamma_0 := |E''(0)|/(g'(m + SP(0)) SP''(0))$: the level field is a local minimum when selection is active and a local maximum when it is not. (d) E and $g \circ SP$ are continuous on the compact $[0, \bar{\Delta}]$, so a maximiser exists.

Case (i): g concave. If $\gamma \leq \gamma_0$ the selection weight is too small to repay any grading – $V'(0) = 0$, $V''(0) \leq 0$ – and the leader keeps the contest level, $\Delta^* = 0$. If $\gamma > \gamma_0$ then $V''(0) > 0$, so $\Delta = 0$ is not a maximiser and V rises as the field first grades. Because office-value satiates, $g'(m + SP(\Delta))$ falls as Δ grows, so the selection gain $\gamma g'(m + SP(\Delta)) SP'(\Delta)$ is bounded and, at a sufficiently graded field, no longer outweighs the persistent effort loss $E'(\Delta) < 0$; taking the field's range wide enough that $V'(\bar{\Delta}) < 0$, continuity of V' gives an interior maximiser $\Delta^* \in (0, \bar{\Delta})$ at which $V'(\Delta^*) = 0$, which is (14). Differentiating that identity, the selection term rises in γ while the effort term steepens in b , so Δ^* increases in γ and decreases in b , with $\Delta^* \downarrow 0$ as $\gamma \downarrow \gamma_0$.

Case (ii): the singular office, g linear ($g' \equiv 1$, calibre never satiating). Now $V'(\Delta) = E'(\Delta) + \gamma SP'(\Delta)$. As the field grades $E'(\Delta) \rightarrow 0$ while $\gamma SP'(\Delta) \rightarrow \gamma/2 > 0$, so $V'(\bar{\Delta}) > 0$: V is increasing at the top and the upper corner beats every neighbouring field. With $V'(0) = 0$ and $V'(\bar{\Delta}) > 0$, any interior stationary point is a local minimum – a trough – so the maximiser is a corner, $\Delta^* \in \{0, \bar{\Delta}\}$. Comparing them, $V(0) - V(\bar{\Delta}) = [E(0) - E(\bar{\Delta})] - \gamma[SP(\bar{\Delta}) - SP(0)]$, with both bracketed differences strictly positive; this is strictly decreasing in γ and vanishes at $\gamma^* = [E(0) - E(\bar{\Delta})]/[SP(\bar{\Delta}) - SP(0)]$. Hence $\Delta^* = 0$ – a level tournament of equals – for $\gamma < \gamma^*$, and $\Delta^* = \bar{\Delta}$ – the maximally-graded elevation of a lone heir – for $\gamma > \gamma^*$. The same corner conclusion holds for any office bounded away from satiation, $g' \geq g_{\min} > 0$, with $\gamma/2$ replaced by $g_{\min}\gamma/2$. \square

Proof of Lemma 3 and Corollary 1. For any pair, $z_i - z_j = (\theta_i - \theta_j) + (\eta_i - \eta_j)$ carries no u , so the posterior over the relative competence $\theta_i - \theta_j$ is free of the common fortune for every realisation and every n , its only noise the idiosyncratic η . For the average, $\bar{z} = \bar{\theta} + u + \bar{\eta}$ identifies only the sum $\bar{\theta} + u$; with $\text{Var}(\bar{\theta} - \bar{\mu}) = \sigma^2/n$ and $\text{Var}(\bar{\eta}) = \sigma_\eta^2/n$, both tending to zero, $\bar{z} - \bar{\mu} \rightarrow u$ as $n \rightarrow \infty$, so a large bench recovers the fortune and a small one cannot. A purely common disturbance changes every z_i by the same amount, hence moves no contrast and leaves every relative posterior – and so the entire promotion and dismissal ranking – unchanged, which is Corollary 1. \square

References

- Ashworth, S. (2005). Reputational Dynamics and Political Careers. *Journal of Law, Economics, and Organization* 21(2), 441–466.
- Ashworth, S., E. Bueno de Mesquita, and A. Friedenberg (2017). Accountability and Information in Elections. *American Economic Journal: Microeconomics* 9(2), 95–138.
- Baye, M. R., D. Kovenock, and C. G. de Vries (1996). The All-Pay Auction with Complete Information. *Economic Theory* 8(2), 291–305.
- Dewan, T. (2026a). Talent in Reserve. Working note.
- Dewan, T. (2026b). What a Coalition Can See. Working note.
- Dewan, T., and D. P. Myatt (2007). Scandal, Protection, and Recovery in the Cabinet. *American Political Science Review* 101(1), 63–77.
- Dewan, T., and D. P. Myatt (2012). Political Career Concerns. Unpublished note.
- Fearon, J. D. (1999). Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance. In A. Przeworski, S. C. Stokes, and B. Manin (Eds.), *Democracy, Accountability, and Representation*, pp. 55–97. Cambridge University Press.
- Green, J. R., and N. L. Stokey (1983). A Comparison of Tournaments and Contracts. *Journal of Political Economy* 91(3), 349–364.
- Holmström, B. (1982). Moral Hazard in Teams. *Bell Journal of Economics* 13(2), 324–340.
- Holmström, B. (1999). Managerial Incentive Problems: A Dynamic Perspective. *Review of Economic Studies* 66(1), 169–182.
- Kartik, N., E. Lipnowski, and D. Pei (2025). Replacement and Reputation. Working paper.
- Lazear, E. P., and S. Rosen (1981). Rank-Order Tournaments as Optimum Labor Contracts. *Journal of Political Economy* 89(5), 841–864.
- Mattozzi, A., and A. Merlo (2015). Mediocracy. *Journal of Public Economics* 130, 32–44.
- Meyer, M. A., and J. Vickers (1997). Performance Comparisons and Dynamic Incentives. *Journal of Political Economy* 105(3), 547–581.
- Shleifer, A. (1985). A Theory of Yardstick Competition. *RAND Journal of Economics* 16(3), 319–327.