Designing Promotion and Hiring Procedures with Biased Evaluators

Christopher Avery and Margaret Meyer

Harvard and Oxford

May 29, 2009

Introduction (1)

Recommendations play an important role in hiring and promotion decisions in a variety of organizations and networks:

- civil service; university faculties; judiciary; professional service firms; military; "old boy networks"

Why? The objective (verifiable) information about past performance that is available may be of limited usefulness for hiring/promotion decisions.

- team production obscures individual contributions
- effort and ability hard to disentangle
- capacity for further development or interpersonal skills hard to discern

So current supervisor is often *better informed* than decision-maker about candidate's attributes.

Introduction (2)

But current supervisor's incentives may not match those of decision-maker: supervisor may be *biased in favor of* the candidate.

- predisposition in favor of those like us
- easier to hold or express *positive* opinions of co-workers
- supervisors may bask in "reflected glory"
- anticipation of future favors from those we help

What are consequences of decision-makers' using evaluators' past recommendations to decide (optimally) how much to rely on their current advice?

Can biased evaluators be effectively disciplined by the reputational incentives thus generated?

Introduction (3)

When evaluators' biases are privately known, tracking past recommendations $\Longrightarrow$

- Decision-maker is more informed about evaluators' objectivity in later periods (even without having learned early candidates' true abilities) $\Longrightarrow$ in later pds, decision-maker relies less on those believed to be more biased.
- Very biased evaluators will make less-biased recommendations (be tougher) in early periods.

BUT

- In early periods, relatively unbiased evaluators will use standards that are *too tough*, to preserve their reputation, and even very biased evaluators may be too tough.
- Paradoxically, these reputational incentives may *reduce* early reports' informativeness about evaluators' objectivity— because standards used by different "types" of evaluator may become *more similar* as well as tougher.

Main questions

1. In dynamic settings, in what environments is overall quality of decision-making improved, and in what environments is it worsened, by decision-maker's keeping track of evaluators' past recommendations?
2. Is quality of decision-making improved when reputations are based not only on past recommendations but also on past candidates' actual abilities?

Our analysis highlights the importance of

- factors affecting the *absolute* strength of evaluators' reputational concerns (e.g. weight placed on later decisions relative to early ones)
- factors affecting the *relative* strength of more-biased and less-biased evaluators' reputational concerns
  - degree of consistency of bias over time
  - degree of similarity of average quality of candidates over time

Related Literature

- Cheap talk models: Crawford + Sobel (1982)
  - Privately-known preferences of sender: Morgan + Stocken (2003)
  - Repeated cheap talk: Sobel (1985); Benabou + Laroque (1992); Morris (2001); Jullien + Park (2009)
- Career concerns of experts (decision-makers or advisors): Do career concerns enhance or reduce efficiency?
  - Experts who differ in *"expertise"* (quality of their information): Scharfstein + Stein (1990); Prendergast + Stole (1996); Levy (2004); Prat (2005); Gentzkow + Shapiro (2006)
  - Experts who differ in *motives*: Sobel; Benabou + Laroque; Morris
- Reputational incentives when *all* types are strategic: Morris; Ely and Valimaki (2003); Jullien + Park
- Dynamic *P*-*A* models in which access to more info can reduce *P*'s overall payoff when he cannot commit ex ante how he will use it: Cremer (1995); Meyer + Vickers (1997); Dewatripont et al (1999)
- Political agency: Smart and Sturm (2008) on term limits

The Basic Dynamic Model: Repeated Cheap Talk

- Ability of candidate in pd. t ($t = 1, 2$) is $s_t \sim U[0, 1]$; s_1, s_2 indep.
- In each pd, P must decide whether or not to hire

$$P\text{'s } pd\text{-}t \text{ payoff} = \begin{cases} s_t & \text{if hires} \\ r & \text{otherwise} \end{cases}$$

- Evaluator (A) privately observes s_t at start of pd.t and $c \in \{c_L, c_H\}$ (where $0 \leq c_L < c_H \leq r$) at start of pd.1. Prior prob. p that $c = c_L$.

$$A\text{'s } pd\text{-}t \text{ payoff} = \begin{cases} s_t + c & \text{if candidate is hired} \\ r & \text{otherwise} \end{cases}$$

- r and c remain unchanged over time, so periods 1 and 2 differ only w.r.t. s_t (relaxed later).
- A uses "weighting factor" $\delta \in [0, \infty)$ and P uses $\delta_P \in [0, \infty)$.
- The 5 parameters (r, c_H, c_L, p, δ) are common knowledge.

The Basic Dynamic Model (2)

- In each pd., after observing s_t, A makes a costless but unverifiable report to P. No restrictions on form of A's report.
- Then P chooses optimally whether or not to hire, on basis of updated beliefs.
- P learns nothing more about s_1 before pd-2 decision (relaxed later).
- Payoffs accrue at end of pd 2.

Analyze Perfect Bayesian Eqa. satisfying Assumption ES: the players *never*, at any point in the game, play a continuation eqm that is *strictly worse* for everyone than another eqm.

Static Benchmark: A is consulted for only 1 period

Generically, $\exists$ at most 1 eqm. with informative communication:

1. A_i $(i = H, L)$ uses cutoff $z_i = r - c_i$ and reports "hire" if $s \geq r - c_i$ and "do not hire" if $s < r - c_i$.

2. P's decision rule:
   - If $r < \frac{1}{2}$ (so with no infor, P prefers to hire), then P follows A's advice, $\forall p \in [0, 1]$.
   - If $r > \frac{1}{2}$ (so with no infor, P prefers not to hire), then
     - if $c_H < 1 - r$, P follows A's advice, $\forall p \in [0, 1]$;
     - if $c_L > 1 - r$, P ignores advice and does not hire, $\forall p \in [0, 1]$;
     - if $c_L < 1 - r < c_H$, $\exists\ \overline{p}$ such that for $p > \overline{p}$, P follows advice, and for $p < \overline{p}$, P ignores advice and does not hire.

Reputational Incentives

We say "reputational incentives arise" if, in 1st pd, at least one type of A behaves differently than in static model.

Proposition

Generically necessary conditions for reputational incentives to arise in the dynamic model are

$$\underbrace{r > \frac{1}{2}}_{\substack{\text{prior inclination} \\ \text{not to hire}}} \quad \text{and} \quad \underbrace{c_L < 1 - r < c_H}_{\substack{\text{if } P \text{ were certain of } A\text{'s bias,} \\ \text{he would follow } A_L\text{'s advice and} \\ \text{ignore } A_H\text{'s advice}}}$$

(Otherwise, P's 2nd-pd behavior is independent of posterior beliefs over A's type $\Rightarrow$ A's 1st-pd report has no effect on his 2nd-pd payoff.)

Henceforth, assume that these necessary conditions hold.

Analysis of the Dynamic Model

In a pure-strategy eqm. in which reputational incentives arise:

- In each period, A reports either "hire" or "do not hire".
- P follows A's advice in 1st pd.
- $z_{1H} \leq z_{1L}$, so reporting "hire" ("do not hire") in 1st pd is bad (good) for A's reputation.
- After "hire" in pd 1, P is suffic. pessimistic about c that ignores advice in pd 2 and does not hire.
- After "do not hire" in pd 1, P is suffic. optimistic about c that he follows advice in pd 2.
- $z_{1i} = r - c_i + \delta\Delta_i$, where $\Delta_i = \frac{(1-r+c_i)^2}{2} =$ option value for A_i of having advice followed in 2nd pd. $=$ value of a "good" reputation to A_i.
  - $\Delta_H > \Delta_L$, since being ignored in pd 2 hurts A_H more than A_L.

Existence and Uniqueness

Proposition

Given Assn. ES, generically

1. *there is at most 1 informative eqm in the dynamic model, and in it A reports either "hire" or "do not hire" in each pd;*
2. *reputational incentives arise in 1st pd if and only if* $p \in [p_\ell, p_u)$*, and they induce both types of A to use* **tougher** *standards than in 2nd pd.*

Equilibrium in the Dynamic Model

R: pure-strat. reput. eqm: A can induce P to hire at most 1 of 2 cands.
S: in both pds, P follows advice, and A uses static cutoffs.
B: in both pds, P ignores advice and does not hire.
M_{2y}, M_{2n}, M_1: P acts as in R, except mixes on one branch.

$A_H^{'}$s and $A_L^{'}$s cutoffs rise with δ only in region R and are toughest when $\delta = \overline{\delta}$ and $p = \overline{p}$: at this point, $z_{1H} = z_{1L}$.

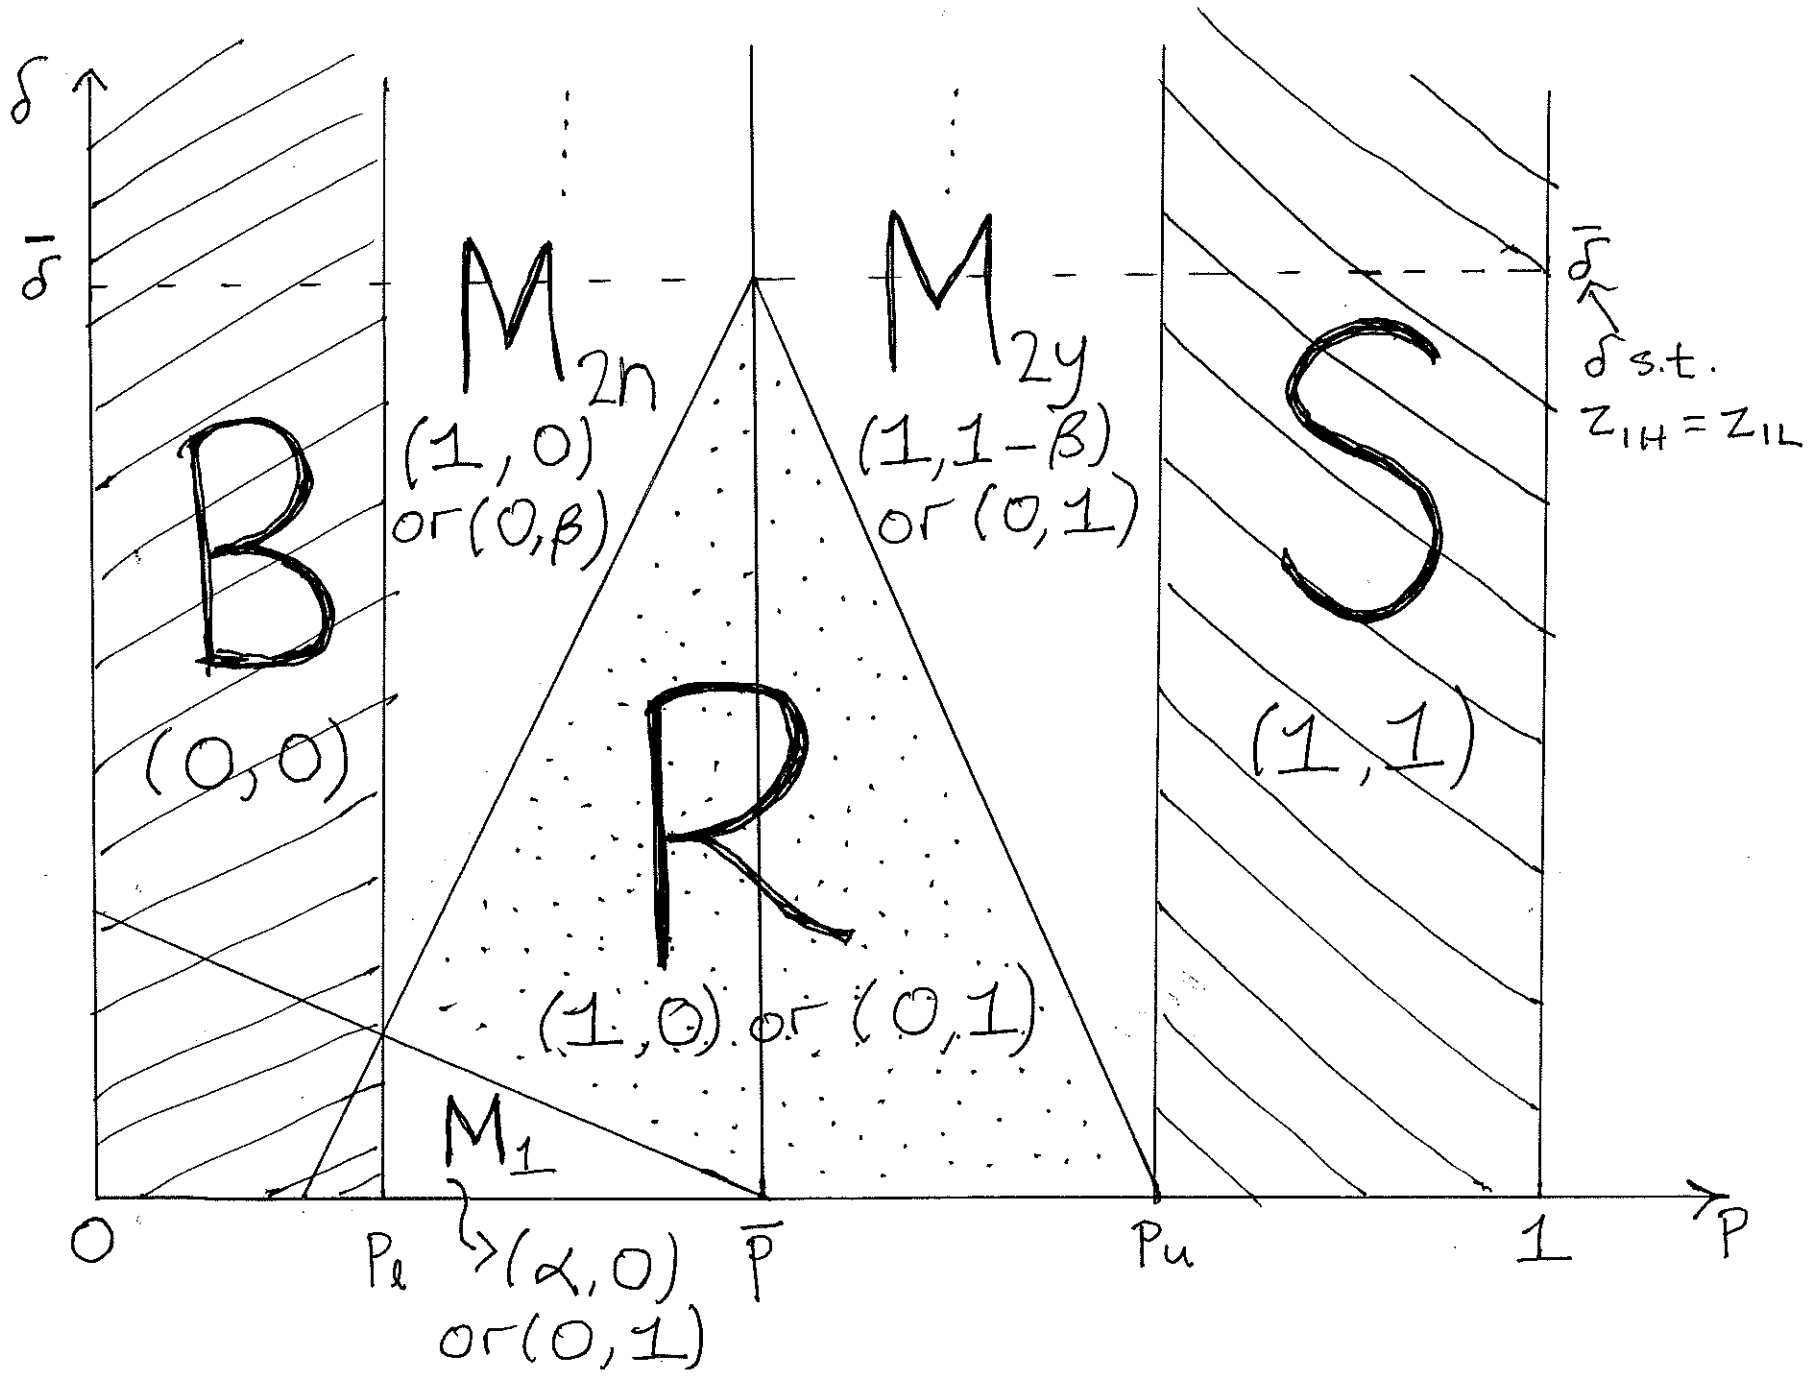
δ
δ̄
B
(0,0)
M2n
(1,0)
or (0,β)
M2y
(1,1−β)
or (0,1)
R
(1,0) or (0,1)
S
(1,1)
δ̄
δ s.t.
z1H = z1L
M1
(α,0)
or (0,1)
0
Pl
P̄
Pu
1
P

Quality of Decision-Making

Compare P's expected payoff in dynamic model to payoff in twice-repeated static setting. Two interpretations:

1. Should P keep track of past evaluations or not?
2. Should P repeatedly consult same evaluator or commit to replace him after period 1?

Not obvious that benefits of inducing reput. incentives outweigh costs:

- Pure-strategy reputational cutoffs for both A types can be too tough ($> r$), e.g. if $c_L = 0$ and δ large.
- Switching from repeated-static to dynamic eqm, P must gain on A_H, even if $z^*_{1H} > r$. But P may lose on A_L, e.g. if $c_L = 0$.
- The period-2 value of the information about A's type from the period-1 report is non-negative, but can be *small* or even *zero,*
  - e.g. as $\delta \to \overline{\delta}$, so $z_{1L} - z_{1H} \to 0$, value of info about A's type $\to 0$.

Reputational Incentives Help in the Basic Dynamic Model

Proposition

$\forall(p,\delta)$, P's payoff in each period in the equilibrium of the basic dynamic model is at least as large as in the equilibrium of the static model.

- Because A_H values a good reputation more than A_L ($\Delta_H > \Delta_L$), and eqm requires $z_{1H} \leq z_{1L}$, reputational incentives are "self-limiting": eqm cutoffs are never tougher than z^*_{1H}, z^*_{1L} such that

$$z^*_{1H} = r - c_H + \delta\Delta_H = r - cL + \delta\Delta_L = z^*_{1L}.$$

- In the basic dynamic model $z^*_{1H} = z^*_{1L} < 1$; hence in 1st pd., even with maximal cutoffs, though P learns nothing about A's bias, he does learn something about s_1.

More General Environments

In more general environments, inducing reputational incentives *may hurt* the principal.

- The cost to P in pd.1 of A using too-stringent standards may outweigh the benefit in pd.2 of knowing more about A's bias.

This happens in (at least) two generalizations of the basic model:

1. A's bias is *imperfectly correlated* across periods/candidates (i.e. $corr(c_1, c_2) < 1$)) and is observed by A period by period.
2. $r_1 > r_2$ (equivalently, 1st-pd candidate's expected ability is lower than that of 2nd-pd candidate).

When Reputational Incentives Worsen Decision-Making

Examples where inducing reputational incentives *hurts* the principal:

δ large; $p > \overline{p}$; $c_L = 0$; and either $corr(c_1, c_2)$ small or $r_1 >> r_2$

- In dynamic eqm., as $\delta \uparrow$, eventually both types of A use cutoffs ≈ 1, i.e. both A types effectively *always reject* pd-1 candidate $\Rightarrow$ P learns nothing in 1st pd, either about s_1 or about c_1.

- Therefore in 1st pd., dynamic $\prec$ static and in 2nd pd, dynamic $\sim$ static.

- $corr(c_1, c_2) \downarrow$ and $(r_1 - r_2) \uparrow$ both $\Longrightarrow$ $\frac{\Delta_L}{\Delta_H}$ rises relative to $\frac{1-r_1+c_L}{1-r_1+c_H}$, i.e. A_L's reputational incentives rise relative to A_H's $\Longrightarrow$ maximal strength of eqm reputational incentives increases.

Imperfect Correlation of Types across Periods

Proposition

Assume c_1 and c_2 are ex ante identically distributed and define $\kappa \equiv corr(c_1, c_2)$. Then $\exists \quad \widehat{\kappa}(r, c_H, c_L) \in (0,1)$ s.t.

1. *for $\kappa \geq \widehat{\kappa}$ and for all (p, δ), P's payoff in each period in the dynamic eqm. is weakly larger than in the eqm. of the static model.*
2. *for $\kappa < \widehat{\kappa}$, we can find, for all δ_p, (p, δ) such that P's overall payoff in the dynamic eqm. is smaller than in the twice repeated version of the static eqm.*

Different Distributions of Candidate Quality across Periods

Proposition

Assume $r_2 > \frac{1}{2}$ *and* $0 \leq c_L < 1 - r_2 < c_H \leq r_2$. *Then*
$\exists \quad \widehat{r}_1(r_2, c_H, c_L) > r_2$ *s.t.*

1. *for* $r_1 \in [c_H, \widehat{r}_1]$ *(i.e. for* r_1 *not too different from* r_2*) and for all* (p, δ)*, P's payoff in each period in the dynamic eqm. is weakly larger than in the eqm. of the static model.*

2. *for* $r_1 > \widehat{r}_1$*, we can find, for all* $\delta_p, (p, \delta)$ *such that P's overall payoff in the dynamic eqm. is smaller than in the twice-repeated version of the static eqm.*

Reput. Incentives when Abilities are Observed Ex Post

How are reputational incentives and quality of decision-making affected by the amount of information P obtains between the 1st- and 2nd-pd decisions about the 1st-pd candidate's ability?

- Suppose that btw. 1st- and 2nd-pd decisions, P observes s_1.
- Suppose A continues to report "hire" or "do not hire" in each period.
- If eqm cutoffs are distinct, it must be that $z_{1H} < z_{1L}$. In this case,
  - If $z_{1H} \leq s_1 < z_{1L}$, then P can deduce A's type from either 1st-pd report.
  - If $s_1 < z_{1H} < z_{1L}$ or $z_{1H} < z_{1L} \leq s_1$, then both types make same report in eqm, so P learns nothing about A's type.
- Focus on $p > \overline{p}$ (so if P learns nothing about A's type in pd 1, will follow advice in pd 2) and $\delta < \frac{c_H - c_L}{\Delta_H}$.

Reputational Incentives when s_1 is Observed Ex Post

Let $p > \bar{p}$. Notation: (1st-pd report, P's 2nd-pd. strategy)

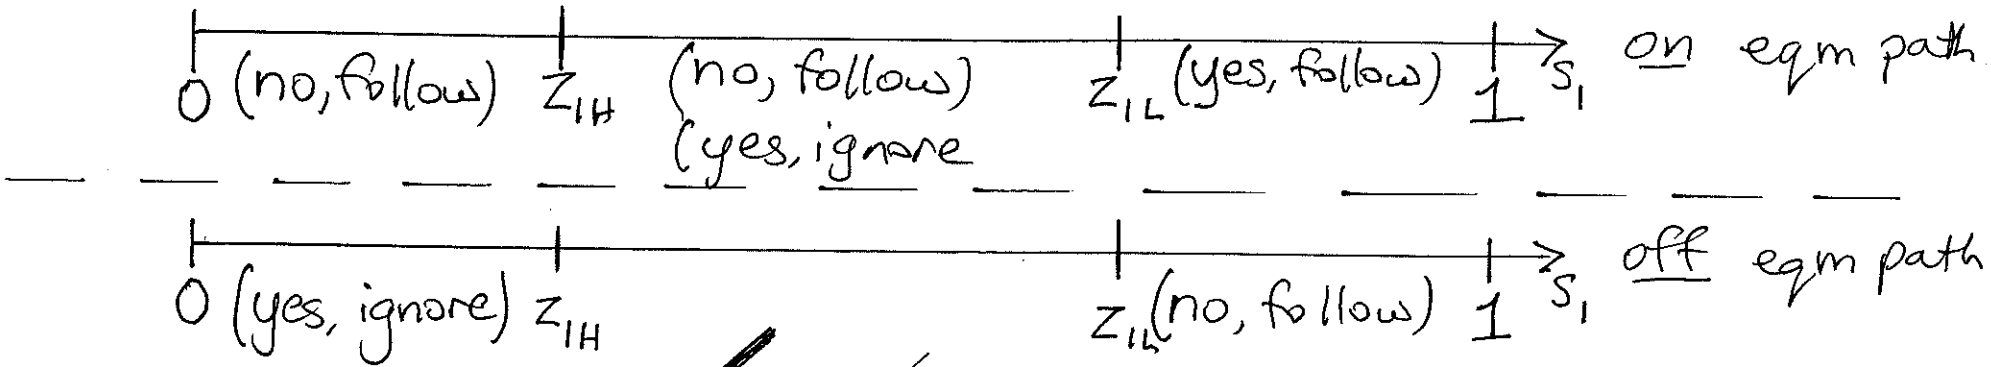

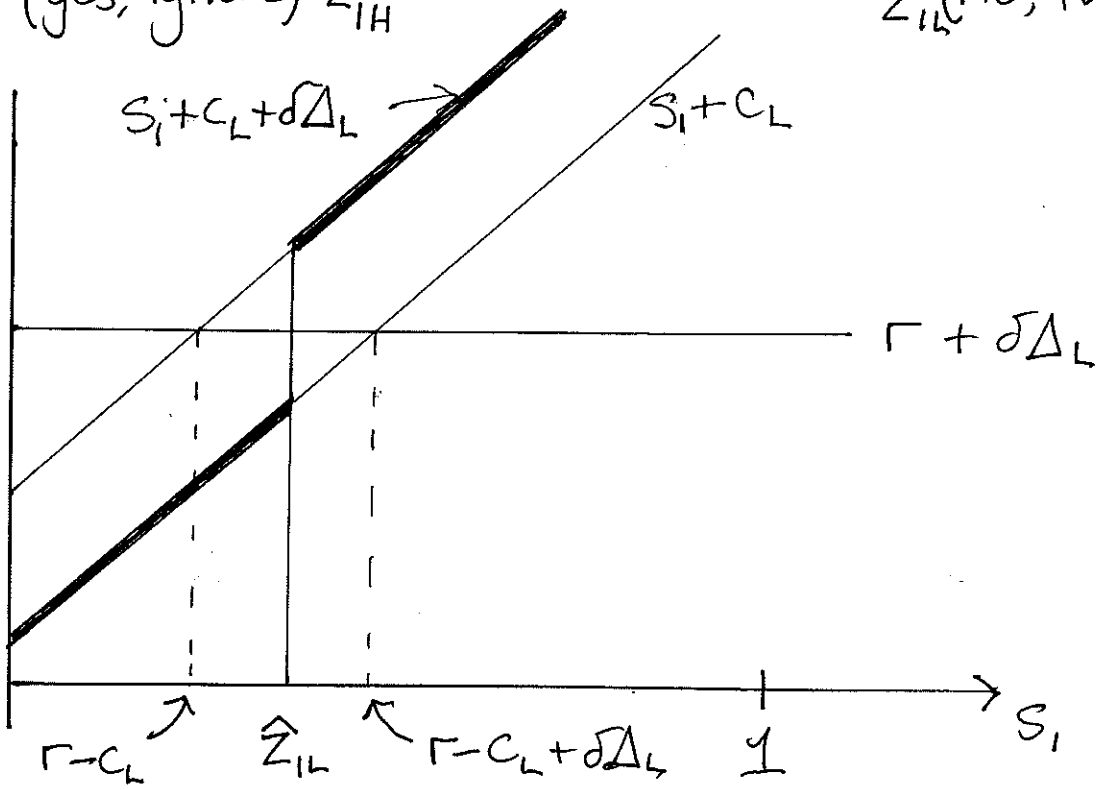

Eqm cutoffs:

$z_{1H} = r - c_H + \delta\Delta_H$

$z_{1L} \in [r - c_L, r - c_L + \delta\Delta_L]$

(NB: We are assuming $r - c_H + \delta\Delta_H < r - c_L$.)

Reput. Incentives when Abilities are Observed Ex Post (2)

Claim: If $p > \overline{p}$ and $\delta < \frac{c_H - c_L}{\Delta_H}$, then in the model where P observes s_1 after the 1st pd, $\exists$ a multiplicity of dynamic eqa:

$$\begin{aligned} z_{1H} &= r - c_H + \delta\Delta_H; \\ z_{1L} &\in [r - c_L, r - c_L + \delta\Delta_L]. \end{aligned}$$

Hence observing s_1 can weaken reputational incentives (when $(p, \delta) \in$ Region R) or strengthen reputational incentives (when $(p, \delta) \in$ Region S).

Claim: $\forall\ \delta_p$, we can find (p, δ) and (r, c_H, c_L) s.t. P's overall payoff in a dynamic eqm is smaller in scenario where he observes s_1 after the 1st pd than in scenario where he does not.

Reput. Incentives when Abilities are Observed Ex Post (3)

Example where observing s_1 after pd 1 lowers P's overall expected payoff: Choose p slightly larger than $\overline{p}$ and δ slightly below $\frac{c_H - c_L}{\Delta_H}$.

- In scenario where P does not observe s_1, we are in Region R:

$$z_{1i}^* = r - c_i + \delta\Delta_i, \qquad i = L, H$$

- In scenario where P observes s_1 after pd 1, select eqm with $z_{1L} = r - c_L$ (and $z_{1H} = r - c_H + \delta\Delta_H$).
- Then for chosen δ, $z_{1H} \approx z_{1L} \Longrightarrow$ from 1st pd, P learns essentially nothing about A's type $\Longrightarrow$ P's 2nd-pd payoff is strictly less in scenario where P observes s_1.
- It is easy to choose $c_L \in (0, c_H)$ s.t. P's 1st-pd payoff from A_L is lower when $z_{1L} = r - c_L$ than when $z_{1L} = r - c_L + \delta\Delta_L$.
- P's 1st-pd payoff from A_H is equal in the two scenarios.

Thus when A anticipates P's access to extra information (s_1), the result can be a worsening of both sorting (how much P learns from pd-1 behavior about A's type) and incentives (quality of pd-1 decisions).

Conclusions

When do *repeated* consultations effectively discipline biased evaluators? Tracking evaluators' past recommendations $\Longrightarrow$ all types of evaluator use **tougher** standards in early periods, but tougher standards may not always benefit decision-maker:

- more-biased A using tougher standards: good
- less-biased A using tougher standards: can be bad

Whether inducing reputational incentives improves overall quality of decision-making depends on

- *absolute* strength of reputational concerns (weight δ placed by A on later decisions)
- strength of good A type's reputational concerns *relative* to bad type's (Δ_L/Δ_H)
  - When both of the above are sufficiently *large*, inducing reputational incentives can *reduce* quality of decision-making.