Designing Hiring and Promotion Procedures When Evaluators are Biased

Christopher Avery

Kennedy School of Government
Harvard University

and

Margaret Meyer

Nuffield College
Oxford University

Examine procedures by which evaluators provide recommendations used for hiring or promotion decisions.

Assume

- evaluators privately informed about candidates' abilities
- each evaluator *biased* in favor of the candidate he is evaluating
- recommendations costless to provide and impossible to verify, i.e. "cheap talk"

Exs.: professional service firms; university faculties; military; "old boy networks"

Other applications:

- interested experts providing policy advice
- managers recommending their "pet projects"

Questions:

How do evaluators' biases affect the amount of information communicated and the quality of decisions?

How can evaluation procedures be designed to mitigate the effects of these biases?

Simplest scenario:

1 evaluator, with publicly known bias, privately observes the ability of a single candidate and then reports to the decision-maker

decision-maker then makes a **binary** decision: whether or not to hire candidate

- simpler than Crawford and Sobel (1982)

Lines of enquiry:

1. If there are several evaluators, each privately informed about their own candidates, how to structure communication?
2. If decision-maker, too, has private information (e.g. about job to be filled), how to structure communication?
3. **If evaluator's degree of bias is private information, what are the consequences of decision-makers' keeping track of evaluators' past recommendations?**

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

- decision-maker is more informed about evaluators' objectivity in later periods
- very biased evaluators will make less-biased recommendations (be tougher) in early periods

BUT

- unbiased evaluators will use standards that are too tough, to preserve their reputation
- even very biased evaluators may be too tough
- paradoxically, these reputational incentives may *reduce* reports' informativeness about evaluators' objectivity
  - because standards may become *more similar* as well as tougher.

Repeated cheap-talk model with 1 principal (P), 1 evaluator (A), and in each of 2 periods, 1 (passive) candidate, whom P must choose whether or not to hire, on basis of A's costless but unverifiable report.

- A privately observes candidate's ability in each period *and* his own bias.
- reputation-building by privately-informed types, *all* of whom are strategic
  - contrast Kreps, Milgrom, Roberts and Wilson (1982), Sobel (1985), Mailath and Samuelson (1998)
  - compare Morris (1997)

Contrast *dynamic eqm.* with *repeated static eqm.*

- *static eqm* – evaluator consulted for only 1 pd.
- *dynamic eqm* – evaluator consulted for 2 pds, and P keeps track of 1st-pd report.

In our static model, $\exists$ generically at most 1 eqm with informative communication.

In the dynamic model, under a natural assn. on eqm selection, $\exists$ generically at most 1 eqm with informative communication.

Welfare comparisons for the principal:

In dynamic model, is P necessarily better off drawing from a less-biased pool of evaluators? No.

How does P's payoff compare in dynamic vs. repeated-static eqa?

Can address other questions in organizational design:

- What is the effect of feedback about 1st-pd candidate's ability, before 2nd-pd decision?

- Is it beneficial for P to retain flexibility about whether to consult or replace A in 2nd pd?

- How does tracking past evaluations compare with other ways of inducing biased evaluators to use tougher standards, e.g. simultaneous evaluations of many candidates, subject to a hiring quota?

The Basic Dynamic Model

Ability of candidate in pd. t $(t = 1, 2)$ is
$s_t \sim U[0, 1]$; s_1 and s_2 independent

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

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

Prior probability that $c = c_L$ is p.

A uses "discount factor" $\delta \in [0, \infty)$ and P uses $\delta_P \in [0, \infty)$.

The 5 parameters (r, c_H, c_L, p, δ) : common knowl.

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 whether or not to hire.

P learns nothing more about s_1 before pd-2 decision.

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 Eqm: A is consulted for only 1 period

A_i's reporting strategy ($i = H, L$) :

$$A_i \text{ reports} \begin{cases} \text{"hire" if } s \geq r - c_i \\ \text{"do not hire" if } s < r - c_i; \end{cases}$$

i.e. A_i uses cutoff $z_i = r - c_i$

P's decision rule:

If $r < \frac{1}{2}$ (so with no information, P prefers to hire), then P follows A's advice, $\forall p \in [0, 1]$

If $r > \frac{1}{2}$ (so with no information, 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 $p \gtrless \overline{p}$, P $\begin{cases} \text{follows advice} \\ \text{ignores advice and does not hire} \end{cases}$

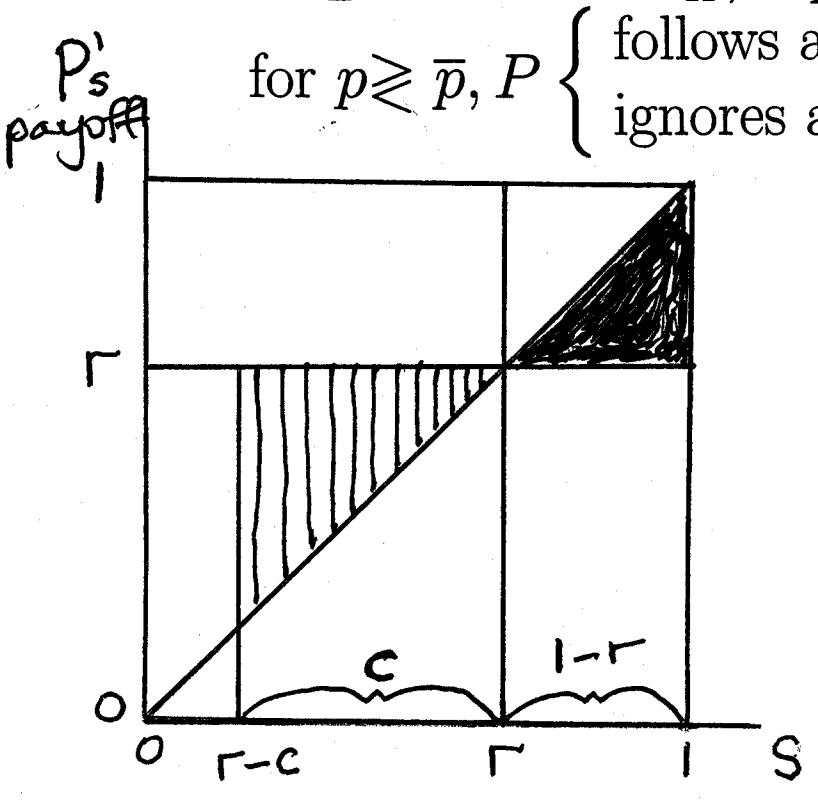

We say "reputational incentives arise" if, in 1st pd, at least one type of A behaves differently than in static model.

Proposition 1 : *Generically necessary conditions for reputational incentives to arise in the dynamic model are*

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

(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.

Dynamic Eqm: A consulted for 2 pds and P keeps track of 1st-pd report

Pure-strategy reputational eqm:

- 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 1st pd, P is pessimistic about c and ignores advice in 2nd pd.
- after "do not hire" in 1st pd, P is optimistic about c and follows advice in 2nd pd.
- $z^*_{1i} = r - c_i + \delta\Delta_i$, where $\Delta_i = \frac{(1-r+c_i)^2}{2}$
  $=$ option value for A_i of having advice followed in 2nd pd. $\equiv$ "reputational gain"
  NB: $\Delta_H > \Delta_L$

Proposition 2 : *Given Assn. ES, generically*

i) there is at most one informative eqm in the dynamic model;

ii) in any informative eqm, A reports either "hire" or "do not hire" in each pd;

iii) reputational incentives arise in 1st pd iff $p \in [p_\ell, p_u)$.

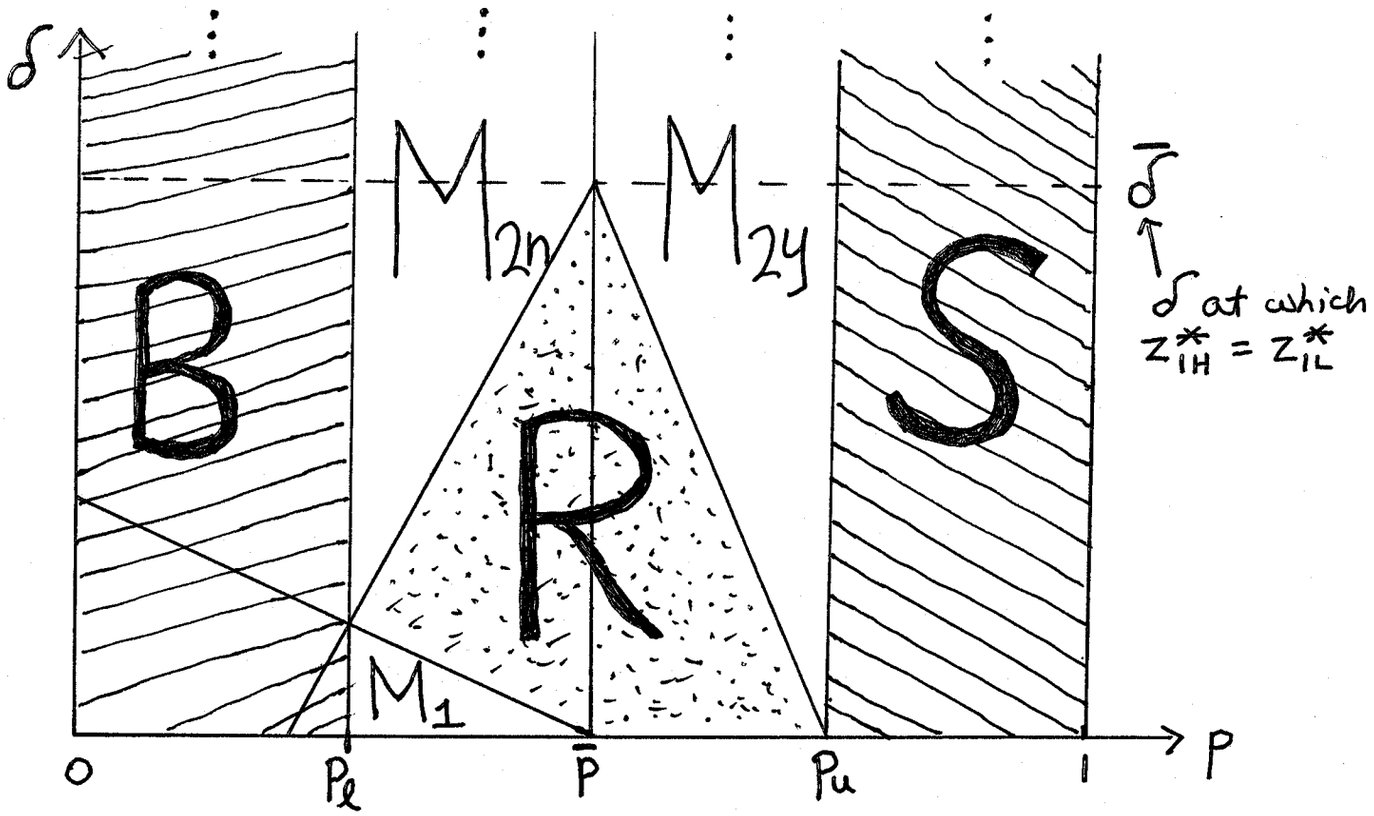

R: pure-strategy reputational eqm: A can induce P to hire at most 1 of the 2 candidates

S: in both periods, P follows advice, and A uses static eqm cutoffs

B: in both periods, P ignores advice and does not hire

M_{2y}, M_{2n}, M_1: mixed-strategy eqa. – P acts as in region R except that in

M_{2y}: P randomizes in pd 2 after "hire" in pd 1

M_{2n}: " " " " "do not hire" "

M_1: P randomizes in pd 1 after "hire"

Welfare Analysis for Principal:

Keeping track of the 1st-pd evaluation is not obviously beneficial:

- pure-strategy reputational cutoffs for both A types can be too tough ($> r$)
  - this can happen if $(1-r)^2 > c_H c_L$, e.g. if $c_L = 0$
- 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 $\to 0$

Proposition 3 : *$\forall (p, \delta)$, P's payoff in each period is at least as large in the dynamic eqm as in the repetition of the static eqm.*

- every MSE is payoff-equivalent for P to some pure-strategy reputational eqm
- in region R, A's cutoffs are toughest when $\delta = \overline{\delta}$ and $p = \overline{p}$
  - at this point, $z^*_{1H} = z^*_{1L} < 1$; hence in 1st pd., though P learns nothing about A's bias, he does learn something about s_1

But more generally, inducing reputational incentives by tracking past evaluations may *hurt* the principal

- P may be made worse off by the too-stringent standards used in pd.1 and gain arbitrarily little or nothing in pd. 2

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

1. allow $r_1 \neq r_2$ (equivalently, allow candidates of *different* expected ability)

or

2. allow A's bias to be *imperfectly* correlated across periods/candidates (i.e. $corr(c_1, c_2) \in (0, 1)$)

Proposition 4: *Assume* $r_2 > \frac{1}{2}$ *and* $0 \leq c_L < 1 - r_2 < c_H \leq r_2$. *Then* $\exists\ \widehat{r}_1(r_2, c_H, c_L) > r_2$ *s.t.*

i) *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 is weakly larger in the dynamic eqm than in the succession of static eqa.*

ii) *for* $r_1 > \widehat{r}_1$*, we can find, for all* δ_p*,* (p, δ) *such that* P*'s overall payoff is smaller in the dynamic eqm than in the succession of static eqa.*

Proposition 5: *Assume* c_1 *and* c_2 *are identically distributed and define* $\kappa \equiv corr(c_1, c_2)$. *Then* $\exists\ \widehat{\kappa}(r, c_H, c_L) \in (0,1)$ *s.t.*

i) *for* $\kappa \geq \widehat{\kappa}$ *and for all* (p, δ) P*'s payoff in each period is weakly larger in the dynamic eqm than in the succession of static eqa.*

ii) *for* $\kappa < \widehat{\kappa}$, *we can find, for all* δ_p, (p, δ) *such that* P*'s overall payoff is smaller in the dynamic eqm than in the succession of static eqa.*

Examples where inducing reputational incentives *hurts* the principal:

Ex.1: $r_1 > \widehat{r}_1$, δ large, p large

Ex. 2: $\kappa < \widehat{\kappa}$, δ large, p close to but above $\overline{p}$

- In both exs., in dynamic eqm., reputational incentives so strong that *both* types of A effectively *reject* 1st-pd candidate *for sure* (i.e. $z^*_{1L} \approx 1$, $z^*_{1H} \approx 1$)

  $\Rightarrow$ P learns *nothing* in 1st-pd, either about s_1 or about c_1.

- p is chosen so that static eqm in 1st pd. involves *informative* communication.

- Therefore in 1st pd., dynamic $\prec$ static

  and in 2nd pd., dynamic $\sim$ static

Other Questions of Organizational Design:

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

- Is it valuable for P to expend resources to observe s_1?

Suppose that btw. 1st- and 2nd-pd decisions, P observes s_1. Suppose A continues to report "hire" or "do not hire".

If $z_{1H} < s_1 < z_{1L}$, then P can deduce A's type from either 1st-pd report.
If $s_1 < z_{1H} < z_{1L}$, then if A reports "do not hire", P learns nothing.
If $s_1 > z_{1L} > z_{1H}$, then if A reports "hire", P learns nothing.

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

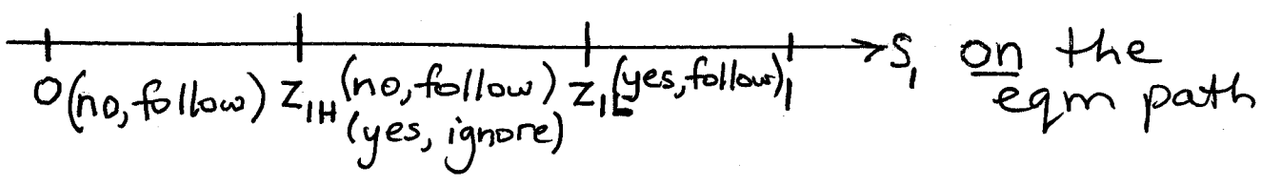

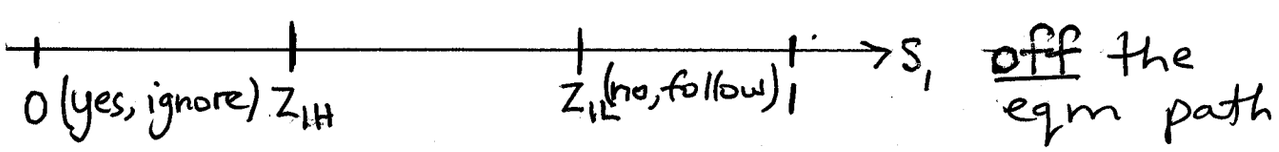

Claim: If $p > \overline{p}$ and $\delta < \frac{c_H - c_L}{\Delta_H}$, then in the model where P observes s_1 after the 1st pd, $\exists$ a multiplicity of dynamic eqa:

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

So 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.

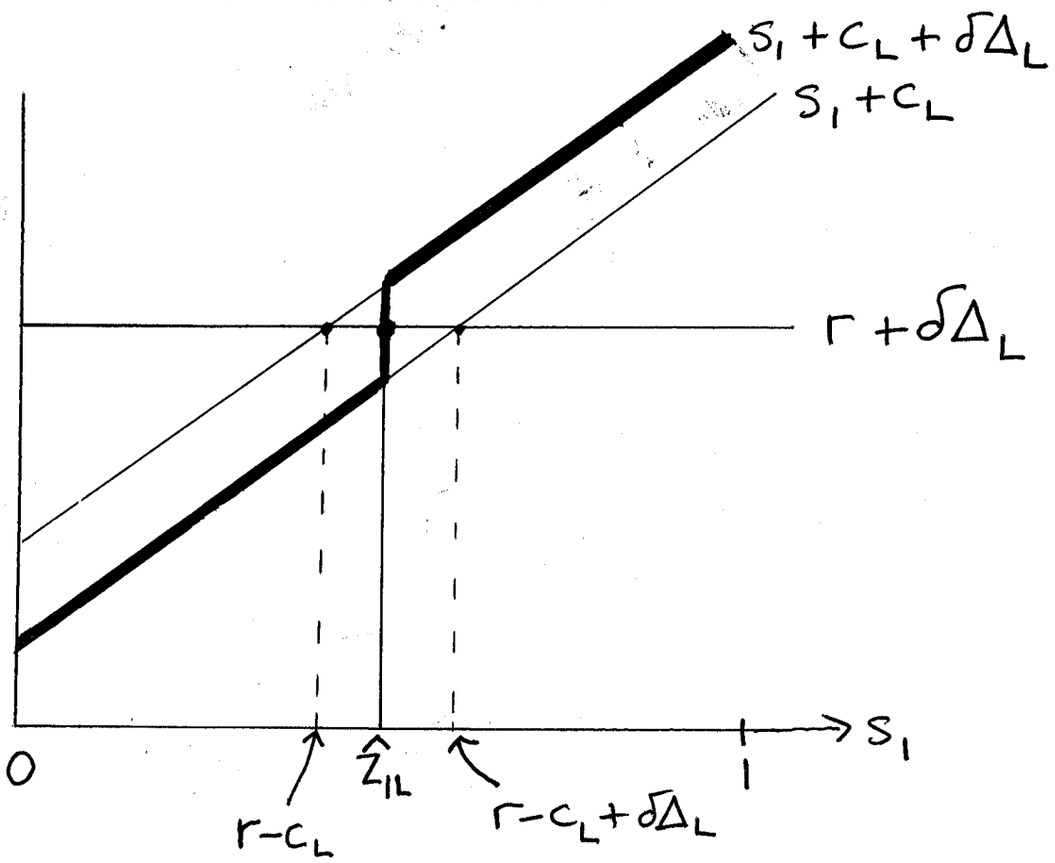

An eqm. in which $z_{1L} = \hat{z}_{1L}$: given that P conjectures A_L uses $\hat{z}_{1L}$, A_L's payoff from reporting "yes" in pd 1 is <u>discontinuous</u> at $\hat{z}_{1L}$

Ex: 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 have a pure-strategy reputational eqm:

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

In scenario where P observes s_1, select eqm with $z_{1L} = r - c_L$. 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 $c_H > c_L > 0$ s.t. P's 1st-pd payoff from A_L is lower when $z_{1L} = r - c_L$ than when $z_{1L} = r - c_L + \delta\Delta_L$.

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)

c.f. Prat (2001) for a result with a similar flavor

Conclusions

Keeping track of evaluators' past recommendations has complex effects on eqm behavior

- by inducing reputational incentives, it affects
  - the quality of the 1st-pd decision and
  - the value of the information generated in pd 1 for the pd-2 decision

We've found sufficient conditions for P to be weakly better off in both pds.

But we've also shown that inducing reputational incentives can hurt P.

Giving P extra information (about s_1) or extra options (flexibility) can also hurt P (relative to the basic dynamic eqm), for the same reasons.