Appendix to "How will growth in Medicare Advantage change the Medicare program's performance?"

The main text considers how shifts of enrollment from TM to MA may affect the performance of TM's ACO programs. This appendix presents a model that formalizes the discussion of two ways in which these enrollment shifts may affect ACOs' willingness to participate in ACO contracts: (1) changing the likelihood that an ACO is assigned a benchmark that it views as unattainable; and (2) changing an ACO's ability to predict its ultimate financial performance.

Model setup

Each ACO has an (unknown) latent spending level Λ drawn from a distribution $N(\lambda, \tau^2)$. It serves a population of N Medicare beneficiaries indexed by i, of whom beneficiaries with $i \leq n \leq N$ are enrolled in TM. There are two time periods of interest: the current period (denoted by $t = 0$) and the prior period (denoted by $t = -1$). Spending by beneficiary i in period t is denoted c_t^i and drawn i.i.d. from the distribution $N(\Lambda, \sigma^2)$. For convenience in what follows, I define the scaled standard deviation $\bar{\sigma}(k) = \sigma \,/\, \sqrt{k}$ and the sample mean $\bar{c}_t(k) = k^{-1} \sum_{i \leq k} c_t^i$.

ACOs face a benchmark $B(k) = w\bar{c}_{-1}(k) + (1 - w)E$, where E is a fixed exogenous component and $w \in [0,1]$ is a fixed weight. Under current policy, $k = n$, so the benchmark reflects spending of TM beneficiaries only, but I also consider scenarios with $k = N$, where MA beneficiaries are also included. The ACO's (gross) savings in the current period are then $S(k) = B(k) - \bar{c}_0(n)$.

ACOs' beliefs at the time of participation decisions

ACOs' participation decisions are likely to depend on their beliefs about the savings they would achieve if they participated, $S(k)$. I assume that ACOs form beliefs based on the outcomes they observe at time $t = -1$. Concretely, I treat each ACO as having an information set of the form $\mathcal{I}(m) = \{c_{-1}^i\}_{i \leq m}$. The formal object of interest is therefore the distribution of $S(k) \mid \mathcal{I}(m)$.

I consider two assumptions about the breadth of information that ACOs observe: the first is that ACOs only observe the spending of their TM patients, corresponding to $m = n$, while the second is that ACOs also observe the spending of their MA patients, corresponding to $m = N$. Throughout, I only consider cases with $k \leq m$. That is, I limit consideration to cases where ACOs have access to at least as much information as is used to set their benchmarks; the rationale is that CMS is likely to make any information it uses to set benchmarks available to ACOs.

Under these assumptions, it is straightforward to show that this distribution has expectation

$$\mathbb{E}[S(k) \mid \mathcal{I}(m)] = B(k) - \mathbb{E}[\Lambda \mid \mathcal{I}(m)]$$

and variance

$$\begin{aligned}\mathrm{Var}[S(k) \mid \mathcal{I}(m)] &= \mathrm{Var}[\bar{c}_0(n) \mid \mathcal{I}(m)] \\ &= \mathbb{E}[\mathrm{Var}(\bar{c}_0(n) \mid \mathcal{I}(m), \Lambda) \mid \mathcal{I}(m)] + \mathrm{Var}[\mathbb{E}(\bar{c}_0(n) \mid \mathcal{I}(m), \Lambda) \mid \mathcal{I}(m)] \\ &= \bar{\sigma}(n)^2 + \mathrm{Var}[\Lambda \mid \mathcal{I}(m)].\end{aligned}$$

These amounts depend in turn on an ACO's beliefs about its latent spending level Λ. Because the c_{-1}^{i} are i.i.d. normal with mean Λ, $\bar{c}_{-1}(m)$ is a sufficient statistic for Λ given the information set $\mathcal{I}(m)$. Now, observe that $\bar{c}_{-1}(m) = \Lambda + \epsilon_{-1}$, where $\epsilon_{-1} \sim N(0, \bar{\sigma}(m)^2)$ and $\epsilon_{-1} \perp \Lambda$. Classic signal extraction results then imply that the distribution of $\Lambda \mid \mathcal{I}(m)$ is normal with

$$\mathbb{E}[\Lambda \mid \mathcal{I}(m)] = \theta(m)\bar{c}_{-1}(m) + \big(1 - \theta(m)\big)\lambda \quad \text{and} \quad \mathrm{Var}[\Lambda \mid \mathcal{I}(m)] = \big(1 - \theta(m)\big)\tau^2$$

where

$$\theta(m) = \frac{\tau^2}{\bar{\sigma}(m)^2 + \tau^2}.$$

Observe that $\theta(m)$ is increasing in m; the more information the ACO observes about prior period spending, the more weight it places on its own experience relative to the overall mean.

Prevalence of "unattainable" benchmarks

I now turn to the question of how often ACOs will receive benchmarks that they view as unattainable. In the language of the model, this is equivalent to how often ACOs will be assigned benchmarks that result in a very low level of expected savings, $\mathbb{E}[S(k) \mid \mathcal{I}(m)]$.

Based on the facts about the distribution of $S(k) \mid \mathcal{I}(m)$ derived above, we obtain

$$\mathbb{E}[S(k) \mid \mathcal{I}(m)] = [1 - w][E - \lambda] + [w - \theta(m)][\bar{c}_{-1}(m) - \lambda] + w[\bar{c}_{-1}(k) - \bar{c}_{-1}(m)]. \qquad (1)$$

The expected savings have three components. The first reflects the difference between the exogenous component of the benchmark and the mean latent spending level. The second reflects the fact that higher spending in the prior period may have a different effect on the ACO's benchmark (as determined by w) than it has on the ACO's subjective beliefs about its latent spending level (as determined by $\theta(m)$). The third reflects the mismatch (if $k \neq m$) between the mean spending used to set the ACO's benchmark and the mean spending observed by the ACO.

It is clear that the *average* value of $\mathbb{E}[S(k) \mid \mathcal{I}(m)]$ across ACOs always equals $[1 - w][E - \lambda]$ regardless of the values of k and m. Thus, to the extent that the likelihood of receiving an unachievable benchmark differs across scenarios, it is because the *dispersion* in ACOs' beliefs and, thus, the prevalence of low expected savings, differs across scenarios.

To assess how the dispersion of beliefs changes, observe that

$$\mathrm{Var}(\mathbb{E}[S(k) \mid \mathcal{I}(m)]) = [w - \theta(m)]^2[\bar{\sigma}(m)^2 + \tau^2] + w^2[\bar{\sigma}(k)^2 - \bar{\sigma}(m)^2], \qquad (2)$$

where this follows from the fact that the second and third terms of equation (1) are independent and some straightforward calculations. There are then three cases to consider:

1. *Benchmarks and ACO beliefs both based on TM enrollees alone* ($k = m = n$): In this case, if $w < \theta(m)$, then falling TM enrollment reduces dispersion. The intuition for this is that falling TM enrollment reduces $\theta(m)$, which brings the effects of prior period spending on the ACO's beliefs and its benchmark into closer alignment.[1] By contrast, the reverse

[1] Formally, observe that, in this case, the right-hand side of equation (2) can be written as $[\theta(m) - w][1 - w/\theta(m)]\tau^2$. Since falling TM enrollment reduces $\theta(m)$, the result follows immediately.

occurs if $w > \theta(m)$. The case with $w < \theta(m)$ is likely the more empirically relevant one since there is likely substantial dispersion in latent spending levels across ACOs, which pushes $\theta(m)$ close to 1, and since policymakers have begun to set $w \ll 1$ in order to give ACOs adequate incentives to reduce spending.

2. *Benchmarks based on TM enrollees, but ACO beliefs based on all enrollees* ($k = n$ and $m = N$): In this case, falling TM enrollment unambiguously increases dispersion. This is because falling TM enrollment no longer changes how prior period spending affects the ACO's beliefs about its latent spending level (reflected in the fact that the first term on the right-hand-side of equation (2) is constant), but does exacerbate mismatches resulting from the fact that the benchmark and the ACO's beliefs reflect different information about prior period spending (reflected in the fact that the second term increases).

3. *Benchmarks and ACO beliefs both based on all enrollees* ($k = m = N$): In this case, dispersion does not depend on the level of TM enrollment, as the first term of equation (2) is not a function of n and the second term vanishes. Dispersion is strictly lower than in case #2, as the second term on the right-hand-side of equation (2) now vanishes, reflecting the fact that mismatches between the data reflected in an ACO's benchmark and the data it observes no longer occur. On the other hand, provided that $w < \theta(n)$, which is likely the empirically relevant case as described above, dispersion is strictly higher than in case #1, reflecting the fact that when the ACO has more information about its latent spending level, it is more likely to perceive its benchmark as unachievable.[2]

Uncertainty of financial performance

I now examine how uncertain an ACO will be about its financial performance under an ACO contract. The facts about the distribution of $S(k) \mid \mathcal{I}(m)$ derived at the outset imply that the ACO's uncertainty about its per capita savings is given by

$$\mathrm{Var}[S(k) \mid \mathcal{I}(m)] = \bar{\sigma}(n)^2 + \big(1 - \theta(m)\big)\tau^2. \tag{3}$$

This uncertainty has two components. The first term reflects the variability in the ACO's average TM spending, while the second reflects the ACO's uncertainty about its latent spending level.

It is unambiguous that the ACO's uncertainty about its per capita savings rises as TM enrollment falls. In particular, the first term on the right side of equation (3) rises in all cases, while the second term rises as well if $m = n$ (since a smaller value of m pushes $\theta(m)$ farther below 1).

It is worth noting, however, that while the ACO's *per capita* savings become more variable as TM enrollment falls, that is not true of the aggregate savings nS:

$$\mathrm{Var}[nS(k) \mid \mathcal{I}(m)] = n\sigma^2 + n^2\big(1 - \theta(m)\big)\tau^2.$$

It is straightforward to show that both terms shrink as TM enrollment shrinks.[3]

---

[2] This can be shown formally through arguments essentially identical to those used in case #1.

[3] This is obvious for the first term. For the second term, it is obvious when $m = N$. When $m = n$, observe that the second term can be written as $n\tau^2\sigma^2[\sigma^2/n + \tau^2]^{-1}$, which is unambiguously increasing in n.