Adjusting Capitation Rates for Risk: An Equivalent Martingale Pricing Approach

Robert T. Maurer

Graduate Program in Health Care Administration
Texas Woman's University
1810 Inwood Rd.
Dallas, Texas 75235

Tel: 214-689-6560
Fax: 214-689-6591
E-mail: rmaurer@twu.edu

Working Paper Version: 6-12-01

Adjusting Capitation Rates for Risk: An Equivalent Martingale Pricing Approach

Abstract

While the use of capitated reimbursement encourages efficiency in the production of health care services, it also creates incentives for health plans to select favorable risks. Incentives to select risks exist when a plan's compensation depends on the financial performance of a risk pool and the plan has an ability, through marketing strategies or other means, to alter the probability distribution of the expected payoff by selectively enrolling favorable risks or by discriminating against unfavorable risks.

Risk selection is an information asymmetry problem in which the health plan exploits its differential ability, relative to the plan sponsor, to predict risk. From the sponsor's perspective, the problem becomes one of eliminating predictable sources of priced risk from the pricing scheme, in order to make the health plan's expected return from the arrangement a random variable. Generally speaking, the normative approach uses risk proxies to detect predictable risk and then groups risks into homogeneous risk categories with risk-adjusted capitation rates. To make this approach work, the sponsor must find observable risk proxies that are also good risk predictors. However, there is some doubt whether observable risk proxies can capture all of the predictable economically-priced risk.

This paper uses financial engineering techniques to formulate a conceptually different approach to the problem that makes the capitation rate a 'fair game' over time. The paper makes three contributions. First, in contrast to the normative static approach, the paper uses an intertemporal framework to capture the evolving nature of the problem over time. Second, since the basic idea underlying risk adjustment is the use of conditioning information, the paper uses a probability measure space framework to emphasize the role of conditioning information in pricing risk. And third, the paper uses the financial theory of martingale pricing processes to derive a risk-adjusted retrospective price as a function of the prospective price conditioned on information revealed by the health plan's actions about changes in the distribution of risk in the population. The paper also shows how this model leads to a straightforward, easily implemented pricing scheme.

Adjusting Capitation Rates for Risk:
An Equivalent Martingale Pricing Approach

1. Introduction

The principal rationale for reimbursing health service plans using risk-based compensation arrangements is to provide the plan with incentives to produce services efficiently. Newhouse (1996) points out, however, that a trade-off exists between desirable incentives that induce the plan to produce efficiently, and undesirable incentives that induce the plan to engage in risk selection.

Risk selection occurs when the health plan uses selective marketing strategies, or other means, to encourage good risks to enroll and to discourage poor risks. The health plan can engage in risk selection only when information asymmetry, over the expected cost of serving enrollees, exists between the sponsor (the agent for the enrollee) and the health plan (van de Ven and Ellis, 1999). In the type of compensation arrangement considered here, the sponsor pays all or most of the enrollee's insurance premium in the form of a capitation fee and the plan's income is a function of the financial performance of the pool of sponsored risks covered by the plan.

In this type of arrangement, the health plan has an incentive to select for good risks when it has the ability to influence the distribution of risks in the risk pool. However, the health plan's estimate of the risk distribution and the estimate used by the sponsor to set the compensation rate must differ. In other words, the health plan must be better informed about the distribution of risks in the covered population than the sponsor. Otherwise, the sponsor could adjust the plan's capitation rate to compensate for the change in the distribution of risks.

From a financial perspective, risk selection looks like an 'arbitrage' strategy in which the health plan takes a short position in the poor risks and an offsetting long position in the good risks. To eliminate the arbitrage opportunity, and discourage the health plan from engaging in risk selection, the sponsor must eliminate predictable sources of priced risk from the pricing scheme. If the sponsor can do so successfully, then the health plan's expected return from serving the sponsor's beneficiaries behaves as a random variable and the incentive diminishes.

Generally speaking, risk adjustment strategies rely on the use of risk proxies to measure unobservable risk-generating state variables in the covered population, and then group the covered risks into homogeneous risk categories with risk-adjusted capitation rates. To make this approach work, the sponsor must find observable risk proxies that are also good risk predictors. Consequently, the academic literature that addresses this problem has focused on the use of clinical and diagnostic information to predict resource consumption (Ash et al., 1989; Ellis et al., 1996; Weiner et al., 1996). One example of this approach is the Medicare adjusted average per-capita cost (AAPCC) methodology. Initially, the AAPCC employed only demographic variables as risk proxies, which capture very little of the predictable variation in costs, but the AAPCC methodology was revised recently to include health status information as well.

There is some doubt, however, whether observable risk proxies capture enough of the predictable economically-priced risk to eliminate selection incentives *ex ante* (Newhouse, 1996; Newhouse, 1998; van de Ven and Ellis, 1999). Since prospective rate adjustment is not likely to eliminate incentives to select risks, Newhouse (1996) suggests that mixed systems containing both prospective and retrospective components have the potential to perform better. This paper uses financial engineering techniques to construct a mixed strategy for adjusting prospective capitation rates for the effects of risk selection.

Since the sponsor's objective is to maintain risk neutral selection incentives, the risk adjustment strategy starts from the assumption that the health plan's expected return from the arrangement should follow a *martingale*, meaning that it remains a 'fair game' over time with respect to the benchmark distribution used to set the prospective capitation rate. Given that the prospective rate provides risk neutral incentives, then the plan's incentive to engage in risk selection remains risk neutral as long as the plan's reward-to-risk ratio remains constant over time. This is because a constant reward-to-risk ratio eliminates the plan's expectation of unpriced future gains from previously selected risks. The basic concept behind the pricing strategy is the use of a risk hedge to adjust the capitation rate periodically for changes in the distribution of risks relative to the benchmark distribution. Since these adjustments maintain a constant reward-to-risk ratio relative to the benchmark ratio, selection incentives remain risk neutral.

This paper makes two contributions to the analysis of the risk selection problem. First, it develops the foundations of the problem using a probability measure space as the mathematical framework. Since a probability measure space framework emphasizes the role of conditioning information in forming expectations, and the basic idea that underlies risk adjustment is the use of conditioning information to price risk, it provides a natural framework to use for a formal analysis of the problem.[1] Second, the paper approaches the problem of risk adjustment from the perspective of capital market pricing theory and the theory of martingale pricing processes in particular. This theory is well suited for this purpose because it treats asset expected returns as martingales under an equivalent risk neutral probability distribution.

The paper is organized as follows. Section 2 develops the mathematical structure for the analysis of the risk selection problem using a probability measure space framework. Section 3 develops the risk selection problem and outlines the theoretical basis for the risk adjustment strategy. Section 4 shows how to take the theoretical idea and implement it in the context of a managed care contract. Section 5 presents concluding comments.

2. Contract Pricing Structure

2.1 Expected Cost Models

In a capitation arrangement, the principal source of financial risk is unexpected variation in the average cost of supplying services. Typically, the expected average cost of supplying services over the life of a *T*-period contract, K_{iT}, is estimated as the sum of the expected average cost in each period,

$$K_{iT} = \sum_{t=1}^{T} \mathrm{E}\left[k_{it} \mid \Phi_t\right], \qquad (1)$$

where k_{it} is the average cost of services consumed by enrollee *i*, or enrollee group *i*, in period *t*, and $\mathrm{E}\left[\ k_{it} \mid \Phi_t\right]$ is the expected average cost of period *t* consumption based on the information Φ_t known about the process that generates service costs in the

[1] Since the mathematics of probability measure theory is very abstract, technical detail with no impact on the development of the argument is omitted from the discussion.

population at the beginning of period t. In general, the average cost expected in each period is estimated using a statistical model of the following type,

$$k_{it} = \mathrm{E}\left[k_{it} \mid \Phi_t\right] + \varepsilon_{it} \tag{2}$$

where ε_{it} is the unexpected component of average cost and is assumed to be random. The unexpected component of the expected average cost accounts for unanticipated changes in average cost related to underlying economic factors (such as cost inflation or the introduction of a new technology), biological factors (reflected in the genetic makeup of the covered population), and ecological factors (such as the emergence of a more virulent strain of the influenza virus) affecting the population.

In practice, the latent risk factors are unobservable and expected cost models must use risk proxies to capture the influence of the risk factors on the cost of consuming services. Frequently, demographic variables (such as age and sex) are used to predict risk, but more sophisticated risk proxies also have been developed using prior expenditure patterns and diagnostic information (Newhouse et al., 1989; Ash et al., 1989; Ellis et al., 1996; Weiner et al., 1996).

In general, the effects of the latent risk factors on the average cost are assumed to be independent. Thus, the overall risk is the sum of the individual risks, and cost prediction models frequently take the form of a linear projection of average cost onto the latent risk factors, as in Newhouse (1996),

$$k_{it} = a + \sum_{j=1}^{M} b_{ij} x_{jit} + u_i + e_{it}\ , \tag{3}$$

where a is the fixed component of average cost and includes routine administrative and service costs, and j indexes M proxies x_{jit} that measure the plan's exposure to the underlying risk factors and M coefficients b_{ij} that project the average cost onto the risk proxies. The random effects term u_i is a disturbance term included under the assumption that group specific fixed costs are distributed randomly across enrollee groups, and e_{it} is a disturbance term distributed randomly across the enrollee groups and over time.

The use of risk proxies to predict the consumption of covered services, however, entails a number of problems (Newhouse, 1998). The most intractable is the information asymmetry that gives rise to risk selection in the first place. Since the health plan always has better information than the sponsor about the expected consumption of services in the covered population, incentives to select favorable risks always exist.

Furthermore, even a prospective risk-adjusted pricing scheme that explains all of the predictable variation in cost is inadequate, because subsequent random events in the population will change the distribution of risk and necessitate pricing adjustments to keep incentives properly aligned. For example, a health plan paid a prospective capitation fee has an incentive to persuade enrollees who later become HIV positive to disenroll. In addition, the information differential required to create an incentive does not need to be large. Since the variance of cost is usually large, and the expected gain from engaging in risk selection typically depends on the standard deviation of cost, even a small amount of additional information can yield substantive gains (Newhouse et al., 1989).

An alternative approach to the risk adjustment problem is to focus on changes in the distribution of risks over time, rather than on uncovering information about the behavior of the latent risk factors. In this case, we are interested in how information induces changes in the expected cost process over time and in statistical models of the following type,

$$\Delta k_{it} = \mathrm{E}\left[\Delta k_{it} \mid \Phi_t\right] + \Delta\varepsilon_{it}, \tag{4}$$

where Δk_{it} is the change in average cost and $\Delta\varepsilon_{it}$ is the change in unexpected cost. This means that we must work with stochastic process models of the expected average cost.

2.2 A State Space Model of the Information Structure

At a fundamental level, risk-based managed care contracts can be thought of as bundles of state contingent claims written on the future consumption of health care services. The expected cost of these future claims is a function of the information available to the health plan, and the sponsor, about the distribution of covered risks in the population. So, in this section, we develop a mathematical model of the information structure and how it evolves over time.

Given a managed care contract, let $(\Omega, \mathbf{F}_t, P)$ be a probability measure space where Ω denotes a sample space of possible outcomes,[2] $\mathbf{F}_t$ denotes the information

[2] The sample space is assumed to be at most countably infinite.

available at the beginning of period t about the outcome (which remains unknown until the contract ends), and P is a probability measure that describes the distribution of outcomes in the sample space. Assume that the contract is subdivided into a series of $t = \{1,\ldots,T\}$ consecutive time periods, and that M linearly independent risk-generating factors determine the period t cost of the services supplied by the health plan. Then, at the end of the contract, an outcome $\omega \in \Omega$ is a M-length vector of risk factor realizations that jointly determine the cost of the contract.[3]

In this framework, an event in the population, such as a plan member presenting with a disease symptom or an outbreak of contagious disease, is represented by an element of the σ-field **E** defined on the outcome space. In mathematical terms, a σ-field is a collection of subsets of the sample space with the following properties: (1) it contains the sample space as a member, (2) it is closed under the formation of complements, and (3) it is closed under the formation of countable unions of the subsets. In practical terms, defining an event as a member of a σ-field ensures that the information structure is closed with respect to the underlying risk factors. Thus, we can think of **E** as the collection of potentially observable, or possible, events.[4]

In this framework, information about the state of the unobserved risk factors is revealed through observations of the average cost of services supplied to the covered population rather than through direct observation of the risk factors themselves. To formalize the concept of an indirect information process, define $\kappa(\omega):\Omega \to \mathbf{R}^+$ to be a nonnegative real-valued function that maps elements of the outcome space onto the

[3] Note that, if the outcome space is spanned by the collection of contingent claims on future resource consumption that define the health plan, and the unit of analysis is an individual plan member, then ω is a covered claim. If the unit of analysis is a group with N plan members, then each outcome ω is represented by an $M \times N$ matrix of factor risk realizations, and ω is a collection of claims that characterize the claims experience of the group as a whole.

[4] To illustrate the basic idea, suppose that we toss a coin once. Then the sample space is $\Omega = \{H,T\}$ and the collection of events is the σ-field $\mathbf{E} = \{\varnothing,\{H\},\{T\},\{H,T\}\}$. Alternatively, suppose that we toss the coin twice. In that case, the sample space is $\Omega = \{HH,HT,TH,TT\}$ and **E** is the much larger collection of all the subsets of Ω.

average cost variable.[5] Then, use this function to partition the outcome space into a collection of observable events and those events that can be excluded from consideration based on the observed realization of the average cost variable.

In formal terms, the average cost k_{it} observed in period *t* reveals information about the state of the underlying risk factors, if and only if the set of risk factor outcomes $\{\omega \in \Omega : \kappa(\omega) = k_{it}\}$ is a member of the set of observable events **E**. Intuitively, this says that an average cost observation is the result of an experiment (in this case, an accounting of the average cost of supplying services to a population of plan members over some period of time), and the information obtained from the experiment partitions the outcome space into a collection of possible outcomes $\{\omega \in \Omega : \kappa(\omega) = k_{it}\}$, and a collection of outcomes which are impossible $\{\omega \in \Omega : \kappa(\omega) \neq k_{it}\}$. Since the observed average cost in period *t* depends on the unknown end-of-contract random outcome ω, average cost is also a random variable.

The information at the beginning of period *t* used to generate expectations of the end-of-period cost, denoted by $\mathbf{F}_t$, is given by the smallest σ-field that can be generated on the sample space by the sequence of cost observations up through period *t*-1. In formal terms, $\mathbf{F}_t$ is generated by the intersection of all the σ-fields on Ω that contain the set of observable events as a subset, and can be thought of as the collection of decidable events when period *t* begins.[6] Note that the sequence of information fields is increasing over time, in the sense that $\mathbf{F}_{t-1} \subseteq \mathbf{F}_t$; consequently, the σ-field $\mathbf{F}_t$ incorporates all of the information in the filtration $\{\mathbf{F}_0, \mathbf{F}_1, \ldots, \mathbf{F}_{t-1}, \mathbf{F}_t\}$ generated by past observations of the average cost variable k_{it}. As cost information accumulates, the information filtration defines an increasingly fine partition on the outcome space until the average cost under the contract is revealed at the end of period *T*.

---

[5] The regression in (3), for instance, defines a function of this type.

[6] Suppose we toss a coin twice and the first toss comes up ‘tails’. Then the information field $\mathbf{F}_1 = \{\varnothing, \{HH, HT\}, \{TH, TT\}, \Omega\}$ consists of events that we know to be either true or false. Note that the event $A = \{HT, TT\}$, ‘tails’ on the second toss, is not decidable and therefore is not a member of $\mathbf{F}_1$.

2.3 Expected Cost Model

Given a model of the information structure at the beginning of time *t*, we need a model for how claims on the underlying events are priced. This requires a probability measure. Let *P* be a probability measure that assigns a probability to any event *A* defined on the outcome space Ω.[7] Suppose a particular outcome $\omega \in \Omega$ is chosen at random, then $P(A)$ is the probability that the outcome ω is included in the collection of outcomes that define the event *A*.

For example, suppose that a health plan member with sub-clinical viral hepatitis (the unknown outcome) presents with a mild fever of unknown origin. Then, there is some probability that on a routine test screen that the individual will test positive with an abnormal hepatic enzyme level (the event). The enzyme test in and of itself, however, is insufficient to distinguish between viral hepatitis, toxic hepatitis, amebic hepatitis, some other disease, or even whether clinical hepatitis will develop at all (the collection of possible outcomes in the event set). The probability measure assigns probabilities to each of the possible outcomes and the triple $(\Omega, \mathbf{F}_t, P)$, together with the cost function $\kappa(\omega)$, defines the contract pricing structure at the beginning of period *t*.

The end-of-period *t* expected average cost, $K_{it} \equiv E\left[k_{it} \,\middle|\, \mathbf{F}_t\right]$, is the conditional expectation of the average cost function with respect to the information available at the beginning of the period. Since expected average cost depends on the unknown outcome ω through the average cost function, it can be defined as follows,

$$\int_A K_{it}\, dP(\omega) = \int_A \kappa(\omega)\, dP(\omega) \quad \text{for all } A \in \mathbf{F}_t\,. \tag{5}$$

This definition is somewhat abstract and depends on the fact that the expected average cost function is measurable on $\mathbf{F}_t$, meaning roughly that the time series leading to the most recent average cost observation k_{it-1} defines a partition of the sample space into disjoint subsets on which the average cost function is constant for each ω in the subset.

---

[7] Formally, the probability measure *P* assigns a number $0 \le P(A) \le 1$ to each event *A* in **E** such that: $P(\varnothing) = 0$, $P(\Omega) = 1$, and the probability assigned to the union of a sequence of disjoint events is the sum of the individual probabilities. The measure *P* is assumed to be stationary over the term of the contract.

Since each event in $\mathbf{F}_t$ is some union of these subsets, the expected cost conditioned on the event *A* is a weighted sum of these subsets that is constant on the event *A*.

For example, a patient with an elevated hepatic enzyme level may develop viral hepatitis, toxic hepatitis, amebic hepatitis, some other disease, or no disease at all. Given a positive test result, the only certain information is that the enzyme levels are elevated and the case has incurred certain costs to this point. Since any forecast of future costs is a probability weighted average of the possible alternative outcomes, the cost forecast is a constant with respect to each of the alternatives used in constructing the forecast.

To define expected average cost operationally, we observe an event *A*. Then, the conditional probability that the outcome ω is responsible for the event *A* is given by

$$P\left[\omega \middle| A \in \mathbf{F}_t\right] = \frac{1}{P(A)} \int_{\Omega} 1_A(\omega)\, dP(\omega)\,, \tag{6}$$

where $1_A(\omega)$ is an indicator function taking on the value one when ω is in the set of outcomes associated with the event *A* and zero otherwise. Then, using the definition of conditional probability in (6), expected average cost is given by the more intuitive expression,

$$K_{it} = \int_A \kappa(\omega)\, dP\left[\omega \middle| A \in \mathbf{F}_t\right]. \tag{7}$$

This is the definition that we will use later. However, note that the definition of expected average cost in (7) is less general since it depends on the event *A*, while the definition in (5) holds generally for all *A* in $\mathbf{F}_t$.

2.3 Compensation Policy

The final element of the pricing strategy is the compensation policy. We focus on the case where the sponsor pays the health plan a capitation fee in each period, but allow the sponsor to vary the plan's fee from one period to the next. Since the fee compensates the health plan for the expected cost of services in each period, we need an expression for the expected average per period cost in each period.

Let $\overline{K}_{it}$ be the expected average per period cost of providing services at the start of period *t* under a *T*-period contract, then

$$\overline{K}_{it} = \frac{1}{T}\left(\sum_{j=1}^{t-1} k_{ij} + \sum_{j=t}^{T} E\left[K_{ij} \mid \mathbf{F}_t\right]\right). \tag{8}$$

The expression in (7) says that the expected average per period cost is a combination of the average costs incurred in past periods plus the expected average costs to be incurred in future periods with the total sum averaged over the full term of the contract.

In general, the prospective compensation rate $C_i(\cdot)$ is a linear combination of the expected per period costs at the beginning of the contract. In this case, for a group *i* with *m* members, the prospective per period compensation rate is given by

$$C_i\left(\overline{K}_{i1}\right) = \frac{1}{mT}\left(\sum_{t=1}^{T} e^{-\delta t}\, E\left[K_{it} \mid \mathbf{F}_1\right] + C_{i0}\right), \tag{9}$$

where δ is a discount rate that adjusts expected future cash flows for the opportunity cost of capital, and C_{i0} captures any indirect costs not included in the service cost estimates. Given that the expected cost estimates in each future period are actuarially fair rates, then the prospective rate is also an actuarial fair rate.

2.4 Expected Cost Model Example

To fix these ideas consider the following example. Suppose a health plan covers only two conditions: viral influenza and pneumonia secondary to influenza. Suppose also that whether one contracts influenza is a completely random event, but once a person has contracted influenza, the probability that one also contracts pneumonia then depends on a biological factor associated with age. Thus, there is a single risk factor in the population for which an individual's age is an observable risk proxy. Assume that the health plan initially enrolls a only single plan member. In this case, the sample space contains three outcomes:

$\Omega = \{\omega_1, \omega_2, \omega_3\}$,

for $\omega_1 =$ no clinical disease,

$\omega_2 =$ influenza,

$\omega_3 =$ pneumonia secondary to influenza,

and the σ-field of observable events defined on Ω is given by

$$\mathbf{E} = \{\varnothing, \{\omega_1\}, \{\omega_2\}, \{\omega_3\}, \{\omega_1, \omega_2\}, \{\omega_1, \omega_3\}, \{\omega_2, \omega_3\}, \Omega\}.$$

Whether an individual contracts influenza is a completely random event, but given that an individual has influenza, then the person's age affects the probability that the person also contracts pneumonia.

Assume that the contract is divided into three sub-periods, and that during the first sub-period the plan member presents with a mild cough, denoted as A_1. Since this information is consistent with all three outcomes in the sample space, it reveals nothing new, so $A_1 = \varnothing$ and the σ-field generated by A_1 is $\mathbf{F}_2 = \mathbf{F}_1 = \{\varnothing, \Omega\}$.[8] Now, in the second period, suppose that the plan member receives some diagnostic tests which rule out a diagnosis of pneumonia. Then, at the end of period 2, the σ-field of decidable events generated by the event $A_2 = \{\omega_1, \omega_2\}$ is given by $\mathbf{F}_3 = \{\varnothing, \{\omega_3\}, \{\omega_1, \omega_2\}, \Omega\}$. Finally, during period 3, suppose that additional diagnostic tests result in a definitive diagnosis of viral influenza. At the end of period 3, then, the σ-field of decidable events generated by the event $A_3 = \{\omega_2\}$ is given by $\mathbf{F}_4 = \mathbf{E}$.

The objective is to construct a forecast of the expected average cost in the next period given the information available at the end of the current period. The probability measure $P(\omega_n)$, the cost of each outcome when the contract terminates $k_T(\omega_n)$, the conditional probabilities $P(\omega_n \mid \mathbf{F}_t)$, and the expected average cost in the next period $E[\kappa(\omega) \mid \mathbf{F}_t]$ are given in Table 1.

Table 1

ω_n	$P(\omega_n)$	$k_T(\omega_n)$	$P(\omega_n \mid \mathbf{F}_2)$	$P(\omega_n \mid \mathbf{F}_3)$	$P(\omega_n \mid \mathbf{E})$
ω_1	0.60	\$100	0.60	0.75	0.00
ω_2	0.20	\$142	0.20	0.25	1.00
ω_3	0.20	\$1,000	0.20	0.00	0.00
$E[\kappa(\omega) \mid \mathbf{F}_t]$			\$96.13	\$36.83	\$47.33

In Table 1, the cost function $\kappa(\omega)$ allocates one third of the expected terminal cost of the contract to each period, which is equivalent to estimating an average cost on the usual per member per month basis. The \$100 cost associated with the no clinical disease outcome

8 Recall that the information field is defined relative to the beginning of period t, thus the information in $\mathbf{F}_t$ is defined with respect to period t-1.

ω_1 represents the fixed costs of maintaining the health plan and includes the cost of routine services provided to a typical plan member. If the final diagnosis is the influenza outcome ω_2, the plan member receives palliative measures, and the plan incurs $42 in additional costs. If the final diagnosis is bacterial pneumonia, the plan member is admitted to the hospital, treated with intravenous antibiotics, and the plan incurs $900 in additional costs.

Using the unconditional probability distribution in Table 1, the plan's expected average per period cost at the beginning of the contract is $96.13. Since the information obtained during period 1 adds nothing to what is already known, the period 1 forecast of the next period average cost conditioned on this information remains $96.13. In period 2, the new information rules out the most expensive outcome (pneumonia) and this reduces the expected next period average cost to $36.83. In period 3, the true state of nature becomes known and the total cost to the plan is revealed to be $142, resulting in an average per period cost of $47.33.

3. The Risk Adjustment Strategy

The economic returns from engaging in risk selection depend on differences in the information available to the health plan and the sponsor about the distribution of risk in the covered population. In the context of a capital market, informational efficiency implies that the equilibrium expected rate of return incorporates all relevant information available to the market, and this in turn implies that the expected rate of return on any asset follows a martingale under the corresponding risk-neutral probability distribution (Harrison and Kreps, 1979).

In this section, we borrow this concept of information efficiency, and use it to require the health plan's expected rate of return follow a martingale with respect to the corresponding risk-neutral probability distribution. However, in the case of the health plan, the expected rate of return is set by the terms of the plan's contract with the sponsor and not through the market. Thus, in this section, we start with the idea that the compensation policy should follow a martingale with respect to a risk-neutral probability distribution (in other words, assume an information-efficient competitive market equilibrium), and then work backwards to show how to construct a compensation policy that behaves as if it was the result of an information-efficient competitive market process.

3.1 Information Asymmetry and Risk Selection

For risk selection to yield an excess return over the return contemplated in the contract, the health plan must have superior information about the probability distribution of average cost in the covered population. Expressed in terms of the model developed above, the health plan has superior information when the probability measure P^+ (that it uses to construct average cost projections) stochastically dominates the probability measure *P* (that was used to construct the compensation policy) in the sense that the following condition is satisfied,[9]

$$K_{it}^{+} \equiv E_{P^{+}}\left[k_{it} \middle| \mathbf{F}_t\right] < E_P\left[k_{it} \middle| \mathbf{F}_t\right] \equiv K_{it}\,, \tag{10}$$

where the subscripts on the expectation operators indicate that the expectations are taken with respect to the given probability measure.

The condition in (10) says that the plan's probability measure results in a lower expected average cost than does the probability measure used by the sponsor to construct the compensation policy, even though the information filtration available to both is identical. Since $\mathbf{F}_1 \subseteq \mathbf{F}_t$, the condition in (10) implies that a health plan that engages in risk selection expects the average cost to be strictly less than the average cost estimate used to set the prospective compensation rate prior to period 1. Thus, an opportunity to profit from risk selection is available. Also, note that, if the compensation policy is actuarially fair in the sense that the provider is compensated for the fair value of the risk assumed, then *P* defines a risk-neutral measure that can be used as a pricing benchmark.

3.2 Fair Compensation and the Martingale Property

A risk-based managed care contract is considered a fair contract if it has the *fair game* property which implies that

---

[9] To be precise, the condition in (10) is satisfied *almost everywhere*. This means that the two expectations in (10) agree only on sets with probability measure zero, which would only be the case for a group that cannot receive services, either because the group is not now covered by the plan, or because none of the group's members are alive.

$$E\left[E\left[k_{it+1} \mid \mathbf{F}_{t+1}\right] - k_{it} \mid \mathbf{F}_t\right] = 0\,. \tag{11}$$

Essentially, the condition in (11) says that a managed care contract is considered to be a fair contract if exchanging k_{it} units of consumption in this period for $E\left[k_{it+1} \mid \mathbf{F}_{t+1}\right]$ units of consumption in the next period results in no expected change in the average cost. The fair game property in turn implies the *martingale* property

$$E\left[E\left[k_{it+1} \mid \mathbf{F}_{t+1}\right] \mid \mathbf{F}_t\right] = K_{it}\,. \tag{12}$$

The condition in (12) says that the best estimate of the expected average cost in the next period, based on currently available information, is the estimate of the expected average cost for the current period.

The martingale property implies that future changes in expected average cost are not predictable. If changes in the expected average cost cannot be predicted, then it is not possible for the health plan to generate reliable excess returns from engaging in risk selection. Thus, a compensation policy is a fair game if the expected average per period cost process $\overline{K}_{it}$ is a martingale over time. That this is the case here is demonstrated in the Appendix at the end of the paper.

3.3 An Equivalent Martingale Pricing Strategy

Now that the basic elements of the model are in place, the problem is to construct a compensation policy that follows a martingale. Since this requires the compensation policy to be a fair game over time, we assume that the contract contains an *ex post* settling-up provision giving the sponsor the right to adjust the compensation rate in the next period for changes in the distribution of risk in the covered population in this period. The objective is to devise a risk-adjustment policy that adjusts the compensation rate so that the plan's expected return in the next period depends on the benchmark measure P, used to negotiate the compensation rate, rather than on the population measure P^+ which can be influenced by the plan.

Information about changes in the distribution of risk in the population is contained in the information filtration $\mathbf{F}_t$. The analytical tool we use for extracting this information is the Radon-Nikodym theorem (Billingsley, 1986, Sect. 32), which guarantees the existence of a random variable $\xi(\omega) = \frac{dP^{+}(\omega)}{dP(\omega)}$ (the Radon-Nikodym derivative) defined by the following decomposition of the plan's expected average cost,

$$K_{it}^{+} = \int_A \kappa(\omega) dP^{+}(\omega) = \int_A \kappa(\omega)\xi(\omega) dP(\omega). \tag{13}$$

This decomposition requires the measures *P* and P^{+} to be *equivalent* in the sense that they agree on sets of measure zero, as is the case here.[10]

Given the Radon-Nikodym derivative of the population measure with respect to the benchmark, it can be used to extract information about changes in the distribution of risks as follows,

$$\begin{aligned} K_{it}^{+} &= E_{P^{+}}\left[k_{it} \mid \mathbf{F}_t\right] \\ &= E_P\left[\xi(\omega) \mid \mathbf{F}_t\right]^{-1} E_P\left[E_P\left[\xi(\omega) \mid \mathbf{F}_{t+1}\right] k_{it} \mid \mathbf{F}_t\right] \\ &= E_P\left[\frac{E_P\left[\xi(\omega) \mid \mathbf{F}_{t+1}\right]}{E_P\left[\xi(\omega) \mid \mathbf{F}_t\right]} k_{it} \mid \mathbf{F}_t\right] \\ &= E_P\left[w_t k_{it} \mid \mathbf{F}_t\right]. \end{aligned} \tag{14}$$

The first line in the decomposition is just the definition of expected average cost. The second line invokes a common strategy[11] for changing the probability measure in an expectation, in this case, from the population measure P^{+} to the benchmark measure *P*. The third line reveals the underlying logic more clearly, showing that the basic strategy involves a rescaling of the of the end-of-period average cost with the ratio of the expected change in the population measure relative to the benchmark measure. The fourth line emphasizes the fact that the scaling factor w_t is a random variable to be estimated.

The decomposition in (14) provides the theoretical foundation for the adjustment method developed in the next section. The method starts from the assumption that the expected average cost used to construct the prospective compensation policy promotes first-best risk sharing between the plan and the sponsor. However, information asymmetry leaves the plan with incentives to select risks. Since the sponsor wants to maintain optimal incentives, the sponsor must adjust the prospective capitation rate at the

---

[10] In other words, the plan sponsor and the plan both agree on the impossible events, but assign different probabilities to the possible events.
[11] See Harrison (1985), Proposition 3, page 10.

beginning of each period for changes in the distribution of risk induced by risk selection on the part of the plan.

Suppose that the plan responds to incentives to select risks, then the sponsor needs an estimate of the health plan's expectation of average cost in the next period $\hat{K}_{it}^{+}$ in order to adjust the capitation rate for the effects of risk selection. Given a cost generating event A, the operational definition of expected average cost implies that

$$\begin{aligned} \hat{K}_{it}^{+} &= E_P\left[w_t k_{it} \mid A \in \mathbf{F}_t\right] \\ &= E_P\left[w_t \mid A \in \mathbf{F}_t\right] E_P\left[k_{it} \mid A \in \mathbf{F}_t\right] \\ &= \hat{w}_t K_{it}, \end{aligned} \tag{15}$$

where the second line in (15) follows from the fact that the conditional expectation of the risk scaling factor is constant on the event A. In the next section, we show how to obtain an estimate of the risk scaling factor $\hat{w}_t$.

4. Implementing the Risk Adjustment Method

The pricing strategy developed in the previous section is closely related to the methods used to price derivative securities in financial markets. In this case, the underlying asset from which the managed care contract derives its value is the expected cost of the investment in a covered life. Since the prospective compensation rate is actuarially fair, given the expected composition of the covered population, the objective of the risk-adjustment procedure is to eliminate the provider's expectation of unearned gains from risk selection. If we view risk selection as a form of 'risk arbitrage' on the part of the health plan, then this is equivalent to enforcing a 'no arbitrage' condition on the plan's expected return from the contract.

The objective is to adjust the prospective rate for changes in outcome risk in the population relative to the benchmark risk measure using the estimated risk scaling factor $\hat{w}_t$. Since the pricing adjustment restores optimal risk sharing relative to the distribution of risk in the population, the contract is again properly priced. In effect, the prospective rate is the price of outcome risk agreed to under the contract and the scaling factor is the change in outcome risk relative to the benchmark. Since provider incentives to select risks cannot be eliminated completely in a contract of finite length, the pricing adjustment is made at the end of each contract period in an effort to maintain an *ex ante* risk-neutral pricing structure.

4.1 Uncertainty over expected cost

Typically, health services contract pricing models introduce uncertainty over cost through the probability distribution of the terminal service costs. Expected cost in our model is a stochastic process, so we introduce uncertainty over the expected average cost of providing services using a continuous-time diffusion process. Since the end-of-period cost of providing services results from the accumulation of a large number of incremental cost events in the covered population over the course of the period, the end-of-period expected average cost can be thought of as the accumulation of incremental expected cost events that build expected average cost over time. The process is represented as follows,

$$K_{it} = \int_{t'}^{t''} \mu_i(s)\,ds + \int_{t'}^{t''} \sigma_i(s)\,dz_i(s)\,, \tag{16}$$

where t' and t'' denote the beginning and ending points of the time interval covered by contract period *t*, $\mu_i(s)$ is the expected increase in cost over a small increment of time *ds*, $\sigma_i(s)$ is the instantaneous standard deviation of the cost process, and $dz_i(s) = \varepsilon_i\sqrt{ds}$ is a standard Weiner process where the random variable ε_i is distributed normally with a zero mean and unit standard deviation.

Since it is more convenient to work with a stochastic integral in differential form, rewrite (16) as follows,

$$dK_{it} = \mu_i(t)\,dt + \sigma_i(t)\,dz_i\,, \tag{17}$$

where dK_{it} denotes the incremental change in expected average cost due to the short interval of time *dt*.[12] Also, since the risk process in the population may have multiple independent risk factors, augment (17) as follows,

$$dK_{it} = \mu_i(t)\,dt + \sum_{j=1}^{M} \sigma_{ij}(t)\,dz_{ij}\,, \tag{18}$$

[12] Recall that *t* indexes contract periods, so time subscripts in (22) are implied and also in the expressions that follow.

where the subscript j indexes the independent risk factors. In general, the parameters of the diffusion can depend on the expected end-of-period average cost and on time, but for simplicity we fix the parameters over each contract period and allow them to vary only across contract periods.

The diffusion process in (18) has two properties that make it a good model for the behavior of uncertainty over the expected average cost through time. First, uncertainty over the expected cost estimates used to set the compensation rate changes proportionally with time,

$$\mathrm{Var}\left(\overline{K}_{it}\right) = \left(1 - \frac{t-1}{T}\right) \mathrm{Var}\left(K_{it}\right), \tag{19}$$

where $1-(t-1/T)$ is the percentage of time remaining in the contract, therefore there is more uncertainty over the cost of the contract early in the contract, as opposed to later in the contract, when more information about the outcome is available. In (18), the variance term is linear in the time variable ($dz_{ij}^2 = \varepsilon_{ij}\, dt$), consequently the variance of the process also changes proportionally with time. Second, the martingale property implies that the current expected average cost is the best possible estimate, therefore the expected average cost process should be non-anticipating. The random variable ε_{ij} is serially uncorrelated over time, consequently the diffusion process is a Markov process, implying that it summarizes all the information about the past behavior of service costs in the current cost estimate.

4.2 Estimating the Scaling Factor

To incorporate multiple subscriber groups in the analysis, let $\mathbf{dK}$ be a vector of expected average cost diffusions, dK_{it}, for i=1 to N subscriber groups,

$$\mathbf{dK} = \boldsymbol{\mu}\, dt + \boldsymbol{\Sigma}\, \mathbf{dz} \tag{20}$$

where $\boldsymbol{\mu}$ is an Nx1 vector of expected incremental costs generated by the contract information structure, $\boldsymbol{\Sigma}$ is an NxM instantaneous variance-covariance matrix, and $\mathbf{dz}$ is a Mx1 vector of standard Weiner processes. Similarly, let $\mathbf{dK}^{+}$ be the process generated by the plan's information structure, where

$$\mathbf{dK}^{+} = \boldsymbol{\mu}^{+}\, dt + \boldsymbol{\Sigma}^{+}\, \mathbf{dz}\,, \tag{21}$$

and the parameters in (21) are defined similar to those in (20). Since the risk generating factors in the population are linearly independent, the elements of $\mathbf{dz}$ are mutually uncorrelated. The process $\mathbf{dK}$ is the pricing benchmark, meaning that the compensation rate is a function of the expected average cost implied by the incremental cost diffusion process. For simplicity, assume a perfectly competitive market for managed care contracts in which the compensation rate is the discounted expected average cost of providing services and that the benchmark process is the least cost process consistent with the desired level of quality.

The pricing strategy depends on the estimation of a risk scaling factor to adjust the expected average cost for changes in the underlying risk measure. Given a diffusion process, as in (20) and (21), changing the probability measure, as in (14), is equivalent to changing the process drift term $\boldsymbol{\mu}\,dt$ (Harrison, 1985, Theorem 9, pg. 10). Thus, to find the risk scaling factor, define an NxM scaling matrix $\mathbf{w}$ such that

$$\mathbf{dK}^{+} = \mathbf{w}\,\mathbf{dK}\,. \tag{22}$$

Since $\mathbf{w}$ scales expected average cost, we need to eliminate terms in $\mathbf{dz}$ from the equation to obtain an expression for $\mathbf{w}$. First, substitute for $\mathbf{dK}$ and $\mathbf{dK}^{+}$ in (22) using (20) and (21), then choose $\mathbf{w} = \boldsymbol{\Sigma}^{+}\boldsymbol{\Sigma}'(\boldsymbol{\Sigma}\boldsymbol{\Sigma}')^{-1}$ to eliminate the terms in $\mathbf{dz}$ and rearrange to obtain the relation

$$\boldsymbol{\mu}^{+} = \boldsymbol{\Sigma}^{+}\boldsymbol{\Sigma}'(\boldsymbol{\Sigma}\boldsymbol{\Sigma}')^{-1}\,\boldsymbol{\mu}\,. \tag{23}$$

As will become clear in the next section, the relation in (23) defines the relationship between the price per unit risk in the contract and the price per unit risk in the covered population.

4.3 Application example

To illustrate the pricing model, consider the following example. Assume that a sponsor enters into an agreement with a health plan to underwrite the cost of providing health services to a homogeneous population whose consumption of health care services is governed by a single latent risk factor. The sponsor reimburses the health plan using a monthly prospective capitation rate and the contract has a term of one year running from January 1 through December 31.

Since the expected cost of supplying services cannot be negative and incremental changes in expected cost per unit time should be proportional to the overall cost, a good model for the expected average cost process is a lognormal diffusion. In this case, the benchmark expected average cost process is given by the diffusion process

$$dK = aK\,dt + sK\,dz\,, \tag{24}$$

and the cost process in the population is given by the diffusion process

$$dK^{+} = a^{+}K^{+}\,dt + s^{+}K^{+}\,dz\,. \tag{25}$$

A lognormal process implies that percentage changes in the expected cost process follow a normal distribution, thus the parameters a, s , and a^{+} , s^{+} , describe the percentage change in the expected average cost that results from extending the contract by a small time interval dt.

In mathematical terms, the diffusion parameters in (24) and (25) are defined for instantaneous changes in expected average cost. Campbell, Lo, and MacKinlay (1997, see pg. 361-363) discuss the problem of obtaining statistical estimates for instantaneous diffusion parameters and show that maximum likelihood estimators can be obtained in this case from the following expressions,

$$\hat{a} = \frac{1}{n\,h}\sum_{i=1}^{n}\log\left(\frac{k_i}{k_{i-1}}\right), \tag{26}$$

$$\hat{s}^2 = \frac{1}{n\,h}\sum_{i=1}^{n}\left(\log\left(\frac{k_i}{k_{i-1}}\right) - \hat{a}h\right)^2, \tag{27}$$

where n is the number of observations, h is the length of the sampling interval, and i is used in this subsection to index the sampling interval.

For example, suppose that the sponsor wants to use information in the monthly distribution of service costs to adjust the capitation rate for selection bias. Given that the sponsor wants to adjust the capitation rate at the end of the 1st Quarter using daily cost observations, then $n = 90$ days and $h = 1$ day. Then, using parameter estimates obtained from (26) and (27) in the pricing relation in (23), we obtain the risk-adjusted benchmark parameter as follows,

$$\hat{a}^{*}_{1Q} = \frac{\hat{s}_{1Q}}{\hat{s}^{+}_{1Q}}\,\hat{a}^{+}_{1Q}\,, \tag{28}$$

which gives us the expected percentage change in cost adjusted for selection bias. The expected cost of supplying services for the month of April is the adjusted percentage change in the benchmark multiplied by the benchmark expected cost: $\hat{a}^*_{1Q}\,\overline{K}_{April}$. Using the compensation policy in (9) to set the capitation rate, and assuming m covered lives in the risk pool, we obtain

$$C^*_{April}\left(\overline{K}_{April}\right) = \frac{1}{m}\left(e^{-\delta(12-3)}\,\hat{a}^*_{1Q}\,\overline{K}_{April} + C_0\right), \tag{29}$$

where δ is the one-month discount rate and C_0 is the estimate of administrative and other non-service related costs.

In the example, the economic rationale behind the expression for the benchmark scaling factor in (23) is easy to see. Rearrange (28), the scaler counterpart to (23), to obtain the equivalence ratios

$$\frac{\hat{a}^*_t}{\hat{s}_t} = \frac{\hat{a}^+_t}{\hat{s}^+_t}, \tag{30}$$

where the numerator in each ratio is the predictable change in cost and the denominator is the uncertainty in the cost estimate. In effect, the ratios measure the expected change in average cost per unit of unpredictable risk in the covered population. The equality in (30) implies that the benchmark change in cost per unit risk should be the same as the change in cost per unit risk in the population. The asterisk on the benchmark parameter indicates that this quantity adjusts to preserve the equality. If the plan selects risks, so that the expected cost-to-risk ratio changes, then the pricing benchmark also changes to preserve the cost-to-risk ratio agreed to under the contract. The expression in (23) is just a more general mathematical statement of this relation.

5. Conclusions

In general, methods for adjusting managed care plan compensation rates for risk selection attempt to replicate the service provider's private information. Although feasible in theory, this approach resembles the sort of cat-and-mouse game that hospitals played with the cost-based reimbursement system. This paper borrows derivative asset pricing techniques from financial economics and applies them to the problem of adjusting the health plan's compensation policy for risk selection. Since the adjustment to the compensation rate depends on the service provider's choice of risks, selection behavior is self-regulating. Of course, there will be random variation in the parameters used to adjust the compensation rate, but for a sufficiently large member base this variation will be small. In addition, it is straightforward to devise statistical confidence limits on when a deviation from the contract benchmark parameters triggers an adjustment in the rate, so it is possible to apply the approach to relatively small member groups.

As Newhouse (1996) advocates, the approach used here combines prospective and retrospective elements. The pricing strategy is prospective in the sense that the benchmark rate defines the contracted price per unit of predictable risk *ex ante*, creating incentives to produce services efficiently. The strategy is retrospective in the sense that the *ex post* adjustments to the price ameliorate incentives to select risks. Finally, the compensation rate adjustments are nonlinear, thus the model takes a more sophisticated view of the interaction between risk and financial return than the traditional model.

6. **References**

Ash, A., F. Porell, L. Gruenberg, E. Sawitz, and A. Beiser. 1989. Adjusting Medicare capitation payments using prior utilization data. *Health Care Financing Review* 10(4): 17-29.

Billingsley, P. 1986. *Probability and Measure, 2nd ed.* John Wiley & Sons. New York, New York.

Campbell, J.Y., Lo, A.W., and MacKinlay, A.C. 1997. *The Econometrics of Financial Markets*. Princeton University Press. Princeton, New Jersey.

Ellis, R. P., G. C. Pope, L. I. Iezzoni, J. Z. Ayanian, D. W. Bates, H. Burstin, and A. S. Ash. 1996. Diagnosis-Based Risk Adjustment for Medicare Capitation Payments. *Health Care Financing Review* 17(3): 101-128.

Ehrlich, I., and Becker, G.S. 1972. Market insurance, self-insurance, and self-protection. *Journal of Political Economy* 80: 623-648

Harrison, J. M. 1985. *Brownian Motion and Stochastic Flow Systems*. John Wiley & Sons, New York.

Harrison, J. M., and D. M. Kreps. 1979. Martingales and arbitrage in multiperiod markets. *Journal of Economic Theory* 20: 381-408.

Holmström, B. 1979. Moral hazard and observability. *Bell Journal of Economics* 10(1): 74-91.

Lambert, R. A. 1983. Long-term contracts and moral hazard. *Bell Journal of Economics* 14(2): 441-452.

LeRoy, S.F. 1989. Efficient capital markets and martingales. Journal of Economic Literature 27: 1583-1621.

Newhouse, J.P., Manning, W.G., Keeler, E.B., and Sloss, E.M. 1989. Adjusting capitation rates using objective health measures and prior utilization. *Health Care Financing Review* 10(3): 41-53.

Newhouse, J.P. 1996. Reimbursing health plans and health providers: Efficiency in production versus selection. *Journal of Economic Literature* 34(3): 1236-1263.

Newhouse, J.P. 1998. Risk adjustment: Where are we now? *Inquiry* 35(2): 122-131.

Pauly, M.V. 1968. The economics of moral hazard: Comment. *American Economic Review* 58: 531-537.

Shavell, S. 1979. Risk sharing and incentives in the principal and agent relationship. *Bell Journal of Economics* 10(1): 55- 73.

van de Ven, W. P. M. M., and R. P. Ellis. 1999. Risk adjustment in competitive health plan markets. Chapter 17 in *Handbook of Health Economics*, eds. A. J. Culyer and J. P. Newhouse.

Weiner, J. P., A. Dobson, S. L. Maxwell, K. Coleman, B. H. Starfield, and G. F. Anderson. 1996. Risk-Adjusted Medicare Capitation Rates Using Ambulatory and Inpatient Diagnoses. *Health Care Financing Review* 17(3): 77-99.

7. Appendix

In this section, we demonstrate that the compensation policy follows a martingale with respect to the expected average cost process. First, we need to demonstrate that $\overline{K}_{it} = \{\overline{K}_{i1}, \ldots, \overline{K}_{it-1}, \overline{K}_{it}\}$, the expected average cost process used to set the compensation policy, is a martingale with respect to a probability measure Q and the filtration $\mathbf{F}_t$ generated by the sequence of average cost observations. This requires that the cost process satisfy four conditions (Billingsley, 1986, Sect. 35), the first three of which are essentially technical in nature. Two of these conditions have been addressed already. The first condition requires an increasing information field which is satisfied by construction. The second condition requires $\overline{K}_{it}$ to be measurable on $\mathbf{F}_t$ which follows from the fact that $\overline{K}_{it}$ is linear in K_{ij} for $j = t, \ldots, T$. The third condition, $E_Q\left[\overline{K}_{it}\right] < \infty$ for all *t*, requires the expected average per period cost to be finite. Since the contract's term is limited to *T* periods, and the sponsor's financial resources must be finite in any period, this condition is satisfied by assumption also. The final condition, that $E_Q\left[\overline{K}_{is} \mid \mathbf{F}_t\right] = \overline{K}_{it}$ for all $t < s \leq T$, is the martingale property. By the definition of K_{it}, and using the property of iterated expectations, we have that

$$\mathrm{E}_Q\left[K_{is} \mid \mathbf{F}_t\right] = \mathrm{E}_Q\left[\mathrm{E}_Q\left[k_{is} \mid \mathbf{F}_s\right] \mid \mathbf{F}_t\right] = \mathrm{E}_Q\left[k_{is} \mid \mathbf{F}_t\right] = K_{it}, \tag{A.1}$$

and, because the realized end-of-period costs up to period *t* in (8) are known, substituting (A.1) in (8) implies that

$$\mathrm{E}_Q\left[\overline{K}_{is} \mid \mathbf{F}_t\right] = \mathrm{E}_Q\left[\frac{1}{T}\left(\sum_{j=1}^{t-1} k_{ij} + \left(T-(t-1)\right)K_{it}\right) \Big| \, \mathbf{F}_t\right] = \overline{K}_{it}, \tag{A.2}$$

which is the martingale property. Thus, the expected average per period cost process is a martingale with respect to the probability measure Q and the filtration $\mathbf{F}_t$ if the process incorporates available information about the distribution of future costs. And, since the compensation policy in (9) is a linear combination of the expected average cost in each period, the compensation policy must inherit the martingale property. Thus, the compensation rate follows a martingale if the expected average cost process follows a martingale.