Hellwig, Martin F. und Schmidt, Klaus M.:

Discrete-Time Approximations of the Holmström-Milgrom Brownian-Motion Model of Intertemporal Incentive Provision

Munich Discussion Paper No. 2001-11

Department of Economics
University of Munich

Volkswirtschaftliche Fakultät
Ludwig-Maximilians-Universität München

Online at https://doi.org/10.5282/ubm/epub.22

Discrete-Time Approximations of the Holmström-Milgrom Brownian-Motion Model of Intertemporal Incentive Provision[1]

Martin F. Hellwig*

University of Mannheim

Klaus M. Schmidt**

University of Munich and CEPR

This draft: July 16, 2001

Abstract: This paper studies the relation between discrete-time and continuous-time principal-agent models. We derive the continuous-time model as a limit of discrete-time models with ever shorter periods and show that optimal incentive schemes in the discrete-time models approximate the optimal incentive scheme in the continuous model, which is *linear in accounts.* Under the additional assumption that the principal observes only cumulative total profits at the end and the agent can destroy profits unnoticed, an incentive scheme that is *linear in total profits* is shown to be approximately optimal in the discrete-time model when the length of the period is small.

Keywords: Principal-agent problems, linear incentive schemes, intertemporal incentive provision, Brownian motion.

JEL classification numbers: C61, D82, J33.

[1]For helpful comments and discussions we are grateful to Darell Duffie, Oliver Hart, Florian Herold, Bengt Holmström, Nobuhiro Kiyotaki, Paul Milgrom, John Moore, Holger Müller, Jae Sung, three anonymous referees and the editor, Drew Fudenberg. We are particularly grateful to Sven Rady for his thorough reading of the paper and valuable suggestions. The first author gratefully acknowledges research support from the Schweizerischer Nationalfonds, the Deutsche Forschungsgemeinschaft, and the Taussig Chair at Harvard University. The second author is grateful for research support from the Deutsche Forschungsgemeinschaft through grants SCHM1196/2-1 and /4-1 and for the hospitality enjoyed at the economics department of Stanford University.

*Department of Economics, University of Mannheim, D-68131 Mannheim, email: hellwig@pool.uni-mannheim.de

**Department of Economics, University of Munich, Ludwigstr. 28 (Rgb.), D-80539 Munich, Germany, email: klaus.schmidt@lrz.uni-muenchen.de

1 Introduction

This paper studies the relation between multi-period discrete-time and continuous-time principal-agent models in the seminal paper of Holmström and Milgrom (1987). For certain intertemporal contracting problems, Holmström and Milgrom have argued that optimal incentive schemes for an agent with constant absolute risk aversion controlling a stationary technology must be linear. We propose to elucidate the structure underlying their linearity results by studying the relation between the discrete-time and the continuous-time models in their paper. We provide an explicit construction establishing the continuous-time model as a limit of discrete-time models with ever shorter periods and establish a suitable continuity result for optimal incentive schemes in the transition between the two. We also clarify the distinction between two kinds of linearity of optimal incentive schemes and provide a discrete-time version of the economically more interesting "linearity in profits" that Holmström and Milgrom establish only for the continuous-time model.

Our work is motivated by two concerns. First, in the introduction of their paper, Holmström and Milgrom argue that the nonlinearities or even discontinuities of optimal incentive schemes that are typical for static principal-agent models are unlikely to be robust to changes in these models that provide the agent with more scope to manipulate information. However, much of their actual analysis does not involve this notion of manipulation at all. Manipulation of information plays no role in either their discrete-time models or their continuous-time model in which the principal observes the entire realization of the outcome process and yet a linearity result is obtained. Manipulation of information does play a role in a version of the continuous-time model, in which the principal is restricted to observing a linear aggregate of the outcome process, where a stronger linearity result is obtained. However, this result has no immediate counterpart in the discrete-time models.

Our second concern relates to the distinction between different kinds of linearity. Most of the linearity results of Holmström and Milgrom refer to counting variables indicating the frequencies of different kinds of outcomes rather than to economic variables such as revenues, costs or profits. To understand the difference, consider their discrete-time models. In the static, i.e. one-period model, the agent chooses a vector of probabilities over a finite set of possible states of the world. The payment the agent receives depends on the outcome, i.e., the realization of the state of the world, according to the incentive scheme that has been agreed on. The m-period model is an m-fold repetition of the static model. In each period, the agent chooses a vector of probabilities over a finite set of possible states of the world for that period; the associated effort cost is assumed to be stationary. Under the assumption that the principal and the agent have constant absolute risk aversion (CARA), a simple m-fold repetition of an optimal incentive scheme for the one-period model is optimal for this m-period model. The overall payment the agent receives then is the sum over the possible period outcomes of the associated one-period incentive payments multiplied by the number of times the outcomes have been realized. The overall incentive scheme is thus a linear function of the frequencies with which the different possible outcomes are observed. The order in which the different outcomes arise plays no role.

This linearity in accounts listing the frequencies of incidence of different outcomes does not in general provide for linearity in outcome variables such as overall revenues, costs or profits. There is no fundamental reason why, say, in a 2-period model, two profit realizations of one should be treated as being equivalent to one profit realization of zero and one profit realization of two. The dependence of payments on outcomes in any period reflects the action the principal wants to implement in that period and the agent's effort cost function, but may have little to do with the revenues, costs or profits that are associated with the different outcomes.

The distinction between the two notions of linearity concerns the continuous-time model as well. In the continuous-time analysis of Holmström and Milgrom, the agent controls the drift rate vector of a multi-dimensional Brownian-motion with an exogenous diffusion process. The different dimensions of this process correspond to the different states of the world in the static model. The process itself may be interpreted as a continuous-time version of the process indicating the evolution of the vector of frequencies of the different outcomes from period to period in the m-period model. As in the discrete-time models, CARA utility functions and stationarity in the technology ensure that the optimal incentive scheme can be represented as a linear function of the cumulative change in this continuous-time analogue of the frequency accounts of the static model. Again though, this linearity in accounts has little to do with economic variables such as revenues, costs and profits.

Within their continuous-time model, Holmström and Milgrom obtain the linearity with respect to outcome variables that they are really after by imposing the additional assumption that the principal observes only the time path of a linear aggregate of the different "accounts", i.e., of the different dimensions of the Brownian motion whose drift rate vector the agent controls. The optimal incentive scheme is then linear in the final value of the aggregate.

However, this result has no immediate discrete-time analogue. To see this consider a multi-period discrete-time model and suppose that the principal observes only the time path of, say, total profits. By substracting total profits at time $t-1$ from total profits at time t, he can easily compute the increment in profits in period t. This increment perfectly reveals the state of the world in this period. Hence, in the discrete-time model the principal is always able to recover the information about the different "accounts" from observing the time path of the aggregate. Again we only get linearity in "accounts" - no matter how short the time period is.

In the following we first develop a unified framework in which the continuous-time Brownian-motion model is explicitly obtained as a limit of discrete-time models. To see that this is a nontrivial matter, note that in a discrete-time model the process indicating how many times each outcome has been observed up to time t is nondecreasing in t, negative instances of observation being out of the question. For convergence of such processes to Brownian motion, one must be comparing the accounts process to some norm so that a negative change can be interpreted as a shortfall from the norm. The question is where this norm comes from and how it is specified. Holmström and Milgrom do not say anything about this. Nor do they say anything about the dependence of outcome variables upon the periodization. We fill this gap, and as a by-product provide a characterization of the class of Brownian-motion models that can be approximated in our setting.

Secondly, we show that in our discrete-time approximations there does exist an analogue of the Holmström-Milgrom result providing for "linearity in profits" in the continuous-time model when the principal observes only the time path of profits as a linear aggregate of the different dimensions of the Brownian motion. For this purpose we assume that in the discrete-time model, the principal does not observe the time path of the outcome process at all. He only observes total profits at the end of the final period. Moreover, by the time the principal observes the aggregate, the agent may have destroyed some returns that had actually been realized. With these assumptions about asymmetric information concerning outcomes, we prove that the linear incentive scheme that is optimal in the continuous model is approximately optimal in the discrete-time models if the period length is sufficiently short. The argument relies on the continuity of certain features of the model in the transition between discrete and continuous time. Stationarity of solutions to discrete-time models, which is central for Holmström and Milgrom, plays no role because under the given information assumption such stationarity is not obtained.

Holmström and Milgrom also argue that the optimal incentive scheme is linear in aggregates if the agent's effort cost depends only on a linear aggregate of the drift rate vector he controls, e.g. on expected total profits. However, in Hellwig and Schmidt (2001) we show that in the discrete-time models and in any continuous-time model that can be obtained as a limit of discrete-time models, this effort cost function is not compatible with their assumption that the principal always wants to induce the agent to choose his actions in the interior of the admissible set.

The next section develops the framework for the analysis by introducing one-period and multi-period models and specifying the role of the "length of the period" as a parameter of these models. Section 3 considers a sequence of control paths in the discrete models and shows that if this sequence converges to a well defined control path in the continuous model, then the incentive schemes that implement the control paths in the discrete models converge to an incentive scheme that implements the limit path in the continuous model. Whereas this result takes the convergence behavior of control paths as given, Section 4 shows that this convergence behaviour is actually obtained for suitable subsequences of optimal control paths and incentive schemes. Section 5 considers the optimality of incentive schemes that are linear in aggregate profits. In this section we show that in a discrete-time model with many short periods, a linear incentive scheme is approximately optimal if the principal only observes cumulative total profits at the end, possibly after some profits have been destroyed by the agent. Section 6 concludes with a discussion of related literature and of the implications of different modeling alternatives in specifying the "length of the period". Appendix A gives the formal proofs of all results in the text. Certain supplementary results are given in Appendix B.

2 A Discrete Multi-Period Model

We start with a static model involving one period. At the beginning of the period the agent chooses an action which gives rise to a random profit $\tilde{\pi} \in \{\pi_0, \ldots, \pi_N\}$, $\pi_0 < \ldots < \pi_N$. Following Holmström and Milgrom (1987) we assume that the agent chooses the probability distribution p over possible profit levels $\pi_i \in I\!R$ directly at personal cost $c(p) \geq 0$. Thus, the agent's action is $p = (p_0, \ldots, p_N) \in P$ where P is the N-dimensional simplex. Throughout the paper we assume the following:

Assumption 1 *The effort cost $c(p)$ of an action $p \in P$ is given by a function $c(\cdot)$ on $I\!R^{N+1}$ which is strictly convex and continuously differentiable on some open set that contains P.*

The agent is assumed to have a constant coefficient of absolute risk aversion $r > 0$. Given an incentive scheme associating the payment s_i with the outcome π_i, he chooses action $p \in P$ so as to maximize his expected utility

$$-\sum_{i=0}^{N} p_i e^{-r(s_i - c(p))} \ . \tag{1}$$

The principal is assumed to be risk neutral. His payoff from implementing an action p by an incentive scheme $s = \{s_0, \ldots, s_N\}$ is given by:

$$\sum_{i=0}^{N} p_i(\pi_i - s_i) \ . \tag{2}$$

In the static model, the contracting problem is to choose an action p and an incentive scheme $(s_0, \ldots, s_N)$ so as to maximize (2) subject to incentive compatibility and to the individual-rationality condition that the certainty equivalent of the contract to the agent be nonnegative.[1]

An m-period model is simply an m-fold repetition of the one-period model. In this model, the agent chooses a new action $(p_0, \ldots, p_N) \in P$ in each period. This action determines outcome probabilities for that period. Overall outcomes in the m-period model are simply sums of period outcomes, i.e., sums of realized incentive payments, effort costs, and profits.

As we look at models with varying numbers of periods m, we want to adapt the "length of one period" so as to keep the total length of the time interval considered fixed. We normalize this time interval to one and suppose that in an m-period model, the individual period has "length" $\Delta = \frac{1}{m}$. The "length of one period" is to be understood as a parameter which enters the data of the static model in such a way that a model with m periods of length $\frac{1}{m}$ has in some sense a similar economic structure as a model with one period of length one. As m goes to infinity, presumably such models converge to a continuous-time model in which the agent adjusts his choices continuously. For the purpose of approximating the Holmström-Milgrom Brownian motion model we want the continuous-time model to be one where the agent's choices concern the drift rate, but not the covariance matrix of a Brownian motion process. In the Concluding Remarks at the end of this paper, we discuss other continuous-time processes and other approximations including some that allow for an endogenous covariance matrix and some that are not diffusions.

The data of the static model consist of (i) the $N+1$ profit levels π_i associated with the outcomes $i = 0, 1, \ldots, N$, and (ii) the effort cost function $c(\cdot)$. We need to specify how these data depend on $\Delta = \frac{1}{m}$. For this purpose, we assume that in a model with period length Δ, the profit levels π_i^Δ and effort cost function $c^\Delta(p^\Delta)$ are given by:

$$\pi_i^\Delta = \pi_i \Delta^{\frac{1}{2}} \quad \forall\, i \in \{0, \ldots, N\}\ . \tag{3}$$

and

$$c^\Delta(p^\Delta) \;\equiv\; \Delta \cdot c\left(\hat{p}_0 + \frac{p_0^\Delta - \hat{p}_0}{\Delta^{\frac{1}{2}}}, \ldots, \hat{p}_N + \frac{p_N^\Delta - \hat{p}_N}{\Delta^{\frac{1}{2}}}\right)\ . \tag{4}$$

[1]We assume that the certainty equivalent of the agent's utility if he does not work for the principal is 0. By Theorem 2 of Holmström and Milgrom (1987), this is no more than a normalization.

where $\pi_0, \ldots, \pi_N$ and $c(\cdot)$ are given as data of the static model with period length one, and $\hat{p} \in P$, $\hat{p} \gg 0$ is some exogenously given reference "standard". It will become clear in Section 3 that this standard is closely related to the disturbance of the stochastic profit process in the continuous-time limit of our model. Without loss of generality, we assume that profits are normalized so that

$$\sum_{i=0}^{N} \hat{p}_i \pi_i^{\Delta} \;=\; \sum_{i=0}^{N} \hat{p}_i \pi_i \;=\; 0 \;. \tag{5}$$

for all Δ.[2]

To see the point of this specification, it is useful to introduce auxiliary variables $\mu_1^{\Delta}, \ldots, \mu_N^{\Delta}$, such that for $i = 1, 2, \ldots, N$

$$\mu_i^{\Delta} \;=\; k_i \frac{p_i^{\Delta} - \hat{p}_i}{\Delta^{\frac{1}{2}}} \;, \tag{6}$$

with $k_i = \pi_i - \pi_0$, or equivalently,

$$p_i^{\Delta}(\mu^{\Delta}) \;=\; \hat{p}_i + \mu_i^{\Delta} \frac{\Delta^{\frac{1}{2}}}{k_i} \quad \forall\, i \in \{1, \ldots, N\} \;, \tag{7}$$

$$p_0^{\Delta}(\mu^{\Delta}) \;=\; 1 - \sum_{i=1}^{N} p_i^{\Delta} \;=\; \hat{p}_0 - \sum_{i=1}^{N} \mu_i^{\Delta} \frac{\Delta^{\frac{1}{2}}}{k_i} \;. \tag{8}$$

From (7) and (8), we see that there is a one-to-one correspondence between probability vectors p^{Δ} and vectors μ^{Δ} of the auxiliary variables, provided the latter satisfy the inequalities

$$-\frac{k_i \hat{p}_i}{\Delta^{\frac{1}{2}}} \;\leq\; \mu_i^{\Delta} \;\leq\; \frac{k_i(1 - \hat{p}_i)}{\Delta^{\frac{1}{2}}} \tag{9}$$

for all $i \in \{1, \ldots, N\}$ and

$$-\frac{1 - \hat{p}_0}{\Delta^{\frac{1}{2}}} \;\leq\; \sum_{i=1}^{N} \frac{\mu_i^{\Delta}}{k_i} \;\leq\; \frac{\hat{p}_0}{\Delta^{\frac{1}{2}}} \;. \tag{10}$$

We may therefore think of the agent choosing a vector μ^{Δ} of auxiliary variables satisfying (9) and (10), with associated probability vector $p = p^{\Delta}(\mu^{\Delta})$ given by (7) and (8), rather

[2] To see that this is indeed a normalization, suppose that $\sum_{i=0}^{N} \hat{p}_i \pi_i = \hat{\pi} \neq 0$. Let $\tilde{\pi} = \pi_i - \hat{\pi}$. Now we can define $\pi_i^{\Delta} = \Delta\hat{\pi} + \tilde{\pi}_i \Delta^{\frac{1}{2}}$. It is easy to check that for any time path of actions chosen by the agent expected total profits in the m-period model are given by equation (11) below plus the constant term $\hat{\pi}$. Furthermore, this constant term has no effect on the disturbance process.

In Holmström and Milgrom (1987, p. 318) the dependence of π_i^{Δ} on Δ is implicit in their appeal to the central limit theorem. For a more systematic discussion of the role played by this dependence, see the Concluding Remarks.

than choosing p^Δ as such. If his choice is formulated in these terms, then in a model with period length Δ, the choice μ^Δ entails expected profits equal to

$$\sum_{i=0}^{N} p_i^\Delta \pi_i^\Delta \;=\; \sum_{i=0}^{N} (p_i^\Delta - \hat{p}_i)(\pi_i - \pi_0)\Delta^{\frac{1}{2}} \;=\; \Delta \sum_{i=0}^{N} (\pi_i - \pi_0) \frac{p_i^\Delta - \hat{p}_i}{\Delta^{\frac{1}{2}}} \;=\; \Delta \sum_{i=1}^{N} \mu_i^\Delta \,. \tag{11}$$

and effort costs equal to

$$\begin{aligned} c^\Delta(p^\Delta) \;&=\; \Delta \cdot c\left(\hat{p} + \frac{p^\Delta(\mu^\Delta) - \hat{p}}{\Delta^{\frac{1}{2}}}\right) \;=\; \Delta \cdot c\left(\hat{p}_0 - \sum_{i=1}^{N} \frac{\mu_i^\Delta}{k_i}, \hat{p}_1 + \frac{\mu_1^\Delta}{k_1}, \ldots, \hat{p}_N + \frac{\mu_N^\Delta}{k_N}\right) \\ &\equiv\; \Delta \cdot \hat{c}(\mu^\Delta). \end{aligned} \tag{12}$$

Both the effort costs and the expected profits with a given vector μ^Δ are thus simply proportional to the “length” of the period. If the agent were to choose a constant vector $\mu^\Delta = \mu$ in each of the m periods having length $\Delta = \frac{1}{m}$, total expected profits would be $\sum_{i=1}^{N} \mu_i$ and total effort cost would be $\hat{c}(\mu)$, both independent of m and $\Delta = \frac{1}{m}$. The tradeoff between expected profits and effort costs is thus not affected by the choice of m and $\Delta = \frac{1}{m}$. This is precisely the invariance property across multi-period models that we need. More generally, if in the m-period model with period length $\Delta = \frac{1}{m}$, the agent chooses $\mu^{\Delta,\tau} = (\mu_1^{\Delta,\tau}, ..., \mu_N^{\Delta,\tau})$ in period $\tau = 1, ..., m$, then expected gross profits are equal to $\Delta \sum_{\tau=1}^{m} \sum_{i=1}^{N} \mu_i^{\Delta,\tau}$, and effort costs are equal to $\Delta \sum_{\tau=1}^{m} \hat{c}(\mu^{\Delta,\tau})$. In particular, if $\mu^{\Delta,\tau} = \mu$ regardless of τ and Δ, expected gross profits and effort costs are equal to $\sum_{i=1}^{N} \mu_i$ and $\hat{c}(\mu)$, regardless of Δ.

However, the probability vector $p^\Delta(\mu)$ that is associated with a given μ depends on Δ. Indeed as Δ goes to zero, for any finite μ, $p^\Delta(\mu)$ converges to the exogenously given vector $\hat{p}$. Underlying this convergence is the following structure: As indicated by (9) and (10), we can think of the agent's actions in terms of departures from the reference vector $\hat{p}$. Consider an action p^Δ that shifts probability mass from some low-profit outcome to some high-profit outcome as compared to the action $\hat{p}$. With profit levels proportional to $\Delta^{\frac{1}{2}}$ this raises expected profits by an order of magnitude $\Delta^{\frac{1}{2}}$. However the agent's cost of this shift in probability mass is made to depend on Δ in such a way that the agent will keep the order of magnitude of such shifts ordinarily to $\Delta^{\frac{1}{2}}$. That is, if the isoquants of the cost function $c(\cdot)$ exhibit nonzero curvature, then the deviation of the action p^Δ that is actually chosen from the zero-expected-profits action $\hat{p}$ will be on the order of $\Delta^{\frac{1}{2}}$, and expected profits per period will be on the order of Δ.

In the m-period model with period length Δ, the principal's problem is to choose a time path of actions $\{\mu^{\Delta,\tau}\}_{\tau=1,2,...,m}$ and an incentive scheme so as to maximize

$$\sum_{\tau=1}^{m}\sum_{i=0}^{N} p_i^{\Delta}(\mu^{\Delta,\tau})(\pi_i^{\Delta} - s_i^{\Delta}) = \Delta \sum_{\tau=1}^{m}\sum_{i=1}^{N} \mu_i^{\Delta,\tau} - \sum_{\tau=1}^{m}\sum_{i=0}^{N} p_i^{\Delta}(\mu^{\Delta,\tau}) s_i^{\Delta} \quad (13)$$

subject to incentive compatibility and individual rationality. This problem depends not only on expected profits and effort costs but also on the variance and the other higher moments of the distribution of profits, all of which depend on Δ. However, as will be shown in the next section, when Δ is small and the agent's action is close to the standard $\hat{p}$, these higher moments are close to being independent of μ, and the incentive payments that are required to implement $p^{\Delta}(\mu)$ admit a simple approximation.

3 Approximation of the Brownian Model

We will consider the continuous-time model as a limit of the discrete-time models when Δ goes to 0. The following definitions will be useful. Let $\tilde{A}_i^{\Delta,\tau} \in \{0,1\}$ be a random variable such that $\tilde{A}_i^{\Delta,\tau} = 1 \Leftrightarrow \tilde{\pi}^{\Delta,\tau} = \pi_i^{\Delta}$, $i \in \{0,\dots N\}$, $\Delta \in \{\frac{1}{m}\}_{m=1,...}$, $\tau \in \{1,\dots,m\}$. Clearly, $\text{Prob}(\tilde{A}_i^{\Delta,\tau} = 1) = p_i^{\Delta,\tau}$, and $\tilde{A}_i^{\Delta,\tau} = 1$ implies $\tilde{A}_j^{\Delta,\tau} = 0$ for all $j \neq i$. Thus, each $\tilde{A}_i^{\Delta,\tau}$ has a Bernoulli distribution, but $\tilde{A}_i^{\Delta,\tau}$ and $\tilde{A}_j^{\Delta,\tau}$ are not stochastically independent.

We consider the incentive payments s_i^{Δ}, $i = 0,1,...,N$ that the principal needs to implement a given action p^{Δ} of the agent if the period length is Δ. For now we assume that the action p^{Δ} belongs to the relative interior of the set of feasible actions in the model with period length Δ. Under this assumption, the first-order conditions for the agent's maximization problem imply that an incentive scheme $s^{\Delta} = (s_0^{\Delta},\dots,s_N^{\Delta})$ which implements p^{Δ} with certainty equivalent 0 in the model with period length Δ must satisfy:

$$s_i^{\Delta} \;=\; c^{\Delta}(p^{\Delta}) - \frac{1}{r}\ln\left(1 - r c_i^{\Delta} + r\sum_{j=0}^{N} p_j^{\Delta} c_j^{\Delta}\right), \quad (14)$$

for $i = 0,1,...,N$, where, as usual, c_i^{Δ} refers to the partial derivative of the effort cost function c^{Δ} with respect to p_i^{Δ}, see Theorem 3 of Holmström and Milgrom (1987). In terms of the vector μ^{Δ}, with $p^{\Delta} = p^{\Delta}(\mu^{\Delta})$, one can use (12) to rewrite s_i^{Δ} as:

$$s_i^{\Delta} \;=\; \Delta\hat{c}(\mu^{\Delta}) - \frac{1}{r}\ln\left(1 - r\hat{c}_i k_i \Delta^{\frac{1}{2}} + r\sum_{j=0}^{N} p_j^{\Delta}\hat{c}_j k_j \Delta^{\frac{1}{2}}\right)\;, \quad (15)$$

where $\hat{c}_i = \frac{c_i - c_0}{k_i}$ is the partial derivative of $\hat{c}$ with respect to μ_i^Δ, $\hat{c}_0 = 0$, and use has been made of the fact that by (4) $c_i^\Delta = \Delta^{\frac{1}{2}} c_i$ for all i.[3]

Using a Taylor series expansion of the logarithmic term in (15), the incentive payment s_i^Δ is approximated by:

$$s_i^\Delta = \Delta \hat{c}(\mu^\Delta) + \left[\hat{c}_i k_i - \sum_{j=0}^{N} p_j^\Delta(\mu^\Delta) \hat{c}_j k_j \right] \Delta^{\frac{1}{2}} + \frac{r}{2} \left[\hat{c}_i k_i - \sum_{j=0}^{N} p_j^\Delta(\mu^\Delta) \hat{c}_j k_j \right]^2 \Delta + O\left(\Delta^{\frac{3}{2}}\right) \quad (16)$$

If the principal wants to implement the time path of actions $\{\mu^{\Delta,\tau}\}_{\tau=1,\ldots,m}$,[4] the total renumeration that has to be offered is given by[5]

$$\begin{aligned} \tilde{s}^\Delta \;=\; & \Delta \sum_{\tau=1}^{m} \hat{c}(\mu^{\Delta,\tau}) + \sum_{\tau=1}^{m} \sum_{i=0}^{N} \hat{c}_i(\mu^{\Delta,\tau}) \left(\tilde{A}_i^{\Delta,\tau} - p_i^\Delta(\mu^{\Delta,\tau}) \right) k_i \Delta^{\frac{1}{2}} \\ & + \frac{r}{2} \sum_{\tau=1}^{m} \sum_{i=0}^{N} \tilde{A}_i^{\Delta,\tau} \left[\hat{c}_i(\mu^{\Delta,\tau}) k_i - \sum_{j=0}^{N} \hat{p}_j \hat{c}_j(\mu^{\Delta,\tau}) k_j \right]^2 \Delta + O(\Delta^{\frac{1}{2}}) \qquad (17) \end{aligned}$$

Equation (17) admits a simple interpretation:[6] The first term reflects the total effort cost of the agent from choosing $\{\mu^{\Delta,\tau}\}$, $\tau = 1, \ldots, m$. The second term gives the appropriate incentives to the agent to actually choose this time path of actions. If outcome i is realized in period τ, this raises the agent's overall incentive payment by an amount $\hat{c}_i(\mu^{\Delta,\tau}) k_i \Delta^{\frac{1}{2}}$, reflecting the marginal cost of shifting probability mass towards outcome i. The expected value of this payment, $p_i^\Delta(\mu^{\Delta,\tau}) c_i^\Delta(\mu^{\Delta,\tau}) k_i \Delta^{\frac{1}{2}}$, is subtracted again because in expected-value terms the principal's payments to the agent depend only on individual-rationality considerations. Thus the second term in (17) is proportional to the difference between the actual realization of this profit level and the expected realization given that the agent chooses $p_i^\Delta(\mu^{\Delta,\tau})$. The third term is the risk premium that has to be paid to the agent

[3] Clearly the convexity and differentiability assumption on $c(\cdot)$ imply that $\hat{c}(\cdot)$ is strictly convex and continuously differentiable on the set of all vectors $\mu = (\mu_1, \ldots, \mu_N)$ for which $\hat{p}_0 > \sum_{i=1}^{N} \mu_i / k_i$ and $\hat{p}_i > -\mu_i / k_i$ for $i = 1, \ldots, N$.

[4] We restrict attention to deterministic time paths of actions that do not depend on the history of profit realizations up to time τ. Holmström and Milgrom have shown that the agent will be induced by the optimal contract to take a constant deterministic action in each period, so there is no need to consider stochastic controls (that may be history dependent) explicitly at this stage. However, we deal with such control paths in Section 5 and in Appendix B.

[5] Here we made use of the fact that $\sum_{i=0}^{N} \tilde{A}_i^{\Delta,\tau} = 1$ for all $\tau \in \{1, \ldots, m\}$. Furthermore, we substituted $p_j^\Delta(\mu^{\Delta,\tau})$ in the squared term by (7) and (8), used $\hat{c}_0 = 0$, and rearranged by putting the appropriate terms into $O(\Delta^{\frac{1}{2}})$.

[6] See also Schättler and Sung (1993, p. 337).

to compensate him for the randomness of the second term. The last term reflects the approximation we are using; it vanishes as Δ goes to zero.

Given the way the per-period deviations $\tilde{X}_i^{\Delta,\tau} = k_i \Delta^{\frac{1}{2}} \left(\tilde{A}_i^{\Delta,\tau} - p_i^{\Delta}(\mu^{\Delta,\tau}) \right)$ of realized from expected profits depend on the length of the period, the cumulative deviations

$$\tilde{X}_i^{\Delta}(\tau) \;=\; k_i \Delta^{\frac{1}{2}} \sum_{\tau'=1}^{\tau} [\tilde{A}_i^{\Delta,\tau'} - p_i^{\Delta}(\mu^{\Delta,\tau'})] \tag{18}$$

from period one to period τ are more convenient to work with. For $i = 1, \ldots, N$, $\tilde{X}_i^{\Delta}(\tau)$ can be thought of as the cumulative deviation (up to τ) of realized profits under outcome i from the expected value of these profits under the given policy of the agent. $\tilde{X}_i^{\Delta}(\tau)$ will be called the "stochastic process of cumulative deviations from the mean". As we restrict attention to deterministic time paths of actions $\{\mu^{\Delta,\tau}\}_{\tau=1,\ldots,m}$, the increments of $\tilde{X}_i^{\Delta}(\tau)$ are stochastically independent.

To embed the discrete-time model in continuous time, we use a linear interpolation to represent the process $\tilde{X}_i^{\Delta}(\tau)$ by a continuous-time process $X_i^{\Delta}(t)$ such that for $t \in [0,1]$:

$$X_i^{\Delta}(t) \;=\; \left(1 - \frac{t}{\Delta} + \left[\frac{t}{\Delta}\right]\right) \tilde{X}_i^{\Delta}\left(\left[\frac{t}{\Delta}\right]\right) + \left(\frac{t}{\Delta} - \left[\frac{t}{\Delta}\right]\right) \tilde{X}_i^{\Delta}\left(\left[\frac{t}{\Delta}\right] + 1\right) \tag{19}$$

where $[\frac{t}{\Delta}]$ denotes the greatest integer less than or equal to $\frac{t}{\Delta}$. Note that $X_i^{\Delta}(t)$ is a random function taking values in $C = C[0,1]$, the space of continuous functions on $[0,1]$. For any deterministic time path of actions $\{\mu^{\Delta,\tau}\}_{\tau=1,\ldots,m}$ in the discrete-time model we also use a continuous-time representation $\mu^{\Delta}(\cdot)$ where

$$\mu^{\Delta}(t) \;=\; \mu^{\Delta,[t/\Delta]} \,. \tag{20}$$

Theorem 1 *Consider a sequence of discrete models with period length Δ, $\Delta = 1, \frac{1}{2}, \frac{1}{3}, \ldots$. Suppose that, as $\Delta \to 0$, the time path of actions $\mu^{\Delta}(t)$ converges uniformly to some continuous function $\mu(t)$, $t \in [0,1]$ such that $\hat{p}_0 > \sum_i \mu_i(t)/k_i$ and $\hat{p}_i > -\mu_i(t)/k_i$ for all t and $i = 1, \ldots, N$. Then, as $\Delta \to 0$,*

(a) the stochastic process of cumulative deviations from the mean $X^{\Delta}(\cdot) = (X_1^{\Delta}(\cdot), \ldots, X_N^{\Delta}(\cdot))$ converges in distribution to a process $X(\cdot)$ which is a

driftless N-dimensional Brownian motion with starting point $X(0) = 0$ *and covariance matrix* $\mathbf{\Sigma} = (\sigma_{ij})$ *where, for* $i = 1, ..., N,$ *and* $j = 1, ..., N,$

$$\sigma_{ij} = k_i^2 \hat{p}_i \delta_{ij} - k_i k_j \hat{p}_i \hat{p}_j \tag{21}$$

with $\delta_{ij} = 1$ *if* $i = j$ *and* $\delta_{ij} = 0$ *if* $i \neq j$.

(b) *the total cost to the agent converges to* $\int_0^1 \hat{c}(\mu(t))dt$*;*

(c) *the joint distributions of the incentive payment* $\tilde{s}^{\Delta}$ *and the cumulative-deviations process* $X^{\Delta}(\cdot)$ *converges to the distribution of the pair* $(\tilde{s}, X(\cdot))$*, where* $X(\cdot)$ *is the driftless Brownian motion specified in (a) and, with probability one,*

$$\tilde{s} \;=\; \int_0^1 \hat{c}(\mu(t))dt + \int_0^1 \hat{c}'(\mu(t))dX + \frac{r}{2}\int_0^1 \hat{c}'(\mu(t))\mathbf{\Sigma}[\hat{c}'(\mu(t))]^T dt \tag{22}$$

where $\hat{c}'(\cdot) = (\hat{c}_1(\cdot), \ldots, \hat{c}_N(\cdot))$.

Theorem 1 is closely related to the results of Holmström and Milgrom (1987) and Schättler and Sung (1993). Equation (22) can be rewritten in the form

$$\tilde{s} = s(Z(\cdot), \mu(\cdot)) \tag{23}$$

where $Z(\cdot)$ is the stochastic process satisfying

$$Z(t) = \int_0^t \mu(t')dt' + X(t) \tag{24}$$

for all t and, for any N-dimensional process $\hat{Z}(\cdot)$ of the form

$$d\hat{Z} = \hat{\mu}\; dt + dX, \tag{25}$$

and any adapted process $\mu(\cdot)$,

$$s(\hat{Z}(\cdot), \mu(\cdot)) = \int_0^1 \hat{c}(\mu(t))dt + \int_0^1 \hat{c}'(\mu(t))d\hat{Z} - \int_0^1 \hat{c}'(\mu(t))\mu(t)dt + \frac{r}{2}\int_0^1 \hat{c}'(\mu(t))\mathbf{\Sigma}[\hat{c}'(\mu(t))]^T dt. \tag{26}$$

In the continuous-time analysis of Holmström and Milgrom or Schättler and Sung, the agent controls the drift rate process process $\hat{\mu}(\cdot)$ of the N-dimensional Brownian motion $\hat{Z}(\cdot)$; the control process $\hat{\mu}(\cdot)$ must be adapted to the disturbance process $X(\cdot)$, which is taken as given. The principal observes the realizations of the process $\hat{Z}(\cdot)$, but not the control process $\hat{\mu}(\cdot)$ or the disturbance process $X(\cdot)$. Under the additional assumption that controls are restricted to a bounded set K, they show that an incentive scheme implements a desired control process $\mu(\cdot)$ with values in the interior of K if and only if it takes the form $s(\hat{Z}(\cdot), \mu(\cdot))$ given by (26).[7] Theorem 1 shows that this implementation condition on incentive schemes and policies in the Brownian model can be interpreted as a limit of implementation conditions on incentive schemes and policies in approximating discrete-time models.

In the continuous-time limit the agent controls only the drift rate while the higher moments of the cumulative output process are fixed and fully determined by the fixed vector $\hat{p}$. In contrast, he does control the entire distribution in the discrete-time models. However, when Δ is small, implementation of a given drift rate process involves actions $p^{\Delta}(\mu(\cdot))$ close to the constant action $\hat{p}$, the difference in any period being of order $\Delta^{\frac{1}{2}}$. Because of this convergence of the path of actions to the constant path with value $\hat{p}$, the second moments of the cumulative-returns process in the limit are fully determined by $\hat{p}$, and the higher moments vanish. The critical action vector $\hat{p}$ thus determines the "noise" of the agency problem in continuous time.

The vector $\hat{p}$ also matters for the interpretation of the process $Z(\cdot)$, by which the principal assesses the agent's performance in the continuous-time model. Discrete-time approximations of the process (24) are obtained by specifying

$$\tilde{Z}_i^{\Delta}(\tau) \;=\; k_i \Delta^{\frac{1}{2}} \left[\sum_{\tau'=1}^{\tau} \tilde{A}_i^{\Delta,\tau'} - \hat{p}_i \tau \right] \tag{27}$$

as the cumulative contribution to total profits stemming from outcome i in excess of some standard that is given by the probability $\hat{p}_i$. Using linear interpolation as before, one obtains continuous-time representations $Z^{\Delta}(\cdot)$ of these discrete-time approximations. By a simple corollary to Theorem 1, the pairs $(\tilde{s}^{\Delta}, Z^{\Delta}(\cdot))$ converge to the pair $(s(Z(\cdot), \mu(\cdot)), Z(\cdot))$ that is given by (26) and (24). For any τ, by (3) and (5), the sum $\sum_{i=1}^{N} \tilde{Z}_i^{\Delta}(\tau)$ is just equal to the sum $\sum_{\tau'=1}^{\tau} \tilde{\pi}^{\Delta,\tau'}$ of profits from period 1 to period τ.

Theorem 1 takes the discrete model as given and shows how to obtain a Brownian model as the limit of a sequence of these discrete models. We could have proceeded the

[7]Theorem 6 of Holmström and Milgrom shows that (26) is necessary to implement $\mu(\cdot)$. In a more general model, Theorems 4.2 and Corollary 4.1 of Schättler and Sung show that if the effort cost function $\hat{c}(\cdot)$ is convex, then (26) is sufficient as well as necessary to implement $\mu(\cdot)$.

other way round, asking whether it is possible to approximate a given Brownian model with some sequence of discrete models. The following result, which is an immediate corollary to Theorem 1, gives an answer to this question.

Corollary 1 *Let B be a given N-dimensional Brownian motion with zero drift and covariance matrix $\mathbf{\Sigma}$. Suppose that if the agent chooses a control process $\mu(\cdot)$, then the agent's cumulative costs are given by the differential equation*

$$dc = \hat{c}(\mu(t))dt \tag{28}$$

and the stochastic process $X(\cdot)$ of cumulative deviations of actual profits from mean profits induced by $\mu(\cdot)$ is B. If there exist real numbers $k_1, \ldots, k_N$ and positive real numbers $\hat{p}_1, \ldots, \hat{p}_N$ with $\sum_{i=1}^N \hat{p}_i \leq 1$ such that $\mathbf{\Sigma} = (\sigma_{ij})$ can be written as in (21), then to this continuous-time model there corresponds a discrete-time model with period length Δ in which

- *there are $N+1$ possible profit levels in each period which are given by*

$$\pi_0^\Delta = -\Delta^{\frac{1}{2}} \sum_{i=1}^N \hat{p}_i k_i \quad \text{and} \quad \pi_i^\Delta = \Delta^{\frac{1}{2}} \left(k_i - \sum_{i=1}^N \hat{p}_i k_i \right) \tag{29}$$

- *the probability $p_i^\Delta(\mu^\Delta)$ is given by (7) and (8), with $\mu^{\Delta,[t/\Delta]} = \mu([t/\Delta])$,*
- *and the cost to the agent in each period is given by $c^\Delta(p^\Delta(\mu^\Delta)) = \Delta \hat{c}(\mu^\Delta)$,*

such that the continuous model is the limit of this discrete model (in the sense of Theorem 1) as Δ approaches 0.

Corollary 1 requires that the covariance matrix $\mathbf{\Sigma} = (\sigma_{ij})$ of the N-dimensional Brownian motion can be written as in (21). If $N = 1$, this is always possible. If $N > 1$, this condition implies a restriction on the set of Brownian models that can be approximated by a sequence of discrete models. This restriction stems from the fact that in the discrete model only one state of the world can materialize in each period. This implies that the accounts $\tilde{A}_i^{\Delta,\tau}$ and thus also the accounts $\tilde{X}_i^{\Delta,\tau}$ have a special covariance structure. For example, it is impossible that the $\tilde{X}_i^{\Delta,\tau}$ are stochastically independent. Therefore this restriction is very natural.

In the entire analysis here, the dimension N of the Brownian motion has referred to the number of different profit levels that can be distinguished. Holmström and Milgrom (1987, p. 322) offer a second interpretation according to which N refers to different activities of the agent. For example, if $N = 2$, account X_1 could be a measure of revenues, while X_2 could be a measure of costs. Thus, μ_1 reflects the agent's effort to increase revenues, while μ_2 reflects his effort to reduce costs. With this interpretation the Brownian model could have any variance-covariance structure. It is possible to show that any N-dimensional Brownian motion model (with an arbitrary covariance matrix $\mathbf{\Sigma}$) can be approximated by a discrete model, if we extend the dimension of the discrete model to $2^N - 1$. However, we do not want to go into the details of this approximation here.[8]

4 The Convergence of Optimal Control Paths and Incentive Schemes

So far, we have only been concerned with the behaviour of the agent. We now consider the behaviour of the principal. Holmström and Milgrom (1987, Theorem 5) have shown that in the repeated discrete-time model, the principal's problem always has a solution in which he induces the agent to take the same action in each period, regardless of prior history. For $m = 1, 2, ...$ and $\Delta = \frac{1}{m}$, let $\mu^{\Delta *}(\cdot)$ be such a control path with constant value $\mu^{\Delta *}$ that is optimal for the principal in the discrete-time model with period length Δ and consider the behaviour of the sequence $\{\mu^{\Delta *}(\cdot)\}$ as Δ goes to zero. Our claim that the discrete-time specifications considered in this paper provide "appropriate" approximations of the continuous-time Brownian-motion model rests on the proposition that the solutions to the principal's problems in the discrete-time models are close to solutions to the principal's

[8]To sketch the basic idea suppose that the Brownian model is two-dimensional and that μ_1 is the drift rate of revenues while μ_2 is the drift rate of costs. To approximate this model by a sequence of discrete models we need at least two different levels of revenues $R \in \{R_1, R_2\}$ and two levels of costs $C \in \{C_1, C_2\}$. Thus, in each period there are four different possible outcomes $\{(R_1, C_1), (R_1, C_2), (R_2, C_1), (R_2, C_2)\}$, which means that there are $2^2 - 1 = 3$ different accounts. With this three-dimensional model we can approximate any variance-covariance matrix in the two-dimensional continuous model by choosing the correlation between the Bernoulli distributions over $\{R_1, R_2\}$ and $\{C_1, C_2\}$ appropriately. If we want to allow for, say, M_r levels of revenues and M_c levels of costs, we need a model with $M_r \cdot M_c$ possible outcomes and hence $M_r \cdot M_c - 1$ different accounts. The problem of getting from these $M_r \cdot M_c - 1$ different accounts to incentive schemes that are defined in terms of just revenue and cost aggregates is then the same as the problem of aggregation across accounts that is studied in Sections 5 and 6.

problem in the continuous-time model if the length of the period is short and there are many periods. This requires that the sequence $\{\mu^{\Delta *}(\cdot)\}$ of optimal control paths in the discrete-time models have a limit path $\mu^*(\cdot)$ and that this limit path in turn be an optimal control path in the continuous-time model.

Three issues arise: First, under what conditions does the sequence $\{\mu^{\Delta *}(\cdot)\}$ have a limit path $\mu^*(\cdot)$ at all? Second, if it does have a limit path, can Theorem 1 - which is based on first-order conditions for interior controls - be applied to $\mu^*(\cdot)$ and the subsequence $\{\mu^{\Delta' *}(\cdot)\}$ of $\{\mu^{\Delta *}(\cdot)\}$ that converges to $\mu^*(\cdot)$? Third, does $\mu^*(\cdot)$ provide a solution to the principal's problem in the continuous-time model? We consider each issue in turn.

First, to ensure that the sequence $\{\mu^{\Delta *}(\cdot)\}$ of constant optimal control paths in the discrete-time models have a limit path $\mu^*(\cdot)$, we assume that controls are restricted to a compact set $\hat{K} \subset I\!R^N$. This guarantees that the sequence $\{\mu^{\Delta *}\}$ of values taken by the paths $\mu^{\Delta *}(\cdot)$ has a subsequence $\{\mu^{\Delta' *}\}$ that converges to a limit μ^*. The corresponding subsequence $\{\mu^{\Delta' *}(\cdot)\}$ of constant optimal control paths converges uniformly to the control path $\mu^*(\cdot)$ with constant value μ^*.

Second, Theorem 1 with its reliance on first-order conditions for interior choices of actions by the agent can be applied to the sequence $\{\mu^{\Delta' *}(\cdot)\}$ and its limit $\mu^*(\cdot)$ if these paths all take values in the interior of the feasible set $\hat{K}$ so that the first-order necessary conditions presumed in Theorem 1 are in fact appropriate. Therefore we impose an additional assumption which ensures that any controls which the principal wants to implement in the continuous-time model or in the discrete-time models when Δ is close to zero must lie in the interior of the set $\hat{K}$.[9]

We start from the idea that there is some minimum payoff level, which the principal can guarantee himself when Δ is small, and that any action in the boundary of the feasible set provides the principal with a payoff strictly below this guaranteed minimum. To specify the guaranteed minimum, we consider the price that the agent would be willing to pay (which may be negative) if he could have the entire project to himself and choose

[9]Alternatively we could allow for control paths with values on the boundary of $\hat{K}$. In Hellwig and Schmidt (1998), the conclusions of Theorems 2 and 3 were obtained regardless of whether optimal control paths take values in the interior or on the boundary of the feasible set, see also Hellwig and Schmidt (2001). However, the formalism is rather more messy when boundary controls have to be considered.

the time path of actions he likes best. Selling the project to the agent for a price q is equivalent to imposing the incentive scheme

$$\tilde{s}^{\Delta} = \sum_{\tau=1}^{m} \tilde{\pi}^{\Delta,\tau} - q. \tag{30}$$

This is acceptable to the agent if q is no greater than the certainty equivalent of the difference between the cumulative total profits $\sum_{\tau=1}^{m} \tilde{\pi}^{\Delta,\tau}$ and the effort cost $\sum_{\tau=1}^{m} c^{\Delta}(p^{\Delta,\tau}) = \Delta \sum_{\tau=1}^{m} \hat{c}(\mu^{\Delta,\tau})$ under the time path of actions $\{\mu^{\Delta,\tau}\}$ that he chooses. If the agent is free to choose $\mu^{\Delta,\tau} = 0$ (or $p^{\Delta,\tau} = \hat{p}$) for $\tau = 1, \ldots, m$, this is no less than the certainty equivalent of the difference $\sum_{\tau=1}^{m} \tilde{\pi}^{\Delta,\tau} - \hat{c}(0)$ between the sum of returns and effort costs at the constant action $\mu^{\Delta,\tau} = 0$. Conditions (3) and (5) imply that as Δ converges to zero, the distribution of the sum of returns $\sum_{\tau=1}^{m} \tilde{\pi}^{\Delta,\tau}$ that is induced by the constant action $\mu^{\Delta,\tau} = 0$ for $\tau = 1, ..., m$ converges to a normal distribution with mean zero and variance $\sigma^2 := \sum_{i=0}^{N} \hat{p}_i \pi_i^2$, and the certainty equivalent of the difference $\sum_{\tau=1}^{m} \tilde{\pi}^{\Delta,\tau} - \hat{c}(0)$ to the agent at the constant action $\mu^{\Delta,\tau} = 0$ converges to $-\frac{r}{2}\sigma^2 - \hat{c}(0)$. If $q < -\frac{r}{2}\sigma^2 - \hat{c}(0)$, then for Δ sufficiently small, the incentive scheme (30) is acceptable to the agent and provides the principal with the payoff q. This is also true in the continuous-time model because $-\frac{r}{2}\sigma^2 - \hat{c}(0)$ is also the certainty equivalent for the agent of the difference $\sum_{i=1}^{N} Z_i(1) - \hat{c}(0)$ under the distribution induced by the constant action $\mu(t) \equiv 0$.[10] Thus we obtain:

Lemma 1 *Assume that $\mu = 0$ is a feasible action. Then for any $\eta > 0$ and any sufficiently small $\Delta > 0$, in the discrete-time model with period length Δ and in the continuous-time model, the principal can guarantee himself a payoff no less than $-\frac{r}{2}\sigma^2 - \hat{c}(0) - \eta$.*

Lemma 1 motivates the following assumption:

Assumption 2 *In the discrete-time models and in the continuous-time model, controls are constrained to lie in a compact convex set $\hat{K} \subset I\!R^N$. Moreover, there exists a compact subset K of the interior of $\hat{K}$ such that $0 \in K$ and for some $\eta > 0$,*

$$\sum_{i=1}^{N} \mu_i - \hat{c}(\mu) \;<\; -\frac{r}{2}\sigma^2 - \hat{c}(0) - \eta \tag{31}$$

[10] By (24) and (5), (30) can also be written as $\tilde{s}^{\Delta} = \sum_{i=1}^{N} \tilde{Z}_i^{\Delta}(m) - q$, so the scheme $\tilde{s} = \sum_{i=1}^{N} Z_i(1) - q$ is just the continuous-time version of (30).

whenever $\mu \notin K$.

Note that Assumption 2 concerns the values of the cost function *not* its derivatives. It does not require any derivative to go out of bounds, e.g., near the boundary of the set K. Unboundedness of derivatives would be incompatible with Assumption 1, which has been used extensively in the proof of Theorem 1(c).

The left hand side of (31) would be the expected net profit of the principal of implementing $\mu \notin K$ if he could do so without compensating the agent for bearing the risk that is necessary to induce him to choose μ. The assumption states that this is strictly less than and bounded away from the minimum payoff that, according to Lemma 1, the principal can guarantee himself in the discrete-time model with a sufficiently small period length Δ and in the continuous-time model. For small Δ, it can therefore not be optimal for the principal to induce the agent to choose controls outside of K.

Lemma 2 *Impose Assumption 2 as well as Assumption 1. For* $\Delta = 1, \frac{1}{2}, \ldots$ *and* $m = \frac{1}{\Delta}$, *let* $\mu^{\Delta *} \in I\!R^N$ *be a control vector such that the control path* $\{\mu^{\Delta *}(\cdot)\}$ *with the constant value* $\mu^{\Delta *}$ *solves the principal's problem in the m-period problem with period length* Δ *when controls are restricted to lie in* $\hat{K}$. *Then* $\mu^{\Delta *} \in K$ *for any sufficiently small* Δ.

Because K is a subset of the interior of the feasible set $\hat{K}$, this justifies the reliance on first-order conditions for interior solutions in Theorem 1 and settles the second issue.

Let us now turn to the third and most difficult issue. As mentioned above, the assumption that in the m-period problem with period length Δ, a control path with a constant value $\mu^{\Delta *}$ is optimal for the principal is without loss of generality, see Holmström and Milgrom (1987, Theorem 5). With controls constrained to lie in the compact set $\hat{K}$, the sequence $\{\mu^{\Delta *}\}$ has a subsequence $\{\mu^{\Delta' *}\}$ which converges to a limit μ^*; by Lemma 2, μ^* actually belongs to K and, hence, to the interior of $\hat{K}$. The corresponding subsequence $\{\mu^{\Delta' *}(\cdot)\}$ of constant optimal control paths converges to the path $\mu^*(\cdot)$ with constant value μ^*. We want to show that this control path is itself optimal for the principal in the continuous-time model.

Technically, we need a continuity property of maximizers in the transition from discrete-time to continuous-time models. This requires an application of the maximum theorem of Berge (1959), by which a limit of a sequence of solutions to different maximization problems is itself a solution to the maximization problem "in the limit" provided the maximization problems along the sequence are sufficiently "similar" to the maximization problem "in the limit". The question is whether the principal's maximization problems in the multi-period discrete-time models are in fact "similar" - in the sense required for the application of Berge's theorem - to his problem in the continuous-time model.

Theorem 1 - in combination with the results of Schättler and Sung - provides one part of the answer to this question: If a subsequence $\{\mu^{\Delta'*}(\cdot)\}$ of optimal control paths with constant values $\mu^{\Delta'*}$ in the discrete-time models converges to a control path $\mu^*(\cdot)$, then the path $\mu^*(\cdot)$, with constant value $\mu^* = \lim \mu^{\Delta'*}$, is implementable in the continuous-time model. Moreover, the incentive payments $\tilde{s}^{\Delta'}$ that are associated with the implementation of the paths $\mu^{\Delta'*}(\cdot)$ in the discrete-time models converge in distribution to the incentive payment $\tilde{s}$ that is associated with the implementation of the path $\mu^*(\cdot)$ in the continuous-time model. The expected cost to the principal of implementing $\mu^*(\cdot)$ in the continuous-time model is not significantly different from the expected cost of implementing $\mu^{\Delta'*}(\cdot)$ in the discrete-time model with period length Δ' when Δ' is close to zero. This ensures that the principal's maximum payoff in the continuous-time model is no worse than the limit of the maximum payoffs in the discrete-time models.

To infer that $\mu^*(\cdot)$ is in fact optimal for the principal in the continuous-time model, we must also show that his maximum payoff in this model is no better than the limit of the maximum payoffs in the discrete-time models. For this purpose we want to argue that in any discrete-time model with short periods the principal is able to implement - at similar costs - a control path which is close to an optimal path for the continuous-time model. In contrast to the analysis so far, which has exclusively focused on first-order necessary conditions for implementation, this requires some account of sufficient conditions for the implementation of control paths in the discrete-time models.

The importance of the distinction between first-order necessary and sufficient conditions for the implementation of an action by an incentive scheme is well known from the work of Mirrlees (1999) or Grossman and Hart (1983). There does not seem to be any direct approach to obtaining sufficient conditions in the multi-period discrete-time models. However we can use an indirect approach: For the continuous-time model, the convexity assumptions on the effort cost function and on the feasible set imply that the first-order necessary conditions for implementation of a stipulated control path by an incentive scheme are also sufficient, see Schättler and Sung (1993, Theorem 4.2). Indeed our assumption of strict convexity of the effort cost function ensures that when the agent is faced with a given incentive scheme in the continuous-time model, his choice of control path is unique. By continuity considerations (another application of Berge's maximum theorem, this time to the agent), it follows that, in a discrete-time model with short periods, any control path that the agent is willing to choose when faced with an incentive scheme close to the scheme that implements an optimal control path in the continuous-time model must itself be close to the optimal control path in the continuous-time model. This ensures that the principal's maximum payoff in the continuous-time model cannot be strictly better than the limit of his maximum payoffs in the discrete-time models.

This result is stated formally in the following theorem which proves our claim that our discrete-time models approximate the Brownian-motion model of Holmström and Milgrom.

Theorem 2 *Impose Assumption 2 as well as Assumption 1. For* $\Delta = 1, \frac{1}{2}, \ldots$ *and* $m = \frac{1}{\Delta}$*, let* $\mu^{\Delta *} \in I\!R^N$ *be a control vector such that the control path* $\{\mu^{\Delta *}(\cdot)\}$ *with the constant value* $\mu^{\Delta *}$ *solves the principal's problem in the m-period problem with period length* Δ *when controls are restricted to lie in* $\hat{K}$*. If* μ^* *is any limit point of the sequence* $\{\mu^{\Delta *}\}$*, then* $\mu^* \in K$*. Moreover the control path* $\mu^*(\cdot)$ *with the constant value* μ^* *solves the principal's problem in the continuous-time model when controls are restricted to* $\hat{K}$*; in terms of the control process* $\mu^*(\cdot)$ *the optimal incentive scheme* s^* *has the representation*

$$s^*(\cdot) \; = \; \hat{c}(\mu^*) + \sum_{i=1}^{N} \hat{c}_i(\mu^*) X_i(1) + \frac{r}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \hat{c}_i \sigma_{ij} \hat{c}_j, \tag{32}$$

with (σ_{ij}) *given by (21).*

We conclude this section with a brief discussion of optimal incentive schemes. The schemes

$$\tilde{s}^{\Delta *} = \sum_{\tau=1}^{m}\sum_{i=0}^{N} \tilde{A}_i^{\Delta,\tau} s_i^{\Delta *}, \qquad (33)$$

which implement the optimal control paths $\mu^{\Delta *}(\cdot)$ in the discrete-time models, can be rewritten in the linear form:

$$\tilde{s}^{\Delta *} = \alpha_0^{\Delta} + \sum_{i=1}^{N} \alpha_i^{\Delta} Z_i^{\Delta}(1) \ , \qquad (34)$$

where, for $i = 1, \ldots, N$, $\alpha_i^{\Delta} = [s_i^{\Delta *} - s_0^{\Delta *}]/k_i \Delta^{\frac{1}{2}}$, $\alpha_0^{\Delta} = s_0^{\Delta *} - \sum_{i=1}^{N} \alpha_i^{\Delta} \mu^{\Delta *}$, and $Z_i^{\Delta}(\cdot)$ is the continuous-time representation of the process $\tilde{Z}_i^{\Delta}(\cdot)$ defined in (27). Under the assumptions of Theorem 2 one easily sees that for any subsequence $\{\mu^{\Delta' *}\}$ of $\{\mu^{\Delta *}\}$ that converges to a limit μ^*, the corresponding subsequence of coefficients $\{\alpha_0^{\Delta'}, \alpha_1^{\Delta'}, \ldots, \alpha_n^{\Delta'}\}$ in (34) converges to a limit $(\alpha_0, \alpha_1, \ldots, \alpha_N)$ such that $\alpha_i = \hat{c}_i(\mu^*)$ for $i = 1, \ldots, N$, and the incentive scheme

$$s(Z(\cdot)) = \alpha_0 + \sum_{i=1}^{N} \alpha_i Z_i(1) = \alpha_0 + \sum_{i=1}^{N} \hat{c}_i(\mu^*) Z_i(1) \qquad (35)$$

implements the control path $\mu^*(\cdot)$ with constant value μ^* in the continuous model. Moreover, (35) and (32) are equivalent in the sense that $s^* = s(Z(\cdot))$ if $Z(\cdot)$ is the stochastic process satisfying $Z(t) \equiv \mu \cdot t + X(t)$.[11]

5 Linearity in Aggregate Profits

So far the optimal incentive schemes in the discrete multi-period and in the corresponding Brownian model are linear only in accounts, not in total gross profits. For total incentive payments in (35) to be representable as a linear function of the principal's cumulative total gross profits $\sum_{i=1}^{N} Z_i(1)$, we should require that

$$\hat{c}_1(\mu^*) = \ldots = \hat{c}_N(\mu^*). \qquad (36)$$

As Holmström and Milgrom point out, (36) is trivially true if $N = 1$ and the static model involves only two possible outcomes ("success" and "failure"). In this case, the principal's

[11] In terms of the processes $X^{\Delta}(\cdot)$ and $X(\cdot)$ that are defined in (18) and (19) and Theorem 1, (34) and (35) become $\tilde{s}^{\Delta} = \hat{\alpha}_0^{\Delta} + \sum_{i=1}^{N} \alpha_i^{\Delta} X_i^{\Delta}(1)$, and $\tilde{s} = \hat{\alpha}_0 + \sum_{i=1}^{N} \alpha_i X_i(1)$, where, for $i = 1, \ldots, N$, α_i^{Δ} and α_i are the same as before and $\hat{\alpha}_0^{\Delta} = \alpha_0^{\Delta} + \sum_{i=1}^{N} \alpha_i^{\Delta} \mu_i^{\Delta *}$, $\hat{\alpha}_0 = \alpha_0 + \sum_{i=1}^{N} \alpha_i \mu_i^*$.

cumulative total gross profit is just $Z_1(1)$, and (35) itself represents $\tilde{s}$ as a linear function of $Z_1(1)$. However, this case is very special.

Holmström and Milgrom also observe that for $N > 1$ (36) is again trivially true if the effort cost function takes the form $\hat{c}(\mu) = g(E(\mu))$, where $E(\mu) := \sum_{i=1}^{N} \mu_i$. In this case (35) yields

$$\tilde{s} = \alpha_0 + g'\left(\sum_{i=1}^{N} \mu_i\right) \sum_{i=1}^{N} Z_i(1) , \tag{37}$$

so again, $\tilde{s}$ is represented as a linear function of the principal's cumulative total gross profits.

However, the specification $\hat{c}(\mu) = g(E(\mu))$ is inconsistent with our Assumptions 1 and 2 as well as Holmström and Milgrom's assumption that in the static model the principal wants to implement an action in the relative interior of the simplex P. In Hellwig and Schmidt (2001) we show that in a discrete-time model with this effort cost specification, the principal never wants to implement an action in the relative interior of the feasible set.[12] Any limit of a sequence of optimal control paths in the discrete-time models must therefore lie on the boundary of the feasible set; the corresponding optimal incentive scheme in the continuous-time model has a representation of the form

$$\tilde{s} = \alpha_0 + \sum_{i=1}^{N} \left[g'\left(\sum_{i=1}^{N} \mu_i\right) - d_i \right] Z_i(1) \tag{38}$$

where $(d_1, ..., d_N)$ is a vector of shadow values to the agent of being constrained at the boundary of the feasible set, and, typically d_i differs across dimensions. On the boundary of the feasible set, incentive constraints are less demanding; there is no need to discourage deviations that would lead out of the feasible set. This can be used by the principal to reduce the agent's salary risk without adverse incentive effects; the variables $d_1, ..., d_N$ in (38) serve this very purpose.[13]

In their most interesting approach to obtaining linearity of incentive schemes in total

[12]The intuition for this is fairly simple. For any given level of expected total profits, the overall surplus increases if the riskiness of the agent's salary is reduced. Starting from any path of actions in the interior of the feasible set, in the discrete-time models, such a reduction of the agent's salary risk can be obtained by shifting probability mass in any period to the two outcomes with profit levels just adjacent to the desired target for expected profits for that period.

[13]Holmström and Milgrom (1987) assume that in the continuous-time model, the set of feasible actions for the agent is open. For the given effort cost specification, this yields (37), but then the continuous-time model has no discrete-time counterparts.

profits, Holmström and Milgrom impose the additional assumption that the principal cannot observe the time paths of the individual accounts $Z_i(t), i = 1, ..., N$, contributing to total profits, but only the time path of the aggregate $z(t) = \sum_{i=1}^{N} Z_i(t)$ and hence of total profits. Given that in the continuous-time model, the disturbance terms are independent of the agent's actions, in this case, any change of the control process $\mu(\cdot)$ that leaves the drift rate process $(\sum_{i=1}^{N} \mu_i(\cdot))$ of the accounting aggregate $z(\cdot)$ unchanged will remain undetected by the principal and will leave the distribution of payments to the agent unchanged. Therefore the agent will always choose the control process $\mu(\cdot)$ so as to achieve a given target $E(\cdot)$ for the aggregate drift rate process $(\sum_{i=1}^{N} \mu_i(\cdot))$ at minimal cost to himself. At an interior solution, the first-order condition for this minimization problem requires that the marginal costs of raising expected gross profits be the same in all accounts, which is just condition (36). As shown in Theorem 8 of Holmström and Milgrom (1987), the optimal incentive scheme in the continuous-time model must then be linear in cumulative total profits.

We state this result formally. As before, let $\hat{K} \subset I\!R^N$ be the set of admissible control vectors. Let $\underline{E}, \overline{E}$ be the minimum and the maximum of $E(\mu)$ on $\hat{K}$. For any $E \in [\underline{E}, \overline{E}]$, let $\gamma(E)$ be the minimum of $\hat{c}(\mu)$ for $\mu \in \hat{K}$ satisfying $E(\mu) = E$, and let $\mu(E) \in \hat{K}$ be such that $\hat{c}(\mu(E)) = \gamma(E)$ and $E(\mu(E)) = E$. By standard arguments, $\gamma(\cdot)$ is a continuously differentiable, strictly convex function on $[\underline{E}, \overline{E}]$. The derivative of $\gamma(\cdot)$ is denoted as $\gamma'(\cdot)$.

Proposition 1 *Impose Assumption 2 as well as Assumption 1. Suppose also that E^* maximizes the expression*

$$E - \gamma(E) - \frac{r}{2}\gamma'(E)^2\sigma^2 \tag{39}$$

over $[\underline{E}, \overline{E}]$. Then $\mu(E^*) \in K$ and, for all i, $\hat{c}_i(\mu(E^*)) = \gamma'(E^*) \geq 0$. Moreover, the control path $\mu^{**}(\cdot)$ with the constant value $\mu^{**}(t) = \mu(E^*)$ solves the principal's problem in the continuous-time model when controls are restricted to $\hat{K}$ and the principal observes the aggregate process $z(\cdot) = \sum_{i=1}^{N} Z_i(\cdot)$, but not its individual components. In terms of the cumulative profits process $Z(\cdot)$, the optimal incentive scheme s^{**} has the representation

$$s^{**} = \gamma(E^*) + \gamma'(E^*)\left(\sum_{i=1i}^{N} Z_i(1) - E^*\right) + \frac{r}{2}\gamma'(E^*)^2\sigma^2. \tag{40}$$

In the remainder of this section, we consider the relevance of this result for our discrete-time approximations of the continuous-time model. We begin by noting that there is no immediate discrete-time analogue. To see this, note that in any period $\tau \in \{1, \ldots, m\}$ the principal observes $\sum_{\tau'=1}^{\tau} \pi^{\Delta,\tau'}$. As the profit in period τ is simply $\pi^{\Delta,\tau} = \sum_{\tau'=1}^{\tau} \pi^{\Delta,\tau'} - \sum_{\tau'=1}^{\tau-1} \pi^{\Delta,\tau'}$, the principal can compute the individual accounts $\tilde{Z}_i^{\Delta,\tau}$ even if he observes the time path of total profits only. Since this information can be used to improve the incentive scheme, the principal will use it as before. Thus we again get linearity in accounts only. As this is true for any $\Delta > 0$, there seems to be a discontinuity in the transition from discrete to continuous time.

Indeed, such a discontinuity may appear even if the principal observes only the final profit aggregate $z^\Delta = \sum_{i=1}^{N} Z_i^\Delta(1)$. For suppose that in the static model, there are three different profit levels, $\pi \in \{-1, \frac{1}{e}, +1\}$, where e is Euler's number. For any Δ then, the decimal expansion of the aggregate z^Δ reveals how often the profit level $\frac{1}{e}\Delta^{\frac{1}{2}}$ has been realized; this in turn makes it possible to infer how often the profit levels $-\Delta^{\frac{1}{2}}$ and $+\Delta^{\frac{1}{2}}$ have been realized, so again the principal can infer the quantities $Z_i^\Delta(1)$ that he needs to know in order to implement the scheme (35).

To avoid such discontinuities in the transition from discrete to continuous time, we need to impose a stronger form of information asymmetry. Going beyond Holmström and Milgrom, we assume that the principal does not observe the time path of profits and, moreover, that the agent is able destroy profits before he reports them to the principal. The principal observes only the undestroyed profit $\hat{z} \leq \sum_{i=0}^{N} Z_i^\Delta(1)$ at date 1. He might ask the agent to report not only undestroyed total profits but also the time path of profits.[14] Even more generally, he could ask for a message w from some general message space and make the payment $S^\Delta(w, \hat{z})$ depend on this message as well as the level $\hat{z}$ of undestroyed total profits. However, the revelation principle implies that there is nothing to be gained by such a procedure. This is the point of the following result.

Proposition 2 *Suppose that in the discrete-time model with period length Δ the principal can implement a (possibly stochastic or history dependent)*

[14] Like Holmström and Milgrom (1987) we rule out the possibility that the principal requires the agent to continuously report profits at any points in time $t \in [0, 1]$. The idea is that the principal cannot monitor the agent continuously but only at exogenously given discrete points in time.

control path $\mu^{\Delta}(\cdot)$ with an incentive scheme $S^{\Delta}(\cdot,\cdot)$, which makes the payment to the agent, $S^{\Delta}(w,\hat{z})$, depend on some message w as well as undestroyed total profits $\hat{z}$. Then there exists a nondecreasing incentive scheme $s^{\Delta}(\cdot)$ such that, when faced with the scheme $s^{\Delta}(\cdot)$, the agent is willing to choose the control path $\mu^{\Delta}(\cdot)$ and truthfully report total profits as $\hat{z} = \sum_{i=0}^{N} Z_i^{\Delta}(1)$. The resulting payoff for the agent is the same and for the principal is no less than under the implementation of $\mu^{\Delta}(\cdot)$ through the scheme $S^{\Delta}(\cdot,\cdot)$.

Proposition 2 shows that, without loss of generality, we can restrict attention to incentive schemes that depend only on reported total profits. Moreover we may suppose that these incentive schemes are nondecreasing, inducing the agent to report total profits truthfully as $\hat{z} = \sum_{i=0}^{N} Z_i^{\Delta}(1)$, without any destruction. These are in fact the only additional restrictions that the information asymmetry introduces. If an incentive scheme implements a control path $\mu^{\Delta}(\cdot)$ under the information assumptions of Section 2 and if this incentive scheme can be represented by a nondecreasing function of cumulative total profits, then it also implements $\mu^{\Delta}(\cdot)$ under the information assumptions in this section. To see this, note that the incentive scheme is feasible for the principal because it depends only on total profits. Moreover, as the incentive scheme is nondecreasing, the agent is willing to report total profits truthfully. Finally, when the agent chooses the control path $\mu(\cdot)$, he faces exactly the same incentives under the information assumptions of Section 2 and of this section.

The assumption that the agent can destroy profits unnoticed and, therefore, that the incentive scheme must be monotonic has strong implications. For example, any nonlinear incentive scheme that is *stationary* violates monotonicity if the number of periods is sufficiently large.[15] For the model of Section 2, Holmström and Milgrom (1987, Theorem 5) have shown that the optimal incentive scheme is stationary. However, in the model

[15]A simple example illustrates this point. Suppose that there are three different profit levels, $\pi_0 < \pi_1 < \pi_2$, and three corresponding incentive payments, $s_0 < s_1 < s_2$, in the one period model, and normalize $\pi_0 = s_0 = 0$. Suppose the incentive scheme is not linear, i.e. $\frac{s_1}{\pi_1} \neq \frac{s_2 - s_1}{\pi_2 - \pi_1}$. Consider the case where $\frac{s_1}{\pi_1} < \frac{s_2 - s_1}{\pi_2 - \pi_1}$. Then there exist numbers $m_1, m_2 \in I\!N$ such that $m_1\pi_1 > m_2\pi_2$ and $m_1 s_1 < m_2 s_2$. If the number of periods, m, satisfies $m > \max\{m_1, m_2\}$ and if the incentive scheme is stationary, monotonicity is violated. To see this consider two time paths of profit realizations: In path A there are m_1 realizations of π_1 and $m - m_1$ realizations of π_0. In path B there are m_2 realizations of π_2 and $m - m_2$ realizations of π_0. Hence, total profit after path A is larger than after path B, but the agent's total renumeration is smaller.

considered in this section, where the agent's renumeration has to be a nondecreasing function of total profits, the optimal incentive scheme need not be stationary. A nonlinear incentive scheme might enable the principal to neutralize at least some of the effects of having to condition on total profits rather than on each outcome separately.

Nevertheless, if the incentive scheme has to be nondecreasing, the scope for doing so is limited when there are many short periods. To see where monotonicity makes a difference, go back to the example with profit levels $\pi \in \{-1, \frac{1}{e}, 1\}$ in the static model, and suppose that, in any discrete-time model with period length Δ, the principal would like to penalize outcomes $\pi^\Delta = \frac{1}{e}\Delta^{\frac{1}{2}}$ and to reward outcomes $\pi^\Delta = \Delta^{\frac{1}{2}}$ in any period. Without the monotonicity assumption the corresponding incentive schemes as a function of cumulative total profits would fluctuate ever more wildly as Δ becomes small; such a sequence of incentive schemes has no convergent subsequence. In contrast, if the agent can destroy output unnoticed, it is impossible to penalize, e.g., an outcome involving ten realizations of $\pi^\Delta = \frac{1}{e}\Delta^{\frac{1}{2}}$ as opposed to an outcome involving six realizations of $\pi^\Delta = \Delta^{\frac{1}{2}}$ and four realizations of $\pi^\Delta = -\Delta^{\frac{1}{2}}$. Indeed, as shown in the following result, monotonicity automatically implies the existence of a convergent subsequence. It is this property that is crucial to establish Theorem 3 below.

Proposition 3 *Suppose that in the discrete-time model with period length Δ, the principal can implement a control process $\mu^\Delta(\cdot)$ with a nondecreasing incentive scheme $s^\Delta(\cdot)$. Suppose also that the means of the resulting final payments $s^\Delta(\sum_{i=0}^N Z_i^\Delta(1))$ are uniformly bounded. Then there exists a nondecreasing, upper semi-continuous function $s(\cdot)$ such that, for some subsequence $\{s^{\Delta'}(\cdot)\}$ of the sequence $\{s^\Delta(\cdot)\}$,*

$$\lim_{\bar{\Delta}\to 0} \sup_{\Delta' \leq \bar{\Delta}} s^{\Delta'}(z) \leq s(z) \text{ for all } z, \tag{41}$$

and

$$\lim_{\Delta' \to 0} s^{\Delta'}(z) = s(z) \text{ for all } z \text{ at which } s(\cdot) \text{ is continuous.} \tag{42}$$

Given a convergent subsequence of incentive schemes, another application of Berge's maximum theorem shows that as the length of the period goes to zero, the control processes chosen by the agent converge to the control process which he chooses when he is faced with the limit scheme in the continuous-time model. The argument is much more complex than in Section 4 because for nonlinear incentive schemes, the solutions to the agent's optimization problems will typically involve nonstationary, time-dependent and history-dependent controls. Even so, the maximum theorem can be applied, and we can infer that the restriction to nondecreasing incentive schemes eliminates certain discontinuities in the transition from discrete to continuous time.

Convergence of control processes is discussed in terms of cumulative controls. Instead of processes $\mu^{\Delta}(\cdot)$, $\mu(\cdot)$, which refer to the time paths of actions taken by the agent, we look at the associated cumulative processes $M^{\Delta}(\cdot)$ and $M(\cdot)$ where, for any $t \in [0,1]$

$$M^{\Delta}(t) = \int_0^t \mu^{\Delta}(t')dt' \text{ and } M(t) = \int_0^t \mu(t')dt'. \tag{43}$$

We think of the processes $M^{\Delta}(\cdot)$ and $M(\cdot)$ as random functions taking values in the space C^N where C is again the space of continuous real-valued functions on $[0,1]$.

Proposition 4 *Impose Assumption 2 as well as Assumption 1. Suppose that in the discrete-time model with period length Δ, the principal can implement a control process $\mu^{\Delta}(\cdot)$ with a nondecreasing incentive scheme $s^{\Delta}(\cdot)$. Suppose also that the processes $\mu^{\Delta}(\cdot)$ take values in the interior of the set $\hat{K}$ and that there exists a nondecreasing, upper semi-continuous function $s(\cdot)$ such that (41) and (42) are satisfied for some subsequence $\{s^{\Delta'}(\cdot)\}$ of the sequence $\{s^{\Delta}(\cdot)\}$. Then there exists a control process $\mu(\cdot)$ in the continuous-time model such that (i) the joint distributions of the cumulative control processes $M^{\Delta'}(\cdot)$ associated with $\mu^{\Delta'}(\cdot)$ and the processes $X^{\Delta'}(\cdot)$ of cumulative deviations from the mean converge to the joint distribution of the cumulative control process $M(\cdot)$ associated with $\mu(\cdot)$ and the driftless Brownian motion $X(\cdot)$ that is specified in Theorem 1, and (ii) the incentive scheme $s(\cdot)$ implements $\mu(\cdot)$ in the continuous-time model. Moreover, as Δ' converges to zero, the expected incentive payment $Es^{\Delta'}(\sum_{i=0}^{N} Z_i^{\Delta}(1))$ of the principal to the agent under the control process $\mu^{\Delta'}(\cdot)$ in the discrete-time model with period length Δ' converges to the expected incentive payment $Es(\sum_{i=0}^{N} Z_i(1))$to the agent under the control process $\mu(\cdot)$ in the continuous-time model.*

Propositions 3 and 4 imply that whatever the principal can implement in a discrete-time model with many short periods will not be very different from something he can implement in the continuous-time model. Given that in the continuous-time model, a linear incentive scheme is optimal, it follows that not much is to be gained from a nonlinear incentive scheme, i.e., a suitably chosen linear incentive scheme is approximately optimal in a discrete-time model when the number of periods is large and the length of each period is small. The linear incentive schemes that we consider take the form

$$s^{\Delta**}(\hat{z}, \eta) = \gamma(E^*) + \gamma'(E^*)\,[\hat{z} - E^*] + \frac{r}{2}\gamma'(E^*)^2\sigma^2 + \eta \tag{44}$$

where E^* is the optimal value of the aggregate drift rate in Proposition 2 and $\eta > 0$ is a sweetener that we introduce to avoid running afoul of the agent's participation constraint in the discrete-time model.

Theorem 3 *Impose Assumption 2 as well as Assumption 1, and suppose that controls are restricted to lie in the interior of $\hat{K}$. If the principal is restricted to using incentive schemes that can be represented as nondecreasing functions of cumulative total gross profits, then for any $\varepsilon > 0$, there exist $\underline{\Delta} > 0$ and $\eta > 0$ such that for any $\Delta < \underline{\Delta}$, the principal's payoff from using the linear incentive scheme $s^{\Delta**}(\cdot, \eta)$ that is given by (44) is no more than ε below his payoff from the optimal incentive scheme in the discrete-time model with period length Δ.*

Theorem 3 shows that the optimal incentive scheme in the continuous model (plus an arbitrarily small constant η is ε-optimal in the discrete model if Δ is sufficiently small. It is our discrete-time counterpart to the optimality of linear incentive schemes in continuous-time models that is established in Theorem 8 of Holmström and Milgrom (1987) or in Proposition 1 above. Like those results, Theorem 3 reflects the basic intuition that nonlinear incentive schemes are vulnerable to manipulation. First, because profits are aggregated over time and across accounts and because the agent can destroy profits unnoticed, the principal has to use an incentive scheme that is a nondecreasing function of total profits. As the number of periods increases, this severely restricts his ability to infer the actual profit realizations by exploiting the discrete structure of the game. Second, as the principal cannot distinguish between profits generated by different accounts, he cannot prevent the agent from spreading his effort across accounts so that the cost to him of generating a given time path of expected profits is minimal; the equality (36) that permits the aggregation across accounts in (35) is derived from this behaviour of the agent rather than any structural property of the cost function. Because of these limitations on implementability, the choice of an optimal incentive scheme is eventually driven by the stationarity of preferences and technologies that push the principal to providing the agent with stationary incentives that are independent of time and histories; as mentioned above, a monotonic incentive scheme that provides for stationary incentives is necessarily linear.

Note that the principal's inability to obtain the relevant information about individual accounts stems not only from the fact that he observes only an aggregate across accounts, but also from the aggregation over time that results from his observing only cumulative totals at the end and from the agent's ability to destroy profits unnoticed. Aggregation over time also plays a role in the continuous-time analysis of Holmström and Milgrom. When the principal is said to "observe" the time path of a Brownian motion $Z(\cdot) = (Z_1(\cdot), ..., Z_N(\cdot))$ or $z(\cdot) = \sum_{i=1}^{N} Z_i(\cdot)$, he is not actually observing any incremental changes that could be interpreted as "rates of growth" of cumulative profits at any given instants. After all, with probability one, the sample paths of the processes $Z(\cdot)$ and $z(\cdot)$ are almost nowhere differentiable, and the stochastic differential equation $dz = \sum_{i=1}^{N} dZ_i$ is meaningful only in terms of its integral version

$$z(t_2) - z(t_1) = \sum_{i=1}^{N} \int_{t_1}^{t_2} dZ_i,$$

involving aggregation over time as well as across accounts.

As for the assumption that the agent can destroy profits unnoticed, we suspect that this can be replaced by something else. This assumption restricts the principal to non-decreasing incentive schemes, which in turn guarantees the existence of a convergent subsequence of such schemes. It seems likely that other assumptions on the information available to the principal can serve the same purpose.

This should also be true of the restriction to interior controls. In the present setting, in contrast to Section 4, we have not been able to endogenize this condition, which we use to ensure that as we vary the length of the period the agent's payoffs are uniformly integrable and there is no discontinuous change in expected payoffs. It should be possible though to obtain the desired uniform integrability of payoffs in some other way.

6 Concluding Remarks

The unified framework developed in this paper makes it possible to directly compare the discrete-time and the continuous-time models of Holmström and Milgrom and to explicitly derive the latter as a limit of the former with ever shorter periods. For a variant of the discrete-time model, in which the agent can shift profits over time and across accounts and in which he may destroy profits unnoticed, we have shown that an incentive scheme that is linear in total profits is approximately optimal if the length of the period is sufficiently small. This result reflects the intuition that a linear incentive scheme is desirable because it is robust against manipulation of information by the agent.

The Holmström-Milgrom paper has given rise to a large literature, but only a few papers are concerned with the methodological and mathematical underpinnings of the analysis. Most papers in the literature just appeal to their results to justify the use of linear incentive schemes in applications; for an example see Holmström and Milgrom (1991). Among the more method-oriented papers, Schättler and Sung (1993) develop a general mathematical framework for the study of agency problems when the agent controls the drift of a Brownian motion in continuous time; we heavily rely on their results, which strengthen and greatly extend the continuous-time results of Holmström and Milgrom (1987). Sung (1995) further extends the analysis by allowing for moral hazard with respect to risk choices, more precisely, choices concerning the diffusion parameters of a Brownian motion. Sung (1997) develops the corresponding analysis for continuous-time agency problems concerning jump processes. Bolton and Harris (1997) also consider continuous-time problems concerning jump processes as well as diffusions. However they are concerned with first-best rather than second-best problems, considering optimal risk sharing and optimal actions without concern for incentive compatibility and showing that for arbitrary preference specifications risk-sharing considerations will typically call for nonlinear contracts in a first-best setting.

The above-mentioned papers all work directly in continuous time and do not discuss the relation between static or multi-period discrete-time and continuous-time agency models. Multi-period discrete-time agency models and their relation to continuous-time models are studied by Schättler and Sung (1997) and by Müller (2000). Unlike Holmström and Milgrom (1987) these papers consider multi-period models as m-fold repetitions of a static model in which the agent's effort choice determines the mean of a normally distributed random variable. Müller (2000) shows that, as in Mirrlees (1974), in a model of this type, the principal's problem typically does not have a solution because a first-best allocation can be approximated (but not reached) by a sequence of incentive schemes using penalties for low outcome realizations to discourage shirking, the incidence of penalties becoming ever rarer and the penalties themselves becoming ever more severe as one goes along the sequence.[16] In Schättler and Sung (1997), existence of a solution to the principal's problem in discrete time is also a problem. To get around it the authors impose a restriction on the class of incentive schemes they admit; with this restriction, they find that conditions for implementing a given strategy of the agent are similar in an m-period discrete-time and in a continuous-time Brownian-motion model. Neither paper asks how a continuous-time Brownian-motion model would be approximated by discrete-time models with finitely many possible outcomes in any one period or what is the relation between linearity of incentive payments in "accounts" and linearity of incentive payments in outcome aggregates, e.g. in profits.

Finally, we would like to point out that our procedure for linking discrete-time and continuous-time models is not the only one to be considered. Different approaches to what one means by the "length of the period" yield different approaches to the relation between discrete-time and continuous-time models; for some of these the continuous-time limits may have very different properties. In the analysis here, the length of the period was introduced through a particular parameterization of the agent's effort cost function $c^{\Delta}(\cdot)$ and of the profit levels π_i^{Δ} associated with different outcomes. As an alternative to the approach considered here, Hellwig (2001) considers the effort cost specification

$$c^{\Delta}(p) = \Delta \bar{c}\left(\frac{E(p)}{\Delta}, \frac{V(p)}{\Delta}\right), \tag{45}$$

where $E(p)$ is the mean and $V(p)$ is the variance of profits in the given period under the action p. He also considers a profit level specification $\pi_i^{\Delta} = \pi_i$, regardless of Δ, in

[16] In Holmström and Milgrom (1987), this difficulty was avoided by the assumption that the set of states of the world in the static model is finite.

addition to the specification (3), $\pi_i^\Delta = \pi_i \Delta^{\frac{1}{2}}$ that was considered here. The effort cost specification (45) is a generalization of the specification considered in Section 5. The difficulties considered there are avoided if $\lim_{V\to 0} \bar{c}(E,V) = \infty$. With this assumption the discrete-time models have well-defined continuous-time limits. However, depending on the specification of profit levels at different outcomes, the qualitative properties of the continuous-time limits are quite different: If $\pi_i^\Delta = \pi_i$, regardless of Δ, the continuous-time limit of optimal outcome processes in the discrete-time models is a jump process; in this case optimal incentive schemes are bounded away from being linear. If $\pi_i^\Delta = \pi_i \Delta^{\frac{1}{2}}$, the continuous-time limit is a Brownian motion with drift, this time with endogenous diffusion parameters. In this case, optimal incentive schemes in the continuous-time limit are linear and a version of Proposition 1 is true even without an additional restriction on the principal's information.

These results suggest that there may be several "canonical" continuous-time agency models all of which embody the notion that the environment is stationary and the agent takes frequent decisions, being all the time fully aware of the history up to the present. As one would expect from the theory of infinitely divisible processes, continuous-time specifications involving jump processes can be approximated just as specifications involving Brownian motion, but whenever a specification like $\pi_i^\Delta = \pi_i \Delta^{\frac{1}{2}}$ is available to guarantee that sample paths in the continuous-time limit are almost surely continuous, the continuous-time specification will involve Brownian motion. Among such specifications the one considered by Holmström and Milgrom stands out for the exogeneity of the diffusion parameters. The effort cost parametrization (4) provides the basis for approximating just this class of models.

Appendix A

Proof of Theorem 1: Part (a) is a special case of Proposition B.1 in Appendix B. To prove (b), suppose that the agent takes the time path of actions $\mu^{\Delta,\tau}$, $\tau = 1, \ldots, m$, where $\mu^{\Delta}(t)$ converges to $\mu(t)$ uniformly in t as Δ goes to zero. Since $\mu(t)$ is continuous, there exists a compact set $K \subset I\!R^N$ such that for any sufficiently small Δ one has $\mu^{\Delta}(t) \in K$ for all t. Since $\hat{c}(\cdot)$ is continuous - and hence bounded on K - we find for the agent's cost in the limit as Δ goes to 0:

$$\lim_{\Delta\to 0} \sum_{\tau=1}^{m} \Delta \hat{c}(\mu^{\Delta,\tau}) \;=\; \lim_{\Delta\to 0} \int_{\Delta}^{1} \hat{c}(\mu^{\Delta}(t))dt \;=\; \int_0^1 \hat{c}(\mu(t))dt \;. \tag{A.1}$$

As for (c), substituting $\tilde{X}_i^{\Delta,\tau}$ in (17) yields:

$$\begin{aligned}\sum_{\tau=1}^{m} \tilde{s}^{\Delta,\tau} \;&=\; \sum_{\tau=1}^{\Delta} \Delta \hat{c}(\mu^{\Delta,\tau}) + \sum_{\tau=1}^{m}\sum_{i=1}^{N} \hat{c}_i(\mu^{\Delta,\tau})\tilde{X}_i^{\Delta,\tau} \\ &\quad + \frac{r}{2}\sum_{\tau=1}^{m}\sum_{i=0}^{N} \Delta \tilde{A}_i^{\Delta,\tau}\left[\hat{c}_i(\mu^{\Delta,\tau})k_i - \sum_{j=1}^{N}\hat{p}_j\hat{c}_j(\mu^{\Delta,\tau})k_j\right]^2 + O(\Delta^{\frac{1}{2}})\end{aligned} \tag{A.2}$$

From (b), the first term in (A.2) converges to $\int \hat{c}(\mu(t))dt$ as Δ converges to zero. As for the second term, the same argument as in (b) implies that for any sufficiently small Δ, $\mu^{\Delta,\tau}$ belongs to a compact set K for all τ, and $\hat{c}_i(\mu^{\Delta,\tau})$ is bounded, uniformly in τ. From part (a) therefore, one finds that for any sufficiently small Δ', the pairs

$$\left(X^{\Delta}(\cdot), \sum_{i=0}^{N}\sum_{\tau=1}^{m} \hat{c}_i(\mu^{\Delta',[\tau\Delta/\Delta']})\tilde{X}_i^{\Delta,\tau}\right) \;=\; \left(X^{\Delta}(\cdot),\; \sum_{i=0}^{N}\sum_{\tau=1}^{m} \hat{c}_i(\mu^{\Delta',\tau})\left[X_i^{\Delta}(\tau\Delta') - X_i^{\Delta}(\tau\Delta' - \Delta')\right]\right) \tag{A.3}$$

converge in distribution to

$$\left(X(\cdot), \sum_{i=0}^{N}\sum_{\tau=1}^{m} \hat{c}_i(\mu^{\Delta',\tau})\left[X_i(\tau\Delta') - X_i(\tau\Delta' - \Delta')\right]\right) \;=\; \left(X(\cdot), \sum_{i=0}^{N}\int_0^1 \hat{c}_i(\mu^{\Delta'}(t))dX_i(t)\right)\;, \tag{A.4}$$

as Δ goes to zero, uniformly in Δ'. One also has $\lim_{\Delta'\to 0}\hat{c}_i(\mu^{\Delta'}(t)) = \hat{c}_i(\mu(t))$ uniformly in t; hence $\sum_{i=0}^{N}\int \hat{c}_i(\mu(t))dX_i(t)$ converges in square mean to $\int \hat{c}_i(\mu^{\Delta'}(t))dX_i(t)$ as Δ' converges to zero (see, e.g. Gikhman and Skorokhod (1972), p. 233). Therefore the joint distributions of the processes $X^{\Delta}(\cdot)$ and the sums $\sum_{i=0}^{N}\sum_{\tau}\hat{c}_i(\mu^{\Delta,\tau})X_i^{\Delta,\tau}$ in the second term of (A.2) converge to the joint distribution of the process $X(\cdot)$ and the stochastic integral $\sum_{i=0}^{N}\int \hat{c}_i(\mu(t))dX_i(t)$ as Δ goes to zero.

Turning to the third term in (A.2), we rewrite this in the form

$$\frac{r}{2}\sum_{\tau=1}^{m}\sum_{i=0}^{N} \; \Delta \; \left(\tilde{A}_i^{\Delta,\tau} - p_i^{\Delta}(\mu^{\Delta,\tau})\right)\left[\hat{c}_i(\mu^{\Delta,\tau})k_i - \sum_{j=1}^{N}\hat{p}_j\hat{c}_j(\mu^{\Delta,\tau})k_j\right]^2$$

$$+ \; \frac{r}{2}\sum_{\tau=1}^{m}\sum_{i=0}^{N} p_i^{\Delta}(\mu^{\Delta,\tau})\left[\hat{c}_i(\mu^{\Delta,\tau})k_i - \sum_{j=1}^{N}\hat{p}_j\hat{c}_j(\mu^{\Delta,\tau})k_j\right]^2 \tag{A.5}$$

By the law of large numbers, the first of these terms converges to zero almost surely as Δ goes to zero. The second term converges to the integral $\frac{r}{2}\int_0^1 \hat{c}'(\mu(t))\mathbf{\Sigma}[c'(\mu(t))]^T dt$, a nonrandom constant. Thus, if the time paths of actions that the principal wants to implement converge to $\mu(\cdot)$, then the joint distributions of the processes $X^{\Delta}(\cdot)$ and the incentive payments $\tilde{s}^{\Delta}$ converge to the joint distribution of the process $X(\cdot)$ and

$$\tilde{s} \;=\; \int_0^1 \hat{c}(\mu(t))dt + \int_0^1 \hat{c}'(\mu(t))dX + \frac{r}{2}\int_0^1 \hat{c}'(\mu(t))\mathbf{\Sigma}[c'(\mu(t))]^T dt \;. \tag{A.6}$$

Q.E.D.

Lemma 1 follows from the argument in the text and does not require any further argument.

<u>Proof of Lemma 2:</u> For any Δ, let $s^{\Delta*} = (s_0^{\Delta*}, ..., s_N^{\Delta*})$ be the incentive scheme implementing $\mu^{\Delta*}$ in the static model with period length Δ, and consider the expected payoff

$$U^{P\Delta}(\mu^{\Delta*}) = \frac{1}{\Delta}\sum_{i=0}^{N} p_i^{\Delta}(\mu^{\Delta*})[\pi_i^{\Delta} - s_i^{\Delta*}] \tag{A.7}$$

that the principal gets from implementing $\mu^{\Delta*}$ in $m = \frac{1}{\Delta}$ periods of length Δ. The agent's participation constraint implies

$$-\sum_{i=0}^{N} p_i^{\Delta}(\mu^{\Delta*})\; e^{-r[s_i^{\Delta*} - c^{\Delta}(p^{\Delta}(\mu^{\Delta*}))]} \geq -1, \tag{A.8}$$

hence, by the concavity of the CARA utility function,

$$\sum_{i=0}^{N} p_i^{\Delta}(\mu^{\Delta*})\; s_i^{\Delta*} \geq c^{\Delta}(p^{\Delta}(\mu^{\Delta*})). \tag{A.9}$$

Upon combining (A.9) and (A.7) and using (12) and (13), one finds that the principal's expected payoff from implementing $\mu^{\Delta*}$ in $m = \frac{1}{\Delta}$ periods of length Δ is bounded above by

$$U^{P\Delta}(\mu^{\Delta*}) \leq \sum_{i=1}^{N} \mu_i^{\Delta*} - \hat{c}(\mu^{\Delta*}). \tag{A.10}$$

By Lemma 1 and the optimality of $\mu^{\Delta *}$ for the principal in the multi-period model with period length Δ, it must also be the case that

$$-\frac{r}{2}\sigma^2 - \hat{c}(0) - \eta \leq U^{P\Delta}(\mu^{\Delta *}) \tag{A.11}$$

whenever Δ is close to zero. Thus, by (A.10) and (A.11) and by the definition of K in Assumption 2, $\mu^{\Delta *}$ must belong to K whenever Δ is close to zero. *Q.E.D.*

The proof of Theorem 2 relies on some properties of the continuous-time model under Assumptions 1 and 2. The following proposition states these properties formally.[17]

Proposition A.1 *Impose Assumption 2 as well as Assumption 1. Let μ^{**} maximize the expression*

$$U^P(\mu) = \sum_{i=1}^N \mu_i - \hat{c}(\mu) - \frac{r}{2}\sum_{i=1}^N\sum_{j=1}^N \hat{c}_i(\mu)\sigma_{ij}\hat{c}_j(\mu)$$

over the set $\hat{K}$. Then $\mu^{**} \in K$. Moreover the control path $\mu^{**}(\cdot)$ with the constant value μ^{**} is a solution to the principal's problem in the continuous-time model when controls are restricted to $\hat{K}$. When faced with the incentive scheme $s(\hat{Z}(\cdot), \mu^{**}(\cdot))$ that is given by (26) for $\mu(\cdot) = \mu^{**}(\cdot)$, the agent strictly prefers the control path $\mu^{**}(\cdot)$ to any alternative path that is distinct from it.

Proof of Proposition A.1: We first show that $\mu^{**} \in K$. For this purpose we claim that there exists a control $\hat{\mu} \in K$ such that

$$U^P(\hat{\mu}) \geq -\frac{r}{2}\sigma^2 - \hat{c}(0). \tag{A.12}$$

From the definitions of $U^P(\cdot)$ and μ^{**}, we have

$$\sum_{i=1}^N \mu_i^{**} - \hat{c}(\mu^{**}) > U^P(\mu^{**}) \geq U^P(\hat{\mu}), \tag{A.13}$$

so (A.12) implies $\sum_{i=1}^N \mu_i^{**} - \hat{c}(\mu^{**}) > -\frac{r}{2}\sigma^2 - \hat{c}(0)$ and hence $\mu^{**} \in K$.

[17] The key property is that under Assumption 2, any solution to the principal's problem necessarily takes values in the interior of the admissible set $\hat{K}$. In Schättler and Sung (1993), this is simply assumed, with some indication of the tension between this assumption and the requirement that the admissible set be bounded.

The desired $\hat{\mu}$ is obtained as the solution to the problem

$$\max_{\mu \in \hat{K}} \left[\sum_{i=1}^{N} \mu_i - \hat{c}(\mu) \right]. \tag{A.14}$$

Because $0 \in K$, we have $\sum_{i=1}^{N} \hat{\mu}_i - \hat{c}(\hat{\mu}) \geq -\hat{c}(0) > -\frac{r}{2}\sigma^2 - \hat{c}(0)$. Therefore $\hat{\mu} \in K$, i.e., the solution to the maximization problem (A.14) lies in the interior of $\hat{K}$ and satisfies the first-order conditions $\hat{c}_i(\hat{\mu}) = 1$ for $i = 1, ..., N$. Thus

$$U^P(\hat{\mu}) \;=\; \sum_{i=1}^{N} \hat{\mu}_i - \hat{c}(\hat{\mu}) - \frac{r}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \sigma_{ij} \;\geq\; -\hat{c}(0) - \frac{r}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \sigma_{ij}\ . \tag{A.15}$$

From (21) and (5), one easily verifies that $\sum_{i=1}^{N} \sum_{j=1}^{N} \sigma_{ij} = \sum_{i=0}^{N} \hat{p}_i \pi_i^2 = \sigma^2$, so (A.12) follows. As mentioned above, this implies $\mu^{**} \in K$.

Next, we show that the process $\mu^{**}(\cdot)$ with the constant value μ^{**} solves the principal's problem in the continuous-time model. Consider any admissible[18] control process $\mu(\cdot)$ taking values in $\hat{K}$ and incentive scheme $\tilde{s}$ that implements $\mu(\cdot)$ relative to $\hat{K}$. By a result of Schättler and Sung (1993, Theorem 4.1), there exist adapted processes $q_i(\cdot)$, $i = 1, ..., N$, taking values in $I\!R$ such that in terms of the control process $\mu(\cdot)$, the implementing incentive scheme has the representation

$$\tilde{s} = \int_0^1 \hat{c}(\mu(t))dt + \int_0^1 \sum_{i=1}^{N} q_i(t)\ dX_i(t) + \frac{r}{2} \int_0^1 \sum_{i=1}^{N} \sum_{j=1}^{N} q_i(t)\sigma_{ij} q_j(t)\ dt; \tag{A.16}$$

moreover $q_i(t) = \hat{c}_i(\mu(t))$ whenever $\mu(t)$ belongs to the interior of $\hat{K}$. This yields

$$E\left[\int_0^1 \sum_{i=1}^{N} \mu_i(t) dt - \tilde{s} \right] = E \int_0^1 \left[\sum_{i=1}^{N} \mu_i(t) - \hat{c}(\mu(t)) - \frac{r}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} q_i(t)\sigma_{ij} q_j(t) \right] dt \tag{A.17}$$

for the principal's expected payoff from implementing the process $\mu(\cdot)$. Consider the integrand on the right-hand side of (A.17). For any t and any event in which $\mu(t)$ belongs to K, by the definition of μ^{**}, we have

$$\begin{aligned} \sum_{i=1}^{N} \mu_i(t) \;&-\; \hat{c}(\mu(t)) - \frac{r}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} q_i(t)\sigma_{ij} q_j(t) \\ &=\; \sum_{i=1}^{N} \mu_i(t) - \hat{c}(\mu(t)) - \frac{r}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \hat{c}_i(\mu(t))\sigma_{ij} \hat{c}_j(\mu(t)) \\ &=\; U^P(\mu(t)) \leq U^P(\mu^{**}). \end{aligned}$$

[18]For a precise definition of of admissibility of a control process, as well as other details of the principal's problem in continuous time, we refer the reader to Schättler and Sung (1993).

For any t and any event in which $\mu(t)$ belongs to $I\!R^N \backslash K$, by the definition of K in combination with (A.12) and (A.13), we have

$$\begin{aligned}\sum_{i=1}^{N} \mu_i(t) \quad &- \quad \hat{c}(\mu(t)) - \frac{r}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} q_i(t) \sigma_{ij} q_j(t) \\ &\leq \quad \sum_{i=1}^{N} \mu_i(t) - \hat{c}(\mu(t)) < -\frac{r}{2}\sigma^2 - \hat{c}(0) \leq U^P(\mu^{**}) \ .\end{aligned}$$

The integrand on the right-hand side of (A.17) is thus bounded above by $U^P(\mu^{**})$ for all t almost surely. Therefore

$$U^P(\mu^{**}) \geq E\left[\int_0^1 \sum_{i=1}^{N} \mu_i(t)dt - \tilde{s}\right] \tag{A.18}$$

for any admissible process $\mu(\cdot)$ and incentive scheme $\tilde{s}$ that implements $\mu(\cdot)$. From (26) applied to the process $\mu^{**}(\cdot)$, one also has

$$U^P(\mu^{**}) = E\left[\int_0^1 \sum_{i=1}^{N} \mu_i^{**}(t)dt - \tilde{s}\right] \tag{A.19}$$

when $\tilde{s} = s(Z(\cdot), \mu^{**}(\cdot))$ is the scheme that implements $\mu^{**}(\cdot)$. Together, (A.19) and (A.18) show that the control process $\mu^{**}(\cdot)$ solves the principal's problem in the continuous-time model when controls are restricted to $\hat{K}$.

To conclude the proof of Proposition A.1, we observe that by the results of Schättler and Sung (1993, Theorem 4.1 and formula (25), any alternative control process $\hat{\mu}(\cdot)$ that also solves the agent's problem at the incentive scheme $s(\cdot, \mu^{**}(\cdot))$ that is given by (26) applied to the process $\mu(\cdot) = \mu^{**}(\cdot)$ must have the property that for almost all $t \in [0, 1]$, with probability one, the expression $\sum_{i=1}^{N} \mu_i \hat{c}_i(\mu^{**}) - \hat{c}(\mu)$ is maximized at $\mu = \hat{\mu}(t)$. Because $\hat{c}(\cdot)$ is strictly convex, it follows that any such alternative solution to the agent's problem must satisfy $\hat{\mu}(t) = \mu^{**}$ for almost all $t \in [0, 1]$ almost surely, i.e., the process $\hat{\mu}(\cdot)$ must be indistinguishable from $\mu^{**}(\cdot)$. *Q.E.D.*

Proof of Theorem 2: Relying on the compactness of $\hat{K}$, let $\{\mu^{\Delta' *}\}$ be a convergent subsequence of $\{\mu^{\Delta *}\}$. For any Δ', let $s^{\Delta' *} = (s_0^{\Delta' *}, ..., s_N^{\Delta' *})$ be the incentive scheme implementing $\mu^{\Delta' *}$ in the static model with period length Δ', and let

$$\tilde{s}^{\Delta' *} = \sum_{\tau=1}^{m} \sum_{i=0}^{N} \tilde{A}_i^{\Delta', \tau} s_i^{\Delta' *} \tag{A.20}$$

be the associated cumulative total payment to the agent. By Lemma 2, for any sufficiently small Δ', $\mu^{\Delta'*}$ belongs to the interior of $\hat{K}$ and so does the limit μ^*. By Theorem 1, as Δ' goes to zero, the random variables $\tilde{s}^{\Delta'*}$ converge in distribution to

$$\tilde{s}^* = \hat{c}(\mu^*) + \sum_{i=1}^{N} \hat{c}_i(\mu^*)X_i(1) + \frac{r}{2}\sum_{i=1}^{N}\sum_{j=1}^{N} \hat{c}_i(\mu^*)\sigma_{ij}\hat{c}_j(\mu^*). \tag{A.21}$$

One easily verifies that the principal's expected payoffs,

$$U^{P\Delta'}(\mu^{\Delta'*}) = \frac{1}{\Delta'}\sum_{i=0}^{N} p_i^{\Delta'}(\mu^{\Delta'*})[\pi_i^{\Delta'} - s_i^{\Delta'*}] = \sum_{i=1}^{N} \mu_i^{\Delta'*} - E\tilde{s}^{\Delta'*} \tag{A.22}$$

also converge to

$$U^P(\mu^*) = \sum_{i=1}^{N} \mu_i^* - \hat{c}(\mu^*) - \frac{r}{2}\sum_{i=1}^{N}\sum_{j=1}^{N} \hat{c}_i(\mu^*)\sigma_{ij}\hat{c}_j(\mu^*). \tag{A.23}$$

By a result of Schättler and Sung (1993, Theorem 4.2) and the convexity of $\hat{c}(\cdot)$, the incentive scheme $\tilde{s}^*$ actually induces the agent to choose the control path $\mu^*(\cdot)$ with the constant value μ^*, providing the principal with the payoff $U^P(\mu^*)$. Therefore the principal's maximum payoff in the continuous-time model cannot be less than the limit of the payoffs $U^{P\Delta'}(\mu^{\Delta'*})$ in the discrete-time models.

It remains to be shown that is actually optimal for the principal to implement the control path $\mu^*(\cdot)$ in the continuous-time model. If this is not optimal, we must have

$$U^P(\mu^*) < U^P(\mu^{**}) \tag{A.24}$$

where μ^{**} is the optimal constant control that is given by Proposition A.1. For any Δ', consider an incentive scheme with cumulative total payments

$$\tilde{s}^{\Delta'**} = \sum_{\tau=1}^{m}\sum_{i=0}^{N} \tilde{A}_i^{\Delta',\tau} s_i^{\Delta'**} \tag{A.25}$$

to the agent where, for $i = 0, 1, ..., N$,

$$s_i^{\Delta'**} = \Delta'\hat{c}(\mu^{**}) - \frac{1}{r}\ln\left(1 - r\hat{c}_i(\mu^{**})k_i\Delta'^{\frac{1}{2}} + r\sum_{j=1}^{N} p_j^{\Delta'}(\mu^{**})\hat{c}_j(\mu^{**})k_j\Delta'^{\frac{1}{2}}\right). \tag{A.26}$$

The following facts are easily verified:

- When faced with the incentive scheme (A.25) in the discrete-time model with period length Δ', the agent can assure himself an expected payoff equal to -1 simply by choosing the constant action $p^{\Delta'}(\mu^{**})$. Therefore it is individually rational for him to accept the principal's proposing this incentive scheme.

- If in the discrete-time model with period length Δ', the agent is free to choose control vectors $\mu_\tau \in \hat{K}$, inducing actions $p^{\Delta'}(\mu_\tau) \in P$, for $\tau = 1, ..., m = \frac{1}{\Delta'}$, then in face of the incentive scheme (A.25), an optimal strategy for the agent is to choose a history-independent constant control $\mu_\tau = \mu^{\Delta'} \in \hat{K}$ for all τ, where $\mu^{\Delta'}$ maximizes the expected payoff

$$U^{A\Delta'}(\mu) = -\sum_{i=0}^{N} p_i^{\Delta'}(\mu)\; e^{-r(s_i^{\Delta'**} - \Delta'\hat{c}(\mu))} \tag{A.27}$$

of the agent in the model with period length Δ' and incentive payments $s_i^{\Delta'**}$, $i = 0, 1, ..., N$.

- From (A.26), one may rewrite (A.27) as

$$\begin{aligned} U^{A\Delta'}(\mu) &= -e^{-r\Delta'[\hat{c}(\mu^{**})-\hat{c}(\mu)]}\left(1 - r\sum_{i=0}^{N}[p_i^{\Delta'}(\mu) - p_i^{\Delta'}(\mu^{**})]\hat{c}_i(\mu^{**})k_i\Delta'^{\frac{1}{2}}\right) \\ &= -\left[e^{-r[\hat{c}(\mu^{**})-\hat{c}(\mu)]}\left(1 - r\sum_{i=0}^{N}[\mu_i - \mu_i^{**}]\hat{c}_i(\mu^{**})\Delta'\right)^{\frac{1}{\Delta'}}\right]^{\Delta'}. \end{aligned} \tag{A.28}$$

- As Δ' goes to zero, $(-U^{A\Delta'}(\mu))^{\frac{1}{\Delta'}}$ converges to $\exp(-r[\hat{c}(\mu^{**}) - \hat{c}(\mu) + r\sum_{i=0}^{N}[\mu_i - \mu_i^{**}]\hat{c}_i(\mu^{**})]$, uniformly for $\mu \in \hat{K}$.

- For any Δ', maximization of $U^{A\Delta'}(\mu)$ over $\hat{K}$ is equivalent to minimization of $(-U^{A\Delta'}(\mu))^{\frac{1}{\Delta'}}$ over the set $\hat{K}$, so by Berge's maximum theorem, any limit point of the sequence $\{\mu^{\Delta'}\}$ is a minimizer of $\exp(-r[\hat{c}(\mu^{**}) - \hat{c}(\mu) + r\sum_{i=0}^{N}[\mu_i - \mu_i^{**}]\hat{c}_i(\mu^{**})]$ and hence of $\hat{c}(\mu) - r\sum_{i=0}^{N}\mu_i\hat{c}_i(\mu^{**})$ over the set $\hat{K}$.

- Because $\hat{c}(\cdot)$ is strictly convex, μ^{**} is the unique minimizer of $\hat{c}(\mu) - r\sum_{i=0}^{N}\mu_i\hat{c}_i(\mu^{**})$ over the set $\hat{K}$. Therefore any limit point of the sequence $\{\mu^{\Delta'}\}$ must be equal to μ^{**}, and the sequence $\{\mu^{\Delta'}\}$ itself must be converging to μ^{**}.

- The control vectors $\mu^{\Delta'}$ belong to the interior of $\hat{K}$ whenever Δ' is sufficiently close to zero.

- By Theorem 1, applied to the control processes $\mu^{\Delta'}(\cdot)$ with constant values $\mu^{\Delta'}$, as Δ' goes to zero, the incentive payments $\tilde{s}^{\Delta'**}$ that are associated with the implementation of $\mu^{\Delta'}$ converge in distribution to

$$\tilde{s}^{**} = \hat{c}(\mu^{**}) + \sum_{i=1}^{N} \hat{c}_i(\mu^{**}) X_i(1) + \frac{r}{2} \sum_{i=1}^{N} \sum_{j=1}^{N} \hat{c}_i(\mu^{**}) \sigma_{ij} \hat{c}_j(\mu^{**}), \tag{A.29}$$

and the principal's payoff expectations $U^{P\Delta'}(\mu^{\Delta'})$ converge to $U^P(\mu^{**})$.

Given the convergence of $U^{P\Delta'}(\mu^{\Delta'})$ to $U^P(\mu^{**})$ and of $U^{P\Delta'}(\mu^{\Delta'*})$ to $U^P(\mu^*)$, the inequality (A.24) would imply that $U^{P\Delta'}(\mu^{\Delta'}) > U^{P\Delta'}(\mu^{\Delta'*})$ for any sufficiently small Δ'. As this contradicts the optimality of $\mu^{\Delta'*}$ for the principal in the discrete-time problem with period length Δ', it follows that the inequality (A.24) must be false, and $\mu^*(\cdot)$ must indeed be optimal for the principal in the continuous-time model. *Q.E.D.*

Proof of Proposition 1: To prove that $\mu(E^*) \in K$, one uses the same argument as for the claim in Proposition A.1 that $\mu^{**} \in K$. The maximizer $\hat{\mu}$ of $E(\mu) - \hat{c}(\mu)$ over $\hat{K}$ is also a minimizer of $\hat{c}(\mu)$ over $\hat{K}$ under the constraint $E(\mu) = E(\hat{\mu})$, and we have $\hat{c}(\hat{\mu}) = \gamma(E(\hat{\mu}))$ and $\gamma'(E(\hat{\mu})) = 1$. Then (A.15) implies

$$E(\hat{\mu}) - \gamma(E(\hat{\mu})) - \frac{r}{2}\gamma'(E(\hat{\mu}))^2\sigma^2 \geq -\hat{c}(0) - \frac{r}{2}\sigma^2. \tag{A.30}$$

By the definition of E^*, it follows that

$$E^* - \gamma(E^*) > E^* - \gamma(E^*) - \frac{r}{2}\gamma'(E^*)^2\sigma^2 \geq -\hat{c}(0) - \frac{r}{2}\sigma^2, \tag{A.31}$$

hence, $\mu(E^*) \in K$. By standard arguments, it follows that the first-order conditions for the agent's cost minimization problem take the form $\hat{c}_i(\mu(E^*)) = \gamma'(E^*)$ for $i = 1, ..., N$. To prove that $\gamma'(E^*) \geq 0$, we note that as $\gamma(\cdot)$ is convex, $\gamma'(\cdot)$ is nondecreasing. If $\gamma'(E^*)$ were negative, then at $E = E^*$, $\gamma'(E)^2$ would be nonincreasing, and $E - \gamma(E) - \frac{r}{2}\gamma'(E)^2\sigma^2$ would be strictly increasing, contrary to the definition of E^*.

To show that the control path $\mu^{**}(\cdot)$ with the constant value $\mu(E^*)$ solves the principal's problem under the given information assumption,we again follow the line of argument in the proof of Proposition A.1. Consider any admissible control process $\mu(\cdot)$ taking values in $\hat{K}$ and any incentive scheme $\tilde{s}$ that implements $\mu(\cdot)$ when the principal only observes the aggregate process $z(\cdot) = \sum_{i=1}^{N} Z_i(\cdot)$. By Theorem 4.1 of Schättler and Sung, there exist adapted processes $q_i(\cdot), i = 1, ..., N$, taking values in $I\!R$ such that the incentive scheme $\tilde{s}$ has the representation

$$\tilde{s} = \int_0^1 \hat{c}(\mu(t))dt + \int_0^1 \sum_{i=1}^{N} q_i(t)dZ_i(t) - \int_0^1 \sum_{i=1}^{N} q_i(t)\mu_i(t)dt \qquad \text{(A.32)}$$
$$+\frac{r}{2}\int_0^1 \sum_{i=1}^{N}\sum_{j=1}^{N} q_i(t)\sigma_{ij}q_j(t)dt.$$

If the principal observes only the process $z(\cdot) = \sum_{i=1}^{N} Z_i(\cdot)$, the processes $q_i(\cdot)$ must satisfy $q_1(t) = ... = q_N(t)$ almost surely, for all t. By Theorem 4.1 of Schättler and Sung, it must also be the case that for any t, with probability one, $\sum_{i=1}^{N} q_i(t)\mu_i - \hat{c}(\mu)$ reaches a maximum over $\hat{K}$ at $\mu = \mu(t)$; with $q_1(t) = ... = q_N(t)$, this implies that for any t, with probability one, under the constraint $E(\mu) = E(\mu(t))$, $\hat{c}(\mu)$ reaches a minimum over $\hat{K}$ at $\mu = \mu(t)$. This in turn yields $\hat{c}(\mu(t)) = \gamma(E(\mu(t)))$ and, by a standard envelope argument, $q_1(t) = ... = q_N(t) = \gamma'(E(\mu(t)))$. Then (A.32) can be rewritten as:

$$\tilde{s} = \int_0^1 \gamma(E(\mu(t)))dt + \int_0^1 \gamma'(E(\mu(t))) \sum_{i=1}^{N} dZ_i(t) - \int_0^1 \gamma'(E(\mu(t)))E(\mu(t))dt$$
$$+\frac{r}{2}\int_0^1 \gamma'(E(\mu(t)))^2\sigma^2 dt. \qquad \text{(A.33)}$$

Conditional on the process $E(\mu(\cdot))$, the principal's net expected payoff from implementing the control process $\mu(\cdot)$ through the incentive scheme $\tilde{s}$ is then almost surely equal to

$$\int_0^1 E(\mu(t))dt - \int_0^1 \gamma(E(\mu(t)))dt - \frac{r}{2}\int_0^1 \gamma'(E(\mu(t)))^2\sigma^2 dt. \qquad \text{(A.34)}$$

By the definition of E^*, this is no larger than $E^* - \gamma(E^*) - \frac{r}{2}\gamma'(E^*)^2\sigma^2$, the principal's net expected payoff from implementing the control process $\mu^{**}(\cdot)$ with constant value $\mu^{**}(t) = \mu(E^*)$. This confirms that under the given information assumption it is optimal for the principal to implement the process $\mu^{**}(\cdot)$. The representation (40) of the incentive scheme that implements $\mu^{**}(\cdot)$ is obtained by substituting for $E(\mu(t)) = E^*$ in (A.33) and using the fact that $\int_0^1 \sum_{i=1}^{N} dZ_i(t) = \sum_{i=1}^{N}[Z_i(1) - Z_i(0)] = \sum_{i=1}^{N} Z_i(1)$. *Q.E.D.*

<u>Proof of Proposition 2:</u> Given the incentive scheme $S^{\Delta}(\cdot,\cdot)$, let $(w(\cdot), \hat{z}(\cdot))$ be the optimal reporting strategy of the agent, and define

$$s^{\Delta}(\cdot) = S^{\Delta}(w(\cdot), \hat{z}(\cdot)) . \qquad \text{(A.35)}$$

If the true value of cumulative total profits is z, he prefers to choose the pair $(w(z), \hat{z}(z))$ as opposed to destroying profits unnoticed down to $z' < z$ and reporting $(w(z'), \hat{z}(z'))$. Therefore it must be the case that $S^{\Delta}(w(z), \hat{z}(z)) \geq S^{\Delta}(w(z'), \hat{z}(z'))$, whenever $z > z'$, so $s^{\Delta}(\cdot)$ is nondecreasing. Given that $s^{\Delta}(\cdot)$ is nondecreasing, an optimal reporting strategy $\hat{z}^*(\cdot)$ for the agent when faced with $s^{\Delta}(\cdot)$ is to tell the truth, i.e., set $\hat{z}^*(z) \equiv z$.

To show that under the given information assumptions the incentive scheme $s^{\Delta}(\cdot)$ implements the control path $\mu^{\Delta}(\cdot)$, we note that since $S^{\Delta}(\cdot, \cdot)$ implements $\mu^{\Delta}(\cdot)$, with reports given by $(w(\cdot), \hat{z}(\cdot))$, we have

$$EU^A(S^{\Delta}(w(\cdot), \hat{z}(\cdot)), \mu^{\Delta}(\cdot)) \;\geq\; EU^A(S^{\Delta}(\bar{w}(\cdot), \bar{z}(\cdot)), \hat{\mu}^{\Delta}(\cdot)) \tag{A.36}$$

for any alternative reporting strategy $(\bar{w}(\cdot), \bar{z}(\cdot))$ that is feasible and any alternative control process $\hat{\mu}^{\Delta}(\cdot)$ where $EU^A(S^{\Delta}(\bar{w}(\cdot), \bar{z}(\cdot)), \hat{\mu}^{\Delta}(\cdot))$ denotes the agent's expected payoff from choosing the control process $\hat{\mu}^{\Delta}(\cdot)$ and reporting strategy $(\bar{w}(\cdot), \bar{z}(\cdot))$ when the incentive scheme is $S^{\Delta}(\cdot, \cdot)$. In particular,

$$EU^A(S^{\Delta}(w(\cdot), \hat{z}(\cdot)), \mu^{\Delta}(\cdot)) \;\geq\; EU^A(S^{\Delta}(w(\cdot), \hat{z}(\cdot)), \hat{\mu}^{\Delta}(\cdot)) \tag{A.37}$$

and hence

$$EU^A(s^{\Delta}(\hat{z}^*(\cdot)), \mu^{\Delta}(\cdot)) \;\geq\; EU^A(s^{\Delta}(\hat{z}^*(\cdot)), \hat{\mu}^{\Delta}(\cdot)) \tag{A.38}$$

for any alternative control path $\hat{\mu}^{\Delta}(\cdot)$, where $\hat{z}^*(\cdot)$ is again the truthtelling strategy. Since $s^{\Delta}(\cdot)$ is nondecreasing, it follows that

$$EU^A(s^{\Delta}(\hat{z}^*(\cdot)), \mu^{\Delta}(\cdot)) \geq EU^A(s^{\Delta}(\bar{z}(\cdot)), \hat{\mu}^{\Delta}(\cdot)) \tag{A.39}$$

for any alternative control path $\hat{\mu}^{\Delta}(\cdot)$ and reporting strategy $\bar{z}(\cdot)$ under the incentive scheme $s^{\Delta}(\cdot)$.

Finally, since $s^{\Delta}(\cdot)$ implements the same time path of actions and yields the same payments to the agent, his expected utility is the same under $s^{\Delta}(\cdot)$ with the truthtelling strategy $\hat{z}^*(\cdot)$ as under $S^{\Delta}(\cdot, \cdot)$ with the optimal reporting strategy $(w(\cdot), \hat{z}(\cdot))$. As for the principal, he benefits when there is no destruction of profits, so on top of the payoff that he would get under $S^{\Delta}(\cdot, \cdot)$ with the reporting strategy $(w(\cdot), \hat{z}(\cdot))$ he obtains the (expected value of) the difference $z - \hat{z}(z)$, which is destroyed under the reporting strategy $(w(\cdot), \hat{z}(\cdot))$. *Q.E.D.*

Proof of Proposition 3: Proposition 3 is essentially a version of Helly's selection theorem, see, e.g. Billingsley (1968, p. 227). To apply Helly's argument, we need to show that for any $z \in I\!R$, there exists a compact interval $[\underline{s}(z), \bar{s}(z)] \subset I\!R$ such that $s^\Delta(z) \in [\underline{s}(z), \bar{s}(z)]$ for any sufficiently small Δ.

For any Δ, the fact that the incentive scheme $s^\Delta(\cdot)$ serves to implement the control process $\mu^\Delta(\cdot)$ implies that the participation constraint

$$-E \exp\{-r[s^\Delta(z^\Delta) - \int_0^1 \hat{c}(\mu^\Delta(t))dt]\} \geq -1 \tag{A.40}$$

is satisfied, where $z^\Delta := \sum_{i=1}^N Z_i^\Delta(1) = \sum_{i=1}^N [M_i^\Delta(1) + X_i^\Delta(1)]$. Given that $\mu^\Delta(t) \in \hat{K}$ with probability one for all t, we have $\underline{E} + x^\Delta \leq z^\Delta \leq \bar{E} + x^\Delta$, where, as before, $\underline{E}$ and $\bar{E}$ are the minimum and the maximum of $\sum_{i=1}^N \mu_i$ on $\hat{K}$, and $x^\Delta := \sum_{i=1}^N X_i^\Delta(1)$. By the monotonicity of $s^\Delta(\cdot)$, (A.40) implies

$$-E \exp\{-rs^\Delta(\bar{E} + x^\Delta)\} \geq -\exp\{-r\underline{c}\}, \tag{A.41}$$

where $\underline{c}$ is the minimum of $\hat{c}(\mu)$ on $\hat{K}$. By the monotonicity of $s^\Delta(\cdot)$, it follows that, for any fixed $z \in I\!R$,

$$-\exp\{-rs^\Delta(z)\}\text{prob}\{x^\Delta \leq z - \bar{E}\} \geq -\exp\{-r\underline{c}\}. \tag{A.42}$$

Proposition B.1 in Appendix B implies that, as Δ goes to zero, $\text{prob}\{x^\Delta \leq z - \bar{E}\}$ converges to $\Phi(z - \bar{E})$, where $\Phi(\cdot)$ is the cumulative normal distribution function with mean zero and variance σ^2. If we fix some $\varepsilon \in (0, \Phi(z - \bar{E}))$ and set

$$\underline{s}(z) := r\underline{c} + \ln(\Phi(z - \bar{E}) - \varepsilon), \tag{A.43}$$

we may conclude that $s^\Delta(z) \geq \underline{s}(z)$ for any sufficiently small Δ.

To obtain a uniform upper bound $\bar{s}(z)$ for $s^\Delta(z)$, we note that, by assumption, there exists a bound $\bar{S}$ such that $Es^\Delta(z^\Delta) \leq \bar{S}$ and, hence,

$$E[s^\Delta(z^\Delta)|\, z^\Delta \leq z]\ \text{prob}\{z^\Delta \leq z\} + s^\Delta(z)(1 - \text{prob}\{z^\Delta \leq z\}) \leq \bar{S}. \tag{A.44}$$

for all Δ. From (A.40), we also have

$$-E[\exp\{-rs^\Delta(z^\Delta)\}|\, z^\Delta \leq z]\ \text{prob}\{z^\Delta \leq z\} \geq -\exp\{-r\underline{c}\},$$

hence, since the agent is risk averse,

$$E[s^{\Delta}(z^{\Delta})|\ z^{\Delta} \leq z] \geq r\underline{c} + \ln \text{prob}\{z^{\Delta} \leq z\}.$$

Thus (A.44) implies

$$s^{\Delta}(z)(1 - \text{prob}\{z^{\Delta} \leq z\}) \leq \bar{S} - [r\underline{c} + \ln \text{prob}\{z^{\Delta} \leq z\}]\ \text{prob}\{z^{\Delta} \leq z\} \tag{A.45}$$

and hence

$$\begin{aligned} s^{\Delta}(z) &\leq \max\{\bar{S},\ 0,\ -[r\underline{c} + \ln \text{prob}\{z^{\Delta} \leq z\}]\} \\ &\leq \max\{\bar{S},\ 0,\ -[r\underline{c} + \ln \text{prob}\{x^{\Delta} \leq z - \underline{\text{E}}\}]\}. \end{aligned} \tag{A.46}$$

By Proposition B.1 in Appendix B, $\text{prob}\{x^{\Delta} \leq z - \underline{\text{E}}\}$ converges to $\Phi(z - \underline{\text{E}})$ as Δ converges to zero. If for some $\varepsilon > 0$, we set

$$\bar{s}(z) = \max\{\bar{S},\ 0,\ -[r\underline{c} + \ln \text{prob}\{x^{\Delta} \leq z - \underline{\text{E}}\}]\} + \varepsilon,$$

we may conclude that $s^{\Delta}(z) \leq \bar{s}(z)$ for any sufficiently small Δ.

Now for any $z \in I\!R$, the fact that $s^{\Delta}(z) \in [\underline{s}(z), \bar{s}(z)]$ for any sufficiently small Δ implies that the sequence $\{s^{\Delta}(z)\}$ has a convergent subsequence. More generally if $\{z_1, z_2, ...\}$ be a countable and dense subset of $I\!R$, then, by a standard diagonalization argument, there exists a subsequence of incentive schemes, $\{s^{\Delta'}(\cdot)\}$, such that for any i, the sequence $\{s^{\Delta'}(z_i)\}$ converges to a limit $s(z_i)$. This defines a nondecreasing function $s(\cdot)$ on the set $\{z_1, z_2, ...\}$. Given that the set $\{z_1, z_2, ...\}$ is dense in $I\!R$, an extension of $s(\cdot)$ to $I\!R$ is defined by setting

$$s(z) = \lim_{z_i \downarrow z} s(z_i) \tag{A.47}$$

where the limit is taken with respect to any sequence in $\{z_1, z_2, ...\}$ that converges to z from above. Since $s(\cdot)$ is nondecreasing on $\{z_1, z_2, ...\}$, for any $z \in I\!R$, the limit in (A.47) is well defined. Moreover, from (A.47), one easily sees that $s(\cdot)$ is everywhere nondecreasing and right-continuous, hence upper semi-continuous. From the mononicity of the incentive schemes $s^{\Delta'}(\cdot)$, one also finds that for any $\varepsilon > 0$ and any z,

$$s(z - \varepsilon) \leq \liminf s^{\Delta'}(z) \leq \limsup s^{\Delta'}(z) \leq s(z).$$

Hence $\lim_{\Delta' \to 0} s^{\Delta'}(z) = s(z)$ for any $z \in I\!R$ at which $s(\cdot)$ is continuous. *Q.E.D.*

Sketch of the Proof of Proposition 4:[19] As discussed in the text, for any control processes $\mu^{\Delta}(\cdot)$, $\mu(\cdot)$ in the discrete-time models and the continuous-time model, the associated cumulative control processes can be thought of as random variables that take values in the space C^N of continuous-functions from the time interval $[0,1]$ into $I\!R^N$. If controls are restricted to the subset $\hat{K}$ of $I\!R^N$, the realizations of the cumulative control processes $M^{\Delta}(\cdot)$ and $M(\cdot)$ must actually belong to the subspace $C^N_{\hat{K}}$ of functions F on $[0,1]$ that have a representation

$$F(t) = \int_0^t f_F(t')dt' \tag{A.48}$$

for some function $f_F(\cdot)$ from $[0,1]$ into $\hat{K}$. For any such function F, we can define

$$\Gamma(F) = \int_0^1 \hat{c}(f_F(t))dt \tag{A.49}$$

as the associated cumulative total effort cost. Given that the set $\hat{K}$ is compact, the functions in $C^N_{\hat{K}}$ are uniformly bounded and equicontinuous, and the function space $C^N_{\hat{K}}$ is precompact. Indeed, because $\hat{K}$ is convex and the effort cost function $\hat{c}(\cdot)$ is convex, one can show that $C^N_{\hat{K}}$ is actually compact and the cost function $\Gamma : C^N_{\hat{K}} \to I\!R$ that is defined by (A.49) is lower semi-continuous.

Let $\{(M^{\Delta'}(\cdot), X^{\Delta'}(\cdot))\}$ be the sequence of cumulative-control and cumulative-deviations processes that corresponds to the convergent subsequence of incentive schemes. Compactness of the space $C^N_{\hat{K}}$ is used to show that any subsequence $\{(M^{\Delta''}(\cdot), X^{\Delta''}(\cdot))\}$ of the sequence $\{(M^{\Delta'}(\cdot), X^{\Delta'}(\cdot))\}$ has a further subsequence which converges in distribution to a pair $(M(\cdot), X(\cdot))$ such that $X(\cdot)$ is the driftless Brownian motion specified in Theorem 1. The distributions $\Phi^{\Delta''}$ of the joint processes $(M^{\Delta''}(\cdot), X^{\Delta''}(\cdot))$ are supported by the product $C^N_{\hat{K}} \times C^N$. By Proposition B.1 in Appendix B, the sequence of marginal distributions of the cumulative-deviations processes $X^{\Delta''}(\cdot)$ converges to the distribution of the process $X(\cdot)$, so by Prohorov's Theorem (Billingsley, 1968, p. 240), this sequence is tight. Given that $C^N_{\hat{K}}$ is compact, this implies that the sequence $\{\Phi^{\Delta''}\}$ itself is tight. By Prohorov's Theorem again, it follows that the sequence $\{\Phi^{\Delta''}\}$ has a subsequence which converges to a distribution Φ^{∞} on $C^N_{\hat{K}} \times C^N$. The marginal distribution on C^N that is induced by Φ^{∞} must obviously coincide with the distribution of the process $X(\cdot)$.

[19]The full proof is available upon request or can be downloaded directly from http://www.vwl.uni-muenchen.de/ls_schmidt/research/disc/disc.htm.

The limit process $(M(\cdot), X(\cdot))$ has the additional property that for any t, the history of the process $(M(\cdot), X(\cdot))$ up to t and the continuation of the process $X(\cdot)$ from t on are conditionally independent given $X(t)$, i.e., the cumulative-control process $M(\cdot)$ does not require the agent ever to rely on information about the continuation of the process $X(\cdot)$ that he cannot be presumed to have. One easily verifies that for any $t \in (0,1)$, the argument in the proof of Proposition B.1 in Appendix B remains valid if the distributions of the processes $X^{\Delta}(\cdot) = (X_1^{\Delta}(\cdot), ..., X_N^{\Delta}(\cdot))$ are replaced by the (regular) conditional distributions of the continuations $\{X^{\Delta}(t')\}_{t' \in (t.1]} = \{(X_1^{\Delta}(t'), ..., X_N^{\Delta}(\cdot))\}_{t' \in (t.1]}$ given the histories of the joint processes $\{(M^{\Delta}(\cdot), X^{\Delta}(\cdot))\}$ up to time $\Delta[t/\Delta]$.[20] By an argument given in Hellwig (1996), this convergence property of conditional distributions is sufficient for the desired conditional-independence property in the limit.

Given that any subsequence $\{(M^{\Delta''}(\cdot), X^{\Delta''}(\cdot))\}$ of the sequence $\{(M^{\Delta'}(\cdot), X^{\Delta'}(\cdot))\}$ has a subsequence which converges in distribution, convergence of the sequence $\Big\{(M^{\Delta'}(\cdot), X^{\Delta'}(\cdot))\Big\}$ itself is assured if all convergent subsequences have the same limit. For this it suffices to show that the control process $\mu(\cdot)$ which corresponds to the cumulative control process $M(\cdot)$ is actually an optimal control process for the agent in the continuous-time model with incentive scheme $s(\cdot)$. This is so because, with a strictly convex cost function, by the results of Schättler and Sung (1993, Theorems 4.1 and 4.2), *regardless of the incentive scheme,* the solution to the agent's problem in the continuous-time model is unique in the sense that any two optimal control processes must have the same values almost surely for almost all $t \in [0, 1]$.

In an abuse of notation, we write $\{(M^{\Delta'}(\cdot), X^{\Delta'}(\cdot))\}$ for whichever convergent subsequence is to be considered, and we show that if $\{(M^{\Delta'}(\cdot), X^{\Delta'}(\cdot))\}$ converges in distribution to a pair $(M(\cdot), X(\cdot))$, then the limiting cumulative control process $M(\cdot)$ must almost surely coincide with the process $M^*(\cdot)$ which is optimal for the agent in the continuous-time model with incentive scheme $s(\cdot)$. The argument proceeds in three steps.

Step 1: The agent's expected payoff from the limit pair $(M(\cdot), X(\cdot))$ in the continuous-time model with incentive scheme $s(\cdot)$ is not significantly worse than his expected payoffs

[20] These regular conditional distributions themselves can be regarded as measure-valued random variables, i.e., functions from the appropriate spaces of histories into spaces of measures on continuations. In the absence of a common underlying probability space, convergence in distribution of these measure-valued random variables then is the appropriate convergence concept; for details see Hellwig (1996).

from the pairs $\{(M^{\Delta'}(\cdot), X^{\Delta'}(\cdot))\}$ in the discrete-time models with incentive schemes $s^{\Delta'}$ when Δ' is small. Formally,

$$\begin{aligned}
&\lim_{\bar{\Delta}\to\infty} \sup_{\Delta'\leq\bar{\Delta}} [-E\exp\{-r[s^{\Delta'}(z^{\Delta'}) - \Gamma(M^{\Delta'}(\cdot))]\}] \\
\leq \quad & -E\exp\{-r[s\left(\sum_{i=1}^{N}(M_i(1) + X_i(1))\right) - \Gamma(M(\cdot))]\},
\end{aligned} \tag{A.50}$$

where, as before, $z^{\Delta'} = \sum_{i=1}^{N} Z_i^{\Delta'}(1) = \sum_{i=1}^{N}[M_i^{\Delta'}(1) + X_i^{\Delta'}(1)]$.

Step 2: The agent's payoff from the limit pair $(M(\cdot), X(\cdot))$ in the continuous-time model with incentive scheme $s(\cdot)$ is no better than his maximal payoff in the continuous-time model with this incentive scheme. Formally,

$$\begin{aligned}
& -E\exp\{-r[s\left(\sum_{i=1}^{N}(M_i(1) + X_i(1))\right) - \Gamma(M(\cdot))]\} \\
\leq \quad & -E\exp\{-r[s\left(\sum_{i=1}^{N}(M_i^*(1) + X_i(1))\right) - \Gamma(M^*(\cdot))]\},
\end{aligned} \tag{A.51}$$

where $M^*(\cdot)$ is the optimal cumulative control process in the continuous-time model with incentive scheme $s(\cdot)$.

Step 3: The agent's maximal payoff in the continuous-time model with the incentive scheme $s(\cdot)$ is not significantly better than his maximal payoffs in the discrete-time models with incentive schemes $s^{\Delta'}(\cdot)$ when Δ' is small. Formally,

$$\begin{aligned}
& -E\exp\{-r[s\left(\sum_{i=1}^{N}(M_i^*(1) + X_i(1))\right) - \Gamma(M^*(\cdot))]\} \\
& \quad \leq \lim_{\bar{\Delta}\to\infty} \inf_{\Delta'\leq\bar{\Delta}} [-E\exp\{-r[s^{\Delta'}(z^{\Delta'}) - \Gamma(M^{\Delta'}(\cdot))]\}].
\end{aligned} \tag{A.52}$$

Given that

$$\begin{aligned}
& \lim_{\bar{\Delta}\to\infty} \inf_{\Delta'\leq\bar{\Delta}} [-E\exp\{-r[s^{\Delta'}(z^{\Delta'}) - \Gamma(M^{\Delta'}(\cdot))]\}] \\
& \quad \leq \lim_{\bar{\Delta}\to\infty} \sup_{\Delta'\leq\bar{\Delta}} [-E\exp\{-r[s^{\Delta'}(z^{\Delta'}) - \Gamma(M^{\Delta'}(\cdot))]\}],
\end{aligned}$$

the three inequalities (A.50) - (A.52) imply that

$$\begin{aligned}
& -E\exp\{-r[s\left(\sum_{i=1}^{N}(M_i(1) + X_i(1))\right) - \Gamma(M(\cdot))]\} \\
& \quad = -E\exp\{-r[s\left(\sum_{i=1}^{N}(M_i^*(1) + X_i(1))\right) - \Gamma(M^*(\cdot))]\},
\end{aligned} \tag{A.53}$$

i.e., that the limit pair $(M(\cdot), X(\cdot))$ provides the agent with the maximal payoff that he can obtain in the continuous-time model with incentive scheme $s(\cdot)$. By the results of Schättler and Sung mentioned above (1993, Theorems 4.1 and 4.2), this in turn implies that $M(\cdot) = M^*(\cdot)$ almost surely. The three inequalities (A.50) - (A.52) are therefore sufficient to establish Proposition 4.

Turning to the arguments required to prove these inequalities themselves, the proof of (A.50) is routine. Given the convergence in distribution of the sequence $\{(M^{\Delta'}(\cdot), X^{\Delta'}(\cdot))\}$, one uses (41), the upper semi-continuity of $s(\cdot)$ and the lower semi-continuity of $\Gamma(\cdot)$ as well as Skorokhod's theorem, the change-of-variables formula, and Fatou's lemma (Hildenbrand (1974), pp. 50 and 46). The proof of (A.51) involves a version of Kuhn's theorem to the effect that, with the independence property specified above, the limit pair $(M(\cdot), X(\cdot))$ is payoff-equivalent to a mixed strategy in the sense of a prior randomization of pure strategies of the agent, where a *pure strategy* is defined as a (suitably adapted) function from the space C^N of realizations of the process of cumulative deviations from the mean to the space $C^N_{\hat{K}}$ of realizations of the cumulative control process. Any pure strategy in the support of the mixed strategy corresponding to $(M(\cdot), X(\cdot))$, is dominated by the agent's optimal pure strategy in the continuous-time model with incentive scheme $s(\cdot)$. This implies that the mixed strategy itself - and with it the pair $(M(\cdot), X(\cdot))$ - is so dominated, which is just (A.51).

The proof of (A.52) starts from the optimal control process $\mu^*(\cdot)$ at the incentive scheme $s(\cdot)$ in the continuous-time model. The idea is to construct approximations of this process for the different discrete-time models so that for small period lengths expected payoffs from these approximations are not significantly worse than expected payoffs from choosing the process $\mu^*(\cdot)$ in the continuous-time model with incentive scheme $s(\cdot)$. Trivially, this implies that for small period lengths expected payoffs from *optimal* control processses in the discrete-time models are also not significantly worse than the agent's payoff from choosing the process $\mu^*(\cdot)$ in the continuous-time model with incentive scheme $s(\cdot)$. Construction of the desired approximations is routine except for the need to take account of possible discontinuities in the (pure) strategy that generates the process $\mu^*(\cdot)$ in the continuous-time model and in the incentive scheme $s(\cdot)$. Relying on Lusin's theorem one avoids the former by first approximating the process $\mu^*(\cdot)$ in the continuous-time model by some other process $\mu^{*\eta}(\cdot)$ generated by a *continuous* strategy. To avoid the latter one has the agent add just a little $\varepsilon > 0$ of effort so that the approximations we are considering lead him "to the right" of whatever discontinuity point the limiting incentive scheme may have.

Once the construction of approximating control processes is set up, the remainder of the argument is routine, along the same lines as in the proof of (A.50). However, where the proof of (A.50) can do with Fatou's lemma, the corresponding argument here needs Lebesgue's convergence theorem, hence uniform integrability of the agent's payoffs under the approximating control processes. Specifically, it must not be the case that under the approximating control processes the incentive schemes $s^{\Delta'}(\cdot)$ involve ever larger penalties with ever smaller probabilities affecting expected payoffs along the sequence, but not in the limit. Under the assumption that the original control processes $\mu^{\Delta'}(\cdot)$ take values in the interior of the feasible set, this uniform-integrability condition is satisfied because the first-order conditions for interior control choices of the agent relate payments to (marginal) costs, and these are uniformly bounded on the compact set $\hat{K}$. *Q.E.D.*

Proof of Theorem 3: We first consider the implications of the incentive schemes $s^{\Delta**}(\cdot, \eta)$, $\eta > 0$, in the discrete-time model with period length Δ. Suppose first that controls can be chosen from the boundary as well as the interior of $\hat{K}$. Then one easily verifies that, when faced with this scheme, the agent is willing to choose a control path $\mu^{\Delta**}(\cdot)$ with a constant value $\mu^{\Delta**}$, which is independent of η. Moreover, as Δ goes to zero, $\mu^{\Delta**}$ converges to $\mu(E^*)$ as specified in Proposition 1. Since $\mu(E^*)$ belongs to the interior of $\hat{K}$, for any sufficiently small Δ, $\mu^{\Delta**}$ also belongs to the interior of $\hat{K}$, and the process $\mu^{\Delta**}(\cdot)$ is optimal for the agent even if boundary controls are ruled out. By Theorem 1, the certainty equivalents of the agent's payoffs converge to $\eta > 0$, so for any sufficiently small Δ, the participation constraint is also satisfied. This shows that for any $\eta > 0$ and any sufficiently small Δ, by offering the incentive scheme $s^{\Delta**}(\cdot, \eta)$, the principal obtains the expected payoff

$$\sum_{i=1}^{N} \mu_i^{\Delta**} - \gamma(E^*) - \gamma'(E^*)(E(\mu^{\Delta**}) - E^*) - \frac{r}{2}\gamma'(E^*)^2\sigma^2 - \eta.$$

As Δ goes to zero, this converges to $E^* - \gamma(E^*) - \frac{r}{2}\gamma'(E^*)^2\sigma^2 - \eta$, so for any $\eta > 0$ and any sufficiently small Δ, by offering the incentive scheme $s^{\Delta**}(\cdot, \eta)$, the principal can assure himself an expected payoff which is no less than $E^* - \gamma(E^*) - \frac{r}{2}\gamma'(E^*)^2\sigma^2 - 2\eta$.

Now suppose that the theorem is false. Then there exist $\varepsilon > 0$ and a sequence $\{\Delta'\}$ of period lengths such that for every Δ', there exists an incentive scheme $s^{\Delta'}(\cdot)$ implementing a control process $\mu^{\Delta'}(\cdot)$ with values in the interior of $\hat{K}$ so that the principal obtains an expected payoff equal to at least $E^* - \gamma(E^*) - \frac{r}{2}\gamma'(E^*)^2\sigma^2 + \frac{\varepsilon}{2}$. Given this uniform bound on the principal's expected payoffs, expected incentive payments satisfy

$$Es^{\Delta'}(z^{\Delta'}) \leq \bar{E} - [E^* - \gamma(E^*) - \frac{r}{2}\gamma'(E^*)^2\sigma^2 + \frac{\varepsilon}{2}] \tag{A.54}$$

and are thus uniformly bounded. By Proposition 3, it follows that there exists a non-decreasing, upper semi-continuous function $s(\cdot)$ such that some subsequence of the sequence $\{s^{\Delta'}(\cdot)\}$, converges to $s(\cdot)$ in the sense of (41) and (42). By Proposition 4, it follows that the pairs $(M^{\Delta'}(\cdot), X^{\Delta'}(\cdot))$ of cumulative control and deviations processes associated with the processes $\mu^{\Delta'}(\cdot)$ converge in distribution to a pair $(M(\cdot), X(\cdot))$ of cumulative control and deviations processes associated with the optimal control process $\mu(\cdot)$ in the continuous-time model with incentive scheme $s(\cdot)$. Moreover the expected incentive payments $Es^{\Delta'}(\sum_{i=1}^{N}[M_i^{\Delta'}(1) + X_i^{\Delta'}(1)])$ converge to the expected incentive payment $Es(\sum_{i=1}^{N}[M_i(1) + X_i(1)])$ in the continuous-time model. Then we must have

$$\lim_{\Delta' \to 0} E\left[\sum_{i=1}^{N} M_i^{\Delta'}(1) - s^{\Delta'}(\sum_{i=1}^{N}[M_i^{\Delta'}(1) + X_i^{\Delta'}(1)])\right] = E\left[\sum_{i=1}^{N} M_i(1) - s(\sum_{i=1}^{N}[M_i(1) + X_i(1)])\right] \tag{A.55}$$

However, by the definition of E^* and the optimality of $\mu(\cdot)$ in the continuous-time model with incentive scheme $s(\cdot)$, we also have

$$E\left[\sum_{i=1}^{N} M_i(1) - s(\sum_{i=1}^{N}[M_i(1) + X_i(1)])\right] \leq E^* - \gamma(E^*) - \frac{r}{2}\gamma'(E^*)^2\sigma^2. \tag{A.56}$$

From (A.55) and (A.56) together, one concludes that for any sufficiently small Δ'

$$E\left[\sum_{i=1}^{N} M_i^{\Delta'}(1) - s^{\Delta'}(\sum_{i=1}^{N}[M_i^{\Delta'}(1) + X_i^{\Delta'}(1)])\right] \leq E^* - \gamma(E^*) - \frac{r}{2}\gamma'(E^*)^2\sigma^2 + \frac{\varepsilon}{4}. \tag{A.57}$$

This contradicts the assumption that the principal's expected payoffs from using the incentive schemes $s^{\Delta'}(\cdot)$ in the discrete-time models are bounded below by $E^* - \gamma(E^*) - \frac{r}{2}\gamma'(E^*)^2\sigma^2 + \frac{\varepsilon}{2}$. The assumption that the theorem is false has thus led to a contradiction, which proves the theorem. *Q.E.D.*

Appendix B

In this appendix we discuss the relationship between discrete-time and continuous-time strategies and outcome processes when the agent's choice at any one time may depend on the history of the process up to this time. In the m-period discrete-time model with period length $\Delta = \frac{1}{m}$, the agent chooses a sequence $\{\tilde{\mu}^{\Delta,\tau}\}_{\tau=1}^{m}$ of possibly history-dependent controls. This choice generates a sequence $\{\tilde{\pi}^{\Delta,\tau}\}_{\tau=1}^{m}$ of random profit levels. We assume that for each τ, the agent's control choice $\tilde{\mu}^{\Delta,\tau}$ is given by a - possibly degenerate - function of profit realizations prior to τ, i.e., that we can write

$$\tilde{\mu}^{\Delta,\tau} \;=\; \hat{\mu}^{\Delta,\tau}(\tilde{\pi}^{\Delta,1}, ..., \tilde{\pi}^{\Delta,\tau-1}) \tag{B.1}$$

for all τ, and we identify the agent's strategy with the sequence of functions $\{\hat{\mu}^{\Delta,\tau}\}_{\tau=1}^{m}$. The range of each function $\hat{\mu}^{\Delta,\tau}$ is taken to be the compact set $\hat{K} \subset I\!R^N$. Given the control strategy $\{\hat{\mu}^{\Delta,\tau}\}_{\tau=1}^{m}$ the profit sequence $\{\tilde{\pi}^{\Delta,\tau}\}_{\tau=1}^{m}$ is assumed to satisfy the condition that for each $\tau \in \{1, ..., m\}$, $\tilde{\pi}^{\Delta,\tau}$ takes values in $\{\pi_0\Delta^{\frac{1}{2}}, ..., \pi_N\Delta^{\frac{1}{2}}\}$ and, for given $\pi^{\Delta,1}, ..., \pi^{\Delta,\tau-1}$

$$\begin{aligned}\text{Prob}(\{\tilde{\pi}^{\Delta,\tau} = \pi_i\Delta^{\frac{1}{2}}\} \quad &| \quad \tilde{\pi}^{\Delta,1} = \pi^{\Delta,1}, ..., \tilde{\pi}^{\Delta,\tau-1} = \pi^{\Delta,\tau-1}) \\ &= \; p_i^{\Delta}(\hat{\mu}^{\Delta,\tau}(\pi^{\Delta,1}, ..., \pi^{\Delta,\tau-1}))\end{aligned} \tag{B.2}$$

As in the text, a given control strategy $\{\hat{\mu}^{\Delta,\tau}\}_{\tau=1}^{m}$ and associated profit sequence $\{\tilde{\pi}^{\Delta,\tau}\}_{\tau=1}^{m}$ are used to define counting variables $\tilde{A}_i^{\Delta,\tau}, i = 0, 1, ..., N, \tau = 1, ..., m, m = \frac{1}{\Delta} = 1, 2, ...$ such that $\tilde{A}_i^{\Delta,\tau} = 1 \Leftrightarrow \tilde{\pi}^{\Delta,\tau} = \pi_i\Delta^{\frac{1}{2}}$ and $\tilde{A}_i^{\Delta,\tau} = 0 \Leftrightarrow \tilde{\pi}^{\Delta,\tau} \neq \pi_i\Delta^{\frac{1}{2}}$. From (B.2) one obviously has:

$$E[\tilde{A}_i^{\Delta,\tau}|\tilde{\pi}^{\Delta,1} \;=\; \pi^{\Delta,1}, ..., \tilde{\pi}^{\Delta,\tau-1} \;=\; \pi^{\Delta,\tau-1}] \;\equiv\; p_i^{\Delta}(\hat{\mu}^{\Delta,\tau}(\pi^{\Delta,1}, ..., \pi^{\Delta,\tau-1})) \tag{B.3}$$

It follows that if we write

$$\tilde{\mathbf{X}}_i^{\Delta}(\tau) \;=\; k_i\Delta^{\frac{1}{2}}(\tilde{A}_i^{\Delta,\tau} - p_i^{\Delta}(\hat{\mu}^{\Delta,\tau}(\pi^{\Delta,1}, ..., \pi^{\Delta,\tau-1})) \tag{B.4}$$

for the cumulative deviations of realized from expected profits in dimension i, we have

$$E[\tilde{\mathbf{X}}_i^{\Delta}(\tau)|\tilde{\pi}^{\Delta,1} = \pi^{\Delta,1}, ..., \tilde{\pi}^{\Delta,\tau-1} \;=\; \pi^{\Delta,\tau-1}] \equiv 0 \tag{B.5}$$

i.e., for any Δ the process $\{(\tilde{X}_1^\Delta(\tau), ..., \tilde{X}_N^\Delta(\tau))\}_{\tau=1}^m$ of cumulative deviations from the means is a martingale. Again using a linear interpolation to obtain a continuous-time representation of this process, we write for any i and any $t \in [0,1]$:

$$X_i^\Delta(t) \;=\; \left(1 - \frac{t}{\Delta} + \left[\frac{t}{\Delta}\right]\right) \tilde{X}_i^\Delta\left(\left[\frac{t}{\Delta}\right]\right) + \left(\frac{t}{\Delta} - \left[\frac{t}{\Delta}\right]\right) \tilde{X}_i^\Delta\left(\left[\frac{t}{\Delta}\right] + 1\right) \qquad \text{(B.6)}$$

where again $[\frac{t}{\Delta}]$ denotes the greatest integer less than or equal to $\frac{t}{\Delta}$. For any Δ, $(X_1^\Delta(\cdot), ..., X_N^\Delta(\cdot))$ is a random function taking values in the space C^N of continuous functions from $[0,1]$ into $I\!R^N$. The distribution of $(X_1^\Delta(\cdot), ..., X_N^\Delta(\cdot))$ depends on the chosen control strategy of the agent. Taking C^N to be endowed with the topology of uniform convergence, we obtain:

Proposition B.1 *For $m = 1, 2, \ldots,$ and $\Delta = \frac{1}{m}$, let $\{\hat{\mu}^{\Delta,\tau}\}_{\tau=1}^m$ be a control strategy taking values in the compact set $\hat{K}$, and consider the induced process $(X_1^\Delta(\cdot), ..., X_N^\Delta(\cdot))$ of cumulative deviations from the means. As m goes out of bounds and Δ converges to zero, the processes $(X_1^\Delta(\cdot), ..., X_N^\Delta(\cdot))$ converge in distribution to the Gaussian process $X(\cdot)$ with initial value $X(0) = 0$, zero drift, and covariance matrix $\mathbf{\Sigma}$ given by (21) .*

<u>Proof:</u> For any Δ, let $\{\tilde{\pi}^{\Delta,\tau}\}_{\tau=1}^m$ be the random profit sequence that is induced by the control strategy $\{\hat{\mu}^{\Delta,\tau}\}_{\tau=1}^m$. For any Δ and any $t \in [0,1]$, let $\mathcal{F}_t^\Delta$ be the σ-algebra on the underlying probability space that is generated by the random variables $\tilde{\pi}^{\Delta,1}, ..., \tilde{\pi}^{\Delta,[t/\Delta]}$. Fix some vector $q \in I\!R^N$, and consider the stochastic process $X_q^\Delta(\cdot)$ on $[0,1]$ such that for any $t \in [0,1]$,

$$X_q^\Delta(t) \;=\; \sum_{i=1}^N q_i X_i^\Delta(t) \qquad \text{(B.7)}$$

For any t and Δ, (B.5) and (B.6) yield:

$$E[X_q^\Delta([t/\Delta]\Delta + \Delta) - X_q^\Delta([t/\Delta]\Delta)|\mathcal{F}_t^\Delta] \;=\; E[\sum_{i=1}^N q_i \tilde{X}_i^{\Delta,[t/\Delta]+\Delta}|\mathcal{F}_t^\Delta] \;=\; 0 \qquad \text{(B.8)}$$

For any two vectors $q^1, q^2 \in I\!R^N$, one also has:

$$\begin{aligned} E[(X_{q^1}^\Delta([t/\Delta]\Delta + \Delta) \;&-\; X_{q^1}^\Delta([t/\Delta]\Delta)(X_{q^2}^\Delta([t/\Delta]\Delta + \Delta) - X_{q^2}^\Delta([t/\Delta]\Delta)|\mathcal{F}_t^\Delta] \\ &= \; E[\sum_{i=1}^N \sum_{j=1}^N q_i^1 \tilde{X}_i^{\Delta,[t/\Delta]+\Delta} q_j^2 \tilde{X}_j^{\Delta,[t/\Delta]+\Delta}|\mathcal{F}_t^\Delta] \end{aligned}$$

$$
\begin{aligned}
&= E[\sum_{i=1}^{N} q_i^1 q_i^2 \; k_i^2 \; \Delta \; (1 - p_i^{\Delta}(\tilde{\mu}^{\Delta,[t/\Delta]})) \; p_i^{\Delta}(\tilde{\mu}^{\Delta,[t/\Delta]}) | \mathcal{F}_t^{\Delta}] \\
&\quad - E[\sum_{i=1}^{N} \sum_{j \neq i} q_i^1 q_j^2 \; k_i k_j \; \Delta \; p_i^{\Delta}(\tilde{\mu}^{\Delta,[t/\Delta]}) \; p_j^{\Delta}(\tilde{\mu}^{\Delta,[t/\Delta]}) | \mathcal{F}_t^{\Delta}] \\
&= \sum_{i=1}^{N} q_i^1 q_i^2 \; k_i^2 \; \Delta \; \hat{p}_i - \sum_{i=1}^{N} \sum_{j=1}^{N} q_i^1 q_j^2 \; k_i k_j \; \Delta \; \hat{p}_i \, \hat{p}_j + \Delta O(\Delta) \\
&= \sum_{i=1}^{N} \sum_{j=1}^{N} q_i^1 q_j^2 \; \sigma_{ij} \; \Delta \; + \Delta O(\Delta) \qquad \text{(B.9)}
\end{aligned}
$$

where σ_{ij} is the ij-th element of the matrix $\mathbf{\Sigma}$ (see (21) in the text) and $O(\Delta)$ is a term that goes to zero when Δ goes to zero, uniformly in strategies and histories.

Now consider the vector process

$$
X_Q^{\Delta} \; = \; (X_{q^1}^{\Delta}, ..., X_{q^N}^{\Delta}) \qquad \text{(B.10)}
$$

where $q^1, ..., q^N$ are eigenvectors of the matrix $\mathbf{\Sigma}$, normed so that $q^r\,' q^r = \lambda_r^{-1}$ where λ_r is the eigenvalue of $\mathbf{\Sigma}$ that corresponds to q^r. Given that $\mathbf{\Sigma}$ is a symmetric, positive definite matrix, the eigenvectors $q^r, q^s, r \neq s$, are mutually orthogonal. For this choice of $q^1, ..., q^N$, (B.9) yields:

$$
\begin{aligned}
E[(X_{q^r}^{\Delta}([t/\Delta]\Delta + \Delta) \; &- \; X_{q^r}^{\Delta}([t/\Delta]\Delta)(X_{q^s}^{\Delta}([t/\Delta]\Delta + \Delta) - X_{q^s}^{\Delta}([t/\Delta]\Delta) | \mathcal{F}_t^{\Delta}] \\
&= \sum_{i=1}^{N} \sum_{j=1}^{N} q_i^r q_j^s \; \sigma_{ij} \; \Delta \; + \Delta O(\Delta) \\
&= \sum_{i=1}^{N} q_i^r \lambda_s q_i^s \; \Delta + \Delta O(\Delta) \\
&= \delta_{rs} \; \Delta + \Delta O(\Delta) \qquad \text{(B.11)}
\end{aligned}
$$

where $\delta_{ij} = 1$ if $i = j$ and $\delta_{ij} = 0$ if $i \neq j$.

Finally we also have, for any $\varepsilon > 0$, any t and any i,

$$
\text{Prob}\{ |X_{q^i}^{\Delta}([t/\Delta]\Delta + \Delta) - X_{q^i}^{\Delta}([t/\Delta]\Delta| \geq \varepsilon | \mathcal{F}_t^{\Delta} \} \; = \; 0 \qquad \text{(B.12)}
$$

for any $\Delta < [\varepsilon / \max |k_j|]^2$.

From (B.8), (B.11), and (B.12), one easily sees that the processes $X_Q^{\Delta}, \Delta = 1, \frac{1}{2}, \frac{1}{2}, ...$ satisfy the assumptions of Theorems 6 and 7 of Gihman and Skorohod (1979, p.195). It follows that for any t, as Δ converges to zero, the conditional distributions of $\{X_Q^{\Delta}(t')\}_{t' \in [t,1]}$

given $\mathcal{F}_t^\Delta$ converge to the conditional distribution of a standard N-dimensional Brownian motion $B(\cdot)$ on $[t, 1]$ given the "initial" value $B(t)$. In particular the overall distributions of the processes X_Q^Δ on $[0, 1]$ converge to the distribution of a standard Brownian motion on $[0, 1]$ with ininital value $B(0) = 0$.

Given this convergence result, the proposition follows from Theorem 5.1 of Billingsley (1968, p.30) in conjunction with the observation that for any Δ, one has $(X_1^\Delta(\cdot), ..., X_N^\Delta(\cdot)) = (Q')^{-1} X_Q^\Delta$. *Q.E.D.*

References

BERGE, C. (1959): *Espaces Topologiques.* Paris: Dunod.

BILLINGSLEY, P. (1968): *Convergence of Probability Measures.* New York: John Wiley & Sons.

BOLTON, P., AND C. HARRIS (1997): "The Continuous-Time Principal-Agent Problem: First-Best Risk-Sharing Contracts and their Decentralization," *mimeo*, King's College, Cambridge.

GIHMAN, I. AND A.W. SKOROKHOD (1972): *Theory of Stochastic Processes*, Vol. I. Heidelberg: Springer.

GIHMAN, I., AND A. W. SKOROKHOD (1979): *Theory of Stochastic Processes*, Vol. III. Heidelberg: Springer.

GROSSMAN, S. AND O. HART (1983): "An Analysis of the Principal-Agent Problem" *Econometrica* 51, 7-45.

HELLWIG, M. (1996): "Sequential Decisions under Uncertainty and the Maximum Theorem" *Journal of Mathematical Economics* 25, 443-64.

——— (2001): "Aggregation and Linearity in Multi-Period Mean-Variance Incentive Problems," mimeo, University of Mannheim.

HELLWIG, M., AND K. M. SCHMIDT (1998): "Discrete-Time Approximations of the Holmström-Milgrom Brownian-Motion Model of Intertemporal Incentive Provision" *mimeo*, University of Munich.

——— (2001): "On the Role of Boundary Solutions in Intertemporal Principal-Agent Problems with Effort Costs Depending on Mean Returns" *mimeo*, University of Mannheim.

HILDENBRAND, W. (1974): *Core and Equilibria in a Large Economy.* Princeton: Princeton University Press.

HOLMSTRÖM, B., AND P. MILGROM (1987): "Aggregation and Linearity in the Provision of Intertemporal Incentives," *Econometrica*, 55, 303-328.

——— (1991): "Multi-Task Principal Agent Analyses," *Journal of Law Economics, and Organization*, 7 (Sp.), 24-52.

Mirrlees, J. (1974): "Notes on Welfare Economics, Information and Uncertainty," in: *Essays on Economic Behavior Under Uncertainty*, ed. by M. Balch, D. McFadden and S. Wu. Amsterdam: North-Holland, 243-58.

Mirrlees, J. (1999): "The Theory of Moral Hazard and Unobservable Behaviour: Part I" *Review of Economic Studies*, 66, 3-22.

Müller, H. (2000): "Asymptotic Efficiency in Dynamic Principal Agent Problems," *Journal of Economic Theory*, 91, 292-301.

Schättler, H., and J. Sung (1993): "The First-Order Approach to the Continuous-Time Principal-Agent Problem with Exponential Utility," *Journal of Economic Theory*, 61, 331-371.

——— (1997): "On Optimal Sharing Rules in Discrete- and Continuous-Time Principal-Agent Problems with Exponential Utility," *Journal of Economic Dynamics and Control*, 21, 551-574.

Sung, J. (1995): "Linearity with Project Selection and Controllable Diffusion Rate in Continuous-Time Principal-Agent Problems," *RAND Journal of Economics*, 26, 720-743.

——— (1997): "Corporate Insurance and Managerial Incentives," *Journal of Economic Theory*, 74, 297-332.