

Optimal Risk Sharing in the Presence of Moral Hazard under Market Risk and Jump Risk

Takashi Misumi
Hitotsubashi University
Graduate School of Commerce and Management, Hitotsubashi University, 2-1 Naka, Kunitachi,
Tokyo 186-8601, Japan
takashi.misumi@r.hit-u.ac.jp

Hisashi Nakamura*
Hitotsubashi University
hisashi.nakamura@r.hit-u.ac.jp

Koichiro Takaoka
Hitotsubashi University
k.takaoka@r.hit-u.ac.jp

Received: April 10, 2013 Revised July 10, 2013 Accepted July 31, 2014

ABSTRACT

This paper provides a tractable framework to study optimal risk sharing between an investor and a firm with general utility forms in the presence of moral hazard under market risk and jump risk. We show that, for a two-date discrete-time moral hazard model, there exists a continuous-time model that obtains the same optimal result. Moreover, we characterize the optimal risk sharing explicitly, in particular, the structural effect of jump risk on the optimal allocations.

Keywords: Optimal risk sharing, Moral hazard, Market risk, Jump risk.

JEL Classification: D82, D86

* Corresponding author: Graduate School of Commerce and Management, Hitotsubashi University, 2-1 Naka, Kunitachi, Tokyo 186-8601, Japan. Phone: +81-42-580-8826. Fax: +81-42-580-8747. Email: hisashi.nakamura@r.hit-u.ac.jp. We would like to thank Georges Dionne, Hirotaka Fushiya, Michael Hoy, Soichiro Moridaira, Mahito Okura, Harris Schlesinger, Yoshihiko Suzawa, Paul Thistle, Takau Yoneyama, and all seminar participants at American Risk and Insurance Association 2013 Annual Conference (Washington, DC), Asia-Pacific Risk and Insurance Association 17th Annual Conference (St. John's University, New York), and Japanese Society of Monetary Economics 70th Summer Meeting (Hitotsubashi University, Tokyo) for their valuable comments. Financial support from the Grants-in-Aid for Scientific Research (B) is gratefully acknowledged.

1. Introduction

It is well known that the problem of moral hazard deserves careful consideration in economics and finance. In fact, needless to say, moral hazard has been studied a lot in the theoretical literature in economics (e.g., Holmström, 1979; Mas-Colell et al., 1995; and many others). However, surprisingly, theoretical models of moral hazard have not been applied well to the practice in financial engineering such as fixed-income investment, the term structure of interest rates, corporate risk management, and actuarial insurance. There has been such a big gap in the research of moral hazard between the theory and the practice.

The purpose of this paper is to bridge the gap by providing a tractable framework to study optimal risk sharing between an investor and a firm with general utility forms in the presence of moral hazard under market risk and jump risk. We show that, for a two-date discrete-time moral hazard model, there exists a continuous-time model that obtains the same optimal result. Moreover, we characterize explicitly the optimal risk sharing, in particular, the structural effect of the jump risk on the optimal allocations.

In the previous literature on moral hazard, there have been mentioned various reasons for the gap before. They can be categorized into two: the first is about physical problems and the second is about informational ones. Specifically, with regard to the first, physical environments in theoretical models are often too naive for practical applications to finance. In complicated dynamic environments under the moral-hazard problem, in general, it is very difficult to find its solution analytically and numerically. To resolve the difficulty, the physical structures have been oversimplified in much of the literature. For example, in the theoretical literature on dynamic moral hazard in discrete time, many papers assume independent and identically-distributed shocks over time (e.g., Phelan & Townsend, 1991; Spear & Srivastava, 1987), which are of limited use practically. An exception is Fernandes and Phelan (2000), who study history-dependent (in particular, first-order Markovian) income shocks. Still, it is hard to tract more realistic, more complicated shock structures (such as multiple jumps) in those discrete-time models.

Continuous-time models have been lately used for overcoming such technical difficulty in the moral hazard problem, due to their mathematical tractability. This line of research was explored first by the seminal paper of Holmström and Milgrom (1987). They find linearity of an optimal compensation rule by assuming that an agent with constant absolute risk aversion (CARA) controls the drift rate of a profit process. Schättler and Sung (1993) develop the first-order approach to the problem under the CARA assumption and re-derive the linearity result. However, the CARA assumption is restrictive for practical applications to finance.¹ Instead, constant relative risk aversion (CRRA) utility and log utility are more desirable in practice (see e.g. Cox et al., 1985).

On the other hand, Cvitanic and Zhang (2007) and Nakamura and Takaoka (2012) study the case of an agent with a general utility form in continuous time. Still, there are three limitations in those models. First, they define the consumption space as the whole real space \mathbf{R} and thus do not fit the CRRA utility forms defined only on the space of positive consumption. Second, in both papers, the formulation of corporate profit is very simple: they assume only a Brownian motion as market risk. In practice, however, based on recent experiences of catastrophic natural disasters and serious financial crises, rare-event (i.e., jump) risk has been the center of attention. Those models do not give any answer to the attention. Third, economic and financial events occur in discrete time, not in continuous time, in practice. It remains to be proven that their continuous-time models can be applicable to practical discrete-time analyses.

As to the second reason, information environments are often not so well-constructed in the theoretical literature as to be applicable practically to finance. Holmström and Milgrom (1987) and Cvitanic and Zhang (2007) assume that the agent controls the drift rate based on the information set generated only by a history of the profit, not a history of his own observable true shocks.² That is, the information set continues to lose the information of a history of his own

effort (i.e., the drift rate) over time. In other words, they presume that the agent controls the drift rate while continuing to forget how he has controlled it until then.

To fill the gap in the moral hazard problem, our paper is novel in two respects. First, with regard to the informational environments, we assume that the firm controls directly the probability measure, rather than the drift rate, in the spirit of standard discrete-time moral hazard models (e.g., Holmström, 1979; Mas-Colell et al., 1995). In addition, we formulate the effort cost as relative entropy, which is a measure of statistical discrimination between the reference (i.e., original) measure and the controlled probability measure.³ We then characterize the change of the drift rates and the jump intensities in the firm's return process as a result (not a cause) of the twist of the probability measure. Due to these assumptions, we can make clear the information structure in a way that is consistent with both the theory and the practice.

Second, with regard to the physical environments, we consider a dynamic stochastic economy with the firm's linear production technology. The positivity of the production is ensured and is compatible with CRRA utility and log utility. Also, we deal with two types of risk: not only Brownian motions as market risk but also Poisson processes as jump risk. We can then study the effect of the jump risk on the optimal risk sharing in the presence of moral hazard. Regardless of such dynamic complexity of the physical and informational structures, we are successful in characterizing the optimal risk sharing explicitly by utilizing the tractable technique of continuous-time stochastic processes. Also, we show that, for a two-date discrete-time moral hazard model, there exists a continuous-time model that obtains the same optimal result.

Our framework has vast possibilities for its future applications to monetary economics and finance. In fact, a companion work of this paper, namely Misumi et al. (2014), applies this framework to a Lucas-type general-equilibrium asset-pricing model. More specifically, that paper looks at an optimal consumption/investment problem of a representative lender in financial markets when an endowment process is subject to the moral hazard problem whereas, by contrast, this current paper looks only at optimal bilateral contracting between the firm and the lender, neither of whom has access to financial markets. That paper consequently shows that, under the moral hazard problem, a positive premium is stipulated on a riskless rate in equilibrium. It is because the lender demands compensation for a loss caused due to the necessity to give the firm an incentive to avoid his opportunistic misbehavior. The result implies that the risk-free rate puzzle, explored first by Weil (1989), is more serious in the presence of moral hazard. This current paper provides a theoretical foundation to the application.

The rest of this paper is organized as follows. Next section defines the environment of our model. Section 3 solves for optimal risk sharing in the presence of moral hazard. Section 4 characterizes it explicitly. Final section concludes.

2. Environment

We consider a dynamic stochastic economy with two representative players: a firm and an investor on a time interval $[0, T]$ for a finite time $T > 0$. The firm and the investor are indexed by player 1 and player 2, respectively. For convenience, we will use female pronouns for the investor, and male ones for the firm.

Fix a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = \{\mathcal{F}(t)\}_{0 \leq t \leq T}, \mathbb{P})$. $\{B_1(t), \dots, B_n(t)\}_{0 \leq t \leq T}$ are n independent one-dimensional standard \mathbb{F} -Brownian motions on the probability space, i.e., for any t, s satisfying $0 \leq t \leq s$, $B_j(s) - B_j(t)$ is independent of $\mathcal{F}(t)$ and $B_j(0) = 0$. $\{N_1(t), \dots, N_m(t)\}_{0 \leq t \leq T}$ are m independent Poisson processes, each of which is characterized by its intensity $\lambda_i > 0$ ($i = 1, \dots, m$). Let the compensated Poisson process be denoted by $M_i(t) := N_i(t) - \lambda_i t$, which is a \mathbb{P} -martingale. The Poisson processes

are independent of $\{B_j(t); j = 1, \dots, n\}_{0 \leq t \leq T}$ as well. The filtration \mathbb{F} is generated by $\{B_j(t); j = 1, \dots, n\}_{0 \leq t \leq T}$ and $\{N_i(t); i = 1, \dots, m\}_{0 \leq t \leq T}$.

Define a measure \mathbb{Q} that is absolutely continuous w.r.t. \mathbb{P} , written as $\mathbb{Q} \ll \mathbb{P}$, i.e., $\mathbb{P}(A) = 0$ implies $\mathbb{Q}(A) = 0$ for $A \in \mathcal{F}$. Define also

$$Z(t) := \frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{\mathcal{F}(t)}.$$

By the Martingale Representation Theorem (cf. Theorem 5.43 of Medvegyev, 2007), there exist \mathbb{F} -predictable processes θ_j and $\alpha_i \geq -1$ for all $i = 1, \dots, m$ and all $j = 1, \dots, n$ such that

$$dZ(t) = Z(t_-) \left\{ \sum_{j=1}^n \theta_j(t) dB_j(t) + \sum_{i=1}^m \alpha_i(t) dM_i(t) \right\}. \quad (2.1)$$

Note that, once $Z(\tau) = 0$ at some time τ due to a jump, $Z(t) = 0$ for $t \geq \tau$. And, for each $j = 1, \dots, n$,

$$\tilde{B}_j(t) := B_j(t) - \int_0^t \theta_j(s) ds \quad (2.2)$$

is a \mathbb{Q} -Brownian motion, and for each $i = 1, \dots, m$,

$$\tilde{M}_i(t) := N_i(t) - \int_0^t \tilde{\lambda}_i(s) ds$$

is a \mathbb{Q} -local martingale where $\tilde{\lambda}_i(s) := \lambda_i \{\alpha_i(s) + 1\}$ (cf. Theorem 41 of Protter, 2010, Ch.III). Note that $\tilde{B}_j(t)$ and $\tilde{M}_i(t)$ (or $N_i(t)$) are uncorrelated instantaneously for any i, j , i.e., the quadratic variations $d\tilde{B}_j(t) \cdot dN_i(t) = 0$ and $d\tilde{B}_j(t) \cdot d\tilde{M}_i(t) = 0$ for any i, j , but are not necessarily independent under \mathbb{Q} , whereas $B_j(t)$ and $M_i(t)$ (or $N_i(t)$) are independent under \mathbb{P} for any i, j . Therefore,

$$\begin{aligned} Z(t) &= \mathbb{E}^{\mathbb{P}} \left[\frac{d\mathbb{Q}}{d\mathbb{P}} \Big| \mathcal{F}(t) \right] \\ &= \prod_{j=1}^n \exp \left\{ \int_0^t \theta_j(s) dB_j(s) - \frac{1}{2} \int_0^t (\theta_j(s))^2 ds \right\} \\ &\cdot \prod_{i=1}^m \exp \left\{ \sum_{0 < s < t} \log \left(\frac{\tilde{\lambda}_i(s)}{\lambda_i} \right) \Delta N_i(s) + \int_0^t (\lambda_i - \tilde{\lambda}_i(s)) ds \right\}. \end{aligned} \quad (2.3)$$

The firm produces the wealth process X with a linear production technology, which is characterized by the following stochastic differential equation:

$$dX(t) = X(t_-) dR(t), \quad X(0) = x > 0$$

where R denotes the return process that is defined as

$$dR(t) = c dt + \sum_{j=1}^n \sigma_j dB_j(t) + \sum_{i=1}^m z_i dM_i(t), \quad R(0) = a \in \mathbf{R},$$

where $c, \sigma_j, z_i \forall i, j$ are constants, $\sigma_j > 0 \forall j$, $z_i > -1 \forall i$, and $z_{i_1} \neq z_{i_2}$ if $i_1 \neq i_2$. In financial terms, $\{B_j; j = 1, \dots, n\}$ stands for market risk and $\{N_i; i = 1, \dots, m\}$ stands for rare-event (i.e., jump) risk. For each $i = 1, \dots, m$, z_i denotes the size of the jump. $\sum_{i=1}^m z_i M_i(t)$ can be interpreted as a mixed Poisson process with its intensity $\sum_{i=1}^m \lambda_i$. For each $i = 1, \dots, m$, $\frac{\lambda_i}{\sum_{i=1}^m \lambda_i}$ represents the probability of having the jump size z_i when a jump occurs.

The firm can share the outcome of the wealth with the investor at time T . Let $U_k : \mathbf{R} \rightarrow \mathbf{R} \cup \{-\infty\}$ for $k \in \{1, 2\}$ denote player k 's utility function of his or her own wealth,

defined on \mathbf{R} , at time T . For $k = 1, 2$, the utility function U_k is non-decreasing, and, on its effective domain denoted by $\text{dom}U_k := \{x \in \mathbf{R} \mid U_k(x) > -\infty\}$, it is twice continuously differentiable. In particular, for $k = 1, 2$, the utility functions possess standard properties: $U'_k > 0$ and $U''_k \leq 0$ on the effective domain.⁴ The firm is exogenously given a reservation utility, denoted by a constant $r \in \mathbf{R}$, at time 0. If the investor offers to the firm any lower utility than the reservation utility r , the firm does not take the offer.

We assume that the firm can control the probability measure so as to maximize his own expected payoff, in the spirit of the standard moral hazard literature in economics (e.g., Holmström, 1979; Mas-Colell et al., 1995; and many others). More specifically, \mathbb{P} is the original probability measure, that is, the measure when the firm does not control it – called it the reference measure. The firm can change the probability measure from \mathbb{P} into \mathbb{Q} such that $\mathbb{Q} \ll \mathbb{P}$.⁵ Assume that \mathbb{P} is the public information, and that the investor knows the fact that \mathbb{Q} is absolutely continuous w.r.t \mathbb{P} but cannot observe \mathbb{Q} directly, i.e., \mathbb{Q} is the private information of the firm.

We also assume that the firm incurs a utility cost when controlling the probability measure. The cost is represented by relative entropy, denoted by $\mathcal{H}(\mathbb{Q} \parallel \mathbb{P})$, which is defined as:

$$\begin{aligned} \mathcal{H}(\mathbb{Q} \parallel \mathbb{P}) &:= \mathbb{E}^{\mathbb{P}} \left[\frac{d\mathbb{Q}}{d\mathbb{P}} \left(\log \frac{d\mathbb{Q}}{d\mathbb{P}} \right) \mathbf{1}_{\left\{ \frac{d\mathbb{Q}}{d\mathbb{P}} > 0 \right\}} \right] \\ &= \mathbb{E}^{\mathbb{Q}} \left[\left(\log \frac{d\mathbb{Q}}{d\mathbb{P}} \right) \mathbf{1}_{\left\{ \frac{d\mathbb{Q}}{d\mathbb{P}} > 0 \right\}} \right] = \mathbb{E}^{\mathbb{Q}} \left[\log \frac{d\mathbb{Q}}{d\mathbb{P}} \right]. \end{aligned}$$

Assume that $\mathcal{H}(\mathbb{Q} \parallel \mathbb{P}) < \infty$.⁶ Roughly speaking, the relative entropy is a measure of the distance between the probability measures \mathbb{P} and \mathbb{Q} .⁷ From a statistical viewpoint, it represents a measure of the type-I error of rejecting the true probability measure \mathbb{Q} and, instead, assuming \mathbb{P} incorrectly. That is, it stands for the statistical inefficiency of assuming that the probability measure is \mathbb{P} when the true measure is \mathbb{Q} . A low level of the relative entropy means that \mathbb{Q} and \mathbb{P} are not so distant as to significantly discriminate \mathbb{P} against \mathbb{Q} . Thus the relative entropy means how far the true probability measure \mathbb{Q} is twisted from the reference measure \mathbb{P} . In sum, in our model, the utility cost that the firm incurs due to the effort to twist the probability measure is measured by how far the probability measure \mathbb{Q} is changed from the reference measure \mathbb{P} . The cost impedes the firm's adopting the probability measure far away from \mathbb{P} . As we will show below, due to this cost, the investor may infer the true probability measure \mathbb{Q} as a Nash-equilibrium result of a strategic game between the two players, although she cannot observe it directly.

We look at an example of the relative entropy.

Example 2.1 Consider the case of finite scenarios: $\#\Omega = l$, say $\Omega = \{\omega_1, \dots, \omega_l\}$. Define a random variable $Y: \Omega \rightarrow \mathbf{R}$ on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Under \mathbb{P} , the random variable Y is represented by its realizations $\{y_1, \dots, y_l\}$ with the assigned probabilities $\{p_1, \dots, p_l\}$ where $p_s > 0$ for each scenario $s = 1, \dots, l$ and $\sum_{s=1}^l p_s = 1$. When the probability measure is changed into \mathbb{Q} , the new probabilities are $\{q_1, \dots, q_l\}$ for the realizations $\{y_1, \dots, y_l\}$ where $q_s \geq 0$ for each $s = 1, \dots, l$ and $\sum_{s=1}^l q_s = 1$. We then obtain

$$\mathcal{H}(\mathbb{Q} \parallel \mathbb{P}) = \sum_{s=1}^l p_s \frac{q_s}{p_s} \log \frac{q_s}{p_s} = \sum_{s=1}^l q_s \log \frac{q_s}{p_s}.$$

When $\{q_1, \dots, q_l\}$ are more distant from $\{p_1, \dots, p_l\}$, $\mathcal{H}(\mathbb{Q} \parallel \mathbb{P})$ becomes larger. We can easily confirm that $\mathcal{H}(\mathbb{Q} \parallel \mathbb{P}) = 0$ if and only if $p_s = q_s$ for all s . In this case, the effort is to change the probability distribution from $\{p_1, \dots, p_l\}$ to $\{q_1, \dots, q_l\}$ and the effort cost is measured by $\mathcal{H}(\mathbb{Q} \parallel \mathbb{P})$.

Concretely, let us look at the case of $l = 3$, in particular, of $\{p_1, p_2, p_3\} = \left\{\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right\}$ and $\{q_1, q_2, q_3\} = \left\{\frac{2}{3}, \frac{1}{3}, 0\right\}$. We then have $\frac{d\mathbb{Q}}{d\mathbb{P}} = \{2, 1, 0\}$ and $\frac{d\mathbb{P}}{d\mathbb{Q}} = \left\{\frac{1}{2}, 1, n.a.\right\}$. Hence, $\mathbb{E}^{\mathbb{Q}}\left[\frac{d\mathbb{P}}{d\mathbb{Q}}\right] < 1$. In a similar way, for the random variable Y , $\mathbb{E}^{\mathbb{Q}}\left[Y\frac{d\mathbb{P}}{d\mathbb{Q}}\right] = \mathbb{E}^{\mathbb{P}}\left[Y\mathbf{1}_{\left\{\frac{d\mathbb{Q}}{d\mathbb{P}} > 0\right\}}\right] \leq \mathbb{E}^{\mathbb{P}}[Y]$. ■

Note that we will see this example again below when deriving the firm's optimal effort.

In our framework, noting Eq.(2.2) and Eq.(2.3), we can characterize the relative entropy by using θ_j and α_i for all i, j as follows:

$$\begin{aligned} \mathcal{H}(\mathbb{Q}||\mathbb{P}) &= \mathbb{E}^{\mathbb{Q}}\left[\log\frac{d\mathbb{Q}}{d\mathbb{P}}\right] = \mathbb{E}^{\mathbb{Q}}\left[\sum_{j=1}^n \int_0^T \theta_j(s) dB_j(s) - \frac{1}{2} \int_0^T (\theta_j(s))^2 ds + \right. \\ &\quad \left. \sum_{i=1}^m \int_0^T \left\{ \tilde{\lambda}_i(s) \log\left(\frac{\tilde{\lambda}_i(s)}{\lambda_i}\right) + (\lambda_i - \tilde{\lambda}_i(s)) \right\} ds\right] \\ &= \mathbb{E}^{\mathbb{Q}}\left[\sum_{j=1}^n \int_0^T \frac{(\theta_j(s))^2}{2} ds + \sum_{i=1}^m \int_0^T \left\{ \tilde{\lambda}_i(s) \log\left(\frac{\tilde{\lambda}_i(s)}{\lambda_i}\right) + (\lambda_i - \tilde{\lambda}_i(s)) \right\} ds\right]. \quad (2.4) \end{aligned}$$

Recall that $\frac{\tilde{\lambda}_i(s)}{\lambda_i} = \alpha_i(s) + 1$. Suppose that there are no jump terms. We then obtain $\mathcal{H}(\mathbb{Q}||\mathbb{P}) = \mathbb{E}^{\mathbb{Q}}\left[\sum_{j=1}^n \frac{1}{2} \int_0^t (\theta_j(s))^2 ds\right]$, which is equivalent to the first term inside the expectation on the right-hand side of Eq.(2.4). On the other hand, the second term inside the expectation corresponds to the jump terms.

Remark 2.1 The cost function $\mathbb{E}^{\mathbb{Q}}\left[\sum_{j=1}^n \frac{1}{2} \int_0^t (\theta_j(s))^2 ds\right]$ is exactly the same as the one defined in Cvitanić and Zhang (2007) in the case of $n = 1$, although they do not link the cost function to the notion of the relative entropy. Their paper interprets $\theta_j(s)$ as the firm's effort in the sense that, noting $dB_j(s) = \theta_j(s)ds + d\tilde{B}_j(s)$, a higher (lower) costly effort leads to a higher (lower) expected return of the wealth under \mathbb{Q} .

In contrast to Cvitanić and Zhang (2007), however, our current paper does not assume that the firm controls the drift rate θ_j and the jump intensity $\tilde{\lambda}_i$ for each i, j . The reason is as follows. By the Martingale Representation Theorem, we can find the predictable processes θ_j and α_i for all i, j corresponding to the controlled probability \mathbb{Q} , as in Eq.(2.1). However, θ_j and α_i for each i, j are adapted to \mathbb{F} , not to the filtration generated by the controlled $\tilde{B}_j(t)$ and $\tilde{M}_i(t)$ for all i, j , in the weak formulation.⁸ Thus, if we assume that the firm controls θ_j and α_i for each i, j , then it means that the controls would be undertaken based on the information set that continues to lose the information of a history of the controls over time. This seems irrelevant in practice. Instead, we assume that the firm controls the probability measure, neither θ_j nor α_i for i, j . In sum, the predictable processes θ_j and α_i for each i, j are not controlled objects, but rather a result from controlling the probability measure. That is, the change of \mathbb{Q} results in the change of the drift rate θ_j for each j and the change of the (stochastic) jump intensity α_i for each i ; the converse is not true. ■

The firm enters into a contract with the investor and shares the time- T outcome $X(T)$ with the investor according to terms of the contract for insuring against his wealth risk. Specifically, the investor offers a menu of contract payoffs C_T to the firm, and the firm then decides whether or not to accept it. We assume that the firm's wealth allocation C_T takes the form of $C_T: \Omega \rightarrow \mathbf{R}$ as a functional of $X(\cdot)$, i.e., $C_T = C_T(X(\cdot))$. Call $C_T(X(\cdot))$ a contract. Define mathematical regularities for the contracts C_T :

Definition 2.1 Define the set \mathcal{A} of the contracts $C_T(X(\cdot))$ such that

- (i) $C_T \in \text{dom}U_1$ and $X(T) - C_T \in \text{dom}U_2$ a.s.,
- (ii) $C_T(X(\cdot))$ is continuous and is Gâteaux differentiable,⁹
- (iii) $1 < \exists p < \infty$, $\mathbb{E}^{\mathbb{P}}[e^{pU_1(C_T)}] < \infty$ and $\mathbb{E}^{\mathbb{P}}[|U_2(X(T) - C_T)|^q] < \infty$ where $q = \frac{p}{p-1}$.

3. Optimal Risk Sharing

3.1. Firm's Optimization

For $C_T \in \mathcal{A}$, define the firm's expected utility under the controlled probability measure \mathbb{Q} , denoted by V_1 , as:

$$V_1 := \sup_{\mathbb{Q} \ll \mathbb{P}, \mathcal{H}(\mathbb{Q} \parallel \mathbb{P}) < \infty} \left\{ \mathbb{E}^{\mathbb{Q}} \left[U_1 \left(C_T(X(\cdot)) \right) \right] - \mathcal{H}(\mathbb{Q} \parallel \mathbb{P}) \right\}. \quad (3.1)$$

We obtain the following proposition:

Proposition 3.1 For $C_T \in \mathcal{A}$,

$$V_1 = \log \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(X(\cdot)))} \right]. \quad (3.2)$$

The maximizer, denoted by \mathbb{Q}^* , is then characterized by

$$\frac{d\mathbb{Q}^*}{d\mathbb{P}} = \frac{e^{U_1(C_T(X(\cdot)))}}{\mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(X(\cdot)))} \right]}. \quad (3.3)$$

This result and its variants are known in the fields of operations research and mathematical finance: for a literature review, see e.g. the first remark in Section 1 of Delbaen et al. (2002). For the sake of completeness, we present a proof.

Proof: Taking exponential of $\mathbb{E}^{\mathbb{Q}} \left[U_1 \left(C_T(X(\cdot)) \right) \right] - \mathcal{H}(\mathbb{Q} \parallel \mathbb{P})$,

$$\begin{aligned} & e^{\mathbb{E}^{\mathbb{Q}} \left[U_1(C_T(X(\cdot))) \right] - \mathcal{H}(\mathbb{Q} \parallel \mathbb{P})} = e^{\mathbb{E}^{\mathbb{Q}} \left[U_1(C_T(X(\cdot))) - \log \frac{d\mathbb{Q}}{d\mathbb{P}} \right]} \\ & \leq \mathbb{E}^{\mathbb{Q}} \left[e^{U_1(C_T(X(\cdot))) - \log \frac{d\mathbb{Q}}{d\mathbb{P}}} \right] \quad (\text{by Jensen's inequality}) \\ & \leq \mathbb{E}^{\mathbb{Q}} \left[e^{U_1(C_T(X(\cdot)))} \frac{d\mathbb{P}}{d\mathbb{Q}} \right] = \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(X(\cdot)))} \mathbf{1}_{\left\{ \frac{d\mathbb{Q}}{d\mathbb{P}} > 0 \right\}} \right] \\ & \leq \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(X(\cdot)))} \right] \end{aligned}$$

with equality if and only if $U_1 \left(C_T(X(\cdot)) \right) - \log \frac{d\mathbb{Q}}{d\mathbb{P}}$ is a constant, i.e., $\frac{d\mathbb{Q}}{d\mathbb{P}} = \frac{e^{U_1(C_T(X(\cdot)))}}{\mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(X(\cdot)))} \right]}$

Thus \mathbb{Q}^* is obtained. ■

Note that we can extend this model straightforwardly into continuous-time consumption models with time-separable utility and recursive utility (see e.g. Misumi et al., 2014).

We look at the case of the finite scenarios shown above in Example 2.1 again.

Example 3.1 Set $Y = U_1(C_T)$. In the environment of Example 2.1, the firm's optimization problem is written as:

$$V_1 = \max_{\{q_1, \dots, q_l\}} \sum_{s=1}^l y_s q_s - q_s \log \frac{q_s}{p_s}$$

subject to $\sum_{s=1}^l q_s = 1$ and $q_s \geq 0$ ($s = 1, \dots, l$). We can assume that $q_s \geq 0$ is satisfied for each $s = 1, \dots, l$. Let the Lagrangian multiplier associated with $\sum_{s=1}^l q_s = 1$ be denoted by κ . We then obtain the Lagrangian:

$$L(\{q_1, \dots, q_l\}; \kappa) = \sum_{s=1}^l q_s y_s - q_s \log \frac{q_s}{p_s} + \kappa \left(\sum_{s=1}^l q_s - 1 \right).$$

Differentiating with respect to q_s ,

$$y_s - \log \frac{q_s}{p_s} - 1 + \kappa = 0.$$

This is the sufficient and necessary condition for optimality. Hence,

$$\frac{q_s}{p_s} = e^{\kappa-1} e^{y_s}.$$

Plugging this into $\sum_{s=1}^l q_s = 1$,

$$1 = \sum_{s=1}^l q_s = \sum_{s=1}^l p_s e^{\kappa-1} e^{y_s} = e^{\kappa-1} \mathbb{E}^{\mathbb{P}}[e^Y].$$

Hence, $\frac{q_s}{p_s} = \frac{e^{y_s}}{\mathbb{E}^{\mathbb{P}}[e^Y]}$ is confirmed. The optimal probability distribution $\{q_1^*, \dots, q_l^*\}$ is obtained.

■

To ensure that the firm participates in the contract, the investor provides him with no lower utility than his reservation utility, i.e., $V_1 \geq r$ – call it the participation condition. In particular, as usual in hidden action problems, we assume that the participation condition is binding:

$$V_1 = r. \tag{3.4}$$

We impose Condition (3.4) on the set \mathcal{A} of the contracts C_T as follows.

Definition 3.1 Define the set \mathcal{A}' of the contracts $C_T \in \mathcal{A}$ such that C_T satisfies Condition (3.4).

Due to Condition (3.4), from Eq.(3.2) and Eq.(3.3),

$$\frac{dQ^*}{d\mathbb{P}} = e^{-r} e^{U_1(C_T(X(\cdot)))}. \tag{3.5}$$

Thus the investor can implement the optimal Q^* by controlling $C_T(X(\cdot))$. We call Eq.(3.5) the implementability condition.

Due to the characteristics of the Radon-Nikodym derivative (2.3),

Corollary 3.1 For a two-date (i.e., $\{0, T\}$) discrete-time model of moral hazard, there exists a continuous-time model that obtains the same optimal result.

Accordingly, we are successful in filling a gap between the discrete-time moral hazard problem and the continuous-time one.

3.2. Investor's Optimization

We formulate the investor's optimization problem with respect to $C_T \in \mathcal{A}'$ as follows:

$$\sup_{C_T \in \mathcal{A}'} \mathbb{E}^{\mathbb{Q}^*} \left[U_2 \left(X(T) - C_T(X(\cdot)) \right) \right]. \quad (3.6)$$

Although the investor cannot observe the true probability measure \mathbb{Q} directly, she can verify the optimal \mathbb{Q}^* by designing the contract so as to satisfy the implementability condition (3.5). Accordingly, for $C_T \in \mathcal{A}'$, the investor can take her expectation under \mathbb{Q}^* as in Eq.(3.6). Using the implementability condition (3.5), the optimization problem (3.6) is rewritten as

$$\begin{aligned} & \sup_{C_T \in \mathcal{A}'} \mathbb{E}^{\mathbb{Q}^*} \left[U_2 \left(X(T) - C_T(X(\cdot)) \right) \right] = \sup_{C_T \in \mathcal{A}'} \mathbb{E}^{\mathbb{P}} \left[\frac{d\mathbb{Q}^*}{d\mathbb{P}} U_2 \left(X(T) - C_T(X(\cdot)) \right) \right] \\ & = \sup_{C_T \in \mathcal{A}'} \mathbb{E}^{\mathbb{P}} \left[e^{-r} e^{U_1(C_T(X(\cdot)))} U_2 \left(X(T) - C_T(X(\cdot)) \right) \right]. \quad (3.7) \end{aligned}$$

Due to Definition 2.1 (iii), by Hölder's inequality, the integrability is ensured in Eq.(3.7). Define the Lagrangian multiplier associated with (3.4) as μ . Using Conditions (3.4) and (3.5), the constrained optimization problem (3.7) is rewritten into:

$$\sup_{C_T \in \mathcal{A}'} \left\{ e^{-r} \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(X(\cdot)))} \left\{ U_2 \left(X(T) - C_T(X(\cdot)) \right) + \mu \right\} \right] \right\}.$$

A necessary condition for optimality of C_T is:

$$e^{U_1(C_T)} U_1'(C_T) \left\{ \mu - \left(\frac{U_2'(X(T) - C_T)}{U_1'(C_T)} - U_2(X(T) - C_T) \right) \right\} = 0. \quad (3.8)$$

As to the sufficiency of the condition, setting $v = C_T$ and differentiating Eq.(3.8) with respect to v ,

$$\begin{aligned} & e^{U_1(v)} U_1'(v) \left\{ U_1'(v) \left(\mu - \left(\frac{U_2'(X(T) - v)}{U_1'(v)} - U_2(X(T) - v) \right) \right) \left(1 + \frac{U_1''(v)}{(U_1'(v))^2} \right) \right\} \\ & \left\{ - \left(\frac{-U_2''(X(T) - v) U_1'(v) - U_2'(X(T) - v) U_1''(v)}{(U_1'(v))^2} + U_2'(X(T) - v) \right) \right\} \\ & = -e^{U_1(v)} U_1' \cdot \left\{ \frac{-U_2'' U_1' - U_2' U_1''}{(U_1')^2} + U_2' \right\} < 0. \quad (3.9) \end{aligned}$$

Therefore,

Proposition 3.2 *A necessary and sufficient condition for optimality is as follows: There exists some $\mu \in \mathbf{R}$ that satisfies*

$$\frac{U_2'(X(T) - C_T)}{U_1'(C_T)} - U_2(X(T) - C_T) = \mu \quad a.s. \quad (3.10)$$

Directly from Eq.(3.10),

Corollary 3.2 $C_T(X(\cdot)) = C_T(X(T))$.

4. Characterization of Optimal Risk Sharing

We characterize the optimal C_T from Eq.(3.10). First, we obtain the uniqueness of the optimal contract, if it exists.

Proposition 4.1 *If there exists some optimal contract satisfying Eq.(3.10), then it is unique.*

Proof: Fix an optimal Lagrangian multiplier μ satisfying Eq.(3.10). Suppose that there exist two different optimal contracts C_T, C_T' almost everywhere for the Lagrangian multiplier μ . By Eq.(3.9), the left-hand side of Eq.(3.10) is strictly increasing in C_T . Thus, in optimum, for the two contracts C_T, C_T' , the two associated Lagrangian multipliers should not be the same, say μ, μ' . When $\mu > \mu'$, $C_T < C_T'$ almost everywhere. This contradicts the binding participation condition (3.4) for any ω . Therefore, there exists a unique optimal solution satisfying Eq.(3.10). ■

Next, we examine a relationship between the two players' optimal utility levels. By Proposition 4.1, if an optimal solution exists, we can write $C_T = C_T(r)$. Also, similarly write $\mu = \mu(r)$. On Eq.(3.10), for some μ given, the left-hand side is differentiable with respect to C_T . The derivative, denoted by $\frac{d\mu}{dC_T}$, is strictly positive. Since the left-hand side is differentiable and monotonic with respect to C_T , the inverse function $C_T(\mu)$ is also differentiable with respect to μ satisfying $\frac{d\mu}{dC_T} \neq 0$. Suppose that $C_T(r)$ is differentiable with respect to r and that the Leibniz rule for differentiating integrals holds true, i.e., the order of the differential and the integral (i.e., expectation) operators is interchangeable.¹⁰ We claim that, if there exists an optimal contract, the investor's optimal utility level is strictly decreasing in the firm's one r . Differentiating the investor's optimal utility with respect to r ,

$$\begin{aligned} & \frac{d}{dr} \left\{ e^{-r} \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(r))} U_2(X(T) - C_T(r)) \right] \right\} \\ = & -e^{-r} \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(r))} U_2(X(T) - C_T(r)) \right] \\ & + e^{-r} \mathbb{E}^{\mathbb{P}} \left[\frac{d \left\{ e^{U_1(C_T(r))} U_2(X(T) - C_T(r)) \right\}}{dC_T(r)} \cdot \frac{dC_T(r)}{dr} \right]. \end{aligned}$$

With regard to the first term, noting $-U_2(X(T) - C_T(r)) = \mu(r) - \frac{U_2'(X(T) - C_T(r))}{U_1'(C_T(r))} < \mu(r)$ from Eq.(3.10),

$$\begin{aligned} & -e^{-r} \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(r))} U_2(X(T) - C_T(r)) \right] < e^{-r} \mu(r) \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(r))} \right] \\ = & e^{-r} \mu(r) e^r \\ = & \mu(r). \end{aligned}$$

With regard to the second term,

$$\begin{aligned} & e^{-r} \mathbb{E}^{\mathbb{P}} \left[\frac{d}{dC_T} \left\{ e^{U_1(C_T(r))} U_2(X(T) - C_T(r)) \right\} \cdot \frac{dC_T(r)}{dr} \right] \\ = & e^{-r} \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(r))} U_1'(C_T(r)) \left\{ U_2(X(T) - C_T(r)) - \frac{U_2'(X(T) - C_T(r))}{U_1'(C_T(r))} \right\} \frac{dC_T(r)}{dr} \right] \\ = & e^{-r} \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(r))} U_1'(C_T(r)) (-\mu(r)) \frac{dC_T(r)}{dr} \right] \quad (\text{noting Eq.(3.10)}) \\ = & -e^{-r} \mu(r) \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(r))} U_1'(C_T(r)) \frac{dC_T(r)}{dr} \right] \\ = & -e^{-r} \mu(r) \mathbb{E}^{\mathbb{P}} \left[\frac{d e^{U_1(C_T(r))}}{dr} \right] \\ = & -e^{-r} \mu(r) \frac{d}{dr} \mathbb{E}^{\mathbb{P}} \left[e^{U_1(C_T(r))} \right] \\ = & -e^{-r} \mu(r) \frac{d}{dr} e^r \\ = & -\mu(r). \end{aligned}$$

Hence, $\frac{d}{dr} \{e^{-r} \mathbb{E}^{\mathbb{P}} [e^{U_1(C_T(r))} U_2(X(T) - C_T(r))]\} < 0$. In other words, a higher (lower) r leads to a lower (higher) level of the investor's optimal utility, if the optimal contract exists.

Furthermore, we characterize the optimal contract as a function of the outcome $X(T)$. As in Cvitanic and Zhang (2007) and Nakamura and Takaoka (2012), by comparing Eq.(3.10) with the standard Borch rule $\frac{U_2'}{U_1'} = \mu$ (i.e., in the case of no moral hazard), we see that the term $U_2(X(T) - C_T)$ stands for the effect of moral hazard in Eq.(3.10). Also, C_T is non-linear in $X(T)$ in contrast to Holmström and Milgrom (1987) and Schättler and Sung (1993). This result is similar to Cvitanic and Zhang (2007) and Nakamura and Takaoka (2012). Noting $U_1'' \leq 0$,

$$0 < \frac{dC_T}{dX(T)} = 1 - \frac{U_2' U_1''}{U_2'' U_1' + U_2' U_1'' - U_2'(U_1')^2} \leq 1. \quad (4.1)$$

On the other hand, in the case of no moral hazard, from the standard Borch rule,

$$0 < \frac{dC_T}{dX(T)} = 1 - \frac{U_2' U_1''}{U_2'' U_1' + U_2' U_1''} \leq 1. \quad (4.2)$$

From Eq.(4.1) and Eq.(4.2), when $U_1'' < 0$, $\frac{dC_T}{dX(T)}$ is less than one, either with or without moral hazard. In addition, when $U_1'' < 0$, it is higher in Eq.(4.1) than in Eq.(4.2). I.e., when $X(T)$ becomes higher, the larger compensation is required in the moral hazard case due to the necessity to induce the firm to make the optimal efforts.

Finally, let us see a numerical example, in which we will show some condition under which no such optimal contract exists.

Example 4.1 Consider the case of $U_1(z) = \log z$ and $U_2(z) = z$: more precisely, $U_1(z) = \log z$ for $z > 0$ and $U_1(z) = -\infty$ for $z \leq 0$. The optimal risk sharing is characterized explicitly in a closed form as follows. When Eq.(3.10) holds, $C_T(X(T))$ can be written as a function of μ , denoted by $C_T^\mu(X(T))$. I.e., $C_T^\mu = \frac{X(T) + \mu}{2}$. From Condition (3.5),

$$\begin{aligned} e^r &= \mathbb{E}^{\mathbb{P}} [e^{U_1(C_T^\mu)}] = \int_0^\infty e^{U_1(C_T^\mu(X(T)))} \Phi(dX(T)) \\ &= \frac{\mathbb{E}^{\mathbb{P}}[X(T)] + \mu}{2} \end{aligned}$$

where Φ denotes the cumulative distribution function of $X(T)$ under \mathbb{P} . Therefore, $\mu = 2e^r - \mathbb{E}^{\mathbb{P}}[X(T)]$.

Let it be denoted by μ^* . We assume that $\mu^* \geq 0$, or equivalently, $r \geq \log \frac{\mathbb{E}^{\mathbb{P}}[X(T)]}{2}$: it then follows that $C_T^{\mu^*} > 0$ a.s., which is consistent with Definition 2.1 (i). When the condition is violated, there is no optimal contract.

The optimal contract $C_T^{\mu^*}$ is written as:

$$C_T^{\mu^*} = \frac{X(T) + \mu^*}{2} = e^r + \frac{X(T) - \mathbb{E}^{\mathbb{P}}[X(T)]}{2}.$$

The investor's optimal expected utility is then obtained as:

$$\begin{aligned} &e^{-r} \mathbb{E}^{\mathbb{P}} [e^{U_1(C_T^{\mu^*})} U_2(X(T) - C_T^{\mu^*})] \\ &= e^{-r} \mathbb{E}^{\mathbb{P}} \left[\frac{X(T) + \mu^*}{2} \cdot \frac{X(T) - \mu^*}{2} \right] = \frac{1}{4e^r} \mathbb{E}^{\mathbb{P}} [(X(T))^2 - (\mu^*)^2] \end{aligned}$$

$$\begin{aligned}
 &= \frac{\mathbb{E}^{\mathbb{P}} \left[(X(T) - \mathbb{E}^{\mathbb{P}}[X(T)])^2 \right]}{4e^r} + \mathbb{E}^{\mathbb{P}}[X(T)] - e^r \\
 &= \frac{\text{Var}^{\mathbb{P}}[X(T)]}{4e^r} + \mathbb{E}^{\mathbb{P}}[X(T)] - e^r
 \end{aligned}$$

where $\text{Var}^{\mathbb{P}}[X(T)]$ denotes variance of $X(T)$ under \mathbb{P} . Finally, let $\mathbb{E}^{\mathbb{P}}[X(T)]$ and $\text{Var}^{\mathbb{P}}[X(T)]$ be specified explicitly in a closed form as follows. Recalling $dR(t) = cdt + \sum_{j=1}^n \sigma_j dB_j(t) + \sum_{i=1}^m z_i dM_i(t)$,

$$X(t) = x + \int_0^t X(s_-) \left(cds + \sum_{j=1}^n \sigma_j dB_j(s) + \sum_{i=1}^m z_i dM_i(s) \right).$$

Taking the expectations of both sides, we have

$$\begin{aligned}
 \mathbb{E}^{\mathbb{P}}[X(t)] &= x + \mathbb{E}^{\mathbb{P}} \left[\int_0^t X(s_-) cds \right] \\
 &= x + \mathbb{E}^{\mathbb{P}} \left[\int_0^t X(s) cds \right] \\
 &= x + \int_0^t \mathbb{E}^{\mathbb{P}}[X(s)] cds \quad (\text{by Fubini's theorem}) \tag{4.3}
 \end{aligned}$$

and thus $\mathbb{E}^{\mathbb{P}}[X(T)] = xe^{cT}$. Also, by Itô's formula,

$$\begin{aligned}
 (X(t))^2 &= x^2 + 2 \int_0^t X(s_-) dX(s) + [X]_t \\
 &= x^2 + 2 \int_0^t (X(s_-))^2 \left\{ cds + \sum_{j=1}^n \sigma_j dB_j(s) + \sum_{i=1}^m z_i dM_i(s) \right\} \\
 &\quad + \int_0^t (X(s_-))^2 \left\{ \sum_{j=1}^n \sigma_j^2 ds + \sum_{i=1}^m z_i^2 dN_i(s) \right\}.
 \end{aligned}$$

The same argument as Eq.(4.3) gives

$$\mathbb{E}^{\mathbb{P}} \left[(X(T))^2 \right] = x^2 \exp \left\{ \left(2c + \sum_j \sigma_j^2 + \sum_i z_i^2 \lambda_i \right) T \right\}.$$

Therefore,

$$\begin{aligned}
 \text{Var}^{\mathbb{P}}[X(T)] &= \mathbb{E}^{\mathbb{P}} \left[(X(T))^2 \right] - (\mathbb{E}^{\mathbb{P}}[X(T)])^2 \\
 &= x^2 \exp\{2cT\} \left(\exp \left\{ \left(\sum_j \sigma_j^2 + \sum_i z_i^2 \lambda_i \right) T \right\} - 1 \right).
 \end{aligned}$$

Thus the results are obtained explicitly in a closed form. Note that, obviously, the investor's optimal utility is decreasing in r in this example. ■

5. Concluding Remarks

This paper provided a tractable framework to study optimal risk sharing between an investor and a firm with general utility forms in the presence of moral hazard under both market risk and jump risk. Our framework has vast possibilities for its future applications to monetary economics and finance. In fact, a companion work of this paper, namely Misumi et al. (2014),

applies this framework to a Lucas-type general-equilibrium asset-pricing model. More specifically, that paper looks at an optimal consumption/investment problem of a representative lender in financial markets when an endowment process is subject to the moral hazard problem whereas, by contrast, this current paper looks only at optimal bilateral contracting between the firm and the lender, neither of whom has access to financial markets. That paper consequently shows that, under the moral hazard problem, a positive premium is stipulated on a riskless rate in equilibrium. It is because the lender demands compensation for a loss caused due to the necessity to give the firm an incentive to avoid his opportunistic misbehavior. The result implies that the risk-free rate puzzle, explored first by Weil (1989), is more serious in the presence of moral hazard. This current paper provides a theoretical foundation to the application.

NOTES

1. As Kimball and Mankiw (1989) discuss, there exist very few empirical studies of the CARA parameters.
2. To our knowledge, the only exception is the paper of Nakamura and Takaoka (2012), which generalizes the information set so as to be generated by a history of an agent's efforts as well as a history of true shocks.
3. The relative entropy has been lately used as a cost of controlling probability measures in the economics literature. See e.g. Hansen and Sargent (2007), Hansen et al. (2006), Sims (2003). Also, Delbaen et al. (2002) use it as a penalty in hedging contingent claims in the finance literature.
4. The first two derivatives of a function f are denoted by f' , f'' , respectively.
5. Under the absolute-continuity restriction, zero probability is necessarily assigned, under \mathbb{Q} , to the state to which zero probability is assigned under \mathbb{P} . In other words, we do not look at the states that are supposed not to occur under the reference measure \mathbb{P} . We assume the reference measure \mathbb{P} that covers a very wide range of states of nature. Also, in contrast to standard discussions of moral hazard, we do not impose either first-order stochastic dominance of probability distributions or monotone likelihood ratio property.
6. Note that this finiteness assumption is imposed for removing the indeterminacy of the firm's optimal expected utility, denoted by V_1 , defined in Eq.(3.1) below.
7. The relative entropy is always non-negative and is zero if and only if $\mathbb{Q} = \mathbb{P}$. Strictly speaking, it is not a true distance because neither the symmetry nor the triangle inequality are satisfied. However, it is well known that it is useful to regard the relative entropy as a distance between two probability measures. See e.g. Cover and Thomas (2006, p.18) in the statistics literature.
8. Nakamura and Takaoka (2012) generalize the information set, so as to be generated by a history of an agent's efforts as well as a history of true shocks. That paper then solves the optimization problem in the strong formulation rather than in the weak one.
9. Gâteaux differentiability is a generalization of direction differentiability. The definition is as follows. Suppose X and Y are locally convex topological vector spaces, $U \subset X$ is open, and $f: X \rightarrow Y$. The Gâteaux differential $df(u; \phi)$ of f at $u \in U$ in a direction $\phi \in X$ is defined as:

$$df(u; \phi) := \lim_{\tau \rightarrow 0} \frac{f(u + \tau\phi) - f(u)}{\tau} = \left. \frac{df(u + \tau\phi)}{d\tau} \right|_{\tau=0}$$
 if the limit exists. If the limit exists for all directions $\phi \in X$, then f is said to be Gâteaux differentiable at u .
10. It might be desirable mathematically to impose higher-level assumptions to ensure the interchangeability of the order of the two operators. However, it is very technical and out of our scope in this paper. Instead, we assume the interchangeability.

REFERENCES

- Cover, T. M., & Thomas, J. A. (2006). *Elements of information theory* (2nd ed.). Wiley.
- Cox, J., Ingersoll, J., & Ross, S. (1985). A theory of the term structure of interest rates. *Econometrica*, 53, 385-407.
- Cvitanić, J., & Zhang, J. (2007). Optimal compensation with adverse selection and dynamic actions. *Mathematics and Financial Economics*, 1, 21-55.
- Delbaen, F., Grandits, P., Rheinlander, T., Samperi, D. J., Schweizer, M., & Stricker, C. (2002). Exponential hedging and entropic penalties. *Mathematical Finance*, 12, 99-123.
- Fernandes, A., & Phelan, C. (2000). A recursive formulation for repeated agency with history dependence. *Journal of Economic Theory*, 91, 223-247.
- Hansen, L. P., & Sargent, T. J. (2007). *Robustness*. Princeton University Press.
- Hansen, L. P., Sargent, T. J., Turmuhambetova, G., & Williams, N. (2006). Robust control and model misspecification. *Journal of Economic Theory*, 128, 45-90.
- Holmström, B. (1979). Moral hazard and observability. *Bell Journal of Economics*, 10, 74-91.
- Holmström, B., & Milgrom, P. (1987). Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, 55, 303-328.
- Kimball, M. S., & Mankiw, N. G. (1989). Precautionary saving and the timing of taxes. *Journal of Political Economy*, 97, 863-879.
- Mas-Colell, A., Whinston, M. D., & Green, J. R. (1995). *Microeconomic theory*. Oxford University Press.
- Medvedev, P. (2007). *Stochastic integration theory*. Oxford University Press.
- Misumi, T., Nakamura, H., & Takaoka, K. (2014). Moral-hazard premium. Working paper #G-1-7, Hitotsubashi University Center for Financial Research.
- Nakamura, H., & Takaoka, K. (2012). A continuous-time optimal insurance design with costly monitoring. Working paper #140, Graduate School of Commerce and Management, Hitotsubashi University.
- Phelan, C., & Townsend, R. M. (1991). Computing multi-period, information-constrained optima. *Review of Economic Studies*, 58, 853-881.
- Protter, P. E. (2010). *Stochastic integration and differential equations* (Version 2.1). Springer.

Schättler, H., & Sung, J. (1993). The first-order approach to the continuous-time principal agent problem with exponential utility. *Journal of Economic Theory*, 61, 331-371.

Sims, C. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50, 665-690.

Spear, S. E., & Srivastava, S. (1987). On repeated moral hazard with discounting. *Review of Economic Studies*, 54, 599-617.

Weil, P. (1989). The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics*, 24, 401-421.