Moral hazard and continuous-time contract theory

Nizar Touzi

June 11, 2018

1 Introduction

The notion of moral hazard was introduced by Adam Smith as a major source of economic risk which occurs in any situation of interacting agents where one of the participants may take more risks whenever the burden of those risks is affected to the others. As the economic activity is based on social interactions, with agents having different information and different objectives, it is fundamental to avoid the negative effects of moral hazard on the society.

For instance, insurance contracts set conditions under which an economic agent is covered against some risk. Clearly, fully insured drivers do not have incentive to increase caution, full health insurance misleadingly lowers out-of-pocket expenses, and thus increases demand for medical services. Similarly, un-employment insurance may give negative incentive to return to the job market. In order to prevent from the last negative effects of moral hazard, insurance companies propose only partial compensation of the incurred loss, and the corresponding contracts index the insurance premium on the observed performance of the agent, as inferred from the claimed expenses. Contract theory aims at finding the terms of a contract which best incentivizes the insured for a moral and responsible behavior.

The ongoing digital transition introduces numerous examples of decentralized decision mechanisms which are self-managed by means of appropriate online devices opening access to a large number of participants. Such facilities provide the possibility of decision making at a very high frequency which maybe reduced to the continuous time context. Adapting contract theory to this context is crucial in order to solve the regulation and the optimal mechanism design of such devices. In this note, we focus on the Principal-Agent approach for the modeling of moral hazard in continuous time, and we present a systematic solution methodology.

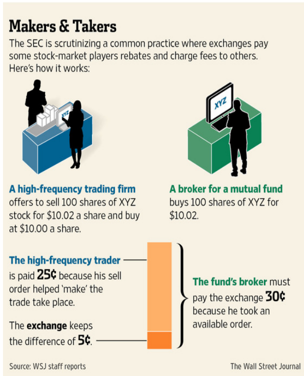

Example 1. The Maker-Taker fees problem of an online platform aims at incentivizing market makers (liquidity makers) to increase the market liquidity so as to attract brokers (liquidity takers), thus increasing the platform revenue accumulated by the fees on the realized transactions. This is turn contributes to increase market liquidity, thus improving the social benefit from the market.
In the situation illustrated by the present image, the broker pays 30 cts, the platform distributes 25 cts compensation to the market maker and makes a revenue of 5 cts.

Example 2. Smart meters (as Linky in France) give access to continuous time information about the electricity consumption profile. This offers to the electricity provider company the possibility of impacting the consumer demand by appropriate tarif signals, so as to improve its management by making use of the demand flexibility.

Example 3. An investor delegates the management of her capital to a fund manager. The latter devotes the effort of choosing the portfolio allocation, possibly in very high frequency, and balances the cost of effort with the compensation agreed on by the contract. The investor optimally chooses the terms of the contract so as to best incentivize the fund manager to work for her interest.

2 Management delegation and optimal contracting

Let us start with the simplest delegation model between a Principal (e.g. employer) and an Agent (e.g. employee). We consider the simplest one period model, with contracting and management decisions taken in the beginning of period, while the outcomes are collected at the end of period.

The Principal owns a production asset, with outcome value at the end of period denoted by X, and delegates its management to the Agent. The terms of this interaction between the Principal and the Agent is an agreement at the starting time on the contract which sets the compensation salary ξ.

Given the contract ξ, the Agent devotes effort α for the project, inducing the following (simplest) impact on the value of the production asset:

$$X := \alpha + X_0, \quad \text{with cost of effort} \quad C(\alpha),$$

for some given random value X_0, and some cost function $C : \mathbb{R} \longrightarrow \mathbb{R}$. Assuming that the Agent and Principal preferences are defined by utility functions U_A and U_P, respectively, their criteria are defined by the expected utilities

$$J_A(\xi, \alpha) := \mathbb{E}\big[U_A(\xi) - C(\alpha)\big] \quad \text{and} \quad J_P(\xi, \alpha) := \mathbb{E}\big[U_P\big(X - \xi\big)\big].$$

Here, we see clearly that a constant salary does not serve the objective of the Principal as the Agent's criterion would not involve the value of the output. In order to force the interest of the Agent in the production output X, it is clear that the contract $\xi = \xi(X)$ should be indexed on the performance of the Agent as deduced from the observation of the output X.

In order to avoid degenerate situations, we assume that contracts are restricted to the following set of acceptable contracts

$$\Xi(\alpha, \rho) \;\; := \;\; \big\{\xi : J_A(\xi, \alpha) \geq \rho\big\},$$

for some reservation utility ρ representing the Agent's acceptance level. This condition allows to avoid that the Principal diverts attention from the production output and takes instead benefit from the Agent's contract !

We now formulate two optimal contracting problems which differ by the power attributed to the Principal.

(FB) First best contracting. We first describe the (unrealistic) situation where the Principal imposes to the Agent the amount of effort to be devoted. Then, the Principal's problem is formulated by the maximization problem:

$$v_p \;\; := \sup_{\alpha, \xi \in \Xi(\alpha, \rho)} J_P(\xi, \alpha). \tag{2.1}$$

Introducing a Lagrange multiplier λ for the Agent's participation constraint, this reduces to

$$v_p \;\; = \;\; \inf_{\lambda \geq 0} \sup_{\alpha, \xi} J_P(\xi, \alpha) + \lambda\big(J_A(\xi, \alpha) - \rho\big) \;\; = \;\; \inf_{\lambda \geq 0} \sup_{\alpha, \xi} \mathbb{E}\big[U_P(X - \xi) + \lambda\big(U_A(\xi) - C(\alpha) - \rho\big)\big].$$

Direct differentiation with respect to ξ provides the so-called Borch rule in risk sharing:

$$\frac{U_P'(X - \xi)}{U_A'(X)} \;\; = \;\; \lambda.$$

In order to further characterize the optimal contract, we assume that both actors have a constant absolute risk aversion (CARA) coefficients $\eta_A, \eta_P > 0$, i.e. $U_A(x) = \eta_A^{-1} e^{-\eta_A x}$ and $U_P(x) = \eta_P^{-1} e^{-\eta_P x}$. Then the optimal contract is given by:

$$\hat{\xi}^{\mathrm{FB}} \;\; = \;\; \frac{\eta_A}{\eta_A + \eta_P} X + \frac{\ln \lambda}{\eta_A + \eta_P}.$$

In other words, the first best optimal contract consists in offering a fixed constant payment and a proportion of the total production. In particular, in the limiting case $\eta_A \longrightarrow \infty$ of risk neutral Agent, the Principal offers the total production against a constant payment.

(SB) Second best contracting. We now assume that the Agent may choose the effort to devote to the project as an optimal response to the contract proposed by the Principal. The Agent's decision problem is then defined by

$$V_A(\xi) \;:=\; \sup_\alpha J_A(\xi,\alpha) \;=\; J_A\big(\xi,\hat\alpha(\xi)\big),$$

where $\hat\alpha(\xi)$ is the optimal response of the Agent to the contract ξ, assumed unique for simplicity. Denoting the corresponding value of output by $\hat X := \hat\alpha(\xi) + X_0$, the Principal optimal contracting problem is:

$$V_P \;:=\; \sup_{(\hat\alpha(\xi),\xi)\in\Xi(\rho)} \mathbb{E}\big[U_P\big(\hat X - \xi(\hat X)\big)\big] \;=\; \sup_{\xi:V_A(\xi)\ge\rho} \mathbb{E}\big[U_P\big(\hat X - \xi(\hat X)\big)\big], \tag{2.2}$$

i.e. the Principal choses optimally the contract ξ given the optimal response of Agent.

No information rent in the one-period setting Clearly $v_P \ge V_P$. The information rent is defined by the difference $R := v_P - V_P$ representing the loss of utility incurred by the Principal by giving up the power of imposing the required effort to the Agent.

- When the Agent is risk neutral $U_A(x) = ax + b$ is affine, then one can see easily that $R = 0$, i.e. the Agent has no benefit from the power given up by the Principal;

- In fact, even in a situation of risk averse Agent, this result turns out to hold true for a large class of distributions for X_0. However, it is shown in the literature that this corresponds to a degenerate situation as V_P can only be realized as a limit of nearly-optimal contracts.

- Because of this, the economic literature focused on special types of parameterized contracts, e.g. affine contracts $\xi(x) = \xi_0 + \xi_1 x$ corresponding to a fixed salary ξ_0 and a proportional component to the realized output.

3 Continuous time contracting for drift management

The continuous-time formulation of the problem, as introduced by Hölstrom & Milgrom 1987, gained an important attention as it induces a possible information rent for a risk averse Agent. In the present setting the output process is defined by the dynamics

$$dX_t \;=\; \alpha_t dt + \sigma dW_t,$$

where the Agent effort α is now adjusted in continuous time, W is a Brownian motion representing an exogenous white noise, and σ is the volatility of the output. The Agent and Principal criteria are now given by

$$J_A(\xi,\alpha) := \mathbb{E}\Big[e^{-rT}U_A(\xi) - \int_0^T e^{-rt}c(\alpha_t)dt\Big] \quad\text{and}\quad J_P(\xi,\alpha) := \mathbb{E}\big[e^{-RT}U_P(X_T-\xi)\big].$$

The first best and second best problems are defined similarly to the one-period model by (2.1) and (2.2) respectively.

The optimal contract is characterized by stochastic control techniques. As a first ingredient, we introduce the so-called Hamiltonian

$$H(z) := \sup_{\alpha}\{\alpha z - c(\alpha)\} = \hat{\alpha}(z)z - c\big(\hat{\alpha}(z)\big),$$

where the maximizer $\hat{\alpha}(z)$ is assumed unique for simplicity. We next introduce a new state Y defined by some initial value Y_0 and a proportional payment process Z:

$$Y_t = Y_0 + \int_0^t e^{r(t-s)}\big(Z_s dX_s - H(Z_s)ds\big).$$

We emphasize that the compensation process Z may depend on the past history of the output process X. The following reduction of the game problem is stated in the context of a one-to-one Agent utility function on $\mathbb{R}$. For othe utility function (as the important CARA case), the result is valid after appropriate adaptation.

Let U_A be one-to-one on $\mathbb{R}$. Then the Principal value of the SB contracting reduces to

$$V_P = \sup_{Y_0 \geq \rho} V(Y_0), \text{ where } V(Y_0) := \sup_Z J_P\big(Y_T, \hat{\alpha}(Z)\big) = \sup_Z \mathbb{E}\big[e^{-RT}U_P(\hat{X}_T - U_A^{-1}(Y_T))\big].$$

Notice that $V(Y_0)$ is the standard stochastic control problem with controlled dynamics

$$d\hat{X}_t = \hat{\alpha}(Z_t)dt + \sigma dW_t \quad \text{and} \quad dY_t = \big[rY_t + c\big(\hat{\alpha}(Z_t)\big)\big]dt + Z_t\sigma dW_t.$$

By the classical tools of stochastic control theory, $V(Y_0) = v(0, X_0, Y_0)$, where v is the solution of the corresponding Hamilton-Jacobi-Bellman equation (HJB):

$$\begin{aligned}
&\partial_t v - Rv - ryv_y + \frac{1}{2}\sigma^2 v_{xx} + h(v_x, v_y, v_{yy}, v_{xy}) = 0, \quad v|_{t=T} = U_P\big(x - U_A^{-1}(y)\big),\\
&\textit{where} \;\; h(v_x, v_y, v_{yy}) := \sup_z \Big\{\hat{\alpha}(z)v_x + c\big(\hat{\alpha}(z)\big)v_y + \frac{1}{2}\sigma^2\big(z^2 v_{yy} + 2zv_{xy}\big)\Big\}.
\end{aligned}$$

In general, one does not hope for an explicit solution of this partial differential equation. Nevertheless, there is a long tradition of numerical approximation techniques ranging from finite differences methods to the most recent Monte Carlo methods for such nonlinear problems.

Finally, we may express the optimal contract in terms of the maximizer Y_0^* of $V(Y_0)$ over the range $Y_0 \geq \rho$, and the maximizer $\hat{z}(t, x, y)$ of h (which depends on v_x, v_y, v_{yy}): differential equation

$$\boxed{\hat{\xi} = Y_T^* \quad \text{with} \quad Y_t^* = Y_0^* + \int_0^t e^{r(t-s)}\big(Z_s^* dX_s - H(Z_s^*)ds\big) \text{ and } \; Z_t^* := \hat{z}(t, X_t, Y_t).}$$

In words, the optimal contract consists of a constant payment Y_0^*, a linear payment with instantaneous decomposition $Z^* dX$, and a substraction of the maximum payment that the rational Agent can hope for.

An explicit example. The last methodology is easily adapted to the case where the agent's criterion is defined by $J_P(\xi, \alpha) = \mathbb{E}\big[-e^{-\eta_A(\xi - \int_0^T c(\alpha_t)dt)}\big]$, where $\eta_A > 0$ is the (constant) absolute risk aversion coefficient of the agent, and the agent's discount factor is $r = 0$. Similarly, let the principal's utility function be defined by a constant risk aversion coefficient $\eta_B > 0$, i.e. $U_P(x) = -e^{-\eta_P x}$, and zero discount factor $R = 0$. Finally, consider the case of a quadratic cost of effort $c(a) := \frac{1}{2}c_0 a^2$. Then, the last methodology leads to the following (constant) optimal effort and optimal contract:

$$\alpha_t^* = \frac{z^*}{c_0} \quad \text{and} \quad \hat{\xi} = \rho + z^*(X_T - X_0) - H(z^*)T, \quad \text{where} \quad z^* := \frac{\frac{1}{c_0} + \eta_P \sigma^2}{\frac{1}{c_0} + (\eta_A + \eta_P)\sigma^2}.$$

For a risk neutral agent $\eta_A = 0$, notice that $z^* = 1$, and we recover the first order optimal contract $\hat{\xi}^{\rm FB}$, in agreement with the comment at the end of Section 2.

Limited liability. In practice, the agent is not expected to have access to infinite amount of cash. Therefore, the Principal cannot offer a contract which may induce any negative payment to the agent, and the optimal contracting problem should be formulated under the addition restriction to the set of contracts which obey to the limited liability restriction. This leads to the following modification of the principal's second best (2.2) to the following problem:

$$V_P^{\rm ll} := \sup_{\xi : V_A(.,\xi) \geq \rho} \mathbb{E}\big[U_P\big(\hat{X} - \xi(\hat{X})\big)\big], \text{ where } V_A(t,\xi) := \sup_\alpha \mathbb{E}_t\Big[e^{-r(T-t)}U_A(\xi) - \int_t^T e^{-r(s-t)}c(\alpha_s)ds\Big],$$

and $\mathbb{E}_t$ denotes the expectation conditional conditional to the available information at time t. In words, for all $t \leq T$, $V_A(t, \xi)$ is the agent dynamic value function defined similar to the initial problem by simply moving the time origin to t.

The previous methodology is easily adapted to this context, and leads to the following characterization:

Under the limited liability constraint, the Principal value of the SB contracting reduces to

$$V_P^{\rm ll} = \sup_{Y \geq \rho} V(Y_0), \text{ where } V(Y_0) := \sup_Z J_P\big(Y_T, \hat{a}(Z)\big) = \sup_Z \mathbb{E}\big[e^{-RT}U_P(\hat{X}_T - U_A^{-1}Y_T)\big].$$

The resulting problem is again a stochastic control problem, with an additional feature of a state constraint $Y_t \geq \rho$ for all $t \in [0, T]$. This additional constraint can be handled by standard tools in stochastic control theory.

4 Optimal contracting for volatility management

The notion of hedging is a key finding at the origin of the development of financial mathematics. This was highlighted by the seminal works of Black & Scholes for derivatives valuation, and Merton for portfolio management. In both contexts, the continuous time updating of the portfolio leads typically to a portfolio value dynamics $dX_t = \Delta_t dS_t$, where Δ represent the portfolio holding of risky asset at each point in time, and S the price process of the risky asset. In the present context, X is the output process, the fund manager in charge of designing the portfolio composition Δ is the Agent, and the Principal is the investor.

Stipulating appropriate dynamics for the process S, we arrive to an output process with dynamics of the form

$$dX_t = \alpha_t dt + \sigma_t dW_t,$$

where the Agent's effort has now two components:

- σ_t representing the position in risky assets (in direct relation with Δ_t),
- α_t may be viewed as the (instantaneous) choice of a subset of assets for the portfolio optimization.

We update the Agent's and the Principal's criteria so as to accounts for the cost of the additional effort σ:

$$J_A(\xi, \alpha, \sigma) := \mathbb{E}\Big[e^{-rT} U_A(\xi) - \int_0^T e^{-rt} c(\alpha_t, \sigma_t) dt\Big] \quad \text{and} \quad J_P(\xi, \alpha, \sigma) := \mathbb{E}\big[e^{-RT} U_P(X_T - \xi)\big].$$

In order to adapt the methodology described in the previous section, we start by introducing the Hamiltonian suitable with the present setting:

$$H(z, \gamma) := \sup_{\alpha, \sigma} \Big\{\alpha z + \frac{1}{2}\sigma^2 \gamma - c(\alpha, \sigma)\Big\} = \hat{\alpha}(z, \gamma) z + \frac{1}{2}\hat{\sigma}(z, \gamma)^2 \gamma - c\big(\hat{\alpha}(z, \gamma), \hat{\sigma}(z, \gamma)\big). \quad (4.1)$$

where the maximizer $(\hat{\alpha}, \hat{\sigma})$ is again assumed to be unique for simplicity. Similar to the previous section, we introduce a new state Y defined by an initial value Y_0 and the dynamics

$$Y_t \;=\; Y_0 + \int_0^t e^{r(t-s)} \Big(Z_s dX_s + \frac{1}{2}\Gamma_s d\langle X\rangle_s - H(Z_s, \Gamma_s) ds\Big). \quad (4.2)$$

Here, $\langle X\rangle$ is the quadratic variation of the output process X, i.e. the squared volatility of X. The quadratic variation $\langle X\rangle$ is an indicator of the riskiness of the value of the output, and can be computed from the observation of the path of X as the following limit:

$$\langle X\rangle_t = \lim_{n\to\infty} \sum_{i=1}^n \big(X_{t_i^n} - X_{t_{i-1}^n}\big)^2 = \int_0^t \sigma_s^2 ds,$$

where $0 = t_0^n < t_1^n < \ldots < t_n^n = t$, with time step $t_i^n - t_{i-1}^n \to 0$ as $n \to \infty$ for all i. Similar to the uncontrolled volatility case, we have:

The SB optimal contracting problem reduces to

$$V_P = \sup_{Y_0 \geq \rho} V(Y_0), \quad \text{where} \quad V(Y_0) := \sup_{Z,\Gamma} J_P\big(Y_T, \hat{a}(Z,\Gamma), \hat{b}(Z,\Gamma)\big).$$

Again, $V(Y_0)$ is the standard stochastic control problem with controlled dynamics

$$d\hat{X}_t = \hat{\alpha}(Z_t,\Gamma_t)dt + \hat{\sigma}(Z_t,\Gamma_t)dW_t \text{ and } dY_t = \big[rY_t + c\big(\hat{\alpha}(z,\gamma),\hat{\sigma}(z,\gamma)\big)\big]dt + Z_t\hat{\sigma}(Z_t,\Gamma_t)dW_t,$$

and with characterization as $V(Y_0) = v(0, X_0, Y_0)$, where v is the solution of the Hamilton-Jacobi-Bellman equation (HJB):

$$\begin{aligned} &\partial_t v - Rv - ryv_y + h(v_x, v_y, v_{xx}, v_{yy}, v_{xy}) = 0, \quad v|_{t=T} = U_P(x-y), \\ &\textit{where} \;\; h(v_x, v_y, v_{xx}, v_{yy}, v_{xy}) := \sup_{z,\gamma} \Big\{ \hat{\alpha}(z,\gamma)v_x + c\big(\hat{\alpha}(z,\gamma),\hat{\sigma}(z,\gamma)\big)v_y \\ &\qquad\qquad\qquad\qquad + \frac{1}{2}\hat{\sigma}(z,\gamma)^2\big(v_{xx} + z^2 v_{yy} + 2v_{xy}\big)\Big\}. \end{aligned}$$

Here again, the solution of the last partial differential equation can be approximated by some numerical method, which also produces as a by-product an approximation of the maximizers $\hat{z}(t,x,y)$ and $\hat{\gamma}(t,x,y)$ of h (which depend on the partial derivatives of v). The last step is to find the maximizer Y_0^* of $V(Y_0)$ over the range $Y_0 \geq \rho$, inducing the optimal contract

$$\begin{aligned} \hat{\xi} = Y_T^* \quad &\text{with} \quad Y_t^* = Y_0^* + \int_0^t e^{r(t-s)}\Big(Z_s^* dX_s + \frac{1}{2}\Gamma_t^* d\langle X\rangle_t - H(Z_s^*, \Gamma_t^*)dt\Big) \\ &\text{and} \quad (Z_t^*, \Gamma_t^*) := (\hat{z},\hat{\gamma})(t, X_t, Y_t). \end{aligned}$$

In the present situation we see that the optimal contract contains the additional infinitesimal contribution $\Gamma_t^* d\langle X\rangle_t$. One typically expects a negative Γ^* so that the optimal contract settled by the Principal discourages the Agent from taking unconsidered risks.

Example Suppose that the principal delegates to the agent the management of a risky project. The agent chooses volatility components of the output process over various risk factors. The principal has no access to the individual components of the quadratic variation, thus inducing a situation of moral hazard with respect to the risk choices of the agent. This model is solved in Cvitanić, Possamaï & Touzi. The optimal contract is shown to be linear in the output and the corresponding quadratic variation. Numerical experiments reveal a significant loss of efficiency if the principal does not use the quadratic variation component of the optimal contract. However, there are parameter values for which the principal rewards the agent for higher values of quadratic variation, thus, for taking higher risk.

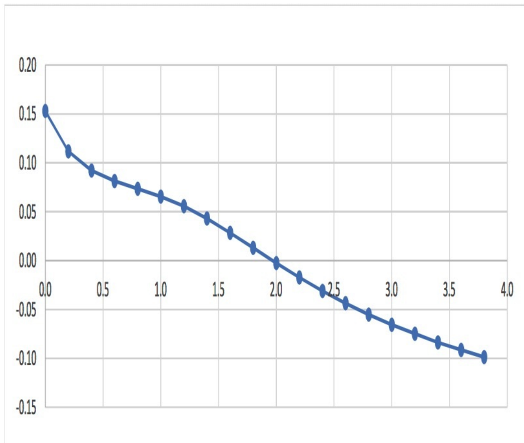

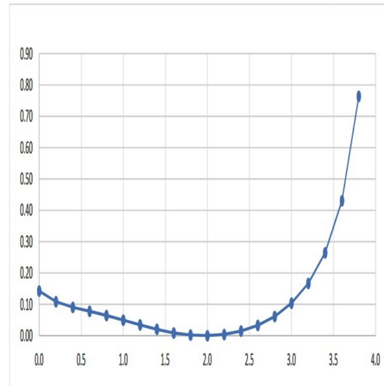

Parameter sensitivity of the percentage loss in principal's certainty equivalent when not using quadratic variation

Sensitivity of optimal contract to quadratic variation

5 The multiple agents case

The previous approach extends naturally to the multiple agents context.

We start with the case of non-interacting finite agents $i = 1, \dots, I$. Each agent i has the delegation to manage an output X^i, with dynamics $dX^i = \alpha^i_t dt + \sigma^i_t dW_t$ for some Brownian motion W, and some effort processes $\nu^i = (\alpha^i, \beta^i)$, and is characterized by the utility index $J^i_A(\xi, \nu) := \mathbb{E}\big[e^{-r^i T} U^i_A(\xi^i) - \int_0^T e^{-r^i t} c(\nu_t) dt\big]$, for some cost function $c^i(\nu^i)$ and discount factor r^i. In this case, we introduce a Hamiltonian H^i together with an additional state Y^i for each agent defined similar to (4.1)-(4.2) by:

$$Y^i_t \;:=\; Y^i_0 + \int_0^t e^{r^i(t-s)} \Big(Z^i_s dX^i_s + \frac{1}{2} \Gamma^i_s d\langle X^i \rangle_s - H^i(Z^i_s, \Gamma^i_s) ds \Big), \quad t \in [0, T].$$

The Principal's problem is defined by the utility index $J_P\big((\xi^i)_{i \in I}, \nu^i\big) := \mathbb{E}\big[e^{-RT} U_P\big(\sum_{i=1}^I X^i_T - \xi^i\big)\big]$, and the contracting problem between the Principal and the group of Agents is defined by a Stackelberg game similar to the previous sections, with participation constraint ρ^i. Denoting by $\hat{\nu}^i(z^i, \gamma^i)$ the maximizer of Agent i's Hamiltonian H^i, and $Y_0 := (Y^1_0, \dots, Y^I_0)$, it follows that the Principal's problem reduces to a standard control problem:

$$V_P = \sup_{Y^i_0 \geq \rho^i} V(Y_0), \quad \text{where} \quad V(Y_0) := \sup_{Z^i, \Gamma^i} J_P\big((Y^i_T)_{i \in I}, \hat{\nu}^i(Z^i, \Gamma^i)\big).$$

Similar to the one-agent case, the last control problem can be characterized by means of the corresponding HJB equation. Moreover , the maximizers $\hat{Y}_0^i, \hat{Z}^i, \hat{\Gamma}^i$ induce an optimal contract

$$\hat{\xi}^i = \hat{Y}_0^i + \int_0^T e^{r^i(T-s)} \Big(\hat{Z}_s^i dX_s^i + \frac{1}{2} \hat{\Gamma}_s^i d\langle X^i \rangle_s - H^i(\hat{Z}_s^i, \hat{\Gamma}_s^i) ds \Big), \quad i = 1, \ldots, I.$$

The case of interacting agents can be dealt with in a similar way. For instance, if the agents behavior is described as a Nash equilibrium, the last methodology is easily adapted to reflect this situation. The case of a continuum of agents can also be addressed by the same method.

Literature overview

Since many decades, economic models accounting for moral hazard have been developed in the economics literature, with real applications in our every day life. We refer to the books by Laffont & Martimort [1] and Bolton & Dewatripont [2] for an overview of the corresponding discrete-time literature. The first continuous-time formulation, as introduced by Holstrom & Milgrom 1985, highlighted the suitability of the continuous-time setting for a convenient resolution of the problem, thanks to the availability of differential calculus. We refer to the book of Cvitanić & Zhang (2008) for a comprehensive exposition of the theory. The seminal work of Holstrom & Milgrom was followed by an important stream of literature, and was even further developed after the very inspiring work of Sannikov 2004. The solution method of the present note was introduced by Cvitanić, Possamaï & Touzi 2015.

References

[1] J.-J. Laffont and D. Martimort. *The Theory of Incentives.* Princeton University Press, 2002.

[2] P. Bolton and M. Dewatripont. *Contract theory.* MIT press, 2005.

[3] J. Cvitanić, D. Possamaï, and N. Touzi. Moral hazard in dynamic risk management. *Management Science*, forthcoming, 2017.

[4] J. Cvitanić, D. Possamaï, and N. Touzi. Dynamic programming approach to principal-agent problems. Finance and Stochastics, to appear.

[5] R. Élie and D. Possamaï. Contracting theory with competitive interacting agents. arXiv:1605.08099 (2016).

[6] B. Holmstrom and P. Milgrom. Aggregation and linearity in the provision of intertemporal incentives, *Econometrica*, 55(2):303–328 (1987).

[7] Y. Sannikov. A continuous–time version of the principal–agent problem. *The Review of Economic Studies*, 75(3):957–984, 2008.