



CENTER FOR RISK MANAGEMENT RESEARCH

Working Paper # 2014 - 04

“Continuous-Time Principal-Agent Problem with Drift and
Stochastic Volatility Control: With Applications to Delegated
Portfolio Management”

Raymond C. W. Leung
University of California, Berkeley
Haas School of Business, Department of Finance

May 4, 2014

University of California
Berkeley

Continuous-Time Principal-Agent Problem with Drift and Stochastic Volatility Control: With Applications to Delegated Portfolio Management

Raymond C. W. Leung *

Current draft: May 4, 2014

Abstract

We study a continuous-time principal-agent problem where the risk-neutral agent can privately and meaningfully choose the drift and volatility of a cash flow, while the risk-neutral principal only continuously observes the managed cash flows over time. Our model contributes a result that is hitherto relatively unexplored in both the continuous-time dynamic contracting and the delegated portfolio management literatures. Firstly, even though there is no direct moral hazard conflict between the principal and the agent on their preferred volatility choices, but to avoid inefficient termination and compensation from excess diffusion, this first best choice is not reached; this is the “reverse moral hazard” effect. Secondly, the dollar incentives the principal gives to the agent critically depends on the volatility choice, endogenous quasi-risk aversion of the principal, and the elasticity to the exogenous factor level; this is the “risk adjusted sensitivity” (RAS) effect. In a delegated portfolio management context, our model suggests outside investors should prefer investment funds with the characteristics: (i) the investment fund has an “internal fund” available only to management, or equivalently, a “flagship fund” that is closed to new outside investors; (ii) the “external fund” for the outside investors closely tracks the investment strategy and value of the internal fund; and (iii) has dynamic incentive fee schemes.

JEL classification: D82, D86, G32, J32

Keywords: continuous-time principal-agent problem, dynamic contracting, delegated portfolio management, stochastic volatility control

*Department of Finance, Haas School of Business, University of California, Berkeley. Email: r_leung@haas.berkeley.edu. I am deeply indebted to and grateful for the countless hours, encouragement, patience and comments of both of my advisers, Robert M. Anderson and Gustavo Manso, on this project and other research endeavors. I also thank Xin Guo, Dmitry Livdan, Alexei Tchisty, and Johan Walden for very helpful discussions. Comments and suggestions from the seminar attendees of the UC Berkeley Center for Risk Management Research, UC Berkeley-Haas Real Estate Student Seminar, UC Berkeley-Haas Finance Student Seminar, and Berkeley-Stanford Joint Finance Student Seminar are gratefully appreciated. All remaining errors are mine and mine alone.

1 Introduction

As of 2009, \$71.3 trillion is invested into managed portfolios worldwide ¹, and the vast majority of these managed portfolios are under active management. Despite the prevalence and importance of delegated portfolio management in the modern capital markets, surprisingly little is known about its optimal contracting characteristics in a dynamic environment. The difficulty of approaching these problems is succinctly captured by a remark in Cuoco and Kaniel (2011):

“A distinctive feature of the agency problem arising from portfolio management is that the agent’s actions (the investment strategy and possibly the effort spent acquiring information about securities’ returns) affect both the drift and volatility of the relevant state variable (the value of the managed portfolio), although realistically the drift and the volatility cannot be chosen independently. This makes the problem significantly more complex than the one considered in the classic paper by Holmström and Milgrom (1987) and its extensions. With a couple of exceptions, as noted by Stracca (2006) in his recent survey of the literature on delegated portfolio management, ‘the literature has reached more negative rather than constructive results, and the search for an optimal contract has proved to be inconclusive even in the most simple settings’ ”.

To emphasize the point, the ability to influence the volatility of a managed cash flow is critical in a delegated portfolio management context. Indeed, numerous papers have recognized that risk shifting ² behavior of the portfolio manager as an important source of moral hazard that is typically not present in traditional principal-agent contexts, such as employer-employee and landlord-tenant relationships.

Here, we present a continuous-time principal-agent model that represents both a first step in the literature in dynamic contracting theory whereby the agent can explicitly and meaningfully privately choose volatility, and equally important, also as a first step into understanding dynamic contracting environment in the context of delegated portfolio management. We consider a dynamic contracting environment in continuous-time with a risk-neutral agent and a risk-neutral principal, whereby the agent can privately choose effort and volatility levels that affect both the mean and overall risk of the cash flows. The principal can continuously observe the cash flows, but not the hidden choices of effort and volatility that the agent chooses. In line with the literature, the agent enjoys a private benefit from exerting low levels of effort (“job shirking”), and we also further assume that the agent enjoys a private benefit from choosing high level volatility control (“lazy quality management”). The expected payoff of the managed cash flows (i.e. the drift) is a “reward function” of both the effort and volatility chosen by the agent. This is to roughly capture the classical “risk-reward trade off” intuition of financial economics, particularly in portfolio choice theory. Here, effort is a binary choice but volatility is chosen from a closed interval. The problem of drift only control, to various degrees of sophistication, has been extensively studied in recent years (see Section 2 for a literature review).

¹Wermers (2011)

²See Stoughton (1993) and Admati and Pfleiderer (1997); see also Stracca (2006) for a summary on how the delegated portfolio management problem presents unique challenges that are not present in standard principal-agent problems.

However, to the best knowledge of this author, continuous volatility control when the principal can continuously observe the cash flows has been given little to no attention in the models commonly used in the literature, and for good reason. Specifically, if the principal can continuously observe the cash flows and the agent directly controls the volatility term of the cash flow, the principal can compute the quadratic variation of the cash flow process, and thereby infer directly the choice of volatility that the agent has been using. Hence, if indeed the agent has deviated from the principal’s prescribed volatility level, the principal could in effect apply a “grim trigger” strategy and punish the agent indefinitely thereon. By using this argument, the agent will never have any incentive to deviate from the principal’s desired and prescribed volatility level, and thus, volatility control by the agent can essentially be abstracted away. Yet economically, the absence of meaningful volatility control by the agent is very unsatisfying. There are several important situations where allowing the agent to influence the volatility of the cash flows is economically significant; for example, the classical considerations of asset substitution in corporate finance ³ and risk shifting in delegated portfolio management. Thus to have any meaningful volatility control by the agent, a richer model of the agent managed cash flows is required.

The key ingredient of our model is to introduce an exogenous *factor level* component to the overall diffusion term of the managed cash flows. In particular, we will allow the instantaneous diffusion of the cash flows be a product of an exogenous factor level process that is completely not managed by the agent, and a component that is directly controlled by the agent. The agent can observe this exogenous factor term, and off equilibrium, this exogenous factor level is unobservable to the principal, even though at equilibrium becomes observable to the principal. Thus when the principal computes the quadratic variation of this cash flow process, the principal can at best observe the product of an unobservable exogenous factor level and a controlled volatility term, but not the two components separately. So economically, even if the principal observes high instantaneous cash flow volatility through the computation of the quadratic variation, the principal cannot disentangle whether this high volatility is due to a high realization of the exogenous factor level, high volatility control by the agent, or both. Clearly, the principal should only punish or reward the agent for the endogenous volatility control by the agent and not for the exogenous component. As well, in line with the models of drift-only control by the agent, the principal must put the agent at risk to induce the agent to work according to the principal’s desired and prescribed plan. However, given that the agent can choose the volatility of the cash flows, the agent can effectively undo or weaken some of the risks that the principal imposes on him. Thus, the incentives involved in a model with combined drift and volatility control are, perhaps understandably, considerably more difficult than a case with only drift control.

Let us now discuss economically the form of the optimal contract. The two relevant state variables here are the agent’s continuation value process and also the exogenous factor level. The use of the agent’s continuation value as a tracker to punish or reward the agent is a standard argument in this literature (see Sannikov (2008)). The principal will terminate the agent when the agent’s continuation value falls below the agent’s outside option, and the principal will receive some recovery value of the firm. In addition, there exists a free boundary that is a function of the factor level,

³Say Leland (1998), among many others.

so namely not a constant, such that when the agent’s continuation value first hits this boundary, a lump sum compensation is made immediately from the principal to the agent. As it is common with principal-agent problems both in discrete or continuous time, the principal needs to incentivize and impose sufficient amount of risk to the agent’s payoff to ensure the agent will make those latent choices that the principal desires. However, since the agent can directly influence the level of uncertainty in this economy, namely through the agent’s choice of volatility, the agent effectively can partially undo or weaken the amount of risk the principal imposes on him. Furthermore, the existence of an exogenous factor level that affects both the principal and the agent further affects the amount of risk the principal can directly impose on the agent. Thus, to this end, we will introduce a concept labeled *risk adjusted sensitivity* (RAS) to represent the amount of risk the principal imposes on the agent, or equivalently, the dollar-per-performance the agent is entitled to. Most notably, RAS is identically a constant in drift-only control models like DeMarzo and Sannikov (2006) and He (2009), and many others, but this is not the case here. Indeed, we can decompose and understand RAS as coming from the volatility precision chosen by the agent, the “endogenous risk aversion” of the principal, and the “elasticity of exogenous factor level”.

Bringing back the discussion and concrete application of delegated portfolio management, when outside investors are seeking an investment firm, of which the prototypical example is a hedge fund, with good corporate governance, our model suggests the outside investors should actively look for investment firms with the following characteristics: (i) the firm has an internal fund available only to management, and an external fund only available to external investors; (ii) the investment strategies of the internal fund and of the external fund are closely correlated to each other; and (iii) the firm has dynamic watermark compensation schemes.

2 Related Literature

This paper contributes to: (i) a growing literature on continuous-time principal-agent problems; and (ii) continuous-time delegated portfolio management problems.

One of the first papers that considered a continuous-time principal agent problem is Holmström and Milgrom (1987). Recent papers in the continuous-time principal-agent problem include DeMarzo and Sannikov (2006) (of which DeMarzo and Fishman (2007) is the discrete-time counterpart), Biais, Mariotti, Plantin, and Rochet (2007), Sannikov (2008), He (2009), Adrian and Westerfield (2009), Hoffmann and Pfeil (2010), Grochulski and Zhang (2011), He (2011), Williams (2011), DeMarzo, Fishman, He, and Wang (2012), He (2012), Szydlowski (2012), Miao and Zhang (2013), Miao and Rivera (2013), Zhu (2013), DeMarzo, Livdan, and Tchisty (2013), Giat and Subramanian (2013) and Hoffmann and Pfeil (2013). We note that Biais, Mariotti, and Rochet (2011), Sannikov (2012a), and Sannikov (2013) all give an excellent survey and overview of the current state in this literature. Please see Table 1 for a selected survey of the models used in the literature; note that even though the table enumerates the agent’s managed cash flow form, these papers often have very different assumptions on the preferences of the agent and the principal, and some also have different assumptions of the timing in which the principal can observe the cash flows.

All the aforementioned papers allow the agent to manage a cash flow in the form of a stochastic

differential equation, of various levels of complexity, but the common setup is that the agent can only exert effort to influence the drift of the cash flow but not its volatility. But in these sorts of papers, the volatility parameter is held constant, known both to the principal and the agent. This is without loss of generality in the case when the noise term of the cash flow is driven exclusively by Brownian motion. DeMarzo, Livdan, and Tchisty (2013) is an interesting example whereby the cash flows have a jump component and the agent can influence the jump, but nonetheless, the agent still does not (and cannot meaningfully) influence the volatility. To our best knowledge, there are some notable exceptions and we will describe these below in Section 2.1 but regardless, none of them allow for meaningful volatility control as in our context.

The setup of this model lends itself naturally to delegated portfolio management problems. As emphasized by Stoughton (1993) and Admati and Pfleiderer (1997), and summarized in Stracca (2006), delegated portfolio management problems present challenges that are not commonly considered in standard principal-agent problems; in particular, the portfolio manager has the ability to influence both the expected return and also volatility of the managed returns or cash flows. While managing expected return part, usually modeled as moral hazard hidden effort selection, is common in standard principal-agent problems, managing volatility is not. Ou-Yang (2003) is one of the key models in the delegated portfolio management literature but modeled in continuous time and we will further discuss this case in Section 2.1.

The problem of “risk shifting”, namely changing volatility of the managed cash flows, is well recognized as a key moral hazard component in the delegated portfolio management literature. Basak, Pavlova, and Shapiro (2007) considers a portfolio manager’s risk taking incentives induced by an increasing and convex relationship of fund flows to relative performance, and how this objective could give rise to risk-shifting incentives. However, it should be noted that the contract in Basak et al. (2007) is exogenously given and there is no explicit principal-agent modeling. Other papers that investigate into risk shifting behavior by portfolio managers include: Chevalier and Ellison (1997), Rauh (2008), Giambona and Golec (2009), Hodder and Jackwerth (2009), Foster and Young (2010) and Huang et al. (2011). Stracca (2006) and Ang (2012) offer excellent recent surveys on the literature in delegated portfolio management.

2.1 Selected important special cases

Several papers in the literature have some level of volatility control in the dynamic principal-agent problem in continuous-time. However, all of them place various levels of restrictive assumptions on the way the agent can control volatility, which is not imposed in our setup.

2.1.1 Sung (1995, 2004)

The papers by Sung (1995) and Sung (2004) are the closest in terms of volatility control but still does not resolve the problem that we have in mind for this paper.

Let’s first review Sung (1995). The author considers a finite time horizon, $[0, 1]$. The cash flows under management is of the form $dY_t = \mu_t dt + \sigma_t dB_t$, where the drift μ_t and volatility σ_t are under

	Cash flow dynamics	References
Brownian motion with controlled drift	$dY_t = \boldsymbol{\mu}_t dt + \sigma dB_t$	Holmström and Milgrom (1987), Biais, Mariotti, Plantin, and Rochet (2007), Sannikov (2008), DeMarzo and Sannikov (2006), Adrian and Westerfield (2009), Zhu (2013)
Geometric Brownian motion with controlled drift	$dY_t = \boldsymbol{\mu}_t Y_t dt + \sigma Y_t dB_t$	He (2009)
General Ito diffusion with controlled drift	$dY_t = f(t, Y, \boldsymbol{\mu}) dt + \sigma(t, Y) dB_t$	Schaettler and Sung (1993)
Brownian motion with controlled drift and controlled jump	$dY_t = (\alpha + \rho \boldsymbol{\mu}_t) dt + \sigma dB_t - D \boldsymbol{\mu}_t dN_t$, α, ρ, D constants	DeMarzo, Livdan, and Tchisty (2013)
Brownian motion with controlled drift via long run incentives	$dY_t = \delta_t dt + \sigma dB_t$, $\delta_t = \int_0^t f(t-s) \boldsymbol{\mu}_s ds$	Sannikov (2012b)
Controlled Poisson intensity	$dY_t = C dN_t$, N has intensity process $\{\boldsymbol{\mu}_t\}_{t \geq 0}$, $C > 0$ constant	He (2012)
Linear Ito diffusion with controlled drift and volatility	$dY_t = f(\boldsymbol{\mu}_t, \boldsymbol{\sigma}_t) dt + \boldsymbol{\sigma}_t dB_t$	(*) Sung (1995, 2004)
Geometric Brownian motion with same control on drift and volatility	$dY_t = [rY_t + \boldsymbol{\mu}_t(\alpha - r)] dt + \boldsymbol{\mu}_t \kappa dB_t$, r, α, κ constants	(*) Ou-Yang (2003)
Geometric Brownian motion with controlled drift and volatility	$dY_t = \kappa Y_t dt + \delta \boldsymbol{\mu}_t dt + \alpha \boldsymbol{\sigma}_t V_t dt + \boldsymbol{\sigma}_t V_t dB_t$, κ, δ, α constants	(*) Cadenillas, Cvitanić, and Zapatero (2004)

Table 1: A selected survey of agent’s managed cash flows in the existing literature. Here, B denotes a standard Brownian motion and N denotes a Poisson process. For consistency, the notations here differ from that of the original papers. The agent’s control (in **bold** for emphasis) is $\boldsymbol{\mu} = \{\boldsymbol{\mu}_t\}_{t \geq 0}$ (and where relevant, $\boldsymbol{\sigma} = \{\boldsymbol{\sigma}_t\}_{t \geq 0}$). For the (*) starred cases where volatility is (seemed to be) under control, please see section 2.1 for discussions of their key caveats.

the agent’s control. A contract is signed between the principal and the agent at time $t = 0$ and the agent is compensated at time $t = 1$. The agent incurs an integrated cost $\int_0^1 c(\mu_t, \sigma_t) dt$ for choosing between the drift μ_t and σ_t over the investment horizon $[0, 1]$. The principal is restricted to compensate the agent according to an exogenously fixed “salary function” S . The author considers two cases. In the first case, the principal observes the entire path of $\{Y_t\}_{t \in [0,1]}$ and hence can also observe σ_t across this path (i.e. via the quadratic variation of Y). Hence, there is no need for the principal to provide incentives to control the volatility σ_t . In the second case, the principal can only observe the ending cash flow value Y_1 . In such a case, the principal cannot know what volatility control σ_t the agent had chosen over the investment horizon $[0, 1]$.

Sung (2004) is similar to Sung (1995), except that the author allows for a more general specification of the cash flow process, and restrict to the second case setup of Sung (1995), whereby the principal can only observe the initial Y_0 cash flow and the terminal Y_1 cash flow. Specifically, Sung (2004)’s specification is of the form $dY_t = f(\mu_t, \sigma_t) + \sigma_t dB_t$, where the agent controls both the drift μ_t and volatility σ_t . The details in the preferences of the agent and principal differ slightly between Sung (1995) and Sung (2004) and we defer the reader to the actual papers for details.

Unlike Sung (1995) and Sung (2004), in this paper we will explicitly allow the principal to observe the agent’s managed cash flow process at all times.

2.1.2 Ou-Yang (2003)

In Ou-Yang (2003), the principal-agent problem is the in form of an investment manager (i.e. agent) has to manage a portfolio for an investor (i.e. principal). Asset returns follow the familiar geometric Brownian motion and together with the risk free asset, it induces a wealth process for the portfolio. The agent can choose the portfolio process and the conflict arises when the investor cannot observe the manager’s chosen portfolio policy ⁴. In this setup, the portfolio choice variable is attached to the diffusion term of the wealth process. But as duly noted in Ou-Yang (2003, Page 178): “If the investor observes both the stock price vector $P(t)$ and the wealth process $W(t)$ of the portfolio continuously, then she can infer precisely the manager’s portfolio policy vector $A(t)$ from the fact that the instantaneous covariance between $W(t)$ and $P(t)$ equals $\text{diag}(P)\sigma\sigma^T A(t)$. Since $\sigma\sigma^T$ is invertible by assumption, the manager’s policy vector $A(t)$ is completely determined. Hence we must assume that the investor does not observe the wealth and stock processes simultaneously.” Hence, through this, rather strong, assumption in restricting what the investor can observe over time, volatility can be controlled without detection by the principal. In this paper, we will not impose such a strong assumption that restricts the principal’s information set.

⁴The wealth process, referring to Ou-Yang (2003, Equation (1)), has the form,

$$dW(t) = [rW(t) + A(t)(\mu - r)]dt + A(t)\sigma dB_t,$$

where r is the risk free rate, and μ is the expected return of risky assets, and $A(t)$ is the portfolio choice policy. Note in particular the choice variable $A(t)$ enters both into the drift and volatility of the wealth process.

2.1.3 Cadenillas, Cvitanić, and Zapatero (2004, 2007)

Cadenillas et al. (2004) does allow the agent to have explicit drift and volatility control but the compensation type is exogenously given. Using the original notation of Cadenillas et al. (2004, Equation (2)), the agent (manager) manages the value of assets V under the agent's management evolves according to,

$$dV_t = \mu V_t dt + \delta \mu_t dt + \alpha v_t V_t dt + v_t V_t dW_t,$$

where u and v are the agent's controls. Moreover, the paper assumes that while effort (drift) u control for the agent is costly, project selection (volatility) v control incurs no cost on the agent, but only implicitly matters to the agent through the principal's compensation. The principal (outside investors) is exogenously allowed to only compensate the agent with stock that becomes vested at a terminal time T . The principal simply needs to choose the number of shares of stock to give to the agent and the level of debt of the firm ⁵. In all, taken in this light, Cadenillas et al. (2004)'s interesting approach of the problem does not have the key ingredients that are present in this paper. Firstly, we do not exogenously fix what the compensation contract the principal must give to the agent, and indeed the compensation structure is dynamic and endogenous. And secondly, the agent does incur private benefits (or equivalently, negative private costs) of controlling the unobservable volatility level of cash flows and so the principal must provide incentives on both effort and volatility.

Cadenillas et al. (2007) allows for the agent to control both the drift and volatility but they explicitly consider a first-best risk-sharing setup whereby agency problems are absent.

2.1.4 Cvitanić, Possamaï, and Touzi (2014)

The closest work in the literature to our paper is the contemporary work by Cvitanić, Possamaï, and Touzi (2014). ⁶ The goal of that paper is also to investigate under what conditions would there be meaningful volatility control by the agent. The authors propose that if the principal can only observe managed cash flows Y continuously overtime, and if the agent controls a *vector* of volatilities $\nu \in \mathbb{R}_+^d$, so that the managed cash flows have the form,

$$Y_t = \int_0^t \nu_s \cdot (b dt + dB_t),$$

where B is a d -dimensional Brownian motion, and $b \in \mathbb{R}^d$ is a common knowledge constant vector. Again, since the principal can continuously observe the managed cash flows Y , then the principal can compute the quadratic variation of Y and obtain an (integrated) matrix,

$$d[Y]_t = \nu_t \cdot \nu_t dt,$$

⁵It should be noted that in a similar spirit, Carpenter (2000) considers a delegated portfolio choice problem of which through the portfolio choice policy, the agent can choose the volatility of the value of the asset portfolio. There are no private benefits or costs to the agent in choosing a particular portfolio choice policy. Exogenously, the principal compensates the agent (only) with a call option with the strike price being the terminal value of the managed portfolio.

⁶This author was not aware of the presence of this paper (with a working date of March 7, 2014) until mid April 2014, but by then, a well working draft of this current paper had already been written and indeed circulated in small private circles.

where, of course, $\nu \cdot \nu'$ is a scalar. But given that the principal can only observe Y , there is no way the principal can decipher the individual managed elements of $\nu \cdot \nu$. Thus, this yields to a setup where there is meaningful volatility control. It should be immediately noted that in the setup of Cvitanić et al. (2014), the dimensionality d plays a critical role. That is, if $d = 1$, then we collapse back to the case where there is no meaningful volatility control. Hence, in their setup, the dimensionality must be $d \geq 2$.

Cvitanić et al. (2014) considers a setup where the agent is only paid at a final deterministic time T and both the principal and the agent have identical CARA utility, and the agent has a quadratic cost in volatility choices. In contrast, our paper considers only risk neutral principal and agent, but allow for intertemporal compensation from the principal to the agent, endogenous termination of the agent, and also private effort and volatility choices by the agent. Moreover, fundamentally, our methods for “hiding” volatility control are fundamentally different in that Cvitanić et al. (2014) relies on a dimensionality argument, whereas this paper relies on reconsidering economically the modeling method of the noise term. It could be an interesting extension for future research to combine both of these approaches.

3 Model Motivation in Discrete-Time

To motivate the continuous-time model of the paper, and also to highlight how existing drift-only control principal-agent models in the literature cannot be used to appropriately model delegation portfolio management problems, we first draw an analogy to a simple discrete-time asset pricing factor models. Suppose in the current period, a group of investors hire a portfolio manager to manage a portfolio that will deliver excess returns $R - r_f$ that is observable to the investor, where r_f is the risk free return. The investors know a priori that the manager has skill so ex-ante the investors are willing to invest into the portfolio manager. What is not known to the investors is whether the manager is exerting sufficient effort to maximize his skills. Suppose further the fund operates like a *mutual fund*, so that the investment holdings of the fund are somewhat transparent, and so the investors know and can observe what is the appropriate market factor, say $R_M - r_f$, to price the portfolio, and moreover, the investors can precisely choose their desired factor loading.⁷ Thus, the portfolio returns are driven by a factor form,

$$R - r_f = \alpha(e) + \beta_M^0 \times (R_M - r_f) + \varepsilon, \quad (3.1)$$

where β_M^0 is the factor loading onto the factor $R_M - r_f$, and ε is the idiosyncratic risk with zero mean, and independent of $R_M - r_f$.⁸ Since β_M^0 had been a priori selected by the principal, there

⁷Indeed, commercial services like Morningstar regularly report the appropriate investment style or factor (i.e. “value”, “growth”, “big cap”, etc.) of the majority of mutual funds available, and they also report a CAPM beta value to the investors.

⁸Another way to view (3.1) is to view the mutual funds types are indexed by their factor loading $\{\beta_{M,j}^0\}_j$. But there the investors can perfectly see the type, and so according to their preferences, select their desired type, say β_M^0 . Once this type has been selected, the investors then proceed to construct contracts to motivate the managers to exert high effort to maximize their skills. But it should be noted that this type selection argument is only valid because the investors know precisely the appropriate market factor is $R_M - r_f$, and hence can compute for themselves the expected risk premium $\mathbb{E}[R_M - r_f]$. If this were unknown, then this argument does not hold.

is no need to condition on this anymore. Thus, conditioning on the effort e the manager will exert, the investors' expected returns from this managed portfolio is,

$$\mathbb{E}[R - r_f | e] = \alpha(e) + \beta_M^0 \mathbb{E}[R_M - r_f], \quad (3.2)$$

where we assumed that effort choice e is independent of the market risk factor R_M . Thus, we see that the expected returns, conditional on effort e is increasing in effort e . Note here that since the investors know the appropriate risk factor $R_M - r_f$, he can also compute the expected risk premium $\mathbb{E}[R_M - r_f]$. That is, the entire term $\beta_M^0 \mathbb{E}[R_M - r_f]$ is common knowledge to both the manager and the investor. Referring to Table 1, the prototypical model in the existing continuous-time principal-agent literature takes the form,

$$dY_t = \mu_t dt + \sigma dB_t, \quad (3.3)$$

where μ_t is a choice that the agent can privately select, and the models in the literature specify μ_t to various degrees of sophistication; see also Section A.1 for further discussion. Mapping (3.3) to the asset pricing model in (3.1), existing drift-only control models can effectively be viewed as,

$$dY_t \approx R - r_f, \quad \mu_t dt \approx \mathbb{E}[R - r_f | e] = \alpha(e) + \beta_M^0 \mathbb{E}[R_M - r_f], \quad \sigma dB_t \approx \beta_M^0 (R_M - r_f) + \varepsilon. \quad (3.4)$$

In particular, we note that for the noise term $\sigma dB_t \approx \beta_M^0 (R_M - r_f) + \varepsilon$, as mentioned earlier, the risk loading β_M^0 is a priori known to the investor, and also the risk factor $R_M - r_f$ is also observable to the investor. Thus, for the overall noise term, the only term that is unobservable to the investor is the idiosyncratic risk ε . And indeed, mapping to the dynamic model, this corroborates well with the notion that the overall noise is simply a Brownian motion dB_t (i.e. continuous-time random walk).

While the viewpoint (3.1), and by extension the continuous-time formulation (3.3) with no volatility control, may be plausible for, say, mutual funds that have fairly transparent investment procedures, this is not the case for numerous other delegated portfolio management practices in the market. Most notably, *hedge funds* and *private equity funds*, unlike mutual funds, are not subject to regulation to reveal their investment positions or trading strategies. And indeed, the investment strategies and positions of these funds are precisely their “secret sauce” or “black box”, of which they are very protective of its details. As such, unlike (3.1), a far more appropriate model here is the form,

$$R - r_f = \alpha(e) + \beta_Z \times (R_Z - r_f), \quad (3.5)$$

where $R_Z - r_f$ is the excess return of an exogenously priced factor that is observable to the manager, but *unknown* to the investor, and β_Z is the factor loading the manager can privately and endogenously control. Note that unlike (3.1), in (3.5) we consider an extreme case and do not write an idiosyncratic term ε . That is because since the factor loading β_Z is endogenously controlled by the manager, in a sense all the idiosyncrasies are due to the manager's decision; and also, this creates a clearer mapping to the continuous-time model to be discussed in Section 4. Concretely speaking, if a hedge fund manager claims that it is a “global macro fund” but does not disclose its actual

positions and trading strategies, there is no way the investor can infer what is the appropriate risk factor to benchmark the fund at. And indeed, even looking at other peer “global macro funds” only give at best a noisy proxy to what the fund in question is actually doing. In particular, that means that unlike the case of mutual funds as per (3.1), the investor cannot a priori view and select the factor loading of the fund. Thus, conditional on the effort e and the factor loading β_Z as chosen by the manager, the expected returns of the managed portfolio are,

$$\mathbb{E}[R - r_f \mid e, \beta_Z] = \alpha(e) + \beta_Z \mathbb{E}[R_Z - r_f], \quad (3.6)$$

where the effort e and factor loading β_Z are independent of the risk factor $R_Z - r_f$. Note that unlike (3.2), the last term $\beta_Z \mathbb{E}[R_Z - r_f]$ is *not* common knowledge to both the investor and the manager. Indeed, even if the investor knew that there is a positive risk premium $\mathbb{E}[R_Z - r_f] > 0$ associated with the risk factor $R_Z - r_f$, the investor still does not know which factor loading β_Z the manager chose.

In all, this means to have a principal-agent model that represents the practices of hedge funds, private equity firms, and other “secret sauce” investment funds, we need at least two additional ingredients, on top of the skill term $\alpha(e)$: (i) exogenous factor term $R_Z - r_f$ observable to the manager but unobservable to the investor; and (ii) endogenous factor loading term β_Z that can be privately controlled by the manager. This represents the starting point of our continuous-time model.

4 Model Outline

Let (Ω, \mathcal{F}, P) be a complete probability space and let $\{B_t\}_{t \geq 0}$ be a standard Brownian motion on this probability space and let $\{\mathcal{F}_t\}_{t \geq 0}$ be the filtration generated by this Brownian motion, suitably augmented. We will write E as the expectation operator under probability measure P . The agent can choose an *action process* $A = \{(e_t, \sigma_t)\}_{t \geq 0}$, where for all times $t \geq 0$, $(e_t, \sigma_t) \in \{e_L, e_H\} \times [\sigma_L, \sigma_H]$, where $e_H > e_L > 0$ and $\sigma_H > \sigma_L > 0$. We will call $\{e_t\}_{t \geq 0}$ the *effort control (process)* and $\{\sigma_t\}_{t \geq 0}$ *volatility control (process)*.⁹¹⁰ Consider a function $\kappa : \{e_L, e_H\} \times [\sigma_L, \sigma_H] \rightarrow [\mu_L, \mu_H]$ that maps both the effort and volatility chosen by the agent to the (expected) return of the (cumulative) *cash flow process* Y ; that is, consider $(e, \sigma) \mapsto \kappa(e, \sigma) = \mu$. We will call κ as the *reward function* and we will discuss further on the assumptions and properties of this function in Section 4.2 below. The cash flow process Y has dynamics that depend on the agent’s action process¹¹,

$$dY_t = \kappa(e_t, \sigma_t)dt + \sigma_t dM_t, \quad Y_0 = y_0 \quad (4.1)$$

$$dM_t = M_t dB_t, \quad M_0 = m_0, \quad (4.2)$$

⁹Of course, strictly speaking in the usual language of stochastic differential equations, we would call $\sigma_t m_t$ as the (stochastic) volatility of Y . However, since in this setup, m_t is an exogenous process, and only σ_t is being directly controlled by the agent, it would be more natural to think and call σ_t as *volatility* that is being managed by the agent.

¹⁰Throughout this article, we will use e to denote the effort parameter / process, and use \mathbf{e} to denote the exponential function.

¹¹Specifically, the action process A is progressively measurable with respect to $\{\mathcal{F}_t\}$.

where $m_0 > 0$ and note we have denoted $\mu_t := \kappa(e_t, \sigma_t)$. Given an action process A , we will call $\{\mu_t\}_{t \geq 0} = \{\kappa(e_t, \sigma_t)\}_{t \geq 0}$ the *drift (process)* of the cash flow Y ¹². The principal cannot observe the agent's action process but can only observe the cash flow Y . The agent can also observe the cash flow Y . Let $\{\mathcal{F}_t^Y\}_{t \geq 0}$ be the (suitably augmented) filtration generated by the cash flow process Y , which represents the principal's information set. The extra term m_t in (4.1) and its dynamics (4.2) is different from the prevailing literature (see Table 1). We shall call $\{M_t\}_{t \geq 0}$ as the *exogenous factor*.

Finally, for illustrative purposes only, Figure 1 plots this cash flow process against some other cash flow processes that have been used in the literature.

Both the principal and the agent are risk neutral. The principal discounts time at rate $r_1 > 0$ and the agent discounts time at rate $r_0 > 0$. As per DeMarzo and Sannikov (2006), we assume that the agent is less patient than investors; so we assume $r_0 > r_1$. The principal needs to compensate the agent and is modeled via the $\{\mathcal{F}_t^Y\}_{t \geq 0}$ -adapted stochastic process $X = \{X_t\}_{t \geq 0}$ and assuming limited liability, we restrict the compensations to be non-negative, so $dX_t \geq 0$. The principal also has the ability to terminate the agent at some \mathcal{F}_t^Y -measurable random time $\tau \in [0, \infty]$. Upon termination, the firm is liquidated for value¹³ $L > 0$ and the agent receives retirement value¹⁴ $R > 0$. A *contract* is the tuple (A, X, τ) , which specifies the recommended action process A , a compensation for the agent X and the termination time τ .

Fix a contract (A, X, τ) and suppose the agent follows the principal's recommended action A . The *agent's payoff* at time $t = 0$,

$$W_0(A) := \mathbb{E}^A \left[\int_0^\tau e^{-r_0 t} \left(dX_t + \left[\phi_e \left(1 - \frac{e_t}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt \right) + e^{-r_0 \tau} R \right], \quad (4.3)$$

where $\phi_e, \phi_\sigma > 0$ are constants known to both the principal and the agent. Here, we denote \mathbb{E}^A as the expectation under the probability measure \mathbb{P}^A induced¹⁵ by the agent's chosen action process A . We will further discuss the properties of the agent's payoffs and incentives in Section 4.3. Likewise, the *principal's payoff* at time $t = 0$ is,¹⁶

$$\mathbb{E}^A \left[\int_0^\tau e^{-r_1 t} (dY_t - dX_t) + e^{-r_1 \tau} L \right] = \mathbb{E}^A \left[\int_0^\tau e^{-r_1 t} \kappa(e_t, \sigma_t) dt - \int_0^\tau e^{-r_1 t} dX_t + e^{-r_1 \tau} L \right]. \quad (4.4)$$

Further discussion of the conflict between the agent and principal is deferred until Section 4.3. We collect some assorted remarks, largely technical in nature, about the model in Section B.

¹² Throughout the paper, we will interchange the notation μ_t and $\kappa(e_t, \sigma_t)$ to denote the drift part of the cash flow process. This is for notational brevity. Thus, given an action process $A = \{(e_t, \sigma_t)\}_{t \geq 0}$, we will also with some abuse of notation, also call and denote $A = \{(\mu_t, \sigma_t)\}_{t \geq 0}$ as the action process, with the understanding that $\mu_t \equiv \kappa(e_t, \sigma_t)$.

¹³We will be more specific about this liquidation value L in Section 5.

¹⁴We will be more specific regarding this retirement value R in Section 6.

¹⁵ See Section B for details.

¹⁶In the second equality, we applied Doob's Optional Stopping Theorem. While τ could have been unbounded (i.e. never terminating the agent) but a standard argument using bounded sequences of stopping times and an usual application of the Dominated Convergence Theorem will also show the result. We will omit these details.

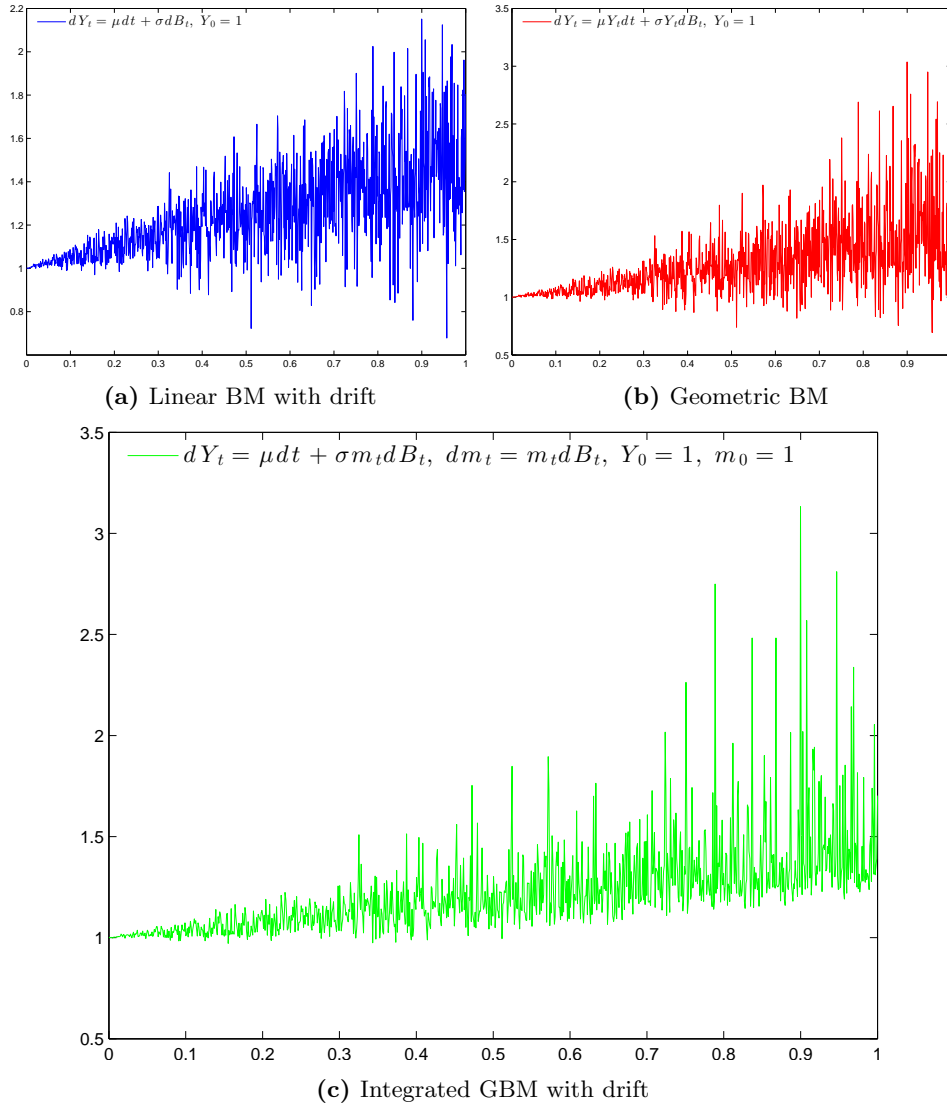


Figure 1: Illustrations of various types of commonly used cash flow processes. Here, the (constant) parameters are chosen to be $\mu = 0.5$, $\sigma = 0.3$, $Y_0 = 1$, $m_0 = 1$, and we simulate over 1000 discrete evenly spaced points over the time interval $t \in [0, 1]$. Subfigure (1a) describes the *linear Brownian motion with drift* cash flow process, $dY_t = \mu dt + \sigma dB_t$, that is used in the models by Holmström and Milgrom (1987), DeMarzo and Sannikov (2006), Sannikov (2008) and several others; see Table 1. Subfigure 1b describes the *geometric Brownian motion* process, $dY_t = \mu Y_t dt + \sigma Y_t dB_t$, that is used by He (2009) (note, He (2009) calls this the *firm value process*). Finally, subfigure (1c) describes an *integrated Geometric Brownian motion with drift*, $dY_t = \mu dt + \sigma dM_t$, $dM_t = M_t dB_t$, as it will be used in this paper. Note also the parameters used in generating this figure are for illustrative purposes only. Unless specified otherwise, these parameters are *not* enforced throughout the paper.

4.1 Mapping back to the discrete-time model

Mapping back to the discrete-time specification in Section 3, and in particular to (3.5), we can map the terms analogously as,

$$\begin{aligned}
 dY_t &\approx R - r_f, \\
 \kappa(e_t, \sigma_t)dt &\approx \mathbb{E}[R - r_f \mid e, \beta_Z] = \alpha(e) + \beta_Z \mathbb{E}[R_Z - r_f], \\
 \sigma_t &\approx \beta_Z, \\
 dM_t &\approx R_Z - r_f.
 \end{aligned}
 \tag{4.5}$$

As discussed, the specification (3.5) and now to the continuous-time extension of (4.1), (4.2), can be viewed as a more appropriate model of delegated portfolio management than existing drift-only control models in the literature.

In particular, in contrast to the mapping in (3.4) where the noise term is mapped like $\sigma dB_t \approx \beta_M^0(R_M - r_f) + \varepsilon$, the mapping in (4.5) for the noise term is mapped like $dM_t \approx R_Z - r_f$. These two forms of mappings represent a fundamental difference in how investors view whether a risk factor is observable or unobservable to them. As explained in Section 3, when the risk factor $R_M - r_f$ is known and observable to the investors, the term $\beta_M^0(R_M - r_f)$ is common knowledge to both the manager and the investor. Thus, the only source of uncertainty to the managed cash flows in that case is the idiosyncratic risk ε , of which its appropriate continuous-time counterpart is a Brownian motion σdB_t , which is essentially a pure noise random walk in continuous-time. In contrast, in the setting where the investors do not know and cannot observe the risk factor $R_Z - r_f$, then the uncertainty is more “informative”. Firstly, the risk loading β_Z is no longer common knowledge. Secondly, the risk factor $R_Z - r_f$, in continuous-time, cannot be appropriately viewed as a random walk. Indeed, in accordance to the tradition of classic asset pricing theory, assets have returns that are log-normally distributed; in particular, the Black-Scholes economy risky asset also models the returns of a risky asset as a geometric Brownian motion. Here, in setting $dM_t \approx R_Z - r_f$, we have simply modeled M as a geometric Brownian motion with zero drift and unit variance; it is straightforward, but algebraically cumbersome, to put in a non-zero drift and non-unit variance.

Remark 4.1. Although slightly beyond the original motivation scope of the paper, it should be noted that in modeling the managed cash flows dY_t , we have that the expected value $\kappa(e_t, \sigma_t)dt$ is dependent on both effort and volatility, and this modeling form has found precedence in the recent empirical and theoretical asset pricing literature. In Buraschi, Kosowski, and Sritrakul (2013), the authors note:

“... [The traditional alpha measure that is independent of beta] raises the question of how well a reduced-form alpha measures the true managerial skill of a hedge fund manager. An answer to this question depends on the determinants of the optimal allocation θ_{At}^* made by that manager. If the optimal allocation is constant and determined exclusively by the risk and return characteristics of the investment opportunity set (as in a traditional Merton model without agency distortions), then reduced-form alpha is an unbiased estimate of managerial skill. However, if the optimal allocation is influenced by

nonlinear agency contracts, then reduced-form alpha is a misspecified estimate of true skill. For instance, a high reduced-form alpha could be the fortunate result of too much leverage as managers aim to maximize their incentive options. Of course, high leverage increases not only the manager’s expected return (because of the call option) but also the likelihood of large negative returns.”

Thus, Buraschi et al. (2013) suggests that the managed portfolio alpha, under the influence of “nonlinear agency contracts”, could depend on leverage and investment opportunities.

In addition, although absent of any agency considerations, Frazzini and Pedersen (2014) also considers an overlapping generations model that implies a factor model structure for risky asset returns, and show that the alpha term could depend explicitly on the factor loading. Also in empirical research of hedge fund performance, Bollen and Whaley (2009) also notes:

“Accurate appraisal of hedge fund performance must recognize the freedom with which managers shift asset classes, strategies, and leverage in response to changing market conditions and arbitrage opportunities. The standard measure of performance is the abnormal return defined by a hedge fund’s exposure to risk factors. If exposures are assumed constant when, in fact, they vary through time, estimated abnormal returns may be incorrect.”

4.2 Reward Function κ

We need to be more specific about the way the agent can control the drift and volatility of the cash flows. It should be noted in drift-control only models, the specification of the drift is usually quite simple (i.e. linear). But in our case, given the volatility control, we must be more careful in modeling and giving economic meaning to link the volatility and drift controls. Note that one possible characterization is to have the agent control drifts and volatilities that are completely unrelated to each other. But this case is not economically meaningful since it destroys the traditional link of risk-return trade-offs of financial economics, particularly that of portfolio choice theory. We will now more specifically define the reward function κ as follows.

Definition 4.1. A strictly positive real valued function $\kappa : \{e_L, e_H\} \times [\sigma_L, \sigma_H] \rightarrow [\mu_L, \mu_H]$, $(e, \sigma) \mapsto \kappa(e, \sigma) = \mu$, that is twice-continuously differentiable in the second argument, is called a *reward (drift) function* if it that has the following properties:

- (a) *Higher effort, higher reward:* $\kappa(e_H, \sigma) > \kappa(e_L, \sigma)$, for all σ ;
- (b) *Higher risk, higher reward but at decreasing rate:* $\kappa_\sigma(e, \sigma) > 0$ and $\kappa_{\sigma\sigma}(e, \sigma) < 0$, for all (e, σ) ;
- (c) *Risk cannot substitute for effort:* $\kappa(e_H, \sigma) > \kappa(e_L, \sigma')$ for all σ, σ' .
- (d) *Range of effort greater than range of risk:* $\frac{\phi_e}{e_H}(e_H - e_L) > \frac{\phi_\sigma}{\sigma_L}\sigma_H$, and $\kappa(e, \sigma) > \phi_e$ for all (e, σ) .

The requirement (a) is natural to interpret; that is, if the agent exerts higher effort to running the project, then the expected payoff should be higher, regardless of the choice of volatility. Requirement (b) is the traditional risk-reward type trade off. One would expect that by choosing a riskier project (higher volatility) over a safer project (lower volatility), it is so that one could enjoy higher expected returns, but we impose that the rate of return from increasing risk is decreasing. Thus, requirements (a) and (b) should have good natural interpretations. Requirement (c) means that exerting high effort always gives a higher return, regardless of the level of risk taken. Effectively, that means that effort and risk are not “substitute goods”; hence, this requirement explicitly rules out a case where the agent can exert low effort and take on a high level of risk such that this return is equal or greater to one with high effort and any level of risk. Note that clearly (c) implies (a) but we write them out separately as (a) is effectively the only requirement assumed in the controlled drift-only models (i.e. when κ is a function only of effort e).¹⁷ More generically, early studies between project selection (viewed as volatility in the current context), risk and effort can be found in Lambert (1986) and Hirshleifer and Suh (1992).

We now give an example that satisfies Definition 4.1 and hence directly showing that the set of reward functions is nonempty.

Example 4.1. Consider the reward function of the form,

$$\kappa(e, \sigma) = \alpha_1 e^{\alpha_0(e - e_L)} \log \sigma,$$

where $\alpha_0, \alpha_1 > 0$ are deterministic constants. Here, we restrict σ_L, σ_H such that $\sigma_L = c$, $\sigma_H \approx 1.763$, where $1 < c < \sigma_H$, and that $\phi_e, \phi_\sigma > 0$ are such that $\frac{\phi_e}{\phi_\sigma} \frac{1}{e_H} (e_H - e_L) > \frac{\sigma_H}{\sigma_L}$. Note here that $\mu_L = \kappa(e_L, \sigma_L)$ and $\mu_H = \kappa(e_H, \sigma_H)$. See Figure 2 for an illustration.

Proof. See the Appendix for proof of this Example and all subsequent proofs throughout this article. \square

4.3 Principal and Agent Conflict

With the reward function specified in Definition 4.1, we are now ready to discuss the sources of conflicts between the principal and the agent. From the agent’s payoff in (4.3), we see that the agent dislikes exerting high effort $e_t = e_H$ and likes to exert low effort $e_t = e_L$, and the agent likes to choose high volatility $\sigma_t = \sigma_H$ and dislikes to choose low volatility $\sigma_t = \sigma_L$. In contrast, from the principal’s payoff in (4.4), and the properties of the reward function as given in Definition 4.1, the principal likes high effort $e_t = e_H$ and dislikes low effort $e_t = e_L$. Moreover, by the properties of the reward function, and also effectively by the risk neutrality of the principal, the principal also seems to like high volatility σ_t . The assumption that the agent likes to job shirk while the principal does not is common in the principal-agent literature.

¹⁷Strictly speaking (d) is not an requirement of the reward function $\kappa(e, \sigma)$ but rather an assumption on the parameters $\phi_e, \phi_\sigma, e_H, e_L, \sigma_H, \sigma_L$, but we will collect this here for subsequent convenient reference.




Figure 2: Illustration of Example 4.1 with the parameters: $\alpha_0 = 0.3, \alpha_1 = 2, e_L = 3, e_H = 5, \sigma_L = 1.01, \sigma_H = 1.763$.

However, in this context, the specification of volatility warrants more discussion. It seems like since both the agent and the principal prefers higher level of volatility, then volatility is not a source of moral hazard conflict between the principal and the agent. However, this is not entirely correct. As it is standard in the principal-agent literature, to incentivize the agent, the principal must put the agent's payoff at risk, and specifically meaning the agent's payoff must be sensitive to the agent's managed cash flows. However, such sensitivity here is also further affected both by the agent's volatility choice σ_t and also the exogenous factor M_t . Hence, even though both the principal and the agent prefers the volatility choices in the same direction, but since volatility choice also affects the overall uncertainty in this economy, this uncertainty indirectly causes the conflict between the agent and the principal. We will have more to say about this important feature of volatility choice in Section 8.2.

5 First Best

Let's begin by characterizing the first best result. At this point, we should further impose a restriction on the recovery value L of the firm upon termination. In line with the literature, we will assume that termination is inefficient so that never terminating $\tau = +\infty$ is indeed optimal in the first best case.

Assumption 5.1. *We assume that termination is inefficient. That is, the salvage value of the firm L is such that,*

$$0 < L < \int_0^{\infty} e^{-r_1 t} \kappa(e_L, \sigma_L) dt = \frac{\kappa(e_L, \sigma_L)}{r_1}. \quad (5.1)$$

Recalling Definition 4.1, the right-hand side of (5.1) is precisely the “worst case” indefinite payoff scenario for the principal.

Suppose the principal can and will operate the firm without the agent. In this case, the principal does not need to pay any compensation, so $X \equiv 0$, nor is there any need for termination, so $\tau \equiv +\infty$. Recalling (4.4), the principal has the optimization problem,

$$b_0^{FB} := \sup_{e, \sigma} \mathbb{E} \left[\int_0^\infty e^{-r_1 t} \kappa(e_t, \sigma_t) dt \right] = \sup_{e, \sigma} \int_0^\infty e^{-r_1 t} \kappa(e_t, \sigma_t) dt \quad (5.2)$$

Proposition 5.2. *Suppose there are no agency conflicts so the principal does not need to hire the agent to run the firm. Then the principal will pay zero compensation, $X \equiv 0$, and never terminate, $\tau \equiv +\infty$. The principal will always exert high effort at all times, so $e_t \equiv e_H$ for all t , and always choose high volatility $\sigma_t = \sigma_H$ for all times. The first best value of the firm b_0^{FB} at time $t = 0$ is,*

$$b_0^{FB} = \int_0^\infty e^{-r_1 t} \kappa(e_H, \sigma_H) dt = \frac{\kappa(e_H, \sigma_H)}{r_1}. \quad (5.3)$$

The first best value of the firm is deterministic and stationary. That is, the time $t = 0$ value b_0^{FB} does not depend on any state variable, and this is economically intuitive. Given that the principal only derives payoff from the reward function $\kappa(e_t, \sigma_t)$, there are no exogenous state variables (namely, say the exogenous factor M) involved.

6 Continuation value and Incentive compatible contracts

Now we proceed to the main focus of the paper. As it is standard in the literature, following the arguments like DeMarzo and Sannikov (2006) and Sannikov (2008), we consider the agent’s continuation value as a state variable to capture the dynamic incentive compatibility constraints. However, again because of the richer volatility setup of (4.1) than the ones in the current literature, we must take more care in deriving the results.

Firstly, we make a trivial substitution that will substantially simplify the problem. Throughout this section, let’s fix an arbitrary contract (A, X, τ) . In particular, note that the action process¹⁸ has the form, $A = \{(e_t, \sigma_t)\}_{t \geq 0}$. As noted in footnote 12, defining $\mu_t \equiv \kappa(e_t, \sigma_t)$, we will also call $A = \{(\mu_t, \sigma_t)\}_{t \geq 0}$ as the action process. Define, the agent’s time t continuation value (or promised value),

$$W_t(A) := \mathbb{E}^A \left[\int_t^\tau e^{-r_0(s-t)} \left(dX_s + \left[\phi_e \left(1 - \frac{e_s}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_s}{\sigma_L} - 1 \right) \right] ds \right) + e^{-r_0(\tau-t)} R \mid \mathcal{F}_t^Y \right]. \quad (6.1)$$

Note here on the left-hand side of (6.1), we have suppressed the notation for the dependence on the payment X and termination time τ , but retained the notation emphasis on the action process A .

¹⁸ We should be clear on the word “fixed” action process here. Although the agent *chooses* an action process that just needs to be $\{\mathcal{F}_t\}$ -adapted, but when the principal fixes a recommended action A to the agent, this recommended action process A is known to the principal and hence A is also $\{\mathcal{F}_t^Y\}$ -adapted. That is, the recommended action must be known to the principal but any general deviation away from the recommended action by agent is not known to the principal.

6.1 Incentive compatible contracts

Definition 6.1. A contract (A, X, τ) at time 0 with expected agent payoff $W_0(A)$ is *incentive compatible* if

(1) (a) $W_t(A) \geq R$ for all times $t \leq \tau$; , where the retirement value $R > 0$ is such that,

$$R > \frac{1}{r_0} \phi_\sigma \left(\frac{\sigma_H}{\sigma_L} - 1 \right); \quad (6.2)$$

(b) $M_t \geq \underline{m}$, for all times $t \leq \tau$, where $\underline{m} > 0$; and

(2) $W_0(A) \geq W_0(A^\dagger)$, for all other action processes A^\dagger .

The *optimal contracting problem* is to find an incentive-compatible contract that maximizes the principal's time 0 expected payoff.

Requirement (1) of Definition 6.1 is also more aptly called agent's *individual participation* (IR) constraints. In particular, (1a) says that the agent's continuation value W_t must be at all times greater than or equal to the agent's reservation value R . This is a standard definition of the IR constraint in the literature. The addition of (1b) warrants slightly more discussion as this is not standard in the literature. Requirement (1b) effectively requires when the agent manages the cash flows dY_t , the agent will only manage it only when the exogenous factor level M_t at any point in time t is not too low, and in particular it must be greater than this lower bound \underline{m} . See further detailed discussions in Remark 6.1, where we also discuss the justification of the retirement value R in (6.2). Finally, requirement (2) is the usual incentive compatibility condition in the literature.

Remark 6.1 (Justification for Definition 6.1(1b) and retirement value R of (6.2)). Let's discuss the economic justification of (1b). If the exogenous factor level M , which again is a geometric Brownian motion so $M > 0$, is too low, say when $M_t \approx 0$ (even though $M_t = 0$ happens on a set of measure zero), then all sources of uncertainty in this economy vanishes. Indeed, suppose in the extreme that we indeed have $M \equiv 0$. And when that happens, the managed cash flows thus become $dY_t = \kappa(e_t, \sigma_t)dt$, without any additional noise term. But this implies the principal, upon observing cash flows dY_t continuously over time, can precisely detect the choice of effort e_t and choice of volatility σ_t that the agent has chosen¹⁹, and clearly then, the principal would instruct the agent to choose the first best effort and volatility choices. However, first best effort and volatility choices are clearly not beneficial for the agent.

Without the presence of uncertainty (so when $M \equiv 0$), the principal no longer needs to compensate the agent, $X \equiv 0$. Mapping back to the context of delegated portfolio management and recalling the discussion in Section 3, when $M \equiv 0$, it is equivalent to saying the outside investor is getting precisely zero premia for the factor exposure of this particular managed fund. In that case, the investor has no particular reason to compensate the manager for management anymore. Indeed, in this case, if the manager does not choose the first best case of highest effort ($e_t \equiv e_H$) and choose

¹⁹This is possible since the reward function κ of Definition 4.1 is bijective and non-crossing in effort e and volatility σ .

the appropriate investment opportunity ($\sigma_t \equiv \sigma_H$), the investor will simply walk away. Anticipating this, the agent is conceivably better off to “walk away” from managing the project before the exogenous factor level M is too low, namely at \underline{m} , and still manage to extract some information rent from the principal. More precisely, the agent’s retirement value R is such that,

$$\begin{aligned}
& \text{Payoff to agent with positive exogenous factor } M \geq \underline{m} > 0, \text{ and arbitrary actions} \\
& \underbrace{\int_0^\tau e^{-r_0 t} dX_t + \int_0^\tau e^{-r_1 \tau} \left[\phi_e \left(1 - \frac{e_s}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_s}{\sigma_L} - 1 \right) \right] dt}_{\geq R} \\
& > \underbrace{\int_0^\infty e^{-r_0 t} \phi_\sigma \left(\frac{\sigma_H}{\sigma_L} - 1 \right) dt}_{\text{Payoff to agent with always zero exogenous factor } M \equiv 0, \text{ and first best actions}} = \frac{1}{r_0} \phi_\sigma \left(\frac{\sigma_H}{\sigma_L} - 1 \right).
\end{aligned}$$

This discussion hence also justifies the retirement value R as specified in (6.2).

Remark 6.2 (Retirement value too low). With Remark 6.1 in mind, we should also consider the counter case. What if the retirement value R is too low? Specifically, what if, unlike (6.2), R is such that,

$$\frac{1}{r_0} \phi_\sigma \left(\frac{\sigma_H}{\sigma_L} - 1 \right) \geq R > 0. \tag{6.3}$$

If the retirement value R is such that (6.3) holds, when the exogenous factor level is identically zero $M \equiv 0$, we will see that it may be *not* optimal for the agent to walk away from the contract. That is, if $M \equiv 0$, again as per the argument in Remark 6.1, the principal can credibly instruct the agent to take on the first best action $e_t \equiv e_H$ and $\sigma_t \equiv \sigma_H$, and pay zero compensation $X \equiv 0$. That is, the agent simply then receives the instantaneous private benefit of $\phi_\sigma \left(\frac{\sigma_H}{\sigma_L} - 1 \right) dt$. Even though the agent knows that there are some positive information rent to extract from the principal if the exogenous factor level is strictly positive, but even at the identically zero exogenous factor level case, his instantaneous private benefit still exceeds the outside retirement value R . So if (6.3) holds, it implies there is a possibility that the principal can give the agent zero compensation and yet the agent will still happily remain employed.

We rule this case out. Specifically, we assume that at time $t = 0$, the agent can anticipate such effects, and negotiate, ex-ante, with the outside labor market to secure a sufficiently high retirement value R that satisfies (6.2), rather than a low retirement value of (6.3). In the context of delegated portfolio management, we may think of a high retirement value R to represent an outside fund management opportunity that’s available to the portfolio manager.

6.2 Continuation value dynamics

The dynamics of the agent’s continuation value is given as follows.

Theorem 6.3. *Fix a contract (A, X, τ) . Then for $t \in (0, \tau)$, the agent’s continuation value $W_t(A)$*

of (6.1) has dynamics,

$$\begin{aligned} dW_t(A) = & r_0 W_t(A) dt - \left(dX_t + \left[\phi_e \left(1 - \frac{e_t}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt \right) \\ & + \beta_t (dY_t - \mu_t dt) + d\epsilon_t^{\perp, A}, \end{aligned} \quad (6.4)$$

where $\epsilon_t^{\perp, A} := \int_0^t e^{r_0 s} dV_s^{\perp, A}$, $\mu_t := \kappa(e_t, \sigma_t)$, and where β_t and $V_t^{\perp, A}$ are given in Proposition D.3.

Let's discuss economic meaning of the dynamics of the agent's continuation value as characterized in (6.4) of Theorem 6.3. Here, β_t represents the sensitivity of the agent's continuation value to output dY_t . When the agent takes the recommended action process A as given in the contract, the term $dY_t - \mu_t dt = dY_t - \kappa(e_t, \sigma_t) dt = \sigma_t M_t dB_t$ is a mean-zero noise term. The term $d\epsilon_t^{\perp, A}$ (explained in more detail below) also has mean zero. Economically and intuitively (though mathematically incorrect), we can view (6.4) in this alternative way:

$$\mathbb{E}_t[r_0 W_t(A) dt] \approx \mathbb{E}_t[dW_t(A)] + \mathbb{E}_t \left[dX_t + \left[\phi_e \left(1 - \frac{e_t}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt \right] \quad (6.5)$$

Hence, viewed in this way, we can think of the expected growth of the agent's continuation value $\mathbb{E}_t[r_0 W_t(A) dt]$, when the agent follows the recommended action process A , can be decomposed into the expected change from the previous continuation value $\mathbb{E}_t[dW_t(A)]$, plus the expected compensation from the principal $\mathbb{E}_t[dX_t]$, and plus the expected benefits from taking not the highest effort ($e_t \neq e_H$) and not the lowest volatility ($\sigma_t \neq \sigma_L$), which yields a strictly positive value $\mathbb{E}_t \left[\left[\phi_e \left(1 - \frac{e_t}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt \right]$. Note that if the agent were to really take the highest effort level ($e_t = e_H$) and the lowest volatility ($\sigma_t = \sigma_L$), then the agent's private benefits vanishes.

The economic interpretation of (6.4) in terms of a logic like (6.5) is similar across models with only drift control, say for instance, DeMarzo and Sannikov (2006) and Sannikov (2008). However, the economic interpretations of the two noise terms $\beta_t (dY_t - \mu_t dt)$ and $d\epsilon_t^{\perp, A}$ warrant more discussion. Firstly, observe that we have two noise terms here, rather than one, as in essentially all the papers with drift-only control. Secondly, while the interpretation of β_t as the sensitivity of the agent's continuation value to output here is still in line with the existing models, the noise term $dY_t - \mu_t dt$ (being multiplied by the sensitivity β_t) is different. Note also (as we will see in subsequent development) β_t still retains the interpretation as the minimal amount of risk the principal wants to subject and incentivize the agent, as in line with the literature. However, since the agent can control the volatility σ_t , and if we read the diffusion term of the agent's continuation value process dW_t as $\beta_t (dY_t - \mu_t dt) = \beta_t \sigma_t M_t dB_t$, then we see that even if the principal can dictate the sensitivity β_t for the agent, the agent still has the ability to "counteract" this dictation by choosing a volatility level σ_t to shift the overall diffusion term $\beta_t \sigma_t M_t$ (recall M is exogenous). Thus, we can already see that in a model where the agent can control volatility σ_t , the principal's tools to incentivize the agent may be weakened, as compared to a model where the agent can only control the drift. This effect is distinctly not present in models without volatility control. Finally, the additional term $d\epsilon_t^{\perp, A}$ is also related to the fact that in this model the agent can control volatility. Recalling that the cash flows are of the form $dY_t = \mu_t dt + \sigma_t M_t dB_t$. If there were no volatility control, so the cash flow

takes on the form $dY_t = \mu_t dt + M_t dB_t$, then uncertainty (as seen by the diffusion term M_t) cannot be dictated by the agent. However, in this current case, the diffusion term in the cash flow is $\sigma_t M_t$, meaning that the agent can actually endogenously change the uncertainty of the cash flows, and as seen from the discussions with regards to the quadratic variation (see Section A), this change of uncertainty cannot be detected by the principal if the agent does not follow an incentive compatible action process. Hence, that is why the term $d\epsilon_t^{\perp, A}$ is there to capture this source of extra (orthogonal) uncertainty (see Proposition D.3). Note that, as shown in Lemma 7.3, when we consider incentive compatible contracts, this term $d\epsilon_t^{\perp, A}$ will become identically zero. The economic intuition is simply that when the principal offers incentive compatible contracts, as opposed to any arbitrary contracts to the agent, the principal knows that the agent will have no incentive deviate from the principal's recommendations. In particular, this also implies the principal can see the instantaneous diffusion of the cash flows $\sigma_t M_t$ and hence there is no source of extra uncertainty that we'd described earlier.

6.3 Incentive compatibility conditions

Now, the following is a necessary and sufficient condition to characterize the incentive compatible contracts in this context.

Lemma 6.4. *Fix a contract (A, X, τ) and consider the process β as given in Proposition D.3. Then we have the following equivalence.*

(i) *The action $A = \{(e_t, \sigma_t)\}_{t \geq 0}$ is such that,*

$$0 \geq -\frac{\phi_e}{e_H}(e' - e_t) + \frac{\phi_\sigma}{\sigma_L}(\sigma' - \sigma_t) + \beta_t(\kappa(e', \sigma') - \kappa(e_t, \sigma_t)) \quad (6.6)$$

for all $(e', \sigma') \in \{e_L, e_H\} \times [\sigma_L, \sigma_H]$.

(ii) *Contract (A, X, τ) is incentive compatible.*

The following corollary is a simple rewriting of Lemma 6.4 but will be useful for the subsequent discussion.

Corollary 6.5. *Under the same setup of Lemma 6.4, if $\beta = \{\beta_t\}$ is a nonnegative process, then a given action process $A = \{(e_t, \sigma_t)\}$ is incentive compatible if and only if for all times t :*

(i) *If $e_t = e_H$,*

$$\beta_t \geq \frac{1}{\kappa(e_H, \sigma_t) - \kappa(e_L, \sigma_H)} \left[\frac{\phi_e}{e_H}(e_H - e_L) + \frac{\phi_\sigma}{\sigma_H}(\sigma_H - \sigma_t) \right]. \quad (6.7)$$

(ii) *If $e_t = e_L$,*

$$0 \leq \beta_t \leq \frac{1}{\kappa(e_H, \sigma_H) - \kappa(e_L, \sigma_t)} \left[\frac{\phi_e}{e_H}(e_H - e_L) + \frac{\phi_\sigma}{\sigma_L}(\sigma_t - \sigma_H) \right] \quad (6.8)$$

Remark 6.6. We should note that in Corollary 6.5, the right hand side of (6.7) is strictly positive for all choices of σ_t . Also, recalling Definition 4.1, the right hand side of the second inequality of (6.8) is also strictly positive for all choices of σ_t .

Remark 6.7. At this point, we can make a direct comparison to the case when only the drift, but not the volatility, is under the agent’s control. That drift only control case has been considered in DeMarzo and Sannikov (2006) and He (2009) but given our current linear cost form, a more direct comparison is with He (2009). It should be noted that in He (2009), the agent manages a geometric Brownian motion (which He (2009) regards as firm value, rather than cash flow). Nonetheless, consider He (2009, Proposition 1) and they derive the analogous necessary and sufficient condition to be,

$$\beta_t \geq \phi_\mu \frac{\sigma}{\mu_H}. \quad (6.9)$$

Note in (6.9), the multiplicative factor by σ is to reflect the fact that the agent managed process in He (2009) is a geometric Brownian motion (with managed drift μ_t and unmanaged constant volatility σ), rather than our linear setup.

What is most striking about the characterization in (6.9) and Lemma 6.4(i) is that on the right-hand side of (6.9), there are no other agent choice variables involved; indeed, the entire incentive-compatible contract is characterized by this single — perhaps rather “static” — inequality. In contrast, on the right-hand side of the inequality in Lemma 6.4(i), there still remains a choice variable by the agent; indeed, this type of characterization is very similar to the one that is provided in Sannikov (2008, Appendix A, Proposition 2), even though in that problem, there is still no volatility control.

Remark 6.8. Despite the addition of volatility control, and in particular that volatility σ_t is chosen from an interval $[\sigma_L, \sigma_H]$ in our model, it might seem surprising that incentive compatibility can still be completely be characterized by *two* inequalities, (6.7) and (6.8) of Corollary 6.5, much alike binary hidden effort or drift choice models of DeMarzo and Sannikov (2006) and He (2009). Economically, it is because the volatility choice σ_t here is not a *direct* source of moral hazard conflict. That is, both the principal and the agent prefer the *same direction* of volatility, even though they may disagree on the level. Hence, the principal need not be concerned with providing direct incentives by altering the sensitivity β_t , and hence the optimal choice of sensitivity β_t should just focus on providing incentives to motivate the correct effort level, and since there are just two effort choices here, this corresponds to the two inequalities. As mentioned earlier, the volatility choice is an *indirect* source of moral hazard conflict. In particular, even though the principal and the agent may agree on the general direction of volatility choices²⁰, the fact that the agent can directly alter the uncertainty of this economy implies the agent’s volatility choice complicates the principal’s task of providing incentives to the agent.

Remark 6.9. At this point, one might step back and ponder about this question: The principal here can only observe a one-dimensional managed cash flow Y , but why is it that it can provide incentives to induce the agent to make the appropriate choices for a two-dimensional moral hazard term (e, σ) , that being effort and volatility choices? The essential explanation lies in the monotonicity of the reward function $\kappa(e, \sigma)$ in both arguments and also the way that the volatility term σ enters linearly into the diffusion term $\sigma_t dM_t$ of the managed cash flows dY_t . For instance, if we had considered an

²⁰A far more difficult characterization happens when the principal and the agent *disagree* on their preferences of the volatility level. This is left for future research.

alternative reward function form, say like $\kappa(e, \sigma m)$, for $M_t = m$, and that the diffusion term is more complicated, like in the form $\sigma_t Y_t dM_t$, then we can see that the above argument will not hold.

7 Principal's Problem

Once the incentive compatible contracts have been characterized as in Lemma 6.4, we are now ready to consider the principal's problem.

7.1 Strengthening the IC condition

If we take the necessary and sufficient IC condition as characterized by Corollary 6.5, it will be difficult to ensure that the resulting principal's value function will be concave in the agent's continuation value w . Hence, we will strengthen the IC condition and consider a sufficient IC condition for Corollary 6.5, and also we will restrict the set of sensitivities to be bounded above.

Assumption 7.1. *Suppose we restrict the set of sensitivities to be,*

$$\mathcal{B} := \{\beta : K \geq \beta \geq \underline{\beta}\}, \quad (7.1)$$

for some sufficiently large $K > 0$, and where we define,

$$\underline{\beta} := \frac{1}{\kappa(e_H, \sigma_L) - \kappa(e_L, \sigma_H)} \left[\frac{\phi_e}{e_H} (e_H - e_L) + \frac{\phi_\sigma}{\sigma_L} (\sigma_H - \sigma_L) \right]. \quad (7.2)$$

Remark 7.2. Note that since $\sigma \mapsto \frac{1}{\kappa(e_H, \sigma) - \kappa(e_L, \sigma_H)} \left[\frac{\phi_e}{e_H} (e_H - e_L) + \frac{\phi_\sigma}{\sigma_L} (\sigma_H - \sigma) \right]$ is monotonically decreasing, it is clear that (7.1) is a sufficient condition that satisfies the IC condition as characterized in Corollary 6.5. Furthermore, also note that $\underline{\beta} > 0$. As well, we impose an upper bound K on \mathcal{B} to ensure that the set \mathcal{B} is compact. If the set \mathcal{B} is not upper bounded, and in particular not compact, then it is conjectured that most of the arguments henceforth will still go through but one might need more sophisticated proof techniques.

7.2 Principal's optimization problem

Henceforth, we will restrict our attention to incentive compatible contracts. And when we write the probability measure \mathbb{P} and expectation \mathbb{E} and other processes where there is dependence on the action process A , we will denote them without the superscript A notation. The following result significantly simplifies the principal's optimization problem.

Lemma 7.3. *Fix an incentive compatible contract (A, X, τ) . Then under the recommended action A , we have that,*

(i) $\{\mathcal{F}_t^Y\}_{t \geq 0} = \{\mathcal{F}_t\}_{t \geq 0}$, where $\{\mathcal{F}_t\}_{t \geq 0}$ is the natural filtration generated by Brownian motion B .

(ii) $d\varepsilon_t^\perp \equiv 0$, \mathbb{P} -a.s.

Thus, with Lemma 7.3(ii) in hand, we are now ready to consider the principal's optimization problem. For the remainder of the discussion, we will only consider the case when the principal wants to induce the agent to choose high effort $e_t = e_H$ at all times t , but the principal still needs to induce the agent to optimally choose the volatility level σ_t . In binary effort models like DeMarzo and Sannikov (2006) and He (2009), the authors also look for an always high effort implementation of the optimal contract. Zhu (2013) considers the model of DeMarzo and Sannikov (2006) and shows that it is possible to induce the agent to choose high effort and switch to low effort shirking at times. We acknowledge this possibility that inducing the agent to shirk could yield a potentially higher payoff for the principal, but for this paper we will only look for an always high effort equilibrium. It should be noted that we do not place such a restriction on the volatility choice; for instance, we do not insist on looking for an equilibrium where the principal induces the agent to always choose high volatility. That is because in our model, as discussed earlier, there is no *direct* moral hazard conflict between the principal and agent's desired direction of volatility. However, there remains an *indirect* moral hazard conflict arising due to volatility choice as the agent can effectively alter the level of uncertainty directly in this economy and thereby making it more difficult for the principal to provide incentives.

Recalling the principal's time $t = 0$ payoff form in (4.4), the *principal's optimization problem*, when the principal desires to induce always high effort $e_t = e_H$, is thus,

$$\tilde{v}(w, m) := \sup_{\sigma, X, \beta, \tau} \mathbb{E} \left[\int_0^\tau e^{-r_1 t} \kappa(e_H, \sigma_t) dt - \int_0^\tau e^{-r_1 t} dX_t + e^{-r_1 \tau} L \right], \quad (\text{P}')$$

subject to state value dynamics,

$$\begin{cases} dW_t = \left[r_0 W_t dt - \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt - dX_t + \beta_t \sigma_t M_t dB_t, & W_0 = w, \\ dM_t = M_t dB_t, & M_0 = m, \end{cases} \quad (\text{S})$$

7.3 Optimal termination time

In (P') the principal maximizes over the set of effort control processes $e = \{e_t\}$ where $e_t \in \{e_L, e_H\}$, volatility control processes $\sigma = \{\sigma_t\}$ where $\sigma_t \in [\sigma_L, \sigma_H]$, compensation processes $X = \{X_t\}$ which is cadlag and nondecreasing, the sensitivity process $\beta = \{\beta_t\}$, where $\beta_t \in \mathcal{B}$, and the termination stopping time τ ; of course, all of the above must be $\{\mathcal{F}_t\}$ -adapted.

At this point, we will note the following. Let's consider the *relaxed principal's optimization problem*, of which we simply remove maximizing over τ in (P'). That is to say, consider,

$$\begin{cases} v(W_0, M_0) := \sup_{e, \sigma, X, \beta} \mathbb{E} \left[\int_0^\tau e^{-r_1 t} \kappa(e_t, \sigma_t) dt - \int_0^\tau e^{-r_1 t} dX_t + e^{-r_1 \tau} L \right], \\ \tau := \inf\{t \geq 0 : W_t \leq R = 0 \text{ or } M_t \leq \underline{m}\}, \end{cases} \quad (\text{P})$$

subject to the state dynamics (S).

Using an argument similar to Cvitanic and Zhang (2012, Chapter 7, Lemma 7.3.2), and also in

accordance to the intuition that the principal would want to hire the agent as long as the agent is getting paid at least his outside option of $R = 0$ (i.e. individual participation constraint), we can show that the problem of (P) subject to (S), and the problem of (P') subject to (S), are equivalent.

7.4 Heuristic HJB

Considering problem (P) subject to (S), this is a stochastic optimal control problem with continuous controls (i.e. the volatility recommendation σ , and the sensitivity β) and singular controls (i.e. compensation process X). Hence, standard results in the optimal control literature suggests the value function v is a solution to the Hamilton-Bellman-Jacobi (HJB) equation,

$$\max \left\{ -r_1 \psi(w, m) + \max_{\sigma} \sup_{\beta} [(\mathcal{L}_{e_H} \psi)(w, m; \sigma, \beta) + \kappa(e_H, \sigma)] - \psi_w(w, m) - 1 \right\} = 0. \quad (7.3)$$

And here, \mathcal{L}_{e_H} is the second order differential operator,

$$\begin{aligned} (\mathcal{L}_{e_H} \xi)(w, m, ; \sigma, \beta) := & \left[r_0 w - \phi_{\sigma} \left(\frac{\sigma}{\sigma_L} - 1 \right) \right] \xi_w(w, m) + \frac{1}{2} m^2 \xi_{mm}(w, m) \\ & + \beta \sigma m^2 \xi_{wm}(w, m) + \frac{1}{2} \beta^2 \sigma^2 m^2 \xi_{ww}(w, m), \end{aligned} \quad (7.4)$$

and where we maximize over $\sigma \in [\sigma_L, \sigma_H]$ and $\beta \in \mathcal{B}$. Also, where not specified, when we write \max_{σ} and \sup_{β} , for notational brevity, it is understood that we are maximizing over $\sigma \in [\sigma_L, \sigma_H]$ and $\beta \in \mathcal{B}$. For convenience, we will also denote the set of admissible controls at initial state (w, m) as $\mathcal{A}_{w,m}$, with a typical control element denoted as $\alpha = (\sigma, X, \beta)$.

Let us denote the state space for the agent's continuation value as $\Gamma_W := (R, \infty)$, and the state space for the exogenous factor as $\Gamma_M := (\underline{m}, \infty)$, and the overall state space be $\Gamma := \Gamma_W \times \Gamma_M$. The appropriate boundary conditions of this problem are:

$$v(w, m) = L, \quad \text{for } (w, m) \in \partial\Gamma. \quad (7.5)$$

7.5 Key illustrations of the value function

Detailed properties of the value function are showed in Section E. The most critical qualitative behaviors of the value function are shown in Figures 3 and 4.

8 Optimal Contract Discussion

In this section we will heuristically discuss the properties of the optimal contract and the implemented actions. The emphasis is on the economic intuition and hence we will suppress the mathematical details in this section. In particular, for the sake of discussion in this section, we will assume outright that the value function is sufficiently smooth such that all the partial derivatives make sense.




Figure 3: Illustration of the state space Γ . The *continuation region* is the set \mathcal{C} , and the *payment condition* is the set \mathcal{D} . Here, the *free (moving) boundary* that separates between the continuation region and the payment condition is $m \mapsto \bar{W}(m)$. It should be noted that the shape of \bar{W} as drawn is only meant to be illustrative.




Figure 4: Illustration of the value function in the w -slice. That is, for each $m \in \Gamma_M$, we plot the value function $w \mapsto v(w, m)$.

8.1 Optimal sensitivity

Let's first begin by discussing the optimal choice of sensitivity β . Fix any $\sigma \in [\sigma_L, \sigma_H]$. From the HJB equation (7.3), when (w, m) is in the no payment region, the optimal choice of sensitivity must thus be,

$$\sup_{\beta \in \mathcal{B}} \beta \sigma v_{wm}(w, m) + \frac{1}{2} \beta^2 \sigma^2 v_{ww}(w, m), \quad (8.1)$$

and recall the definition of \mathcal{B} in (7.1).

Before we proceed to discuss the form of the optimal sensitivity β choice in the optimization problem (8.1), let's first discuss the diffusion term of the agent's continuation value dynamics dW_t in (S), and in particular, highlight how this makes our model significantly different from the drift-only control models. If we recall back to agent's continuation value dynamics dW_t (S), the overall diffusion term is $\beta_t \sigma_t M_t dB_t$. Hence, even focusing on the choices of (β, σ) on the agent's continuation value diffusion term alone, we see several effects at play. On the one hand, the principal wants to provide the cheapest or lowest amount of sensitivity β to induce the agent to adhere to his recommended actions. But on the other hand, the amount of risk (i.e. the diffusion term of the agent's continuation value) is not solely just based on the principal's imposed sensitivity β . It is indeed determined by the product of the sensitivity β , volatility choice σ , and the exogenous factor level $M_t = m$. That is to say, in contrast to drift-only control models, where the total amount of risk (i.e. again, meaning the diffusion term of the agent's continuation value) is of the form $\beta_t dB_t$, so the principal can directly dictate the amount of risk he wants to subject the agent to through the choice of sensitivity β . In contrast, in our case, the principal's choice of sensitivity β is *not* the only source of risk the agent

is facing — the agent faces the *product* $\beta_t \sigma_t M_t$, of which σ_t remains to be a term that the principal wants to recommend and dictate for direct payoff reasons, and M_t is an exogenous factor level not controlled by the agent nor the principal. In all, that is to say when the principal wants to provide incentives through the sensitivity β , the principal must thus take into account providing incentives for an optimal volatility choice σ , and also the exogenous factor level m . It is precisely in this sense, the ability for the principal to provide incentives to the agent to induce the agent to take on the recommended action is weakened, relative to a drift-only control model.

Once we understand the incentive concerns in the diffusion term of the agent’s continuation value dynamics dW_t , we can now be more specific about what the principal needs to consider in choosing the optimal sensitivity β to choose, as in the optimization problem (8.1). Again, we immediately note several effects that are distinctly not present in drift-only control models. The optimal choice of sensitivity β now clearly depends on the volatility choice σ , the cross marginal effect $v_{wm}(w, m)$ of the agent’s continuation value $W_t = w$ and the exogenous factor level $M_t = m$, and the second order effect $v_{ww}(w, m)$ of the agent’s continuation value. Let’s assume that $v_{ww}(w, m) < 0$ in the no payment region, implying that the principal, although is risk neutral, becomes “endogenously quasi risk averse” with respect to the agent’s continuation value. Then the objective function (8.1) is a concave quadratic continuous function in β over a compact convex set \mathcal{B} . Thus, a unique maximizer $\beta^*(\sigma; w, m)$ exists. Let us also define the sets on $[\sigma_L, \sigma_H]$,

$$\mathcal{G}_L(w, m) := \left\{ \sigma \in [\sigma_L, \sigma_H] : -\frac{v_{wm}(w, m)}{\sigma v_{ww}(w, m)} < \underline{\beta} \right\} \quad (8.2a)$$

$$\mathcal{G}_M(w, m) := \left\{ \sigma \in [\sigma_L, \sigma_H] : K \geq -\frac{v_{wm}(w, m)}{\sigma v_{ww}(w, m)} \geq \underline{\beta} \right\} \quad (8.2b)$$

$$\mathcal{G}_H(w, m) := \left\{ \sigma \in [\sigma_L, \sigma_H] : -\frac{v_{wm}(w, m)}{\sigma v_{ww}(w, m)} > K \right\}. \quad (8.2c)$$

Note that $\cup_{j \in \{L, M, H\}} \mathcal{G}_j(w, m) = [\sigma_L, \sigma_H]$, and $\mathcal{G}_j(w, m) \cap \mathcal{G}_k(w, m) = \emptyset$ for $j, k \in \{L, M, H\}, j \neq k$. Then by a usual constrained optimization argument, we see that the optimal sensitivity choice $\beta^*(\sigma; w, m)$ is given by the following.²¹

Proposition 8.1. *The optimal choice of sensitivity associated with the optimization problem (8.1) is,*

$$\beta^*(\sigma; w, m) = \begin{cases} \underline{\beta}, & \text{if } \sigma \in \mathcal{G}_L(w, m) \\ -\frac{v_{wm}(w, m)}{\sigma v_{ww}(w, m)}, & \text{if } \sigma \in \mathcal{G}_M(w, m) \\ K, & \text{if } \sigma \in \mathcal{G}_H(w, m). \end{cases} \quad (8.3)$$

Proof. The proof is immediate by the usual constrained optimization methods via the Kuhn-Tucker conditions. \square

As it is with drift-only control models, the object β is the sensitivity, or “incentives”, that the principal must subject and provide to the agent in order to induce the agent to take the principal’s

²¹We denote $j \in \{L, M, H\}$ for the sets $\mathcal{G}_j(w, m)$ to, respectively, mean “low”, “medium” and “high”. The reason is that if $j = L$ and $\sigma \in \mathcal{G}_L(w, m)$, then the optimal sensitivity $\beta^*(\sigma; w, m)$ is chosen to be the one at the lowest value; and likewise for the other cases of j .

desired action. In this case, we see that the sensitivity β that the principal wants to subject the agent to is underline by two distinct channels: (i) the level of the exogenous factor at $M_t = m$; and (ii) the volatility level $\sigma_t = \sigma$ that should be implemented. For the rest of this discussion, let's hold the recommended volatility level σ as fixed. Also, we recognize that when (w, m) is in the no payment, we only have that $v_w(w, m) \geq -1$. Thus, it is distinctly possible that $v_w(w, m) = 0$. We will suppose that $v_w(w, m) > 0$ for the sake of this economic discussion, but this is not enforced anywhere else. ²²

In the expressions (8.2), we see that the object $-\frac{v_{wm}(w, m)}{\sigma v_{ww}(w, m)}$ plays a significant role to determining the optimal choice of sensitivity in (8.3). To highlight its importance, we will label the term $-\frac{v_{wm}(w, m)}{\sigma v_{ww}(w, m)}$ as *risk adjusted sensitivity* (RAS). Economically, we can decompose RAS as follows:

$$-\frac{v_{wm}(w, m)}{\sigma v_{ww}(w, m)} = \underbrace{\frac{1}{\sigma}}_{\text{(a) Precision of volatility choice}} \times \underbrace{\left(-\frac{v_w(w, m)}{v_{ww}(w, m)}\right)}_{\text{(b) "Risk tolerance"}} \times \underbrace{\left(\frac{v_{wm}(w, m)}{v_w(w, m)}\right)}_{\text{(c) "Elasticity of exogenous factor"}}. \quad (8.4)$$

We see that RAS depends on three different terms: (a) precision of volatility choice; (b) “risk tolerance”; and (c) “elasticity of exogenous factor”.

Let's first discuss the economic channel for which RAS would induce the optimal sensitivity β to be low, that is $\beta^*(\sigma; w, m) = \underline{\beta}$ in (8.3). Noting the form of $\mathcal{G}_L(w, m)$, in order for the set to be nonempty, we see that while $-v_{ww}(w, m) > 0$, given that $\underline{\beta} > 0$, there are no particular sign restrictions on $v_{wm}(w, m)$. Economically, this is the case when the principal does not care or want exposure to the exogenous factor level, that is roughly to say, the principal is relatively “inelastic” to the exogenous factor $M_t = m$. Once the principal does not care about the exogenous factor. then indeed, we return to the perhaps more familiar economic logic of drift-only control models. As well, the precision of volatility choice here must be relatively low, so that the choice of volatility is relatively high. Also, we can infer also here that the principal must have low risk tolerance, which implies the principal wants to achieve the lowest overall volatility of cash flow diffusion, and this is achieved when the principal subjects the agent to the lowest sensitivity $\beta^*(\sigma; w, m) = \underline{\beta}$.

Next, let's discuss the economic channel for which RAS would induce the optimal sensitivity β choice to be high, that is $\beta^*(\sigma; w, m) = K$ in (8.3). Firstly, noting the form of $\mathcal{G}_H(w, m)$ in (8.2c), we see that since $-v_{ww}(w, m) > 0$, if we $v_{wm}(w, m) \leq 0$, then the set $\mathcal{G}_H(w, m) = \emptyset$. So let us suppose and discuss the case when $\mathcal{G}_H(w, m) \neq \emptyset$, which implies $v_{wm}(w, m) > 0$. If $\sigma \in \mathcal{G}_H(w, m)$, then it implies that the precision of volatility choice is high, or that the volatility choice is relatively low. Furthermore, the “risk tolerance” term must also be relatively high, and the “elasticity of exogenous factor” is also relatively high. This is effectively the scenario when the cash flow volatility σ is relatively low, the principal is relatively risk tolerance and so is willing to take on more risk, and hence is willing to let the exogenous factor to give the extra “risk” bump, and so justifying why $v_{wm}(w, m) > 0$. In such a case, the principal wants to put the highest sensitivity or incentives to

²²Indeed, for the rest of this discussion, the sign and value of $v_w(w, m)$ is largely irrelevant. However, including the term $v_w(w, m)$ allows us to identify terms like $-v_w(w, m)/v_{ww}(w, m)$ as “risk tolerance” as it is traditionally defined (when we view the agent's continuation value $W_t = w$ as a “consumption good”) ²³, and view $v_{wm}(w, m)/v_w(w, m)$ as an “elasticity” in the traditional economic sense. But even without this normalization by $v_w(w, m)$, all of the economic reasoning here goes through, except that it may not be appropriate to keep on using the traditional economic labels.

the agent.

Finally, let's discuss the economic channel for which RAS would induce the optimal sensitivity β choice to be medium, that is $\beta^*(\sigma; w, m) = \text{RAS} = -\frac{v_{wm}(w, m)}{\sigma v_{ww}(w, m)}$. And it is through this medium case, which is effectively the interior solution to the optimization problem in (8.1), why we think the label *risk adjusted sensitivity* (RAS) is appropriate. Note that unlike the case when the optimal sensitivity $\beta^*(\sigma; w, m) = p(\sigma)$ is low, and similar to the high case, here for $\mathcal{G}_M(w, m)$ to be nonempty, there must be a sign restriction on $v_{wm}(w, m)$. In particular, we require that $v_{wm}(w, m) > 0$ and also not too high, nor too low. This case of which $\sigma \in \mathcal{G}_M(w, m)$ as in (8.2b) is exactly the “Goldilocks zone” and the optimal sensitivity $\beta^*(\sigma; w, m) = -\frac{v_{wm}(w, m)}{\sigma v_{ww}(w, m)}$ is almost like a “Goldilocks” sensitivity. That is, the precision of volatility choice is neither too high nor too low, the risk tolerance of the principal is neither too high nor too low, and the principal's appetite for the exogenous factor is neither too high nor too low.

In all, the above discussion not only suggests that a model with volatility control differs substantially to drift-only control models on how the optimal sensitivity β should be chosen, but equally important, our model suggests that how it is chosen is through the decomposition of the RAS term in (8.4).

8.2 Optimal volatility

Once the optimal sensitivity has been characterized, as discussed in Section 8.1 and in (8.3), we are now ready to discuss the optimal volatility σ choice. The choice of volatility here also highlights an interesting economic result — while the principal and the agent both desire higher volatility as seen from their direct payoffs, so seemingly there is no moral hazard conflict effect, but there still exist a distinctive presence of a *reverse moral hazard* effect.

Firstly, it should be noted that while we have emphasized and focused the case when the principal implements the high effort $e_t \equiv e_H$ at all times, in the case of volatility control σ , it is a priori unclear whether it is possible to say which fixed volatility level that is prevalent at all time is optimal for both the agent and the principal. While there is indeed some loss of generality in focusing and implementing a high effort choice at all time, and this point is made clear in Zhu (2013) based on the effort-only control model of DeMarzo and Sannikov (2006), the motivation is that high effort at all times is “part of” the first best action. Recall again from the discussion in Section 5, the first best action is indeed to implement high effort $e_t \equiv e_H$ at all times, and also high volatility $\sigma_t \equiv \sigma_H$ at all times.

But it is perhaps difficult to motivate and justify how and why the principal would find it desirable to implement the first best volatility choice, that being the high volatility choice, at all times. It is here that we can pinpoint the source of this *reverse moral hazard* effect. Recalling the payoffs of both the agent in (4.3) and principal in (4.4), it would *appear* that there is no direct moral hazard conflict between the principal and the agent in volatility choice at all. That is, both the principal and agent strictly prefer higher levels of volatility. Thus, it might appear that in the optimal contract, the principal should be able to implement the first best level of volatility, namely setting $\sigma_t \equiv \sigma_H$ to the highest level, at all times. But in light of the discussion in Section 8.1 of the overall risk or

diffusion term of the agent's continuation value dW_t , we can see that fixing at the high volatility choice at all times leads to the overall diffusion term $\beta_t \sigma_H M_t dB_t$. Recall again that in this setup, termination is inefficient. In particular, that implies up to the IR conditions of the agent being met, economically the principal would want to keep the agent employed as long as possible. Then there are now two tensions. While the principal's instantaneous direct payoff $\kappa(e_H, \sigma_t)$ is maximized when choosing $\sigma_t = \sigma_H$ (again, the first best result), if the principal recommends a higher volatility σ_t at time t , it also boosts the probability that the agent's continuation value W_t will hit the termination boundary (i.e. when the first time when $W_t = R = 0$), and thereby the principal will only get the inefficient liquidation value L . This is precisely the *reverse moral hazard* effect. By the IR condition, so $W_t \geq R = 0$, it is clear that it is better for the agent to be employed than to be terminated. But to keep employment, even though the agent desires a higher volatility choice through his private benefits, the agent must at the same time also desire lower volatility to maintain employment. Similarly, while the principal obtains a higher direct payoff from recommending a higher volatility choice, it is endogenously in the interest of the principal to not recommend too high of a volatility choice for fear of terminating the agent and receiving the inefficient liquidation value.

In all, this implies that it is *not necessarily* optimal for the principal to recommend the first best volatility choice at all times. And indeed, the above discussion highly suggests that the *reverse moral hazard* effect will endogenously lead the principal to shade down the choice of volatility. And also, by choosing a higher volatility σ_t , and also observing the decomposition of RAS in (8.3) and the sets (8.2), it also implies that the choice of sensitivity β will tend to be lower. In all, and already suggested in Section 8.1, there is an interplay of effects between the optimal choice of volatility and optimal choice of sensitivity. Let us make this precise below.

With the optimal sensitivity choice $\beta^*(\sigma; w, m)$ characterized in (8.3), we define first the objective function,

$$\begin{aligned}
G(\sigma; w, m) := & \phi_\sigma \left(\frac{\sigma}{\sigma_L} - 1 \right) v_w(w, m) + \beta^*(\sigma; w, m) \sigma m^2 v_{wm}(w, m) \\
& + \frac{1}{2} \beta^*(\sigma; w, m) \sigma^2 m^2 v_{ww}(w, m) + \kappa(e_H, \sigma),
\end{aligned} \tag{8.5}$$

and the optimization problem,

$$\max_{\sigma \in [\sigma_L, \sigma_H] = \cup_j \mathcal{G}_j(w, m)} G(\sigma; w, m). \tag{8.6}$$

Economically, we see that when the principal recommends the volatility choice, there are several effects at play. The first three terms of (8.5) are for the principal to internalize the agent's concerns. The first term $\phi_\sigma \left(\frac{\sigma}{\sigma_L} - 1 \right) v_w(w, m)$ are the direct payoffs to the agent for choosing volatility level σ , multiplied by the weight $v_w(w, m)$. The weight $v_w(w, m) \geq -1$ represents the marginal value of the principal's value function with respect to an increase to the agent's continuation value. So if $v_w(w, m) > 0$, then it is marginally beneficial for the principal to increase the agent's continuation value, and in that case, the principal would prefer to recommend a higher volatility choice, which is also a private benefit again for the agent; vice-versa, if $-1 \leq v_w(w, m) < 0$, then the principal would

want to decrease the agent's continuation value, and this is achieved by picking a lower volatility level that incurs a private benefit cost to the agent. And if $v_w(w, m) = 0$, then the principal is indifferent. The second and third terms $\beta^*(\sigma; w, m)\sigma m^2 v_{wm}(w, m) + \frac{1}{2}\beta^*(\sigma; w, m)\sigma^2 m^2 v_{ww}(w, m)$ of (8.5) capture the sensitivity effects as discussed in Section 8.1, which effectively captures the cost to providing incentives to the agent, except now the exogenous factor level effect $M_t = m$ is now explicitly present. Finally, the last term $\kappa(e_H, \sigma)$ captures the principal's concerns. In all, that means in the problem of choosing and recommending the optimal volatility, the principal must trade off the agent's incentives, the cost and sensitivity to providing correct incentives to the agent, and also the principal's own desired preferences.

At this point, to consider the optimization problem (8.5), we effectively need to partition the volatility control σ into three different regions, according to (8.2) and accordingly change the value of $\beta^*(\sigma; w, m)$ as given in (8.3). While $\mathcal{G}_M(w, m)$ is clearly a compact subset of $[\sigma_L, \sigma_H]$, it is clear that $\mathcal{G}_L(w, m)$ and $\mathcal{G}_H(w, m)$ are just half-open interval subsets of $[\sigma_L, \sigma_H]$. So from an optimization perspective, optimizing over non-compact intervals might have serious non-existence issues. However, this is not a concern in our current case. For instance, recalling (8.2c), if we pick $\sigma \in \mathcal{G}_H(w, m)$ such that $\beta^*(\sigma; w, m) = K$, and if it is indeed the optimizer σ is at the boundary K of the set $\mathcal{G}_H(w, m)$, then that effectively means $-\frac{v_{wm}(w, m)}{\sigma v_{ww}(w, m)} = K$ and hence we are no different from optimizing over the closure $\overline{\mathcal{G}_H(w, m)}$ or evaluating the objective function at $\sigma \in \mathcal{G}_M(w, m)$ such that $-\frac{v_{wm}(w, m)}{\sigma v_{ww}(w, m)} = K$. Either case, the optimal sensitivity is $\beta^*(\sigma; w, m) = K$ and so the overall objective function $G(\sigma; w, m)$ of (8.5) remains the same. Similar arguments apply to the case when we consider $\mathcal{G}_L(w, m)$ of (8.2a). Thus, with this argument in mind, we modify our optimization problem and consider,

$$\max_{\sigma \in \cup_j \mathcal{G}_j(w, m)} G(\sigma; w, m). \quad (8.7)$$

We will consider the optimization each case at a time. When we pick $\sigma \in \overline{\mathcal{G}_j(w, m)}$, for $j = L, M, H$, then the objective function respectively becomes,

$$\begin{aligned} G(\sigma; w, m)|_{\overline{\mathcal{G}_L(w, m)}} &= \phi_\sigma \left(\frac{\sigma}{\sigma_L} - 1 \right) v_w(w, m) + p(\sigma)\sigma m^2 v_{wm}(w, m) \\ &+ \frac{1}{2}p(\sigma)^2 \sigma^2 m^2 v_{ww}(w, m) + \kappa(e_H, \sigma), \end{aligned} \quad (8.8a)$$

$$G(\sigma; w, m)|_{\overline{\mathcal{G}_M(w, m)}} = \phi_\sigma \left(\frac{\sigma}{\sigma_L} - 1 \right) v_w(w, m) - \frac{1}{2} \frac{v_{wm}(w, m)^2}{v_{ww}(w, m)} m^2 + \kappa(e_H, \sigma), \quad (8.8b)$$

$$\begin{aligned} G(\sigma; w, m)|_{\overline{\mathcal{G}_H(w, m)}} &= \phi_\sigma \left(\frac{\sigma}{\sigma_L} - 1 \right) v_w(w, m) + K\sigma m^2 v_{wm}(w, m) \\ &+ \frac{1}{2}K^2 \sigma^2 m^2 v_{ww}(w, m) + \kappa(e_H, \sigma). \end{aligned} \quad (8.8c)$$

From the forms in (8.8), we see that in general there are no closed form and simple analytic expressions of the optimal choice of volatility $\sigma^*(w, m)$. Moreover, one needs to compute the (set of) optimizers $\sigma_j^*(w, m)$ for each case $j = H, M, L$, substitute the optimizer back into the objective function $G(\sigma_j^*(w, m); w, m)|_{\mathcal{G}_j(w, m)}$, and once that is complete, the optimal volatility choice is the set,

$$\sigma^*(w, m) \in \arg \max_{j \in \{L, M, H\}} G(\sigma_j^*(w, m); w, m)|_{\mathcal{G}_j(w, m)}. \quad (8.9)$$

Remark 8.2. As a general remark, it should not be surprising that the optimal volatility choice in the form (8.9) is rather this complicated. Indeed, if one observes the drift-only control model of Zhu (2013), which is based off of the model of DeMarzo and Sannikov (2006), in which the agent has a binary choice of effort, it is readily seen that it is not trivial and indeed rather challenging to characterize the optimal effort choice. Here, we have already simplified matters substantially by concentrating on implementing the always high effort case, but nonetheless, even allowing for volatility to be optimally implemented, the resulting optimal volatility recommendation is nonetheless rather complicated to observe.

Remark 8.3 (Difficulty of direct application of verification theorem). As we conclude the discussion of the optimal choice of sensitivity in Section 8.1 and optimal choice in this Section 8.2, we can now remark the tremendous difficulty in applying the traditional “verification theorem” to conjecture the existence of a smooth solution of the HJB (7.3) that actually coincides with the value function in (P). A classical method is the verification theorem argument, or the “guess and verify” argument, where one conjectures that a PDE that solves HJB equation subject to some well thought out and economically motivated boundary conditions, and from the HJB, one takes the first order conditions to obtain the optimal controls, and substitute these controls back into the original HJB equation. There, one then proceeds to directly solve the PDE by constructing an explicit solution and thereby directly proving existence and also smoothness. Then essentially by Ito’s lemma argument, one can then verify that the HJB is a supersolution of the value function, and hence under the optimal controls, the HJB is the solution to the value function. This type of argument is fairly prevalent in the finance literature, especially in asset pricing theory, and also in continuous-time principal agent problems where there is a single state variable, so the problem is an ODE rather than a PDE.

However, we see here once we substitute the optimal sensitivity $\beta^*(\sigma; w, m)$ in (8.3) and optimal volatility $\sigma^*(w, m)$ in (8.9) back into the HJB equation (7.3), the resulting HJB is sufficiently complex that it is difficult to see how we can indeed obtain the existence of a smooth solution that satisfies the necessary boundary conditions. And this is especially why in the technical proofs, we have to proceed through a more roundabout way via viscosity solutions to show existence, and then “upgrade” our smoothness results.

9 Delegated portfolio management

9.1 Motivation

We are now ready to consider the concrete application of the model in the context of delegated portfolio management. Suppose we regard the principal as outside investors of a managed portfolio, and regard the agent as the portfolio manager. In this context, we will explicitly assume that the portfolio manager has skill and can exert costly effort to search and achieve higher cash flow payoffs in the managed portfolio. That is, the portfolio manager can directly control the drift of the cash flows. Furthermore, we assume the portfolio manager has available assorted tools and financial instruments to engage into hedging and speculating behavior that can change the overall volatility of the cash flows. The portfolio itself is also subject systematically subject to an exogenous market wide or industry wide factor that the portfolio manager cannot control. Hence, portfolio volatility is comprised of a manager specific choice in volatility, reflecting risk management practices, and an exogenous market or industry factor. A delegated portfolio management problem framed in a principal-agent setting has also been considered by Ou-Yang (2003); other recent models that consider delegated portfolio management problems include van Binsbergen et al. (2008), Dybvig et al. (2009) and Cvitanić et al. (2014). In this section, we will relabel and call the agent as the *manager*, and the principal as *investor*.

The main idea here is to have the portfolio manager to have sufficient “skin in the game” through his own investments. Specifically, consider an investment firm whereby the investment manager operates two different investment funds: an *external fund* fund that is available to outside investors and an *internal fund* that is only available to management. Suppose we regard the continuation value W_t as the value of an *internal fund* that is only available to the portfolio manager but not to the outside investor. This form of internal fund that is only available to insiders of the firm, and not outside investors, is also an observed market practice. For instance, numerous banks (at least prior to the Volcker Rule) also run proprietary trading desks, which are effectively internal hedge funds. Several hedge funds also engage into this practice. For example, the hedge fund firm Renaissance Technologies runs three funds that are open to outside investors, but also run a separate fund, the Medallion Fund, that is only open to its employees (see Zuckerman (2013)). Darolles and Gourieroux (2014a) and Darolles and Gourieroux (2014b) also discuss at length the practice of this internal fund in the hedge fund industry. The external fund has value $v(W_t, M_t)$. In particular, this specifically implies that the value of the external fund is explicitly dependent on the value of the internal fund W_t and also the exogenous factor level M_t .

Remark 9.1. Asserting that the portfolio volatility overall is directly influenced by the portfolio manager’s risk management practices and market volatility should be reasonable. However, asserting that the portfolio manager has skill, and moreover that such skill surely translates to higher cash flow payoffs is a strong assumption. Indeed, the question of whether persistent manager skill exists or not is still heavily debated in empirical research.²⁴ As well, what we refer to as skill here differs from the usual interpretation of skill in the empirical research of fund performance. In the

²⁴The literature on this question is too vast to survey here with any justice, but some recent papers here include: Fama and French (2010) and Barras et al. (2010).

empirical literature, skill usually refers to whether the portfolio manager can deliver positive alpha, adjusting for the systematic risk factors the portfolio takes on. In contrast, here we refer to skill as the manager's direct ability to influence the drift of the cash flows (not returns). Thus, recognizing fully that this is a debatable aspect, we will nonetheless make such a strong assumption in this section.

9.2 Interpreting the continuation value

For given sensitivity $\beta = \{\beta_t\}$ and volatility $\sigma = \{\sigma_t\}$, the continuation value dynamics dW_t in (S) in the no payment region so $dX_t \equiv 0$, we can rewrite the expression as,

$$\begin{aligned} dW_t &= \left[r_0 W_t - \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt + \beta_t \sigma_t dM_t \\ &= \left[r_0 W_t - \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt + \beta_t (dY_t - \kappa(e_H, \sigma_t) dt) \\ &= \beta_t dY_t + r_0 W_t dt - \left[\phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) + \beta_t \kappa(e_H, \sigma_t) \right] dt. \end{aligned} \quad (9.1)$$

With the context of delegated portfolio management, we will interpret the continuation value dynamics dW_t via the expression form of (9.1).

Since W is the value of the internal fund, the expression (9.1) suggests that the value of the internal fund is driven by the amount of ownership $\beta_t dY_t$ the agent has of the underlying investment technology, and plus an investment $r_0 W_t dt$ into a riskfree asset that pays off at a rate r_0 . Note here that in this context, we can interpret β_t as the *dynamic incentive fees* of ownership the manager owns of the investment opportunity dY_t , so $\beta_t dY_t$ is the total dollar exposure the manager has to the managed cash flows. The latter two terms of (9.1) represent the cost of implementing an investment strategy σ_t . The term $\phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) dt$ can be thought of as the direct cost of implementing the investment strategy σ_t ; for instance, this could represent the direct trading costs or managerial monitoring costs. The term $\beta_t \kappa(e_H, \sigma_t) dt$ represents the proportional expected return of implementing the strategy σ_t . So the manager's internal fund value only gets a positive bump if $\beta_t (dY_t - \kappa(e_H, \sigma_t) dt) > 0$, and since the ownership amount $\beta_t > 0$, then we have that $\beta_t (dY_t - \kappa(e_H, \sigma_t) dt) > 0$ if and only if $dY_t - \kappa(e_H, \sigma_t) > 0$. That is to say, the manager captures the positive excess returns over the expected return of the investment strategy, only if the investment strategy performs extraordinarily well.

Since we focus on implementing an always high effort action $e_t \equiv e_H$, other than the dollar incentives β_t , the remaining control policy here is the volatility σ_t . Again, we interpret β_t as the dollar incentives of ownership, or *dynamic incentive fees*, the manager owns of the managed investment opportunity dY_t , so $\beta_t dY_t$ is the total dollar exposure the manager has to the managed cash flows. In addition, here we may broadly interpret σ_t as *investment strategies*. To be more specific, once in equilibrium we implement an always high effort $e_t \equiv e_H$ action, then it effectively implies that the manager is already exerting costly skill to find the set investment opportunities with good returns. However, even after exerting skill to find this investment opportunity set, the manager still needs to choose the specific investments from the set, and it is here we interpret σ_t

as the opportunities available to the manager. Specifically, we will let σ_t be effectively a parameter captures both the set of investment opportunities and hedging strategies.

Now using the results in Section 8, under the optimally chosen sensitivity $\beta^*(\sigma; w, m)$ in (8.3) and optimally chosen volatility $\sigma^*(w, m)$ in (8.9), and the set forms $\mathcal{G}_j(w, m)$ in (8.2), we can thus write (9.1) as,

$$dW_t = \sum_{j \in \{L, M, H\}} \left\{ \beta^*(\sigma^*(W_t, M_t); W_t, M_t) dY_t + r_0 W_t dt - \left[\phi_\sigma \left(\frac{\sigma^*(W_t, M_t)}{\sigma_L} - 1 \right) + \beta^*(\sigma^*(W_t, M_t); W_t, M_t) \kappa(e_H, \sigma^*(W_t, M_t)) \right] dt \right\} \mathbf{1}_{\{\sigma^*(W_t, M_t) \in \mathcal{G}_j(W_t, M_t)\}}. \quad (9.2)$$

Economically, the form of (9.2) implies the following contractual implementation. The outside investors offer the manager an initial start up fund value of W_0 at $t = 0$ and in return, the manager commits to the following *dynamic incentive fee* compensation scheme as represented via the optimal dollar incentives $\beta^*(\sigma^*(w, m); w, m)$, viewed as a map from (w, m) . That is to say dependent on the value of the internal fund $W_t = w$ and also the exogenous factor level $M_t = m$, the manager will choose a different investment strategy $\sigma^*(w, m)$. And dependent on this strategy, the internal fund will only get different dollar exposures of the managed cash flows dY_t . For instance, for those investment strategies such that $\sigma^*(w, m) \in \mathcal{G}_L(w, m)$, the manager gets a low dollar incentive $\beta^*(\sigma^*(w, m); w, m) = \underline{\beta}$; for those investment strategies $\sigma^*(w, m) \in \mathcal{G}_M(w, m)$, the manager gets a medium dollar incentive $\beta^*(\sigma^*(w, m); w, m) = -\frac{v_{wm}(w, m)}{\sigma^*(w, m)v_{ww}(w, m)}$; and finally, for those investment strategies $\sigma^*(w, m) \in \mathcal{G}_H(w, m)$, the manager gets a high dollar incentive $\beta^*(\sigma^*(w, m); w, m) = K$.

The critical real-world implication of the above is as follows — investors should contract on the value of the internal fund and the exogenous factor level. Again, take hedge funds as a prototypical example. Hedge fund investment strategies are essentially completely black box, and effectively that means even investors into the fund most often have no idea what types of investment strategies the manager is employing. Hence, that makes directly contracting on investment strategies to be highly unrealistic and impossible. In our framework however, the investor only needs to contract on two things: the value of the internal fund $W_t = w$ and the stochastic factor level $M_t = m$. That is, the investor writes a contract not on the investment strategy that maps to the dollar incentives $\sigma \mapsto \beta^*(\sigma; w, m)$, but rather directly from the value of the internal fund and the stochastic factor level $(w, m) \mapsto \beta^*(\sigma^*(w, m); w, m)$, and we emphasize that $\beta^*(\sigma^*(w, m); w, m)$ only depends on (w, m) , and indeed only has three relatively “small” sets of values, as given by (8.3).

However, while the investors can certainly contract on the value of the internal fund $W_t = w$, it is unclear how the investors can contract on the factor level $M_t = m$. In particular, recall that off equilibrium, the investors cannot observe the exogenous factor level. Thus, to complete the optimal contract implementation, we further require the manager to directly and truthfully report the exogenous factor level to the investors. In practice, that translates to the manager reporting periodically some factor benchmark index to the investors.

Thus, with the internal fund value $W_t = w$ and the exogenous factor level $M_t = m$ known to the investors, the investors can just adjust the level of the dynamic incentive fee $\beta^*(\sigma^*(w, m); w, m)$ accordingly, without knowledge of the actual employed investment strategy $\sigma^*(w, m)$. The advantage

of this implementation is that the manager does *not* need to report to the investors their actual employed investment strategy, which is usually what is observed in practice in the case of hedge funds. Moreover, if the internal fund does well, so when $W_t = w$ hits the payment boundary $\bar{W}(M_t) = \bar{W}(m)$, the external investors will directly compensate manager ²⁵

As a result of the above discussion, the value of the external fund to the investors is $v(w, m)$, when the value of the internal fund is $W_t = w$ and the level of the exogenous factor is $M_t = m$. But economically and conceptually, what does it mean by the value of the external fund is a function of the value of the internal fund and the level of the exogenous factor? Borrowing the language of financial derivatives, we effectively can view the external fund as a *derivative*, where the underlying asset here is written on the value of the internal fund and level of the exogenous factor, with two associated barriers. The lower barrier is the first time $(W_t, M_t) = (w, m)$ hits the level $(w, m) = (R, m)$ or $(w, m) = (w, \underline{m})$; that is, either when the value of the internal fund goes bust (i.e. $W_t = R$), or when the exogenous factor level is sufficiently low (i.e. $M_t = \underline{m}$) that the manager effectively walks away from the firm. The upper barrier is the moving barrier $(\bar{W}(M_t), M_t)$ that determines the optimal capital injection or compensation scheme. However, despite this discussion, a direct implication here is that the investment strategy of the manager still remains a black-box to the investor. Another direct implication here is that the investment strategy of the external fund will closely track that of the internal fund's investment strategy.

Finally, we should observe what is *not* optimal or feasible in this context. Most notably, the perhaps “easier” contractual setup would be that there is a single investment fund, managed by the manager, for which managers and investors commonly invest to. In this context, this is not possible. As a thought experiment, suppose this were true, meaning the internal fund and the external fund are exactly identical. But because the manager can privately select effort and volatility (again, broadly interpreted as investment opportunity), and by limited liability, the manager would effectively have incentives to gamble (i.e. choose the highest volatility) and exert the lowest effort. Focusing on the volatility choice, although by the form of the reward function $\kappa(e, \sigma)$ it may appear that it too is desirable for the investor to choose the highest volatility, the discussion in Section 8.2 argued that this is not the case. As well, and this is perhaps a more cynical view of managers, suppose there does exist only a single common fund, and recalling the black box nature of the investment strategy, how can investors ensure that the managers will not privately squirrel away the best available investment opportunities and leave the subpar investment opportunities to the common fund? However, the establishment of an internal only fund for the manager with an external fund that closely tracks the investment strategy of the internal fund does mitigate this concerns. That is, although in equilibrium, the manager will collect some “information rent”, namely in the form of keeping the best investment strategies still for the internal fund, but if the investment strategy of the external fund commits to following that of the internal fund, the external fund still benefits from the exposure of those good investment strategies.

In all, as a concrete and simple investment policy recommendation, this model highly suggests that outside investors should actively seek funds with the following characteristics: (i) the fund has an

²⁵In this setup, injecting more capital dX_t into the internal fund is equivalent to compensating the manger, as we assume the manager derives all utility from maximizing the value of the internal fund.

internal fund available only to management, and an external fund only available to external investors; (ii) the fund investment strategies of the internal fund and external fund are closely correlated to each other; and (iii) the investment firm has dynamic incentive fee compensation schemes.

10 Conclusion

We studied continuous-time principal-agent problem where the agent can continuously choose the drift and volatility parameters, while the principal continuously observes and receives the resulting controlled cash flows. The key ingredient yielding a meaningful private volatility control by the agent, in that the agent's deviation cannot be easily detected by the principal's computation of the quadratic variation of the cash flows, is via the introduction of an exogenous factor level. Hence effectively, even though the principal can infer the overall instantaneous diffusions of the cash flows, the principal cannot disentangle the component that is due to the agent's endogenous volatility control and the exogenous factor level. As a result, beyond merely hidden drift or effort control, the principal must provide incentives now for both inducing the desired effort and volatility. Most importantly, as a concrete application, our current model provides a first step to considering the dynamic contracting environment in the context of delegated portfolio management.

By introducing this meaningful sense of volatility control, we now open a new economic channel for researchers to study continuous-time principal-agent problems. In particular, there are further questions one can consider between the interplay of effort and volatility. In particular, there are several questions that this framework researchers could consider:

- The current model assumes the principal is risk neutral. However, once we give the agent the meaningful ability to privately select volatility, an immediate and relevant extension is to consider a case when the principal is risk averse.
- In the context of delegated portfolio management, since prices of risky assets jump (i.e. “disaster” states), it will be interesting to pair, say, our current model to that of DeMarzo et al. (2013), where the agent can also influence the likelihood of a disaster state occurring.

In all, we feel that this model is an important step to the growing literature of continuous-time dynamic contracting, and also to our better understanding of delegated portfolio management contracting practices.

Appendices

A The “quadratic variation test”

Given the continuous-time and Brownian noise driven setup, the computation of the quadratic variation turns out to be of utmost importance, especially in the case when we have volatility control. Technical details revolving around the quadratic variation can be found in standard references like Karatzas and Shreve (1991), Protter (2005) and Øksendal (2010), among many others.

In what follows, we will first specialize the discussion to when the agent controls the volatility σ_t so that it is continuous in time t . In the latter section, we will consider the case when the agent does not necessarily need to make the volatility control to be continuous in time.

A.1 Models in the existing literature

In continuous-time principal agent models like Holmström and Milgrom (1987) and the more recent papers by DeMarzo and Sannikov (2006), Biais et al. (2007), Sannikov (2008), He (2009) and others, all consider the case of *constant* volatility. For instance, taking the model of DeMarzo and Sannikov (2006) and Sannikov (2008), cash flows are of the form (D.2). Thus, the agency conflict only arises due to the principal’s unobservability of the agent’s choice of the drift, which is broadly interpreted as hidden effort. However, in those model, it is without loss of generality to consider constant volatility. For instance, even if volatility σ_t (non constant) is a choice variable for the agent in this setting, and since the principal can continuously observe the cash flow process Y , then at any time t , the principal can infer the agent’s choice of volatility with arbitrary precision. To see this, suppose the principal has observed the cash flow process $\{Y_s\}_{0 \leq s \leq t}$ according to (D.2) up to time t (replacing σ by σ_t). Then by computing the quadratic variation ²⁶, and using the fact that noise here is driven by a standard Brownian motion B , and that σ_t is continuous in time,

$$[Y]_t = \int_0^t \sigma_s^2 ds.$$

The right-hand side is a standard Riemann integral, of which the principal can apply the fundamental theorem of calculus and compute that,

$$[Y]_t' = \sigma_t^2, \quad \text{for all } t > 0, \text{ a.s.},$$

and since the left-hand side of the above is observable to the principal, and thus so is the volatility choice σ_t at each point in time ²⁷. Thus, it follows that even if the agent can choose the volatility, given that the principal can perfectly observe such a choice, the principal can thus directly dictate his desired level of volatility and immediately detect deviations. Thus, it is for this reason, choosing volatility $\sigma_t = \sigma$ as a known constant to both the principal and the agent in these types of models is without loss of generality. Indeed, Schaettler and Sung (1993) and Sung (2004) use the same type of dynamics (i.e. noise driven by Brownian motion) in which the agent can control both the drift and volatility; they show that the volatility is naturally optimally chosen to be a constant. Thus indeed, in these types of models, the only variable that could have moral hazard type conflicts is in the choice of the drift.

A.2 Current model

Here, the dynamics are fundamentally different. If we use the same “quadratic variation test”, we immediately see that even if the principal observes the cash flow $\{Y_s\}_{0 \leq s \leq t}$ up to time t of (4.1), and since both σ_t and m_t are continuous in time t , we have that,

$$[Y]_t = \int_0^t \sigma_s^2 m_s^2 ds.$$

²⁶Throughout this paper, we will denote the *quadratic variation* of X as $[X]_t$, and likewise, the *quadratic covariation* between X and Y as $[X, Y]_t$.

²⁷Since we assume that the volatility term is always strictly positive, that means that even if the square σ_t^2 is observed, by taking the square root, we know for sure the value of σ_t .

Again, applying the fundamental theorem of calculus, we see that,

$$[Y]'_t = \sigma_t^2 m_t^2, \quad \text{for all } t > 0, \text{ a.s.}$$

Thus, even if the principal can observe the left-hand side $[Y]'_t$ above, he can at best observe the product $\sigma_t^2 m_t^2$ but not the two components separately, since the process m is not observed by the principal independently. Thus, under the specification of (4.1) and (4.2), the agent can choose volatility without being directly detected by the principal and hence, this is a meaningful moral hazard term.

A.3 Detecting deviations in the current model?

As argued in section A.1, for the models in the existing literature, there is no loss of generality in assuming that the volatility is not under control; whether we allow the non-controlled exogenous volatility to be constant or stochastic then is just a matter of taste. But continuing on with the discussion in section A.2, we now have a legitimate concern of whether the principal can detect deviations from his recommended actions without the need to provide incentives.

A.3.1 An incorrect computation (i)

We will first discuss a type of computation that is seemingly correct and seemingly suggest that even in this current more complex model volatility control is mute, but is indeed faulty. Suppose the principal offers a contract (A, X, τ) with action process $A = \{(\mu_t, \sigma_t)\}_{t \geq 0}$. Of course, we will just consider the case when the contract is not incentive compatible for the agent, else if the contract was to be incentive compatible, the agent will not deviate from the recommendation. Suppose the agent ignores the principal's recommended action A and deviates to $A^\dagger = \{(\mu_t^\dagger, \sigma_t^\dagger)\}_{t \geq 0}$. Under the deviant action A^\dagger , the cash flows will now evolve according to,

$$\begin{aligned} dY_t &= \mu_t^\dagger dt + \sigma_t^\dagger m_t dB_t, & Y_0 &= 0 \\ dm_t &= m_t dB_t, & m_0 &= 1. \end{aligned}$$

Repeating the same quadratic variation test, and for emphasis here, we insert the notation ω to denote this is a sample space path-by-path computation, the principal would compute at time t ,

$$[Y]_t(\omega) = \int_0^t (\sigma_s^{\dagger 2} m_s^2)(\omega) ds. \tag{A.1}$$

One might strongly suggest the following. Since the principal knows his own recommended action A , and if the agent complied and followed the principal's recommended action, the principal would correctly deduce that the cash flows might be evolving according to dynamics, and we denote Y^{conj} to mean "conjecture",

$$\begin{aligned} dY_t^{\text{conj}} &= \mu_t dt + \sigma_t m_t dB_t, & Y_0 &= 0 \\ dm_t &= m_t dB_t, & m_0 &= 1. \end{aligned}$$

And if the principal were to do the exact same quadratic variation computation, the principal would deduce that,

$$[Y^{\text{conj}}]_t(\omega') = \int_0^t (\sigma_s^2 m_s^2)(\omega') ds, \tag{A.2}$$

and note in particular that we do not necessarily have $\omega = \omega'$.

It is tempting, thus, for us to conclude that in order to detect deviations, all the principal needs to do is to observe $[Y]_t$ in (A.1) and consider his conjecture $[Y^{\text{conj}}]_t$ in (A.2) at each point in time t , and take the difference $[Y]_t - [Y^{\text{conj}}]_t$. If this difference is zero, then the agent must have followed the recommendation and if it is not zero, then the agent must have deviated and the principal can then appropriately punish the agent. Thus, this seems like we are back to the case of section A.1 where volatility control is mute.

However, this thinking is *incorrect*. Recall again the crucial assumption that the principal only observes the cash flows Y but *not* the exogenous volatility m . In particular, this means that generically, we cannot even compare $[Y]_t$ and $[Y^{\text{conj}}]_t$ on the same trajectory realization ω ! To make this point, let's apply the fundamental theorem of calculus in (A.1) and (A.2). While in (A.1), the principal, say, observes $10 = [Y]_t(\omega) = (\sigma_t^2 m_t^2)(\omega)$. It remains true that the

principal cannot disentangle $(\sigma_t^2 m_t^2)(\omega)$ separately; more forcibly, that is to say the principal cannot tell whether the observed value 10 is factored as $10 = 2 \times 5$ or is it $10 = 5 \times 2$, or infinitely many other factor combinations. Thus, the principal at most can see the product $(\sigma_t^2 m_t^2)(\omega)$ but not $\sigma_t^2(\omega)$ and $m_t^2(\omega)$ separately. More succinctly, this means the product $\sigma_t^2 m_t^2$ is \mathcal{F}_t^Y -measurable, but the individual components σ_t^2 and m_t^2 are not, and recall the information set of the principal is $\{\mathcal{F}_t^Y\}$.

Continuing onto the conjectured cash flow Y^{conj} , this means even if the principal can compute $[Y]_t^{\text{conj}}(\omega') = \sigma_t^2(\omega') m_t^2(\omega')$, for each ω' , and of course the principal knows his recommendation σ_t evaluated at any ω' and in particular when $\omega' = \omega$, the principal does not know what realization $m_t^2(\omega)$ should be plugged in (again, since m_t is not \mathcal{F}_t^Y -measurable). Thus, the wishful computation of evaluating at $\omega' = \omega$ to have $[Y]_t^{\text{conj}}(\omega) = \sigma_t^2(\omega) m_t^2(\omega)$, is not even possible, again since $m_t^2(\omega)$ is not known to the principal. Thus, the wishful deviation test of simply observing and checking whether the difference $[Y]_t - [Y^{\text{conj}}]_t$ is zero or not is invalid.

A.3.2 An incorrect computation (ii)

The discussion in section A.3.1 above may lead to wonder what the principal can detect deviations with more time points. Could the principal make a further level of inference to detect deviation not based on the quadratic variation but directly via the cash flows?

Suppose the principal makes the same computations as described in section A.3.2 for each point in time t . This implies that if the principal were to use the conjectured cash flows Y^{conj} that follows his recommended action, the principal could indeed compute ²⁸ $m_t^2 = m_t^{\text{conj},2} = [Y^{\text{conj}}]_t / \sigma_t^2$, ω' -by- ω' . And since the principal can do this for each ω' and for each t , and since $m_t^{\text{conj}} > 0$ by nature of the geometric Brownian motion, the principal has thus constructed a process $\{m_t^{\text{conj}}\}_{t \geq 0}$. Thus, that implies at any time t , the principal could construct another cash flow process $Y^{\text{conj}+}$ (i.e. denoting “conjecture plus”),

$$Y_t^{\text{conj}+} = \int_0^t \mu_s ds + \int_0^t \sigma_s m_s^{\text{conj}} dB_s.$$

Comparing it against the true cash flow Y , it is tempting to conclude that the principal simply needs to take the difference $Y_t^{\text{conj}+} - Y_t$ at each point in time and see if it equals zero at each point in time, as if this difference is zero, it would indicate the agent is compliant and if this difference is nonzero, it would indicate the agent had deviated.

Again, however, this thinking is *incorrect*. Taking the difference between $Y_t^{\text{conj}+}$ and Y_t , one finds,

$$Y_t^{\text{conj}+} - Y_t = \int_0^t (\mu_s - \mu_s^\dagger) ds + \int_0^t (\sigma_s m_s^{\text{conj}} - \sigma_s^\dagger m_s) dB_s.$$

We note that even if the agent were to be completely compliant in the drift recommendation so that $\mu_s = \mu_s^\dagger$ for all s , and so the first integral on the right-hand side above vanishes, we see that generically $m_s^{\text{conj}} \neq m_s$ for all s . Hence — however tempting — the conclusion to say that $m_s^{\text{conj}} = m_s$ holds with equality so that in the integrand of the second integral we can factor as $m_s(\sigma_s - \sigma_s^\dagger)$ and need to simply observe whether this is zero or not — is faulty.

Remark A.1. There is a very important case by which $m^{\text{conj}} \equiv m$. This is precisely the case when the principal offers a contract (A, X, τ) that is *incentive compatible* for the agent. Specifically, since the agent is recommended an action for which it is incentive compatible, the agent will not deviate from the principal’s recommended action. We will use this critical observation in a key lemma in the later development.

A.4 The case of volatility controls not continuous in the time path

The discussion of section A.2 and A.3 assumed that the agent must make time-continuous volatility controls; that is, we had restricted the discussion to σ_t to be continuous in time t . The consideration of jumps in the controls of the cash flow process is important, even in the drift-only control models of DeMarzo and Sannikov (2006), say when the cash flow had the form $dY_t = \mu_t dt + \sigma dB_t$, where the drift controls are binary in state values, $\mu_t \in \{0, A\}$. In the drift-only control case, the distinction of whether allowing for the state values to be binary (i.e. $\mu_t \in \{0, A\}$; see DeMarzo and Sannikov (2006) amongst others) or to be in a compact set (i.e. $\mu_t \in [0, A]$; see Sannikov (2008)), up to possibly difficult technical details, is not that important economically.

²⁸This type of computation is used again later in the discussion.

So now, suppose that the volatility control σ_t need not be continuous in time, and so we in particular allow it to have jumps ²⁹, so $\Delta\sigma_t = \sigma_t - \sigma_{t-} \neq 0$. But nonetheless, by standard stochastic integration theory, even if the integrand (i.e. in this case $\sigma_t m_t$) is cadlag, as long as the integrator is continuous (i.e. in this case dB_t), then there exists a version of the stochastic integral $\int_0^t \sigma_s m_s dB_s$ that is continuous in the time path t . Hence working with this version, the principal will still conclude that, by calculating the quadratic variation,

$$[Y]_t = \int_0^t \sigma_s^2 m_s^2 ds.$$

Note however now, the integrand in the Riemann integral $\sigma_t^2 m_t^2$ is not continuous in the time path t since σ_t is cadlag, not continuous. But if we define $g(s, \omega) := \sigma_s^2(\omega) m_s^2(\omega) \mathbf{1}_{[0, t]}(s)$, then it is immediate that for each ω , $g(\cdot, \omega)$ is absolutely integrable on \mathbb{R} . That means for each ω , we can apply the Lebesgue differentiation theorem ³⁰ to conclude that,

$$[Y]'_t(\omega) = (\sigma_t^2 m_t^2)(\omega), \quad \lambda\text{-a.e.},$$

where λ is the Lebesgue measure on \mathbb{R} . Up to this null set difference in the time $t \in \mathbb{R}_{++}$ coordinate, we can apply the same arguments as in sections A.3.1 and A.3.2 to illustrate that it is still not possible for the principal to detect deviations from the agent.

A.5 Economic intuition and summary of the “quadratic variation” test

As noted in section A.1, existing models in the literature does not have room for meaningful volatility control, precisely because they cannot “pass” the quadratic variation test. We assert that our model (4.1) does pass the quadratic variation test. Furthermore, we argue that for any meaningful agency problems associated with volatility control, the first pass must be to pass the quadratic variation test. Although we have not tried this approach yet but if we were to switch out the Brownian noise dB_t attached to the instantaneous controlled volatility $\sigma_t m_t$ with some more general stochastic process, it may also pass this quadratic variation test. It is also possible to compute even higher n th-order variation (i.e. cubic, quartic, etc.) but it is unclear how this will be useful for the principal in detecting deviations.

Economically, what exactly is this “quadratic variation” test trying to tell us and why does it matter so much for stochastic volatility control? Stepping away from the current setup (4.1), let’s consider a standard asset pricing type model. Suppose we have a stock with price S and has dynamics,

$$dS_t = \mu(t, S_t)dt + \sigma(t, S_t)dB_t,$$

where B is again a standard Brownian motion ³¹. Computing the quadratic variation of S and have,

$$[S]_t = \int_0^t \sigma(u, S_u)^2 du,$$

which roughly is precisely the integrated instantaneous variance of the process S . Effectively, the quadratic variation gives us the “sum (over time)” of the variances of the stock price S . And if variance is related to some measure of “risk” ³², this quadratic variation object, in some loose sense, is a time-tracker of “sum of risks”.

A.5.1 Brief comment to the link with financial econometrics

It should be noted that the econometric *estimation* of quadratic variation is an active topic of study in the field of financial econometrics theory. For instance, a recent survey of the developments in this field can be found in Aït-Sahalia and Jacod (2012) and a textbook reference like Kessler, Lindner, and Sorensen (2012) deals with statistical methodologies of SDE’s. The key point is that it is not too “far-fetched” to think of an individual (in this context,

²⁹In particular, we still restrict the trajectories to be cadlag and the process to be adapted.

³⁰See, say, Folland (1999, Chapter 3, Section 3.4).

³¹We omit spelling out the restrictions on μ and σ here to ensure that a solution to this SDE exists. Details can be found in Protter (2005).

³²This claim of “risk” is made precise when we consider individuals with mean-variance preferences or risk averse preferences.

the principal) to in reality employ a battery of statistical tests of this sort to estimate and make inference of the instantaneous volatility via the quadratic variation. Indeed, there is a well developed literature using the so called “realized variance” to estimate the quadratic variation using real (discrete) data; see for instance, among many others, Barndorff-Nielsen and Shephard (2002a), Barndorff-Nielsen and Shephard (2002b), Barndorff-Nielsen and Shephard (2004), Andersen, Bollerslev, Diebold, and Ebens (2001) and Andersen, Bollerslev, Diebold, and Labys (2003).

B Assorted Remarks

B.1 Notations

Throughout this article, we will use enforce the following notations. We will use e to denote the effort parameter / process, and use \mathbf{e} to denote the exponential function. If $\{S_t\}_{t \geq 0}$ is a stochastic process, we will either use S to denote the process or also $\{S_t\}$. We will also define and denote the indicator function as $\mathbf{1}_A(x)$, which equals 1 if $x \in A$ and 0 if $x \notin A$. If E is a set, we will denote E^c as its complement. Given an action process $A = \{e_t, \sigma_t\}$, we will interchange the notation μ_t and $\kappa(e_t, \sigma_t)$ to denote the drift part of the cash flow process. Hence, we will also with some abuse of notation, also call and denote $A = \{(\mu_t, \sigma_t)\}_{t \geq 0}$ as the action process, with the understanding that $\mu_t \equiv \kappa(e_t, \sigma_t)$.

B.2 Induced probability measure

Fix an action process $A = \{(\mu_t, \sigma_t)\}_{t \geq 0}$ and recall that the action process A is $\{\mathcal{F}_t\}$ -adapted. Consider the stochastic process $Y^A =: Y = \{Y_t\}_{t \geq 0}$ with dynamics as given in (4.1) and (4.2), which clearly and explicitly depend on the action process A . Define \mathbb{P}^A to be the *law* of the stochastic process $Y = Y^A$, and likewise, \mathbb{E}^A the expectation under \mathbb{P}^A . Observe that the definition (and indeed existence) of this law \mathbb{P}^A is possible through the deep and highly technical Kolmogorov Extension Theorem (see Durrett (2010, Theorem A.3.1) or Revuz and Yor (2005, Chapter 1, §3) for more details). Note here that we write P to be the probability measure on (Ω, \mathcal{F}) and likewise, E as the expectation under P .

While the law of a stochastic process is deep measure theoretic concept, but the intuition can more or less be grasped when we have a single random variable (i.e. the trivial stochastic process with a singleton time index). Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space and let $X : \Omega \rightarrow \mathbb{R}$ be a real-valued random variable. We often call Ω the *sample space* and \mathbb{R} here as the *state space* of the random variable X . And suppose X is integrable so that the expectation is well defined,

$$\mathbb{E}_{\mathbb{P}}[X] := \int_{\Omega} X(\omega) d\mathbb{P}(\omega).$$

Note in particular that we are integrating over the sample space Ω . However, in a lot of probability applications, it is far easier and more intuitive to work with the state space of the random variable (of which we may know something about its distribution function) than the sample space. By effectively a change of variables, we can integrate over the state space \mathbb{R} of the random variable. That is, we may define the *law* (or more commonly, the *distribution*) \mathbb{P}_X of the random variable X as,

$$\mathbb{P}_X(B) := \mathbb{P}(X^{-1}(B)), \quad B \in \mathcal{B}(\mathbb{R}),$$

where $\mathcal{B}(\mathbb{R})$ is the Borel set of \mathbb{R} . Then under this definition, we see that,

$$\mathbb{E}_{\mathbb{P}}[X] = \int_{\Omega} X(\omega) d\mathbb{P}(\omega) = \int_{\mathbb{R}} x d\mathbb{P}_X(x) = \mathbb{E}_{\mathbb{P}_X}[X].$$

The idea of the law under stochastic process is essentially the above consideration, except that we have to work through with finite dimensional joint distributions and cylinder sets to make the definition of the law of a stochastic process rigorous and precise.

B.3 Single Brownian motion

In (4.1) and (4.2), we use a single Brownian motion B , rather than say two different Brownian motions. In particular, this might be surprising coming from an asset pricing perspective; in asset pricing applications, say like the classical

Heston (1993) model, if S is the price of an asset, then it has say the dynamics,

$$\begin{aligned} dS_t &= \mu_S(t, S_t)dt + \sigma_S(t, \sigma_t, \nu_t)dB_{1t} \\ d\nu_t &= \mu_\nu(t, \nu_t)dt + \sigma_\nu(t, \nu_t)dB_{2t}, \end{aligned}$$

where B_1 and B_2 are correlated Brownian motions (with possibly zero correlation). In this type of specification, we see that the stochastic volatility dynamics of price S are also further driven by the process ν . In contrast, the specification of (4.1) and (4.2) uses the same Brownian motion. Indeed, in our specification, if we were to use two different Brownian motions, then economically, then there is little hope for an equilibrium. Economically, recall here the principal can only observe a single source of information (i.e. the cash flow Y_t over time), but if there are two sources of risk (i.e. two Brownian motions), then the agent has far too much room to deviate from the principal's recommended actions. Indeed, we will see the importance of using a single Brownian motion in Lemma 7.3.

B.4 Stochastic Time Change

We can actually view (4.1) and (4.2) as a time-changed process. We will not use this fact elsewhere in the paper. Since we can write (4.1) and (4.2) as $dY_t = \kappa(e_t, \sigma_t)dt + \sigma_t dm_t$, but since m is a geometric Brownian motion, by the Dambis-Dubins-Schwartz theorem³³ and by expanding the filtration and changing the probability space if needed, there exists a stochastic time change $\{T(t)\}_{t \geq 0}$ given by $T(t) := \inf\{u : \int_0^u \exp\{-v + 2B_v\}dv > t\}$, and another Brownian motion Z , such that $m_t = Z_{T(t)}$. Hence, we can rewrite the cash flows as,

$$dY_t = \kappa(e_t, \sigma_t)dt + \sigma_t dZ_{T(t)}.$$

Indeed, more is true. Since m is a geometric Brownian motion, by Lamperti's relation³⁴, we can further write m as a time-changed squared Bessel process

C Proofs of Section 4

Proof of Example 4.1. Let's quickly verify that the conditions in Definition 4.1 are satisfied.

- (a) $\kappa(e_H, \sigma) = \alpha_1 e^{\alpha_0(e_H - e_L)} \log \sigma > \alpha_1 e^{\alpha_0(e_L - e_L)} \log \sigma = \alpha_1 \log \sigma$, which holds since clearly $\alpha_0(e_H - e_L) > 0$, which implies $e^{\alpha_0(e_H - e_L)} > 1$.
- (b) $\kappa_\sigma(e, \sigma) = \alpha_1 e^{\alpha_0(e - e_L)} \frac{1}{\sigma} > 0$, and $\kappa_{\sigma\sigma}(e, \sigma) = -\alpha_1 e^{\alpha_0(e - e_L)} \frac{1}{\sigma^2} < 0$.
- (c) (i) If $\sigma = \sigma'$, then we return back to case (a);
(ii) If $\sigma > \sigma'$, then $\kappa(e_H, \sigma) = \alpha_1 e^{\alpha_0(e_H - e_L)} \log \sigma > \alpha_1 \log \sigma' = \kappa(e_L, \sigma')$ clearly holds since $e^{\alpha_0(e_H - e_L)} > 1$;
(iii) If $\sigma < \sigma'$, then $\kappa(e_H, \sigma) = \alpha_1 e^{\alpha_0(e_H - e_L)} \log \sigma > \alpha_1 \log \sigma' = \kappa(e_L, \sigma')$ will hold if and only if $e^{\alpha_0(e_H - e_L)} > \frac{\log \sigma}{\log \sigma'}$ holds. And since $\sigma < \sigma'$, clearly $1 > \frac{\log \sigma}{\log \sigma'}$. But since $e^{\alpha_0(e_H - e_L)} > 1$, then we are done.
- (d) Observe that,

$$\begin{aligned} \kappa(e_H, \sigma) - \kappa(e_L, \sigma_H) &< \sigma \kappa_\sigma(e_H, \sigma) \\ \iff \alpha_1 e^{\alpha_0(e_H - e_L)} \log \sigma - \alpha_1 \log \sigma_H &< \alpha_1 e^{\alpha_0(e_H - e_L)} \frac{1}{\sigma} \\ \iff \log \sigma - e^{-\alpha_0(e_H - e_L)} \log \sigma_H &< \frac{1}{\sigma}. \end{aligned}$$

But we note that if for all $\sigma \in [\sigma_L, \sigma_H]$ are such that $1/\sigma \geq \log \sigma$, then clearly,

$$1/\sigma \geq \log \sigma > \log \sigma - e^{-\alpha_0(e_H - e_L)} \log(\sigma_H + K),$$

holds. Note that $1/\sigma > \log \sigma$ is clearly equivalent to $1/\sigma - \log \sigma > 0$, and when we view the left-hand side of the equality as a map on \mathbb{R}_+ , $f(y) := 1/y - \log y$, f is strictly decreasing, continuous, $\lim_{y \rightarrow 0^+} f(y) = +\infty$, and

³³See Revuz and Yor (2005, Chapter V, §1, Theorem 1.6) for the precise statement.

³⁴See Revuz and Yor (2005, Chapter XI, §1, Exercise 1.28).

$\lim_{y \rightarrow +\infty} f(y) = 0$. Thus by the intermediate value theorem, there exists some $\bar{y} \in (0, \infty)$ such that $f(\bar{y}) = 0$. Note also, we can view $1/\sigma > \log \sigma$ equivalently as $1 > (\log \sigma)e^{\log \sigma}$.

Now, we must investigate the equation $z = W(z)e^{W(z)}$, where W is the complex valued *Lambert W function*³⁵, and z is also complex (note, even though the Lambert W function is a map from the complex numbers to the complex numbers, for this example, we just need the real parts). Our case of interest is when $z = 1$ and σ solves $\log \sigma = W(1)$, or that $\sigma = e^{W(1)}$. That is, we need to find the positive real number $W(1) = x$ (it is indeed unique) such that $1 = xe^x$ and numerically, we have that $x = W(1) \approx 0.5671$, which implies $e^x = e^{W(1)} \approx 1.763$.

Hence, if we consider the set of volatility controls as the interval $[\sigma_L, \sigma_H] = [c, e^{W(1)}] \approx [c, 1.763]$, where $\sigma_H > c > 1$, then the condition $1/\sigma - \log \sigma > 0$ will hold for any $\sigma \in [\sigma_L, \sigma_H]$.

(e) Condition (d) of Definition 4.1 holds immediately by hypothesis. □

Remark C.1. One would also wonder why the range $\sigma_H - \sigma_L$ of the volatility control set $[\sigma_L, \sigma_H]$ in Example 4.1 to be relatively numerically “small”. However, this is just due to the choice of the log function in defining $\kappa(e, \sigma)$ for this example. Essentially, the log function grows too slowly in σ , even though it retains all of the required properties in Definition 4.1. If we were to choose another function that satisfies the properties of Definition 4.1 but one that grows much quicker in σ , we will widen the numerical range of the volatility control set $[\sigma_L, \sigma_H]$. As well, we might be concerned with how small numerically the maximum σ_H controlled volatility value can take on. This is especially in contrast to the overall instantaneous diffusion term $\sigma_t m_t$ of dY_t , where m_t could be relatively large. This might beg the question how “meaningful” is volatility control in light of an exogenous volatility term m_t that could “swamp” the endogenous volatility control of the agent. However, by inspection of the reward function form in Example 4.1, we see that we can translate the values of volatility control set $[\sigma_L, \sigma_H]$ as it is written to a form $[\sigma_L + K, \sigma_H + K]$, where $K > 0$ is some deterministic constant, and thereby enlarging the absolute values of the volatility controls.

D Proofs of Section 6

First let’s define,

$$V_t(A) := \mathbb{E}^A \left[\int_0^\tau e^{-r_0 s} \left(dX_s + \left[\phi_e \left(1 - \frac{e_s}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_s}{\sigma_L} - 1 \right) \right] ds \right) + e^{-r_0(\tau-t)} R \mid \mathcal{F}_t^Y \right] \quad (\text{D.1})$$

It is important to note that both (6.1) and (D.1), which represents the continuation value of the agent at time t if the action process A is being taken, are conditioned on the filtration $\{\mathcal{F}_t^Y\}$. That is, the information observable to the principal. In particular, the filtration that is being conditioned on is not a Brownian one. This is a key and important departure from the usual papers in continuous-time principal agent problems. In providing incentives to the agent, since the principal can only observe the cash flows Y , this implies the continuation value of the agent, from the perspective of the principal, can only condition on the information $\{\mathcal{F}_t^Y\}$ generated by the cash flows Y , and hence (6.1) and (D.1) have the correct conditioning.

Remark D.1. Consider the typical setup of DeMarzo and Sannikov (2006) in the form (D.2),

$$dY_t = \mu_t dt + \sigma dB_t. \quad (\text{D.2})$$

For a fixed recommended action $\{\mu_t\}_{t \geq 0}$, and since the principal observes the cash flows $\{Y_t\}_{t \geq 0}$, by simply rearranging terms, we see that,

$$\frac{dY_t - \mu_t dt}{\sigma} = dB_t.$$

Indeed, this is the key step to the analysis of both DeMarzo and Sannikov (2006) and Sannikov (2008). And so, in this case, once the action process $\{\mu_t\}_{t \geq 0}$ is fixed³⁶, then the left-hand side is completely observable by the principal. And thus, in these types of setup, the information set available to the principal $\{\mathcal{F}_t^Y\}_{t \geq 0}$ is exactly identical to the

³⁵For details, see Wright (1959). This W notation is *not* to be confused with our subsequent treatment of the agent’s continuation value process. We retain the notation W for the Lambert W function out of convention.

³⁶More precisely, this means we work under the induced probability measure \mathbb{P}^A , where $A = \{\mu_t\}_{t \geq 0}$.

Brownian information set $\{\mathcal{F}_t\}_{t \geq 0}$, and so we do achieve $\mathbb{E}^A[\cdot | \mathcal{F}_t^Y] = \mathbb{E}[\cdot | \mathcal{F}_t]$. That is, in words, if the principal knows the action process and can observe the cash flows, that means he must also know the Brownian motion noise.

However, in this current setup of (4.1) and (4.2), this is clearly not the case. In particular, even if the agent knows the action process and the cash flows, he does not know the Brownian motion noise, and so we have a clear inequality, $\mathbb{E}^A[\cdot | \mathcal{F}_t^Y] \neq \mathbb{E}^A[\cdot | \mathcal{F}_t]$. To see this, observe that even if we repeat the above rearranging computation,

$$dY_t - \mu_t dt = \sigma_t M_t dB_t,$$

where specifically, we cannot “divide” over the generically not constant over time volatility choice σ_t (i.e. simply write out the above SDE in it’s integrated form). Recall also that m is not observable to the principal. And from this expression, we see that even if an action process $A = \{(\mu_t, \sigma_t)\}_{t \geq 0}$ is known to the principal, his information set \mathcal{F}_t^Y cannot equal to the information set generated by Brownian motion \mathcal{F}_t .

D.1 A trivial rewriting

In light of Remark D.1, it motivates for the following rewriting. The idea is to *not* think of noise driven by Brownian motion but rather driven by a more general continuous martingale process. In particular, observe, trivially, from (4.1) and (4.2), we can write,

$$dY_t = \mu_t dt + \sigma_t dM_t. \tag{D.3}$$

In particular, note that from (4.2), it is a geometric Brownian motion with zero drift and unit volatility. Thus, we have an explicit solution,

$$M_t = M_0 \exp\left\{-\frac{1}{2}t + B_t\right\}, \quad t \geq 0. \tag{D.4}$$

For all the proofs that follow, we take, without loss of generality that $M_0 \equiv 1$; the proof goes through with a generic $M_0 = m_0$ but we just have to carry more algebra.

Proposition D.2. *For a fixed contract (A, X, τ) , the stochastic process $t \mapsto \int_0^t \sigma_s dm_s$ is an $(\mathbb{P}^A, \{\mathcal{F}_t^Y\})$ -martingale.*

Proof. First, from (D.3) and again since the action process A is held fixed, we write,

$$Y_t - \mu_t dt = \sigma_t dM_t. \tag{D.5}$$

From (D.4), we see immediately that M is a true $(P, \{\mathcal{F}_t\})$ -martingale (i.e. the Dolean exponential with respect to Brownian motion). Since σ_t is $\{\mathcal{F}_t\}$ -adapted and since m is a P -square integrable continuous martingale, then this implies that $t \mapsto \int_0^t \sigma_s dM_s$ is also a $(P, \{\mathcal{F}_t\})$ -martingale.

Now, let’s show that $t \mapsto \int_0^t \sigma_s dM_s$ is also a $(P, \{\mathcal{F}_t^Y\})$ -martingale. Observe that $t \mapsto \int_0^t \sigma_s dM_s$ is \mathcal{F}_t^Y -adapted. But this is immediate by viewing the left-hand side of (D.5) and recalling footnote 18. It remains to verify the martingale property. Pick any time $t_1 \geq t_0$. Since $t \mapsto \int_0^t \sigma_s dM_s$ is a martingale, then we immediately have that,

$$E\left[\int_{t_0}^{t_1} \sigma_s dM_s \mid \mathcal{F}_{t_0}\right] = 0.$$

But by applying the Law of Iterated Expectations,

$$E\left[\int_{t_0}^{t_1} \sigma_s dM_s \mid \mathcal{F}_{t_0}^Y\right] = E\left[E\left[\int_{t_0}^{t_1} \sigma_s dM_s \mid \mathcal{F}_{t_0}\right] \mid \mathcal{F}_{t_0}^Y\right] = 0,$$

so the martingale property holds for $\{\mathcal{F}_t^Y\}$. Finally, the fact that $t \mapsto \int_0^t \sigma_s dM_s$ is a $(\mathbb{P}^A, \{\mathcal{F}_t^Y\})$ -martingale follows from the discussion in footnote 15. \square

This trivial rewriting of (D.3) and observation of Proposition D.2 that $\sigma_t dM_t$ is an $(\mathbb{P}^A, \{\mathcal{F}_t^Y\})$ -martingale are the key steps to the beginning of our analysis. Indeed, to emphasize the importance of the observation in Proposition D.2, we will in the subsequent few sections, denote explicitly the $(\mathbb{P}^A, \{\mathcal{F}_t^Y\})$ -martingale as,

$$\sigma_t dM_t^A = Y_t - \mu_t dt. \tag{D.6}$$

D.2 Martingale Representation Theorem

With Proposition D.2 in mind and also the cash flow process in the form (D.3), we can now state the following key proposition.

Proposition D.3. *Fix a contract (A, X, τ) . Then there exist processes $\{\beta_t\}_{t \geq 0}$ and $\{V_t^{\perp, A}\}_{t \geq 0}$ such that the dynamics of $V_t(A)$ in (D.1) can be written as,*

$$dV_t(A) = e^{-r_0 t} \beta_t \sigma_t dM_t^A + dV_t^{\perp, A}, \quad (\text{D.7})$$

where β_t is some predictable process such that $\int_0^t (e^{-r_0 s} \beta_s)^2 d[\int_0^\cdot \sigma_u dM_u^A]_s < \infty$, \mathbb{P}^A -a.s., and $V_t^{\perp, A}$ is continuous and orthogonal to $t \mapsto \int_0^t \sigma_s dM_s^A$ (i.e. meaning, $[V^{\perp, A}, \int_0^\cdot \sigma_s dM_s^A]_t = 0$, \mathbb{P}^A -a.s.), and $V_0^{\perp, A} = 0$.

Proof. It is easy to verify that (D.1) is a $(\mathbb{P}^A, \{\mathcal{F}_t^Y\})$ -martingale (i.e. Doob's martingale). Then the result is an immediate consequence of Proposition D.2 and the (general) martingale representation theorem. For the precise statement for this general martingale representation theorem result, please see Hunt and Kennedy (2004, Chapter 5, Theorem 5.37), Protter (2005, Chapter IV, Section 3, Corollary 1) and Revuz and Yor (2005, Chapter V, Section 4, Lemma 4.2). \square

Remark D.4. In the term $e^{-r_0 t} \beta_t \sigma_t dM_t^A$, the time discount factor term $e^{-r_0 t}$ is merely a convenient normalization; this is also done in Sannikov (2008). Also, strictly speaking, β is clearly dependent on the choice of the action process A but will suppress it for notational convenience when the context is clear. We keep the notation A on the orthogonal process $V_t^{\perp, A}$ as the choice of A will make a meaningful difference in the subsequent discussions.

Compared to the papers like Holmström and Milgrom (1987), DeMarzo and Sannikov (2006), Sannikov (2008), and others, they all invoke a martingale representation theorem for the case when the filtration is generated by Brownian motion (i.e. see Karatzas and Shreve (1991, Chapter 3, Theorem 4.2), among others); recall again the discussion in Remark D.1 on why a Brownian filtration setup here is inappropriate. Specifically in the Brownian case, the orthogonal term, denoted above as $V_t^{\perp, A}$, would be identically zero. It is also worth noting that in this line of continuous-time principal-agent literature, there have been some notable cases where the filtration is not Brownian. For instance, in Sannikov (2007, Proposition 1), the “extra” orthogonal term is interpreted as a public randomization device.

D.3 Dynamics of the agent's continuation value

With Proposition D.3 in mind, the following is an easy application of Ito's lemma.

Proof of Theorem 6.3. From (D.1), we can write,

$$V_t(A) = \int_0^t e^{-r_0 s} \left(dX_s + \left[\phi_e \left(1 - \frac{e_s}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_s}{\sigma_L} - 1 \right) \right] \right) ds + e^{-r_0 t} W_t(A). \quad (\text{D.8})$$

Applying Ito's lemma, we obtain,

$$dV_t(A) = e^{-r_0 t} \left(dX_t + \left[\phi_e \left(1 - \frac{e_t}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt \right) + \underbrace{d(e^{-r_0 t} W_t(A))}_{-r_0 e^{-r_0 t} W_t dt + e^{-r_0 t} dW_t}. \quad (\text{D.9})$$

Equating (D.7) with (D.9), we obtain,

$$\begin{aligned} & e^{-r_0 t} \beta_t \sigma_t dM_t^A + dV_t^{\perp, A} \\ &= e^{-r_0 t} \left(dX_t + \left[\phi_e \left(1 - \frac{e_t}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt \right) - r_0 e^{-r_0 t} W_t(A) dt + e^{-r_0 t} dW_t(A) \end{aligned} \quad (\text{D.10})$$

Defining $e_t^{\perp, A} := \int_0^t e^{r_0 s} dV_s^{\perp, A}$, rearranging, and recalling that $\sigma_t dM_t = Y_t - \mu_t dt$, we obtain,

$$dW_t(A) = r_0 W_t(A) dt - \left(dX_t + \left[\phi_e \left(1 - \frac{e_t}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt \right) + \beta_t (dY_t - \mu_t dt) + de_t^{\perp, A}, \quad (\text{D.11})$$

and we are done. \square

D.3.1 Checking for Deviations

With Theorem 6.3 on hand, we are now ready to give a simple condition (hopefully that is both sufficient and necessary) to pin down the incentive compatibility constraints of the agent. Fix two action processes ³⁷ $A = \{(\mu_t, \sigma_t)\}_{t \geq 0}$ and $A^\dagger = \{(\mu_t^\dagger, \sigma_t^\dagger)\}_{t \geq 0}$.

If the agent plays A , then the agent's time zero continuation value is $W_0(A)$ as in (6.1). But suppose the agent deviates to A^\dagger . Specifically, the cash flow processes under the two different action processes evolve as ³⁸,

$$\text{Under } \mathbb{P}^A: \quad dY_t = \mu_t dt + \sigma_t dM_t^A. \quad (\text{D.12})$$

$$\text{Under } \mathbb{P}^{A^\dagger}: \quad dY_t = \mu_t^\dagger dt + \sigma_t^\dagger dM_t^{A^\dagger}. \quad (\text{D.13})$$

Phrased in this light, this strongly calls for a change-of-measure type analysis. To do so, we need to invoke a stronger version of Girsanov's theorem, which is usually applied in a Brownian setting. Here, we will use the slightly more general Girsanov-Meyer theorem. We will reiterate it here for reference:

Girsanov-Meyer Theorem. *Let P and Q be equivalent measures. Let X be a continuous (classical) semimartingale under P with decomposition $X = M + A$. Then X is also a continuous (classical) semimartingale under Q and has decomposition $X = L + C$, where*

$$L_t = M_t - \int_0^t \frac{1}{Z_s} d[Z, M]_s \quad (\text{D.14})$$

is a Q local martingale, and $C = X - L$ is a Q finite variation process.

Proof. For details of the theorem and its proof, please see Protter (2005, Chapter III, Section 8, Theorem 39). \square

Using the Girsanov-Meyer theorem as a guide, and we can set $P = \mathbb{P}^A$ and $Q = \mathbb{P}^{A^\dagger}$, then it is natural to set,

$$A_t = \int_0^t \mu_s ds, \quad M_s = \int_0^t \sigma_s dM_s^A, \quad (\text{D.15})$$

and,

$$C_t = \int_0^t \mu_s^\dagger ds, \quad L_s = \int_0^t \sigma_s^\dagger dM_s^{A^\dagger}. \quad (\text{D.16})$$

Then, to ensure that the correct change of measure is possible, it remains to identify the process Z .

In particular, we define the Radon-Nikodym derivative as,

$$Z_t = \mathbb{E}^A \left[\frac{d\mathbb{P}^{A^\dagger}}{d\mathbb{P}^A} \mid \mathcal{F}_t^Y \right]. \quad (\text{D.17})$$

To continue the discussion, we will need the following mild technical assumption that we have an appropriate kernel.

Assumption D.5. *Suppose associated with the Radon-Nikodym derivative in (D.17), there exists a square integrable process $\{\varphi_t\}_{t \geq 0}$ such that,*

$$dZ_t = \varphi_t Z_t \sigma_t dM_t^A. \quad (\text{D.18})$$

With Assumption D.5 on hand, and letting $N_t := \varphi_t \sigma_t dM_t^A$, we see that,

$$Z_t = \mathcal{E}(N)_t, \quad (\text{D.19})$$

the Dolean's exponential ³⁹ for the process N . For now, let's suppose that Z is a true $\{\mathcal{F}_t^Y\}$ -martingale but we will

³⁷Again, recalling the notation convention in footnote 12, what we mean here is to fix two action processes $A = \{(e_t, \sigma_t)\}_{t \geq 0}$ and $A^\dagger = \{(e_t^\dagger, \sigma_t^\dagger)\}_{t \geq 0}$, then define $\mu_t \equiv \kappa(e_t, \sigma_t)$ and $\mu_t^\dagger \equiv \kappa(e_t^\dagger, \sigma_t^\dagger)$.

³⁸To be precise, we mean the following and please see Proposition D.2 again. Under \mathbb{P}^A , we have the $(\mathbb{P}^A, \{\mathcal{F}_t^Y\})$ -martingale $\sigma_t dM_t^A = dY_t - \mu_t dt$; and under \mathbb{P}^{A^\dagger} , we have the $(\mathbb{P}^{A^\dagger}, \{\mathcal{F}_t^Y\})$ -martingale $\sigma_t^\dagger dM_t^{A^\dagger} = dY_t - \mu_t^\dagger dt$. Rearrange to obtain the two displayed equations.

³⁹The Dolean's exponential $\mathcal{E}(N)_t$ for a continuous local martingale S is given by,

$$\mathcal{E}(S)_t := \exp \left\{ S_t - S_0 - \frac{1}{2} [S]_t \right\}.$$

verify this in the subsequent discussion ⁴⁰. Thus, it remains to find φ_t . Note in the above, we have two expressions for L , (D.14) as given in the Girsanov-Meyer theorem statement, and also in (D.16). Equating these two expressions for L , we find,

$$L_t = \int_0^t \sigma_s dM_s^A - \int_0^t \varphi_s d \left[\int_0^\cdot \sigma_u dM_u^A \right]_s = \int_0^t \sigma_s^\dagger dM_s^{A^\dagger}, \quad (\text{D.20})$$

and since $dM_t^A = M_t^A dB_t^A$ then we immediately have by the quadratic variation of Brownian motion,

$$d \left[\int_0^\cdot \sigma_u dM_u^A \right]_t = \sigma_t^2 (M_t^A)^2 dt.$$

Rewriting everything in differential form and substituting, we thus have that,

$$\sigma_t dM_t^A - \varphi_t \sigma_t^2 (M_t^A)^2 dt = \sigma_t^\dagger dM_t^{A^\dagger}.$$

But recall again that $\sigma_t dM_t^A = dY_t - \mu_t dt$ and $\sigma_t^\dagger dM_t^{A^\dagger} = dY_t - \mu_t^\dagger dt$, so we substitute,

$$dY_t - \mu_t dt - \varphi_t \sigma_t^2 (M_t^A)^2 dt = dY_t - \mu_t^\dagger dt.$$

Canceling terms and equating, we have thus,

$$\int_0^t \varphi_s \sigma_s^2 M_s^2 ds = \int_0^t (\mu_s^\dagger - \mu_s) ds.$$

But since the integrands on the left-hand side and the right-hand side are well bounded, this immediately implies the integrands must equal. And thus, we have that the Girsanov kernel is,

$$\varphi_t = \frac{\mu_t^\dagger - \mu_t}{\sigma_t^2 (M_t^A)^2}. \quad (\text{D.21})$$

Now in particular, we have the following important relationship.

Lemma D.6. *Fix a contract (A, X, τ) . Suppose the agent considers the recommended action process A and fixes another action process A^\dagger . The agent considers a deviation from A to A^\dagger . Then the noise terms are related in the following manner:*

$$\sigma_t dM_t^A = (\mu_t^\dagger - \mu_t) dt + \sigma_t^\dagger dM_t^{A^\dagger}. \quad (\text{D.22})$$

Proof. Using the Girsanov kernel in (D.21), simply substitute,

$$\begin{aligned} \sigma_t dM_t^A - \varphi_t \sigma_t^2 (M_t^A)^2 dt &= \sigma_t dM_t^A - \frac{(\mu_t^\dagger - \mu_t)}{\sigma_t^2 (M_t^A)^2} \sigma_t^2 (M_t^A)^2 dt \\ &= \sigma_t^\dagger dM_t^{A^\dagger}. \end{aligned}$$

Rearrange, and we get (D.22). □

D.3.2 Change of Measure and Novikov's Criterion

However, an important task remains — we need to verify that the Dolean's exponential $\mathcal{E}(N)$ for N is a true martingale to allow for a valid change of measure. With the Girsanov kernel given in (D.21), and recalling again that $dM_t^A = M_t^A dB_t^A$, we substitute back to see that,

$$dN_t = \frac{\mu_t^\dagger - \mu_t}{\sigma_t (M_t^A)^2} M_t^A dB_t^A = \frac{\mu_t^\dagger - \mu_t}{\sigma_t M_t^A} dB_t^A$$

But this implies N has quadratic variation,

$$[N]_t = \int_0^t \frac{(\mu_s^\dagger - \mu_s)^2}{\sigma_s^2 (M_s^A)^2} ds.$$

⁴⁰Of course, with the equivalence of \mathbb{P}^A and \mathbb{P}^{A^\dagger} , there is no need to specify for which probability measure is Z a martingale.

In all, that means to ensure that $\mathcal{E}(N)$ is a true martingale, a sufficient condition is to ensure that the Novikov's criterion ⁴¹ holds. For what follows, consider a fixed deterministic time horizon $T < \infty$. The *Novikov's criterion* requires,

$$\mathbb{E}^A \left[\exp \left\{ \frac{1}{2} \int_0^T \frac{(\mu_t^\dagger - \mu_t)^2}{\sigma_t^2 (M_t^A)^2} dt \right\} \right] < \infty. \quad (\text{D.23})$$

Lemma D.7. *The Novikov criterion (D.23) holds for finite time horizon $T < \infty$ and infinite time horizon $T = \infty$. And thus, $\mathcal{E}(N)_t$ is a true martingale for all $t \in [0, \infty]$.*

Proof. Recall again that M^A itself is a geometric Brownian motion with zero drift and unit volatility. Thus, we can immediately write $M_t^A = \exp \{-\frac{1}{2}t + B_t^A\}$; we have taken, without loss of generality $M_0 = 1$. Furthermore, since we know that $\mu_t, \mu_t^\dagger \in [\mu_L, \mu_H]$ and likewise $\sigma_t, \sigma_t^\dagger \in [\sigma_L, \sigma_H]$, then we clearly have,

$$|\mu_t^2 - \mu_t^{\dagger 2}| \leq 2\mu_H^2, \quad \frac{1}{\sigma_L^2} \geq \frac{1}{\sigma_t^2}, \quad \frac{1}{\sigma_L^2} \geq \frac{1}{(\sigma_t^\dagger)^2}.$$

Substituting,

$$\begin{aligned} \mathbb{E}^A \left[\exp \left\{ \frac{1}{2} \int_0^T \frac{(\mu_t^\dagger - \mu_t)^2}{\sigma_t^2 (M_t^A)^2} dt \right\} \right] &\leq \mathbb{E}^A \left[\exp \left\{ \frac{1}{2} \frac{2\mu_H^2}{\sigma_L^2} \int_0^T \frac{1}{(M_t^A)^2} dt \right\} \right] \\ &= e^{\mu_H^2/\sigma_L^2} \mathbb{E}^A \left[\exp \left\{ \int_0^T \frac{1}{(M_t^A)^2} dt \right\} \right] \\ &= e^{\mu_H^2/\sigma_L^2} \mathbb{E}^A \left[\exp \left\{ \int_0^T e^{t-2B_t^A} dt \right\} \right] \end{aligned}$$

Thus, the problem reduces now to proving that, $\mathbb{E}^A \left[\exp \left\{ \int_0^T e^{t-2B_t^A} dt \right\} \right] < \infty$. It should be noted that this is a non-trivial problem since we essentially have to show that the expectation of the exponential of an integrated geometric Brownian motion is finite. As noted from Yor (1992), the aforementioned problem is essentially the same as investigating the properties of,

$$\int_0^t e^{aB_s + bs} ds, \quad a, b \in \mathbb{R}.$$

But by scaling properties of Brownian motion ⁴² B , it suffices to consider the process,

$$A_t^{(\nu)} := \int_0^t e^{2(B_s + \nu s)} ds, \quad \nu \in \mathbb{R}. \quad (\text{D.24})$$

Hence, to solve our problem of showing (D.23), it is equivalent to showing that,

$$\mathbb{E} \left[\exp \left\{ A_t^{(\nu)} \right\} \right] < \infty. \quad (\text{D.25})$$

But to show that (D.25) is finite is equivalent to showing that the Laplace transform (moment generating function) of $A_t^{(\nu)}$ is well defined and finite. But using Yor (1992, Equation (7.e)) as pointed out by Kim (2004) (see also Albanese and Lawi (2005)), we are ensured that the aforementioned Laplace transform is well defined and finite.

Thus, this implies that (D.23) does indeed hold for each finite $T < \infty$. However, given in this model, we allow for a termination time τ that could be finite (i.e. terminating the agent at some time) or infinite (i.e. never terminating the agent), considering the deterministic finite time case is insufficient. But invoking Revuz and Yor (2005, Chapter VIII, §1, Corollary 1.16), we can extend the discussion from finite time interval $[0, T]$ for $T < \infty$ to $[0, \infty]$. Thus, this shows that $\mathcal{E}(N)_t$ is a martingale for all times $t \in [0, \infty]$. \square

The following lemma will be useful when we further characterize the effects of deviation in the subsequent discussion.

Lemma D.8. *Fix a contract (A, X, τ) . Consider the setup and the process $\epsilon^{\perp, A}$ as defined in Theorem 6.3. Fix another action process A^\dagger . The stochastic process $\epsilon^{\perp, A}$ is a $\{\mathcal{F}_t^Y\}$ -martingale under both probability measures \mathbb{P}^A and \mathbb{P}^{A^\dagger} .*

⁴¹See Protter (2005, Chapter III, Section 8, Theorem 45), Revuz and Yor (2005, Proposition 1.15, Corollary 1.16) and also Duffie (2001, Appendix D)

⁴²For this discussion, it suffices to suppress the dependence on action process A .

Proof. The fact that $\epsilon^{\perp, A}$ is a $(\mathbb{P}^A, \mathcal{F}_t^Y)$ -martingale is immediate from the fact that by Proposition D.3, $V^{\perp, A}$ is a $(\mathbb{P}^A, \mathcal{F}_t^Y)$ martingale.

Thus, it remains to prove that $\epsilon_t^{\perp, A}$ is a $(\mathbb{P}^{A^\dagger}, \mathcal{F}_t^Y)$ -martingale. Clearly, this is equivalent to showing that $V_t^{\perp, A}$ is an (\mathbb{P}^{A^\dagger}) -martingale. But thanks to Lemma D.7, this is equivalent to showing that for any $s \geq t$, we have,

$$V_t^{\perp, A} = \mathbb{E}^{A^\dagger} \left[V_s^{\perp, A} | \mathcal{F}_t^Y \right] = \mathbb{E}^A \left[V_s^{\perp, A} \frac{d\mathbb{P}^{A^\dagger}}{d\mathbb{P}^A} \Big|_{\mathcal{F}_s^Y} \Big| \mathcal{F}_t^Y \right]$$

Hence, we see that it suffices to prove that the stochastic process as a product, $t \mapsto V_t^{\perp, A} \frac{d\mathbb{P}^{A^\dagger}}{d\mathbb{P}^A} \Big|_{\mathcal{F}_t^Y}$ is an \mathcal{F}_t^Y -martingale.⁴³ But we observe that $V_t^{\perp, A}$ is a square-integrable martingale and likewise for $d\mathbb{P}^{A^\dagger}/d\mathbb{P}^A|_{\mathcal{F}_t^Y} = \mathcal{E}(N)_t$. Then using the notion of strongly orthogonal martingales in, say, Protter (2005, Chapter IV, §3), to show that the product $V_t^{\perp, A} \mathcal{E}(N)_t$ is a martingale, it is equivalent to showing,

$$[V^{\perp, A}, \mathcal{E}(N)]_t \text{ is a uniformly integrable martingale.}$$

But observe that since $dV_t^{\perp, A}$ and $\sigma_t dM_t^A$ are orthogonal (i.e. recall Proposition D.3), then computing the quadratic covariation (here, for convenience, we use the differential notation),

$$\begin{aligned} dV_t^{\perp, A} d\mathcal{E}(N)_t &= dV_t^{\perp, A} \mathcal{E}(N)_t \varphi_t \sigma_t dM_t^A \\ &= \mathcal{E}(N)_t \varphi_t \underbrace{(dV_t^{\perp, A})(\sigma_t dM_t^A)}_{=0} \\ &= 0. \end{aligned}$$

Hence, we have that $[V^{\perp, A}, \mathcal{E}(N)]_t = 0$ (i.e. a constant stochastic process, which is a trivial martingale). Thus, we have that $V^{\perp, A}$ is also a martingale under \mathbb{P}^{A^\dagger} , in addition to being a martingale under \mathbb{P}^A . \square

D.3.3 Characterizing Deviations

Proof of Lemma 6.4. As before, fix a contract (A, X, τ) . And fix another action process A^\dagger . Consider a deviation from A to A^\dagger . Let's do some preliminary computations before showing the equivalence of (i) and (ii). Define,

$$\hat{V}_t := \int_0^t e^{-r_0 s} \left(dX_s + \left[\phi_e \left(1 - \frac{e_s^\dagger}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_s^\dagger}{\sigma_L} - 1 \right) \right] ds \right) + e^{-r_0 t} W_t(A), \quad (\text{D.26})$$

which is the time t expectation of the agent's total payoff if he experienced the cost of effort from the action process A^\dagger before time t , and plans to follow the recommended action process A after time t . Let's write the dynamics of \hat{V}_t under the measure \mathbb{P}^{A^\dagger} . Observe that in differential form, and noting the expression in (D.10), and equating,

$$\begin{aligned} d\hat{V}_t &= e^{-r_0 t} \left(dX_t + \left[\phi_e \left(1 - \frac{e_t^\dagger}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t^\dagger}{\sigma_L} - 1 \right) \right] dt \right) + d(e^{-r_0 t} W_t(A)) \\ &= e^{-r_0 t} \left(dX_t + \left[\phi_e \left(1 - \frac{e_t^\dagger}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t^\dagger}{\sigma_L} - 1 \right) \right] dt \right) \\ &\quad + e^{-r_0 t} \beta_t^A \sigma_t dm_t^A + dV_t^{\perp, A} - e^{-r_0 t} \left(dX_t + \left[\phi_e \left(1 - \frac{e_t}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt \right). \end{aligned} \quad (\text{D.27})$$

⁴³Note by the equivalence in measures, we no longer need to be explicit about the probability measure for which this is a martingale.

But using (D.22) of Lemma (D.6), and collecting terms, we can further rewrite (D.27) as,

$$\begin{aligned}
d\hat{V}_t &= e^{-r_0 t} \left(dX_t + \left[\phi_e \left(1 - \frac{e_t^\dagger}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t^\dagger}{\sigma_L} - 1 \right) \right] dt \right) \\
&\quad + e^{-r_0 t} \beta_t \left\{ [\kappa(e_t^\dagger, \sigma_t^\dagger) - \kappa(e_t, \sigma_t)] dt + \sigma_t^\dagger dm_t^{A^\dagger} \right\} + dV_t^{\perp, A} \\
&\quad - e^{-r_0 t} \left(dX_t + \left[\phi_e \left(1 - \frac{e_t}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt \right) \\
&= e^{-r_0 t} \left[-\phi_e \frac{e_t^\dagger - e_t}{e_H} + \phi_\sigma \frac{\sigma_t^\dagger - \sigma_t}{\sigma_L} + \beta_t [\kappa(e_t^\dagger, \sigma_t^\dagger) - \kappa(e_t, \sigma_t)] \right] dt \\
&\quad + dV_t^{\perp, A} + e^{-r_0 t} \beta_t \sigma_t^\dagger dM_t^{A^\dagger}
\end{aligned} \tag{D.28}$$

Thus, writing in integrated form for both (D.26) and (D.28), and equating,

$$\begin{aligned}
\hat{V}_t &= \hat{V}_0 + \int_0^t e^{-r_0 s} \left[-\phi_e \frac{e_t^\dagger - e_t}{e_H} + \phi_\sigma \frac{\sigma_t^\dagger - \sigma_t}{\sigma_L} + \beta_t [\kappa(e_t^\dagger, \sigma_t^\dagger) - \kappa(e_t, \sigma_t)] \right] ds \\
&\quad + V_t^{\perp, A} + \int_0^t e^{-r_0 s} \beta_s \sigma_s^\dagger dM_s^{A^\dagger}
\end{aligned} \tag{D.29}$$

$$= \hat{V}_0 + \int_0^t e^{-r_0 s} \left(dX_s + \left[\phi_e \left(1 - \frac{e_s^\dagger}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_s^\dagger}{\sigma_L} - 1 \right) \right] ds \right) + e^{-r_0 t} W_t(A). \tag{D.30}$$

Let's consider taking the time 0 expectation of \hat{V}_t in (D.29) above under \mathbb{P}^{A^\dagger} . From Lemma D.8, we have that $\mathbb{E}^{A^\dagger}[V_t^{\perp, A}] = 0$. Furthermore, note that the stochastic process $t \mapsto e^{-r_0 t} \beta_t \sigma_t^\dagger$ is an \mathcal{F}_t^Y square integrable martingale, since the integrator $m_t^{A^\dagger}$ is a square integrable martingale and the terms in the integrand are well bounded; see Protter (2005, Chapter IV, §2, Theorem 11). Thus, in expectation, the last two terms in the sum of (D.29) vanish. Thus, picking any two times $t \geq t_0 \geq 0$ in mind, observe that, under condition (i),

$$\begin{aligned}
&\mathbb{E}^{A^\dagger} \left[\hat{V}_t - \hat{V}_{t_0} \mid \mathcal{F}_{t_0}^Y \right] \\
&= \mathbb{E}^{A^\dagger} \left[\int_{t_0}^t e^{-r_0 s} \left[-\phi_e \frac{e_t^\dagger - e_t}{e_H} + \phi_\sigma \frac{\sigma_t^\dagger - \sigma_t}{\sigma_L} + \beta_t [\kappa(e_t^\dagger, \sigma_t^\dagger) - \kappa(e_t, \sigma_t)] \right] ds \mid \mathcal{F}_{t_0}^Y \right] \\
&\leq 0.
\end{aligned}$$

Specifically, this implies that \hat{V}_t is an \mathcal{F}_t^Y -supermartingale (under both probability measures \mathbb{P}^A and \mathbb{P}^{A^\dagger} , by Lemma D.7).

Now, let's show that (i) \implies (ii). The above supermartingale property implies, $\mathbb{E}^{\varepsilon A^\dagger}[\hat{V}_t] \leq \hat{V}_0 = W_0(A)$. Rewriting the left-hand side of the this inequality, and using (D.30), we have that,

$$\begin{aligned}
&\mathbb{E}^{A^\dagger} \left[\int_0^t e^{-r_0 s} \left(dX_s + \left[\phi_e \left(1 - \frac{e_s^\dagger}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_s^\dagger}{\sigma_L} - 1 \right) \right] ds \right) \right] + \mathbb{E}^{A^\dagger} [e^{-r_0 t} W_t(A)] \\
&\leq \hat{V}_0 \\
&= W_0(A) \\
&= \mathbb{E}^A \left[\int_0^\tau e^{-r_0 s} \left(dX_s + \left[\phi_e \left(1 - \frac{e_s}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_s}{\sigma_L} - 1 \right) \right] ds \right) + e^{-r_0 \tau} R \right].
\end{aligned}$$

Since $t \geq 0$ was arbitrary, set it to $t = \tau$ and note that $W_\tau(A) = R$ a.s. (both in \mathbb{P}^A and in \mathbb{P}^{A^\dagger}), then we have that,

$$\begin{aligned}
&\mathbb{E}^{A^\dagger} \left[\int_0^\tau e^{-r_0 s} \left(dX_s + \left[\phi_e \left(1 - \frac{e_s^\dagger}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_s^\dagger}{\sigma_L} - 1 \right) \right] ds \right) + e^{-r_0 \tau} R \right] \\
&\leq \mathbb{E}^A \left[\int_0^\tau e^{-r_0 s} \left(dX_s + \left[\phi_e \left(1 - \frac{e_s}{e_H} \right) + \phi_\sigma \left(\frac{\sigma_s}{\sigma_L} - 1 \right) \right] ds \right) + e^{-r_0 \tau} R \right]
\end{aligned}$$

Thus, if the recommended action process is A , then it is not optimal for the agent to deviate to A^\dagger . Thus, (i) \implies (ii) holds.

Let's show that (ii) \implies (i) holds. We will prove by contrapositive. Suppose that (i) does not hold on a set of non-zero measure (again, both under \mathbb{P}^A or \mathbb{P}^{A^\dagger}). Let's show that deviating away from A is optimal (i.e. A is suboptimal). On the set of times with non-zero measure such that (6.6) does not hold for some time s and some $(e, \sigma) \in \{e_L, e_H\} \times [\sigma_L, \sigma_H]$. But this implies we can construct an action process \tilde{A} such that there would exist some time t' such that,

$$\mathbb{E}^{\tilde{A}}[\hat{V}_{t'}] > \hat{V}_0 = W_0(A).$$

But since the agent gets utility $\mathbb{E}^{\tilde{A}}[\hat{V}_{t'}]$ if he follows \tilde{A} until time t' and switches to A , the action process A is suboptimal. This shows that (ii) \implies (i). \square

Proof of Corollary 6.5. Consider (6.7). To consider a deviation from the recommended action (e_t, σ_t) to the deviated action (e', σ') , that is $(e_t, \sigma_t) \neq (e', \sigma')$, we have three cases to consider:

- (i) $e_t \neq e', \sigma_t \neq \sigma'$;
- (ii) $e_t = e', \sigma_t \neq \sigma'$; and
- (iii) $e_t \neq e', \sigma_t = \sigma'$.

Case (i): Suppose $e_t \neq e'$ and $\sigma_t \neq \sigma'$. Let's prove the case when $e_t = e_H$. Then we must have that $e' = e_L$. So, we have,

$$0 \geq -\frac{\phi_e}{e_H}(e_L - e_H) + \frac{\phi_\sigma}{\sigma_L}(\sigma' - \sigma_t) + \beta_t [\kappa(e_L, \sigma') - \kappa(e_H, \sigma_t)].$$

But from Definition 4.1(c), we have $\kappa(e_L, \sigma') - \kappa(e_H, \sigma_t) < 0$. Rearranging the above, we have that,

$$\begin{aligned} \beta_t &\geq \frac{1}{\kappa(e_L, \sigma') - \kappa(e_H, \sigma_t)} \left[\frac{\phi_e}{e_H}(e_L - e_H) - \frac{\phi_\sigma}{\sigma_L}(\sigma' - \sigma_t) \right] \\ &= \frac{1}{\kappa(e_H, \sigma_t) - \kappa(e_L, \sigma')} \left[\frac{\phi_e}{e_H}(e_H - e_L) + \frac{\phi_\sigma}{\sigma_L}(\sigma' - \sigma_t) \right], \end{aligned} \quad (\text{D.31})$$

for all $\sigma' \in [\sigma_L, \sigma_H]$. But we observe immediately that the inequality (D.31) holds if and only if (6.7) holds. This shows the equivalence of (6.7) with (6.6) when $e_t = e_H$. The case of when $e_t = e_L$ is proved similarly.

Case (ii): Suppose that $e_t = e' = e$ and $\sigma_t \neq \sigma'$. Suppose first if $\sigma' > \sigma_t$, which implies $\kappa(e, \sigma') - \kappa(e, \sigma_t) > 0$. Then we have,

$$\frac{\phi_\sigma}{\sigma_L}(\sigma_t - \sigma') \geq \beta_t (\kappa(e, \sigma') - \kappa(e, \sigma_t)),$$

which implies that $0 > \beta_t$ — contradiction. That is to say, if the principal's recommended volatility is σ_t , the agent will not deviate to a higher volatility $\sigma' > \sigma_t$. Next, if $\sigma' < \sigma_t$, so $\kappa(e, \sigma') - \kappa(e, \sigma_t) < 0$, then we have,

$$\frac{\phi_\sigma}{\sigma_L}(\sigma_t - \sigma') \geq \beta_t (\kappa(e, \sigma') - \kappa(e, \sigma_t)),$$

which implies β_t is greater than or equal to some strictly negative term. But nonnegativity of β_t , this imposes no restriction on β_t .

Case (iii): Suppose $e_t \neq e'$, and $\sigma_t = \sigma' = \sigma$. Consider first the case when $e_t = e_H$, so $e' = e_L$, which implies $\kappa(e_L, \sigma) - \kappa(e_H, \sigma) < 0$. Then we have,

$$0 \geq -\frac{\phi_e}{e_H}(e_L - e_H) + \beta_t (\kappa(e_L, \sigma) - \kappa(e_H, \sigma)),$$

implying,

$$\beta_t \geq \frac{1}{\kappa(e_H, \sigma) - \kappa(e_L, \sigma)} \frac{\phi_e}{e_H}(e_H - e_L). \quad (\text{D.32})$$

But we also have that $\kappa(e_H, \sigma) - \kappa(e_L, \sigma) \geq \kappa(e_H, \sigma) - \kappa(e_L, \sigma_H)$, which then is equivalent to the following chain of inequalities,

$$\begin{aligned} \frac{1}{\kappa(e_H, \sigma) - \kappa(e_L, \sigma)} \frac{\phi_e}{e_H}(e_H - e_L) &\leq \frac{1}{\kappa(e_H, \sigma) - \kappa(e_L, \sigma)} \left[\frac{\phi_e}{e_H}(e_H - e_L) + \frac{\phi_\sigma}{\sigma_L}(\sigma_H - \sigma) \right] \\ &\leq \frac{1}{\kappa(e_H, \sigma) - \kappa(e_L, \sigma_H)} \left[\frac{\phi_e}{e_H}(e_H - e_L) + \frac{\phi_\sigma}{\sigma_L}(\sigma_H - \sigma) \right], \end{aligned}$$

and hence the condition (6.7) covers Case (iii) when $e_t = e_H$. The case of when $e_t = e_L$ is similar.

□

E Principal's value function

E.1 Properties of the value function

First we obtain some basic properties of the value function \tilde{v} .

Proposition E.1. *The value function \tilde{v} is concave in w . That is, for all $m > \underline{m}$, and $w^1, w^2 > R$ and $\lambda \in [0, 1]$,*

$$\lambda \tilde{v}(w^1, m) + (1 - \lambda) \tilde{v}(w^2, m) \leq \tilde{v}(\lambda w^1 + (1 - \lambda)w^2, m).$$

Proof of Proposition E.1. Pick any $W^j \geq R = 0$ and controls $\sigma^j, X^j, \beta^j, \tau^j$, for $j = 1, 2$, and in particular pick $\tau^1 = \tau^2 \equiv \tau^0$, from the admissible control set. Then from (S), we have the dynamics,

$$dW_t^j = \left[r_0 W_t^j - \phi_\sigma \left(\frac{\sigma_t^j}{\sigma_L} - 1 \right) \right] dt - dX_t^j + \beta_t^j \sigma_t^j dM_t,$$

for any $j = 1, 2$. Fix any $\lambda \in [0, 1]$. Multiplying and summing, we obtain,

$$\begin{aligned} d(\lambda W_t^1 + (1 - \lambda)W_t^2) &= \left[r_0(\lambda W_t^1 + (1 - \lambda)W_t^2) - \phi_\sigma \left(\frac{\lambda \sigma_t^1 + (1 - \lambda)\sigma_t^2}{\sigma_L} - 1 \right) \right] dt \\ &\quad - d(\lambda X_t^1 + (1 - \lambda)X_t^2) + (\lambda \beta_t^1 \sigma_t^1 + (1 - \lambda)\beta_t^2 \sigma_t^2) dM_t. \end{aligned}$$

Now, let us define,

$$\beta_t := \frac{\lambda \beta_t^1 \sigma_t^1 + (1 - \lambda)\beta_t^2 \sigma_t^2}{\lambda \sigma_t^1 + (1 - \lambda)\sigma_t^2}. \quad (\text{E.1})$$

Let's show that $\beta_t \in \mathcal{B}$, as given in (7.1). That is, let's show that,

$$K \geq \beta_t \geq \underline{\beta}. \quad (\text{E.2})$$

The upper bound is clear since $\beta_t^1, \beta_t^2 \leq K$. But the lower bound is also clear since $\beta_t^1, \beta_t^2 \geq \underline{\beta}$. Thus, $\beta_t \in \mathcal{B}$.

Hence, we have that if $(\sigma^j, X^j, \beta^j, \tau^j) \in \mathcal{A}_{w^j, m}$, $j = 1, 2$, then $(\lambda \sigma^1 + (1 - \lambda)\sigma^2, \lambda X^1 + (1 - \lambda)X^2, \beta, \tau) \in \mathcal{A}_{\lambda(w^1, m) + (1 - \lambda)(w^2, m)}$, where β is as constructed in (E.1).

Thus, this implies by optimality, and concavity of $\kappa(e_H, \sigma)$ in σ ,

$$\begin{aligned} &\lambda \mathbb{E} \left[\int_0^{\tau^0} e^{-r_1 t} \kappa(e_H, \sigma_t^1) dt - \int_0^{\tau^0} e^{-r_1 t} dX_t^1 + e^{-r_1 \tau^0} L \right] \\ &+ (1 - \lambda) \mathbb{E} \left[\int_0^{\tau^0} e^{-r_1 t} \kappa(e_H, \sigma_t^2) dt - \int_0^{\tau^0} e^{-r_1 t} dX_t^2 + e^{-r_1 \tau^0} L \right] \\ &\leq \mathbb{E} \left[\int_0^{\tau^0} e^{-r_1 t} \kappa(e_H, \lambda \sigma_t^1 + (1 - \lambda)\sigma_t^2) dt - \int_0^{\tau^0} e^{-r_1 t} d(\lambda X_t^1 + (1 - \lambda)X_t^2) + e^{-r_1 \tau^0} L \right] \\ &\leq \tilde{v}(\lambda w^1 + (1 - \lambda)w^2, m). \end{aligned}$$

Take the supremum to the above over the admissible set of controls, and we obtain, $\lambda \tilde{v}(w^1, m) + (1 - \lambda)\tilde{v}(w^2, m) \leq \tilde{v}(\lambda w^1 + (1 - \lambda)w^2, m)$, as desired. □

Proposition E.1 is not only mathematically important, but also economically critical. In the model of DeMarzo and Sannikov (2006), it was explicitly shown that the principal's value function as a function of the agent's continuation value is concave, and thus, public randomization does not improve the payoff for the principal. Note that public randomization is effectively concavification of the principal's value function. However, in DeMarzo, Livdan, and Tchisty (2013), the authors show that in the case the agent manages a cash flows with a jump component (interpreted as "disasters"), then public randomization does indeed improve the value for the principal. Economically, public

randomization implies the following.⁴⁴ To induce the current agent to work, the manager could effectively flip a coin every morning, and if the coin lands in heads, the manager keeps the current agent employed, but if the coin lands in tails, the manager fires the current agent and finds an identical agent in the labor market to replacement the outgoing agent. This coin flipping act implies that the principal is indifferent to the identity of the agent, as long as there does exist a competitive labor market of identical agents, and the principal can frictionlessly hire and fire agents from this labor market pool. Also equally important, the production technology of the firm is completely independent of the identity of the agent. Effectively, that means the firm is effectively a factory with a fixed production technology, and the agent is simply hired to spend effort (or not) to press a button in the factory; it does not matter who presses that button.

However, in the context of delegated portfolio management, this public randomization argument does not hold. In particular, for investment firms, the technology *is* the agent. Effectively, investment firms, and hedge funds in particular, live and die by the investment manager. Threatening the investment manager via the aforementioned coin flipping exercise is not credible, as the manager knows if he is fired, the firm also collapses with him. Thus, the importance of Proposition E.1 is that the principal does not need to resort to a public randomization device to achieve a better outcome, as if otherwise, this public randomization device is not even feasible.

E.2 Further properties of the value function

Lemma E.2. *For any $(w^i, m^i) \in \Gamma$, $i = 1, 2$, and $\lambda \in [0, 1]$, if $(\sigma^i, \beta^i, X^i) \in \mathcal{A}_{w^i, m^i}$, then there does not exist some β such that $(\lambda\sigma^1 + (1-\lambda)\sigma^2, \beta, \lambda X^1 + (1-\lambda)X^2) \in \mathcal{A}_{\lambda(w^1, m^1) + (1-\lambda)(w^2, m^2)}$.*

Proof of Lemma E.2. We proceed by contradiction. Fix any (w^i, m^i) , $i = 1, 2$ and $\lambda \in [0, 1]$. Without loss of generality, let us pick $m^2 > m^1$. Then there exists some admissible controls (σ^i, β^i, X^i) such that,

$$dW_t^i = \left(r_0 W_t^i - \phi_\sigma \left(\frac{\sigma_t^i}{\sigma_L} - 1 \right) \right) dt - dX_t^i + \beta_t^i \sigma_t^i dM_t^i.$$

In particular, we must have that there exist some point (w, m) such that $\lambda(w^1, m^1) + (1-\lambda)(w^2, m^2) = (w, m)$ and some admissible controls (σ, β, X) associated with the point (w, m) . In particular, multiplying by λ and summing, we must have that,

$$\begin{aligned} d(\lambda W_t^1 + (1-\lambda)W_t^2) &= \left(r_0(\lambda W_t^1 + (1-\lambda)W_t^2) - \phi_\sigma \left(\frac{\lambda\sigma_t^1 + (1-\lambda)\sigma_t^2}{\sigma_L} - 1 \right) \right) dt \\ &\quad - d(\lambda X_t^1 + (1-\lambda)X_t^2) + \lambda\beta_t^1 \sigma_t^1 dM_t^1 + (1-\lambda)\beta_t^2 \sigma_t^2 dM_t^2. \end{aligned} \quad (\text{E.3})$$

But since M^1, M^2 are both geometric Brownian motions on the same underlying Brownian motion term except for different initial conditions, so $M_t^i = m^i e^{-1/2t+Bt}$, the diffusion term above in (E.3) can be rewritten as,

$$\begin{aligned} \lambda\beta_t^1 \sigma_t^1 dM_t^1 + (1-\lambda)\beta_t^2 \sigma_t^2 dM_t^2 &= \lambda\beta_t^1 \sigma_t^1 m^1 e^{-1/2t+Bt} dB_t + (1-\lambda)\beta_t^2 \sigma_t^2 m^2 e^{-1/2t+Bt} dB_t \\ &= (\lambda\beta_t^1 \sigma_t^1 m^1 + (1-\lambda)\beta_t^2 \sigma_t^2 m^2) e^{-1/2t+Bt} dB_t. \end{aligned}$$

But if there exist some admissible control σ associated with (w, m) , then from the $\phi_\sigma \left(\frac{\lambda\sigma_t^1 + (1-\lambda)\sigma_t^2}{\sigma_L} - 1 \right) dt$ term, it means this admissible volatility control σ must be $\sigma_t = \lambda\sigma_t^1 + (1-\lambda)\sigma_t^2$. As well, the admissible compensation must be $X = \lambda X^1 + (1-\lambda)X^2$. Then from this form, it implies the admissible sensitivity β must thus be the form, for $m = \lambda m^1 + (1-\lambda)m^2$,

$$\beta_t [\lambda\sigma_t^1 + (1-\lambda)\sigma_t^2] [\lambda m^1 + (1-\lambda)m^2] = \lambda\beta_t^1 \sigma_t^1 m^1 + (1-\lambda)\beta_t^2 \sigma_t^2 m^2,$$

or that,

$$\beta_t = \frac{\lambda\beta_t^1 \sigma_t^1 m^1 + (1-\lambda)\beta_t^2 \sigma_t^2 m^2}{[\lambda\sigma_t^1 + (1-\lambda)\sigma_t^2] [\lambda m^1 + (1-\lambda)m^2]}. \quad (\text{E.4})$$

If this β_t is admissible, it must thus be in \mathcal{B} . But the lower bound in \mathcal{B} cannot hold for β_t of (E.4). To see this,

⁴⁴I thank Dmitry Livdan for pointing this out.

since $\beta_t^i \in \mathcal{B}$, $i = 1, 2$, we have that,

$$\beta_t \geq \beta \frac{\lambda \sigma_t^1 m^1 + (1 - \lambda) \sigma_t^2 m^2}{[\lambda \sigma_t^1 + (1 - \lambda) \sigma_t^2][\lambda m^1 + (1 - \lambda) m^2]}.$$

Thus, in order for β_t to be admissible, we must thus have that,

$$\frac{\lambda \sigma_t^1 m^1 + (1 - \lambda) \sigma_t^2 m^2}{[\lambda \sigma_t^1 + (1 - \lambda) \sigma_t^2][\lambda m^1 + (1 - \lambda) m^2]} \geq 1. \quad (\text{E.5})$$

Rearranging and after some algebra, (E.5) implies,

$$m^2(\sigma_t^2 - \sigma_t^1) \geq m^1(\sigma_t^2 - \sigma_t^1). \quad (\text{E.6})$$

Recall we had assumed, without loss of generality, $m^2 > m^1$ — contradiction, this is impossible to hold for all choices of $\sigma_t^1, \sigma_t^2 \in [\sigma_L, \sigma_H]$ for all times t . In particular, it suffices to pick those times t and controls such that $\sigma_t^2 < \sigma_t^1$ and the above inequality will imply $m^2 \leq m^1$. Thus, there does not exist an admissible control β associated with the point $(w, m) = \lambda(w^1, m^1) + (1 - \lambda)(w^2, m^2)$. \square

Remark E.3. The significance of Lemma E.2 is that it is not possible that the value function is concave in the coordinate pair of (w, m) . Since Γ is clearly a convex set, that means it must be that for any $(w^i, m^i) \in \Gamma$, $i = 1, 2$ and $\lambda \in [0, 1]$ we can for sure find a point (w, m) such that $(w, m) = \lambda(w^1, m^1) + (1 - \lambda)(w^2, m^2)$. The difficult in making the concavity argument of the value function is that from those controls (σ^i, β^i, X^i) associated with point (w^i, m^i) , can we find or construct a control (σ, β, X) associated with the point (w, m) , which again is a convex combination of (w^i, m^i) . Lemma E.2 shows that we cannot. However, to be clear, that is not to say there does not exist an admissible associated with the point (w, m) . Lemma E.2 merely states that if (w, m) is a convex combination of (w^i, m^i) , $i = 1, 2$, that admissible control associated with the point (w, m) cannot be constructed out of the controls associated with (σ^i, β^i, X^i) .

Finally, we note that Lemma E.2 is not contradicting Proposition E.1. In particular, Proposition E.1 is *not* claiming concavity in the coordinate pair (w, m) , but rather it is claiming that if we hold the exogenous factor level m fixed and look at the w -slice of the state space, then the value function is concave in the w -direction, with respect to the agent's continuation value. We summarize and formalize this below in Corollary E.4.

Corollary E.4. *The value function v is not concave on Γ .*

The next result shows that the value function is decreasing in the exogenous factor level.

Proposition E.5. *The value function is decreasing in the exogenous factor level. That is, for any $w \in \Gamma_W$, $m_1, m_2 \in \Gamma_M$ with $m_2 \geq m_1$, we have,*

$$v(w, m_1) \geq v(w, m_2). \quad (\text{E.7})$$

Proof of Proposition E.5. Fix any $w \in \Gamma_W$ and fix any $m_1, m_2 \in \Gamma_M$ and let's suppose $m_2 \geq m_1$. Pick the admissible control as follows. Pick an arbitrary volatility choice $\sigma = \{\sigma_t\}$ and let $\sigma^1 = \sigma^2 = \sigma$ and also pick an arbitrary sensitivity choice $\beta = \{\beta_t\}$ and let $\beta^1 = \beta^2 = \beta$. For the compensation process, pick an arbitrary compensation process $X = \{X_t\}$, and set X^1, X^2 such that,

$$\begin{aligned} X_t^1 &= X_t, \\ X_t^2 &= X_t \mathbf{1}_{\{t \in (0, \tau_1)\}}. \end{aligned}$$

Then we have the associated state variable dynamics associated with those controls as,

$$\begin{aligned} dW_t^i &= \left[r_0 W_t^i - \phi_\sigma \left(\frac{\sigma_t}{\sigma_L} - 1 \right) \right] dt - dX_t^i + \beta_t \sigma_t dM_t^i \\ dM_t^i &= M_t^i dB_t, \end{aligned}$$

where $(W_0^i, M_0^i) = (w, m_i)$, $i = 1, 2$. Let τ^i be the associated hitting time of the form,

$$\tau^i := \inf \{t \geq 0 : W_t^i \leq R \text{ or } M_t^i \leq \underline{m}\}, \quad (\text{E.8})$$

corresponding to the stopping time form in (P).

But recalling that M^i is a geometric Brownian motion, that implies the diffusion terms of dW_t^i is simply $\beta_t \sigma_t dM_t^i = \beta_t \sigma_t M_t^i dB_t = \beta_t \sigma_t m_i e^{-1/2t+B_t} dB_t = \beta_t \sigma_t m_i dM_t$, where M is a geometric Brownian motion on B with zero drift and unit variance and with initial value $M_0 = 1$. But given that $m_2 \geq m_1$, it implies that the diffusion term of dW_t^2 is weakly greater than that of dW_t^1 . But recalling (E.8), and since $W_0^1 = W_0^2 = w$, it implies that we have $\tau^1 \geq \tau^2$; that is, with a higher diffusion (from dW_t^2), and on the same Brownian path B_t , it is likely the first time to get bumped out of the region in (E.8) comes before that of one with a lower diffusion (from dW_t^1).

Then consider that,

$$\begin{aligned} & \mathbb{E} \left[\int_0^{\tau^1} e^{-r_1 t} \kappa(e_H, \sigma_t) dt - \int_0^{\tau^1} e^{-r_1 t} dX_t^1 + e^{-r_1 \tau^1} L \right] - \mathbb{E} \left[\int_0^{\tau^2} e^{-r_1 t} \kappa(e_H, \sigma_t) dt - \int_0^{\tau^1} e^{-r_1 t} dX_t^2 + e^{-r_1 \tau^2} L \right] \\ &= \mathbb{E} \left[\int_{\tau^2}^{\tau^1} \kappa(e_H, \sigma_t) dt + (e^{-r_1 \tau^1} - e^{-r_1 \tau^2}) L \right] \\ &\geq 0. \end{aligned}$$

But rearranging the above, and recalling the chosen admissible controls were arbitrary, and by optimality, we have that,

$$v(w, m_1) \geq \mathbb{E} \left[\int_0^{\tau^1} e^{-r_1 t} \kappa(e_H, \sigma_t) dt - \int_0^{\tau^1} e^{-r_1 t} dX_t^1 + e^{-r_1 \tau^1} L \right] \geq \mathbb{E} \left[\int_0^{\tau^2} e^{-r_1 t} \kappa(e_H, \sigma_t) dt - \int_0^{\tau^1} e^{-r_1 t} dX_t^2 + e^{-r_1 \tau^2} L \right].$$

And again by arbitrariness of the admissible controls and optimality, we have that,

$$v(w, m_1) \geq v(w, m_2) \geq \mathbb{E} \left[\int_0^{\tau^2} e^{-r_1 t} \kappa(e_H, \sigma_t) dt - \int_0^{\tau^1} e^{-r_1 t} dX_t^2 + e^{-r_1 \tau^2} L \right].$$

This concludes the proof. \square

The next result provides a lower bound on the value function v and directly shows that the value function is positive.

Proposition E.6. For any $(w, m) \in \Gamma$, define the processes \tilde{W}, \tilde{M} , given by

$$\begin{aligned} d\tilde{W}_t &= r_0 \tilde{W}_t dt + \underline{\beta} \sigma_L d\tilde{M}_t, & \tilde{W}_0 &= w \\ d\tilde{M}_t &= \tilde{M}_t dB_t, & \tilde{M}_0 &= m, \end{aligned}$$

Define the hitting time θ as,

$$\theta := \inf \left\{ t \geq 0 : (\tilde{W}_t, \tilde{M}_t) \notin \bar{\Gamma} \right\}.$$

Then,

$$\frac{\kappa(e_H, \sigma_L)}{r_1} - \mathbb{E}[e^{-r_1 \theta}] \left(\frac{\kappa(e_H, \sigma_L)}{r_1} - L \right) \leq v(w, m), \quad (\text{E.9})$$

where $\kappa(e_H, \sigma_L)/r_1 - L > 0$, and holds with equality if and only if $(w, m) \in \partial\Gamma$, in which case $\theta = 0$, and $v(w, m) = L$. Thus, the value function is bounded below by a finite, positive constant.

Proof of Proposition E.6. Fix any $(w, m) \in \Gamma$. Pick the controls (σ, X, β) as $\sigma_t \equiv \sigma_L$, $X \equiv 0$ and $\beta_t \equiv \underline{\beta}$ for all times t . Then the state variables thus becomes,

$$\begin{aligned} d\tilde{W}_t &= r_0 \tilde{W}_t dt + \underline{\beta} \sigma_L d\tilde{M}_t, & \tilde{W}_0 &= w \\ d\tilde{M}_t &= \tilde{M}_t dB_t, & \tilde{M}_0 &= m. \end{aligned}$$

With these choices of controls, the principal's payoff is thus,

$$\begin{aligned} \mathbb{E} \left[\int_0^{\tau} e^{-r_1 t} \kappa(e_H, \sigma_L) dt + e^{-r_1 \tau} L \right] &= \kappa(e_H, \sigma_L) \mathbb{E} \left[\frac{1 - e^{-r_1 \tau}}{r} \right] + L \mathbb{E}[e^{-r_1 \tau}] \\ &= \frac{\kappa(e_H, \sigma_L)}{r_1} - \mathbb{E}[e^{-r_1 \tau}] \left(\frac{\kappa(e_H, \sigma_L)}{r_1} - L \right), \end{aligned}$$

where note that $\kappa(e_H, \sigma_L)/r_1 - L > 0$ by Assumption 5.1 that $\kappa(e_L, \sigma_L)/r_1 > L$ and we have $\kappa(e_H, \sigma_L) > \kappa(e_L, \sigma_L)$ by Definition 4.1. Now, since the stopping time $\tau =: \theta$ is now viewed as,

$$\theta := \inf \{t \geq 0 : W_t \leq 0 \text{ or } M_t \leq 0\} = \inf \{t \geq 0 : (W_t, M_t) \notin \bar{\Gamma}\}.$$

Since $(w, m) \in \Gamma$ was arbitrary, then we clearly have (E.9) as desired. \square

Remark E.7. Economically, the value on the left hand side of the inequality (E.9) represents the following. The term $\kappa(e_H, \sigma_L)/r_1$ is the “second worst” value of the firm, in which the agent effectively chooses the lowest possible volatility $\sigma_t \equiv \sigma_L$ for all times t . We note that it is “second worst” value because the absolute “worst” value of the firm is $\kappa(e_L, \sigma_L)/r_1$, that is when the lowest effort $e_t = e_L$ is exerted at all times, but note in this discussion we are concentrating on implementing the high effort e_H contract. However, the agent is still running the firm and recall from Assumption 5.1 that terminating the firm remains to be inefficient. Hence, the term $\kappa(e_H, \sigma_L)/r_1 - L$ effectively represents the premium the principal has to give up to the agent to operate the firm, even at its “second worst” value. However, to maintain IR constraints of the agent, the principal will only allow the agent to run the firm up until the stopping time θ .

E.3 Comparison Principle

We first establish a comparison principle for the value function v .

Proposition E.8. *Suppose ψ is a smooth solution on Γ that satisfies $\psi(0, m) \geq L$ for all $m > 0$, and also satisfies,*

$$\max \left\{ -r_1 \psi(w, m) + \max_{\sigma} \sup_{\beta} (\mathcal{L}_{e_H} \psi)(w, m; \sigma, \beta) + \kappa(e_H, \sigma), -\psi_w(w, m) - 1 \right\} \leq 0, \quad (\text{E.10})$$

Then we have that,

$$\psi \geq v.$$

Proof of Proposition E.8. Fix an initial state $(w_0, m_0) \in \Gamma$ and select an arbitrary admissible control $\alpha = (\sigma, \beta, X) \in \mathcal{A}_{w_0, m_0}$. Furthermore, for $k, n \in \mathbb{N}$, set $\theta_k := \inf\{t \geq 0 : W_t \geq k \text{ or } W_t \leq 1/k\}$, and $\rho_n := \inf\{t \geq 0 : M_t \geq n \text{ or } M_t \leq 1/n\}$. Then we have that $\theta_k, \rho_n \uparrow \infty$ as $k, n \rightarrow \infty$. Now by Ito’s formula, we have that,

$$\begin{aligned} & e^{-r_1 \tau \wedge \theta_k \wedge \rho_n} \psi(W_{\tau \wedge \theta_k \wedge \rho_n}, m_{\tau \wedge \theta_k \wedge \rho_n}) \\ &= \psi(w_0, m_0) + \int_0^{\tau \wedge \theta_k \wedge \rho_n} e^{-r_1 s} [-r_1 \psi(W_{s-}, M_s) + (\mathcal{L}_{e_H} \psi)(W_{s-}, M_s; \sigma_s, \beta_s) + \kappa(e_H, \sigma_s)] ds \\ & \quad - \int_0^{\tau \wedge \theta_k \wedge \rho_n} e^{-r_1 s} \kappa(e_H, \sigma_s) ds + \int_0^{\tau \wedge \theta_k \wedge \rho_n} e^{-r_1 s} \psi_w(W_{s-}, M_s) \beta_s \sigma_s dM_s + \int_0^{\tau \wedge \theta_k \wedge \rho_n} e^{-r_1 s} \psi_m(W_{s-}, M_s) dM_s \\ & \quad + \sum_{0 \leq s \leq \tau \wedge \theta_k \wedge \rho_n} e^{-r_1 s} (\psi(W_s, M_s) - \psi(W_{s-}, M_s)) - \int_0^{\tau \wedge \theta_k \wedge \rho_n} e^{-r_1 s} \psi_w(W_{s-}, M_s) dX_s^c, \end{aligned}$$

where X^c is the continuous part of X . Using the mean value theorem and since ψ satisfies the variational inequality (E.10), we have that $\psi_w \geq -1$, and moreover since for times $s \in [0, \tau \wedge \theta_k \wedge \rho_n]$, all the integrands in the diffusion terms are bounded, and noting also that the controls are also in a compact set, and so taking expectations and rearrange, we obtain,

$$\psi(w_0, m_0) \geq \mathbb{E} e^{-r_1 \tau \wedge \theta_k \wedge \rho_n} \psi(W_{\tau \wedge \theta_k \wedge \rho_n}, m_{\tau \wedge \theta_k \wedge \rho_n}) + \mathbb{E} \left[\int_0^{\tau \wedge \theta_k \wedge \rho_n} e^{-r_1 s} \kappa(e_H, \sigma_s) ds - \int_0^{\tau \wedge \theta_k \wedge \rho_n} e^{-r_1 s} dX_s \right].$$

Now, take $k, n \rightarrow \infty$ and applying Fatou’s lemma, we obtain, and recalling $\psi(W_\tau, m_\tau) = \psi(0, m_\tau) \geq L$,

$$\psi(w_0, m_0) \geq \mathbb{E} \left[\int_0^\tau e^{-r_1 s} \kappa(e_H, \sigma_s) ds - \int_0^\tau e^{-r_1 s} dX_s + e^{-r_1 \tau} L \right].$$

Since the set of admissible controls \mathcal{A}_{w_0, m_0} were arbitrary, taking the supremum on the right hand side of the inequality above, then we are done. \square

With Proposition E.8 on hand, we can now derive some easy growth conditions on the value function v .

Corollary E.9. *For all $(w, m) \in \Gamma$, the value function v satisfies,*

$$v(w, m) \leq (1 + r_0)w + m + \frac{\kappa(e_H, \sigma_H)}{r_1}.$$

Proof of Corollary E.9. Take $\psi(w, m) := (1 + r_0)w + m + \frac{\kappa(e_H, \sigma_H)}{r_1}$ on Γ and $\psi(w, m) = L$ for $w \leq 0$ and $m > 0$. Then clearly ψ is smooth on Γ and moreover, $\psi_{ww} = \psi_{wm} = \psi_{mm} = 0$, and $\psi_w = 1 + r_0$. Observing (E.10), we have that,

$$-\psi_w(w, m) - 1 = -(1 + r_0) - 1 < 0,$$

and

$$\begin{aligned} & -r_1\psi(w, m) + \max_{\sigma} \sup_{\beta} (\mathcal{A}_{e_H} \psi)(w, m; \sigma, \beta) + \kappa(e_H, \sigma_H) \\ & \leq -r_1 \left((1 + r_0)w + m + \frac{\kappa(e_H, \sigma_H)}{r_1} \right) + r_0w + \kappa(e_H, \sigma_H) \\ & = (r_0 - r_1(1 + r_0))w - r_1m \\ & \leq 0, \end{aligned}$$

since we have $r_0/r_1 < 1 + r_0$. Then this choice of ψ satisfies the hypothesis of Proposition E.8 and we are done. \square

E.4 Viscosity solution

E.4.1 Overview

Unlike the approach by the existing continuous-time principal-agent problem literature where either the value function of the principal is only dependent on one single state variable⁴⁵, namely the agent's continuation value, or there are multiple state variables but can be shown that the value function can be written in such a way that dynamic programming only applies to a single state variable⁴⁶. In particular, because there is only one relevant state variable in considering dynamic programming, the literature can rely on the extensive literature on existence and uniqueness results of ODE theory, and in some cases even compute explicitly the form of the principal's value function from the ODE form.

However in our case, it is not evident or perhaps even possible, to consider a rewriting to reduce the two state variables of the agent's continuation value W and the exogenous factor level M to a single state variable case. As a result the conventional and classical approach of the "verification theorem" does not apply. In particular, it means unlike the extensive results from ODE theory that can ensure existence and uniqueness of smooth solutions, we cannot a priori assume that there will exist a smooth solution (namely $C^2(\Gamma)$) such that we can take the first order conditions in (7.3), substitute the maximizer back into (7.3) and hope that there will exist a C^2 solution that still satisfies the highly nonlinear HJB PDE (7.3). Without existence of such a C^2 solution to the HJB PDE (7.3), a verification theorem to show that the solution to the HJB PDE (7.3) is indeed the value function (P) may likely fail. Thus, we must use more general techniques to understand the value function (P) and the HJB PDE (7.3) and hence we will consider *viscosity solution* methods.

To this end, we will first define the PDE operator F . Let us define,

$$\begin{aligned} F(w, m, u, p, A) := \max \left\{ -r_1u + \max_{\sigma} \sup_{\beta} \left(\left[r_0w - \phi_{\sigma} \left(\frac{\sigma}{\sigma_L} - 1 \right) \right] p_w \right. \right. \\ \left. \left. + \frac{1}{2}m^2 A_{mm} + \beta \sigma m^2 A_{wm} + \frac{1}{2}\beta^2 \sigma^2 m^2 A_{ww} + \kappa(e_H, \sigma) \right) \right. \\ \left. - p_w - 1 \right\}. \end{aligned} \quad (\text{E.11})$$

⁴⁵There are many examples here. Most notably, DeMarzo and Sannikov (2006), Sannikov (2008), He (2009), among many others.

⁴⁶He, Wei, and Yu (2014) is an interesting recent example.

Hence, the HJB PDE in (7.3) is the rewriting,

$$F(w, m, v, Dv, D^2v) = 0, \quad (\text{E.12})$$

where we denote Dv as the gradient vector and D^2v as the Hessian matrix of v , respectively. We will now show that the value function v of (P) can be understood as the viscosity solution to (E.12).

We now give some definitions.

Definition E.1. We say that u is a *viscosity supersolution* of (E.12) in Γ if, for every $(w, m) \in \Gamma$ and $\varphi \in C^2(\bar{\Gamma})$ such that (w, m) is a local minimum of $u - \varphi$ in $\bar{\Gamma}$, then

$$F(w, m, u, D\varphi, D^2\varphi) \geq 0. \quad (\text{E.13})$$

We say that u is a *viscosity subsolution* of (E.12) in Γ if, for every $(w, m) \in \Gamma$ and $\varphi \in C^2(\bar{\Gamma})$ such that (w, m) is a local maximum of $u - \varphi$ in $\bar{\Gamma}$, then

$$F(w, m, u, D\varphi, D^2\varphi) \leq 0. \quad (\text{E.14})$$

We say that u is a *viscosity solution* of (E.12) in Γ if it is both a viscosity supersolution and viscosity subsolution.

It is widely known that it is without loss of generality at the point (w, m) in the definition above to take $v(w, m) = \varphi(w, m)$ and also to replace local optimality with global optimality in the above.

E.4.2 Dynamic Programming Principle (DPP)

We will also assume and state without proof the *Dynamic Programming Principle* (DPP).

Theorem E.10 (Dynamic Programming Principle). *For every initial state $(w, m) \in \Gamma$ and every stopping time θ ,*

$$v(w, m) = \sup_{\alpha \in \mathcal{A}_{w, m}} \mathbb{E}_0 \left[\int_0^{\tau \wedge \theta} e^{-r_1 s} \kappa(e_H, \sigma_s) ds - \int_0^{\tau \wedge \theta} e^{-r_1 s} dX_s + e^{-r_1 \tau \wedge \theta} v(W_{\tau \wedge \theta}, M_{\tau \wedge \theta}) \right].$$

E.4.3 Value function as a viscosity solution

Proposition E.11. *The value function v of (P) is the unique viscosity solution of (E.12) in Γ .*

Remark E.12. In the proof of Proposition E.11, we directly show that v is both a viscosity subsolution and a viscosity super solution of (E.12), and thus by definition, v is a viscosity solution of (E.12).⁴⁷ The proof for uniqueness is lengthy and technical. Hence, on a first pass, we will omit the proof for uniqueness.

Proof to Proposition E.11, Viscosity Subsolution. Fix any $(w, m) \in \Gamma$ and let $\varphi \in C^2(\bar{\Gamma})$ with $v - \varphi$ is a local max in $\bar{\Gamma}$ and $v(w, m) = \varphi(w, m)$. By Theorem E.10, if we pick any $x \in (0, w]$ with $X \equiv x$, then we have that,

$$\varphi(w, m) = v(w, m) \geq v(w - x, m) - x \geq \varphi(w - x, m) - x.$$

Rearrange and take $x \downarrow 0$, then we have

$$\varphi_w(w, m) \geq -1. \quad (\text{E.15})$$

Next, fix any constant $\bar{\beta}, \bar{\sigma}$ in the control set, and set $\beta_t \equiv \bar{\beta}$ and $\sigma_t \equiv \bar{\sigma}$, and let $X_t \equiv 0$ for all times t . Let (W, M) be the state variables with those associated control policies. Fix any $h > 0$. Define $\tau_\rho := \inf \{t \geq 0 : (W_t, M_t) \notin B_\rho(w, m) \cap \Gamma\}$, where for $\rho > 0$ sufficiently small, $B_\rho(w, m)$ is the ball centered at (w, m) with radius ρ . Then from Theorem E.10,

⁴⁷The proof ideas are largely inspired by and inherited from Yong and Zhou (1999), Fleming and Soner (2006), Budhiraja and Ross (2008) and Ly Vath, Pham, and Villeneuve (2008).

and noting that $\tau \wedge \tau_\rho = \tau_\rho$, and applying Ito's lemma, we have that,

$$\begin{aligned}
0 &\geq \mathbb{E} \left[\int_0^{\tau_\rho \wedge h} e^{-r_1 s} \kappa(e_H, \bar{\sigma}) ds + e^{-r_1 \tau_\rho \wedge h} v(W_{\tau_\rho \wedge h}, M_{\tau_\rho \wedge h}) \right] - v(w, m) \\
&\geq \mathbb{E} \left[\int_0^{\tau_\rho \wedge h} e^{-r_1 s} \kappa(e_H, \bar{\sigma}) ds + e^{-r_1 \tau_\rho \wedge h} [v(W_{\tau_\rho \wedge h}, M_{\tau_\rho \wedge h}) - \varphi(W_{\tau_\rho \wedge h}, M_{\tau_\rho \wedge h}) + \varphi(W_{\tau_\rho \wedge h}, M_{\tau_\rho \wedge h})] \right] - v(w, m) \\
&\geq \mathbb{E} \left[\int_0^{\tau_\rho \wedge h} e^{-r_1 s} \kappa(e_H, \bar{\sigma}) ds + e^{-r_1 \tau_\rho \wedge h} \underbrace{[v(w, m) - \varphi(w, m)]}_{=0} + \varphi(W_{\tau_\rho \wedge h}, M_{\tau_\rho \wedge h}) \right] - \underbrace{v(w, m)}_{=\varphi(w, m)} \\
&= \mathbb{E} \left[\int_0^{\tau_\rho \wedge h} e^{-r_1 s} \kappa(e_H, \bar{\sigma}) ds + \varphi(w, m) + \int_0^{\tau_\rho \wedge h} e^{-r_1 s} [-r_1 \varphi(W_s, M_s) + (\mathcal{L}_{e_H} \varphi)(W_s, M_s; \bar{\sigma}, \bar{\beta})] ds \right. \\
&\quad \left. + \int_0^{\tau_\rho \wedge h} e^{-r_1 s} \varphi_w(W_s, M_s) \bar{\beta} \bar{\sigma} dM_s + \int_0^{\tau_\rho \wedge h} e^{-r_1 s} \varphi_m(W_s, M_s) dM_s \right] - \varphi(w, m) \\
&= \mathbb{E} \left[\int_0^{\tau_\rho \wedge h} e^{-r_1 s} [-r_1 \varphi(W_s, M_s) + (\mathcal{L}_{e_H} \varphi)(W_s, M_s; \bar{\sigma}, \bar{\beta}) + \kappa(e_H, \bar{\sigma})] ds \right] \\
&\geq \mathbb{E} \left[\int_0^{\tau_\rho \wedge h} e^{-r_1 s} \inf_{B_\rho(w, m)} [-r_1 \varphi(\tilde{w}, \tilde{m}) + (\mathcal{L}_{e_H} \varphi)(\tilde{w}, \tilde{m}; \bar{\beta}, \bar{\beta}) + \kappa(e_H, \bar{\sigma})] ds \right] \\
&= \mathbb{E} \left[\frac{1 - e^{-r_1 \tau_\rho \wedge h}}{r_1} \right] \inf_{B_\rho(w, m)} [-r_1 \varphi(\tilde{w}, \tilde{m}) + (\mathcal{L}_{e_H} \varphi)(\tilde{w}, \tilde{m}; \bar{\beta}, \bar{\beta}) + \kappa(e_H, \bar{\sigma})].
\end{aligned}$$

Since with $X \equiv 0$, then the state variable process (W, M) are continuous and hence $\tau_\rho > 0$. By dominated convergence theorem, let $h \downarrow 0$ and we have,

$$\mathbb{E} \left[\frac{1 - e^{-r_1 \tau_\rho \wedge h}}{h} \right] \rightarrow r_1.$$

As well, dividing the above inequality by h , and letting $\rho \downarrow 0$ so $\tau_\rho \rightarrow \infty$ and $B_\rho(w, m) \rightarrow \{(w, m)\}$ and $h \downarrow 0$, and recall $v(w, m) = \varphi(w, m)$, we obtain,

$$0 \geq -r_1 v(w, m) + (\mathcal{L}_{e_H} \varphi)(w, m; \bar{\sigma}, \bar{\beta}) + \kappa(e_H, \bar{\sigma}).$$

But since the choice of $\bar{\beta}, \bar{\sigma}$ were arbitrary, the above also implies,

$$0 \geq -r_1 v(w, m) + \max_{\sigma} \sup_{\beta} (\mathcal{L}_{e_H} \varphi)(w, m; \sigma, \beta) + \kappa(e_H, \sigma). \quad (\text{E.16})$$

Putting (E.16) and (E.15) together and we are done. \square

Proof to Proposition E.11, Viscosity Supersolution. Let $\varphi \in C^2(\bar{\Gamma})$ and (\hat{w}, \hat{m}) be a local minimizer of $v - \varphi$ on Γ with $v(\hat{w}, \hat{m}) = \varphi(\hat{w}, \hat{m})$. We need to show that,

$$F(\hat{w}, \hat{m}, \varphi, D\varphi, D^2\varphi) \geq 0 \quad (\text{E.17})$$

For contradiction, suppose not. Then the left hand side of (E.17) is strictly negative and by smoothness of φ , there exists $\delta, \gamma > 0$ satisfying,

$$F(w, m, \varphi, D\varphi, D^2\varphi) \leq -\gamma, \quad (w, m) \in B_\delta(\hat{w}, \hat{m}), \quad (\text{E.18})$$

where $B_\delta(\hat{w}, \hat{m}) := \{(w, m) : \|(w, m) - (\hat{w}, \hat{m})\|_2 < \delta\}$. Since Γ is an open set, by changing δ , if necessary, we may assume that $B_\delta(\hat{w}, \hat{m}) \subset \Gamma$.

Fix an arbitrary control $\alpha = (\sigma, X, \beta) \in \mathcal{A}_{\hat{w}, \hat{m}}$, and let θ be the first exist time of (W, M) from $B_\delta(\hat{w}, \hat{m})$. Since $B_\delta(\hat{w}, \hat{m}) \subset \Gamma$, we have that $\theta < \tau$.

Let W^c, X^c denote the continuous parts of W, X , respectively, and noting that $\Delta W_t := W_t - W_{t-} = -\Delta X_t := -(X_t - X_{t-})$. By the continuity of sample paths, $M_s = M_{s-}$. Now by Ito's lemma and taking expectations, we have

that,

$$\begin{aligned}
& \mathbb{E}e^{-r_1\theta^-}\varphi(W_{\theta^-}, M_{\theta^-}) - \varphi(\hat{w}, \hat{m}) \\
&= \mathbb{E} \int_0^{\theta^-} e^{-r_1s} [-r_1\varphi(W_s, M_s) + (\mathcal{L}_{e_H}\varphi)(W_s, M_s; \sigma_s, \beta_s) + \kappa(e_H, \sigma_s)] ds \\
&\quad - \mathbb{E} \int_0^{\theta^-} e^{-r_1s} \varphi_w(W_s, M_s) dX_s^c + \mathbb{E} \sum_{0 \leq s < \theta^-} e^{-r_1s} [\varphi(W_s, M_s) - \varphi(W_{s-}, M_s)]
\end{aligned} \tag{E.19}$$

But for $0 \leq s < \theta^-$, (E.18) implies,

$$-r_1\varphi(W_s, M_s) + (\mathcal{L}_{e_H}\varphi)(W_s, M_s; \sigma_s, \beta_s) \leq -\gamma, \tag{E.20}$$

$$-\varphi_w(W_s, M_s) - 1 \leq -\gamma. \tag{E.21}$$

And using the mean value theorem and (E.21), we obtain,

$$\varphi(W_s, M_s) - \varphi(W_{s-}, M_s) \leq (1 - \gamma)\Delta X_s. \tag{E.22}$$

Substituting (E.20), (E.21) and (E.22) and noting that $X_t = X_t^c + \Delta X_t$, we obtain,

$$\begin{aligned}
& \mathbb{E}e^{-r_1\theta^-}\varphi(W_{\theta^-}, M_{\theta^-}) - \varphi(\hat{w}, \hat{m}) \\
&\leq -\mathbb{E} \int_0^{\theta^-} e^{-r_1s} \kappa(e_H, \sigma_s) ds + \mathbb{E} \int_0^{\theta^-} e^{-r_1s} (1 - \gamma) dX_s - \mathbb{E} \int_0^{\theta^-} e^{-r_1s} \gamma ds.
\end{aligned} \tag{E.23}$$

Note that while $(W_{\theta^-}, M_{\theta^-}) \in \bar{B}_\delta(\hat{w}, \hat{m})$, we have that (W_θ, M_θ) is either on the boundary $\partial\bar{B}_\delta(\hat{w}, \hat{m})$ or out of $\bar{B}_\delta(\hat{w}, \hat{m})$. However, there exists some random variable $\lambda \in [0, 1]$ such that,

$$(W^\lambda, M^\lambda) := (W_{\theta^-} + \lambda\Delta W_\theta, M_\theta) = (W_{\theta^-} - \lambda\Delta X_\theta, M_\theta) \in \partial\bar{B}_\delta(\hat{w}, \hat{m}). \tag{E.24}$$

And again by the mean value theorem and (E.21), we have that,

$$\varphi(W^\lambda, M^\lambda) - \varphi(W_{\theta^-}, M_{\theta^-}) \leq (1 - \gamma)\lambda\Delta X_\theta. \tag{E.25}$$

Note also that,

$$\begin{aligned}
W^\lambda &= W_{\theta^-} - \lambda\Delta X_\theta \\
&= (W_\theta - \Delta W_\theta) - \lambda\Delta X_\theta \\
&= W_\theta + \Delta X_\theta - \lambda\Delta X_\theta \\
&= W_\theta + (1 - \lambda)\Delta X_\theta.
\end{aligned} \tag{E.26}$$

From (E.26) and properties of the value function, we also have that,

$$v(W^\lambda, M^\lambda) \geq v(W_\theta, M_\theta) - (1 - \lambda)\Delta X_\theta. \tag{E.27}$$

And since $v - \varphi$ is a local min at (\hat{w}, \hat{m}) , with $v(\hat{w}, \hat{m}) = \varphi(\hat{w}, \hat{m})$, so we have that,

$$v(W^\lambda, M^\lambda) \leq \varphi(W^\lambda, M^\lambda). \tag{E.28}$$

So from (E.25), (E.27) and (E.28), we obtain,

$$\begin{aligned}
\varphi(W_{\theta^-}, M_{\theta^-}) &\geq \varphi(W^\lambda, M^\lambda) - (1 - \gamma)\lambda\Delta X_\theta \\
&\geq v(W^\lambda, M^\lambda) - (1 - \gamma)\lambda\Delta X_\theta \\
&\geq v(W_\theta, M_\theta) - (1 - \lambda)\Delta X_\theta - (1 - \gamma)\lambda\Delta X_\theta \\
&= v(W_\theta, M_\theta) - (1 - \lambda\gamma)\Delta X_\theta
\end{aligned} \tag{E.29}$$

Substituting (E.29) into (E.23), and rearranging, we thus have,

$$\begin{aligned} \varphi(\hat{w}, \hat{m}) \geq & \mathbb{E} \left[\int_0^{\theta^-} e^{-r_1 s} \kappa(e_H, \sigma_s) ds - \int_0^{\theta^-} e^{-r_1 s} dX_s - e^{-r_1 \theta^-} \Delta X_\theta + e^{-r_1 \theta^-} v(W_\theta, M_\theta) \right] \\ & + \mathbb{E} \left[\int_0^{\theta^-} e^{-r_1 s} dX_s - \int_0^{\theta^-} e^{-r_1 s} ds + e^{-r_1 \theta^-} \lambda \Delta X_\theta \right]. \end{aligned} \quad (\text{E.30})$$

Now suppose we can show that there exists a constant $c_0 > 0$ such that,

$$\mathbb{E} \left[\int_0^{\theta^-} e^{-r_1 s} dX_s - \int_0^{\theta^-} e^{-r_1 s} ds + e^{-r_1 \theta^-} \lambda \Delta X_\theta \right] \geq c_0. \quad (\text{E.31})$$

Suppose for now that (E.31) is true. Then from (E.30), using (E.31), recalling that the chosen controls were arbitrary so we may take the supremum over all admissible controls, and using Theorem E.10 we have that,

$$\varphi(\hat{w}, \hat{m}) \geq \gamma c_0 + v(\hat{w}, \hat{m}), \quad (\text{E.32})$$

implying that $\varphi(\hat{w}, \hat{m}) - v(\hat{w}, \hat{m}) \geq \gamma c_0 > 0$ — contradiction, since we had assumed that $\varphi(\hat{w}, \hat{m}) = v(\hat{w}, \hat{m})$.

So, the proof is complete once we can prove the existence of the constant $c_0 > 0$ that satisfies (E.31). To this end, let us define the C^2 function,

$$\psi(w, m) := c_0 \left(1 - \frac{\|(w, m) - (\hat{w}, \hat{m})\|_2^2}{\delta^2} \right), \quad (\text{E.33})$$

where,

$$c_0 := C_0 \wedge (\delta/2), \quad (\text{E.34})$$

and C_0 is given by,

$$\begin{aligned} C_0 := & \delta^2 (\kappa(e_H, \sigma_L) - r_1) \\ & \times \inf_{\bar{\sigma} \in [\sigma_L, \sigma_H]} \inf_{(\tilde{w}, \tilde{m}) \in B_\delta(\hat{w}, \hat{m})} \left[2r_0 \tilde{w}(\tilde{w} - \hat{w}) - 2\phi_\sigma \left(\frac{\bar{\sigma}}{\sigma_L} - 1 \right) (\tilde{w} - \hat{w}) + \tilde{m}^2 (\tilde{m} - \hat{m}) + K^2 \sigma_H^2 \right]^{-1}. \end{aligned} \quad (\text{E.35})$$

Then a direct (but somewhat messy) computation will show that for any admissible choice $(\bar{\beta}, \bar{\sigma})$, (E.33) satisfies,

$$\begin{cases} \max \{ 1 - \psi_w, \psi - 1, -[-r_1 \psi + (\mathcal{L}_{e_H} \psi)(\cdot, \cdot; \bar{\sigma}, \bar{\beta})] - 1 \} \leq 0, & \text{on } \bar{B}_\delta(\hat{w}, \hat{m}), \\ \psi = 0, & \text{on } \partial \bar{B}_\delta(\hat{w}, \hat{m}). \end{cases} \quad (\text{E.36})$$

By Itô's lemma applied to $e^{-r_1 \theta^-} \psi(W_{\theta^-}, M_{\theta^-})$, taking expectations and rearranging, we will arrive at (E.31).

This completes the proof. \square

E.4.4 Regularity upgrade

Once we have obtained Proposition E.11 and thus we can understand the value function v as the viscosity solution to the HJB PDE (7.3), we are now ready to “upgrade” our results. First we give a “partial” C^1 result.

Proposition E.13. *The value function v is C^1 in the w -direction; that is, for each $m \in \Gamma_M$, the partial derivative $v_w(w, m)$ exists for all $w \in \Gamma_W$, and is continuous in w .*

Proof of Proposition E.13. Fix any $m_0 \in \Gamma_m$. Define the limits,

$$\nabla_w^+ v(w, m_0) := \lim_{\delta \downarrow 0} \frac{v(w + \delta, m_0) - v(w, m_0)}{\delta}, \quad (\text{E.37a})$$

$$\nabla_w^- v(w, m_0) := \lim_{\delta \downarrow 0} \frac{v(w, m_0) - v(w - \delta, m_0)}{\delta}. \quad (\text{E.37b})$$

By Proposition E.1, the map $w \mapsto v(w, m_0)$ is concave in the w -direction. Thus from standard results in convex analysis (see, for instance, Rockafellar (1970)), the limits (E.37) exist. We want to show that $\nabla_w^+ v(w, m_0) = \nabla_w^- v(w, m_0)$, and hence equals $v_w(w, m_0)$ for $w \in \Gamma_w$. We proceed in three steps.

Step 1. Let's show that $\nabla_w^+ v(w, m_0) \geq \nabla_w^- v(w, m_0)$. For contradiction, suppose there exist some $w_0 \in \Gamma_w$ such that $\nabla_w^+ v(w_0, m_0) < \nabla_w^- v(w_0, m_0)$. Fix any nonzero $q \in (\nabla_w^+ v(w_0, m_0), \nabla_w^- v(w_0, m_0))$, and any $\varepsilon > 0$. Consider the function,

$$\varphi(w, m) := v(w_0, m_0) + q \left(1 + \frac{4m_0^2}{\varepsilon q w_0} + \frac{4K^2 \sigma_L^2 m_0^2}{\varepsilon q w_0} \right) (w - w_0) - \frac{1}{2\varepsilon} (w - w_0)^2 - \frac{1}{2\varepsilon} (m - m_0)^2.$$

Then φ is quadratic and concave in (w, m) , and then clearly (w_0, m_0) is a local maximum of $v - \varphi$, with $v(w_0, m_0) = \varphi(w_0, m_0)$, $\varphi_w(w_0, m_0) = q \left(1 + \frac{4m_0^2}{\varepsilon q w_0} + \frac{4K^2 \sigma_L^2 m_0^2}{\varepsilon q w_0} \right)$, $\varphi_{wm} = 0$, and $\varphi_{ww}(w_0, m_0) = \varphi_{mm}(w_0, m_0) = -\frac{1}{\varepsilon}$. By the viscosity subsolution property of v , and suboptimality,

$$\begin{aligned} 0 &\geq F(w_0, m_0, \varphi, D\varphi, D^2\varphi) \\ &= \max \left\{ -r_1 v(w_0, m_0) + \max_{\sigma} \sup_{\beta} \left\{ \left[r_0 w_0 - \phi_{\sigma} \left(\frac{\sigma}{\sigma_L} - 1 \right) \right] \cdot q \left(1 + \frac{4m_0^2}{\varepsilon q w_0} + \frac{4K^2 \sigma_L^2 m_0^2}{\varepsilon q w_0} \right) \right. \right. \\ &\quad \left. \left. + \frac{1}{2} m_0^2 \cdot \left(-\frac{1}{\varepsilon} \right) + \beta \sigma m_0^2 \cdot 0 + \frac{1}{2} \beta^2 \sigma^2 \sigma^2 m_0^2 \cdot \left(-\frac{1}{\varepsilon} \right) + \kappa(e_H, \sigma) \right\}, \right. \\ &\quad \left. - q \left(1 + \frac{4m_0^2}{\varepsilon q w_0} + \frac{4K^2 \sigma_L^2 m_0^2}{\varepsilon q w_0} \right) - 1 \right\} \\ &\geq -r_1 v(w_0, m_0) + r_0 w_0 q \left(1 + \frac{4m_0^2}{\varepsilon q w_0} + \frac{4K^2 \sigma_L^2 m_0^2}{\varepsilon q w_0} \right) - \frac{m_0^2}{2\varepsilon} - \frac{K^2 \sigma_L^2 m_0^2}{2\varepsilon} + \kappa(e_H, \sigma_L). \end{aligned}$$

Rearranging the above, we will have,

$$0 \geq -\varepsilon r_1 v(w_0, m_0) + \varepsilon r_0 w_0 q + \varepsilon \kappa(e_H, \sigma_L) + m_0^2 \left(4r_0 - \frac{1}{2} \right) + K^2 \sigma_L^2 m_0^2 \left(4r_0 - \frac{1}{2} \right).$$

Take $\varepsilon \downarrow 0$, then the above implies,

$$0 \geq m_0^2 \left(4r_0 - \frac{1}{2} \right) + K^2 \sigma_L^2 m_0^2 \left(4r_0 - \frac{1}{2} \right),$$

— contradiction, as we recall $r_0 \in (0, 1)$. Thus, we have $\nabla_w^+ v(w, m_0) \geq \nabla_w^- v(w, m_0)$ for $w \in \Gamma_w$.

Step 2. Now, it remains to show $\nabla_w^+ v(w, m_0) \leq \nabla_w^- v(w, m_0)$. But since for each $m_0 \in \Gamma_m$, $v(w, m_0)$ is concave in the w -direction, by again standard results from convex analysis (see Rockafellar (1970)), the desired inequality $\nabla_w^+ v(w, m_0) \leq \nabla_w^- v(w, m_0)$ immediately holds.

Step 3. By Step 1 and 2, we have that $\nabla_w^+ v(w, m_0) = \nabla_w^- v(w, m_0) \equiv v_w(w, m_0)$. But furthermore, since we have concavity in the w -direction and also continuity, this implies $v_w(w, m_0)$ is also continuous. Thus, we have C^1 in the w -direction. \square

Now, we give a C^2 regularity upgrade.

Proposition E.14. *Define the sets,*

$$\mathcal{D} := \{(w, m) \in \Gamma : v_w(w, m) = -1\}, \quad (\text{E.38})$$

$$\mathcal{C} := \Gamma \setminus \mathcal{D}. \quad (\text{E.39})$$

Then,

1. v is C^2 in the w -direction on $\mathcal{C} \cup \mathcal{D}^{\circ}$.
2. In the classical C^2 solution sense, we have,

$$-r_1 v + \max_{\sigma} \sup_{\beta} [(\mathcal{L}_{e_H} v)(\cdot, \cdot; \sigma, \beta) + \kappa(e_H, \sigma)] = 0, \quad \text{on } \mathcal{C}. \quad (\text{E.40})$$

Proof of Proposition E.14. Part 1. It is clear that v is C^2 in the w -direction on \mathcal{D}° . That is, for $(w, m) \in \mathcal{D}^{\circ}$, we have $v_w(w, m) = -1$. Note that by Proposition E.13, the expression $v_w(w, m)$ makes sense. But moreover, since the right hand side of $v_w(w, m) = -1$ is a constant, which is trivially differentiable in the w -direction, and so it also implies $v_w(w, m)$ is trivially C^1 , and so v is C^2 in the w -direction.

It remains to prove that v is C^2 in the w -direction on \mathcal{C} .

Part 2. Let's first show that v is a viscosity solution to,

$$-r_1 v(w, m) + \max_{\sigma} \sup_{\beta} [(\mathcal{L}_{e_H} v)(w, m; \sigma, \beta) + \kappa(e_H, \sigma)] = 0, \quad (w, m) \in \mathcal{C}. \quad (\text{E.41})$$

Viscosity supersolution. Indeed, let $(\hat{w}, \hat{m}) \in \mathcal{C}$ and φ be a C^2 function on \mathcal{C} such that (\hat{w}, \hat{m}) is a local minimum of $v - \varphi$ with $v(\hat{w}, \hat{m}) = \varphi(\hat{w}, \hat{m})$. But that means by first order conditions, and again recalling Proposition E.13, it implies we have $0 = \frac{\partial}{\partial w}(v - \varphi)(\hat{w}, \hat{m})$; so in particular, we have that for $(\hat{w}, \hat{m}) \in \mathcal{C}$, $\varphi_w(\hat{w}, \hat{m}) = v_w(\hat{w}, \hat{m}) < -1$. And thus from the viscosity supersolution property of v from Proposition E.11, we have,

$$-r_1 \varphi(\hat{w}, \hat{m}) + \max_{\sigma} \sup_{\beta} [(\mathcal{L}_{e_H} v)(\hat{w}, \hat{m}; \sigma, \beta) + \kappa(e_H, \sigma)] \geq 0.$$

This shows the desired viscosity supersolution property.

Viscosity subsolution. The subsolution property is immediate by the fact v is (at least) a viscosity subsolution to (7.3), as given by Proposition E.11. Thus, v is also a viscosity solution to (E.41).

Now, fix any arbitrarily bounded set $O \subset \mathcal{C}$. Consider the nonlinear Dirichlet boundary value problem,

$$-r_1 \xi + \max_{\sigma} \sup_{\beta} [(\mathcal{L}_{e_H} \xi)(\cdot, \cdot; \sigma, \beta) + \kappa(e_H, \sigma)] = 0, \quad \text{on } O, \quad (\text{E.42a})$$

$$\xi = v, \quad \text{on } \partial O. \quad (\text{E.42b})$$

In particular, we see that for any $a = (a_1, a_2) \in \mathbb{R}^2$, we have that by extracting out the coefficients to the second order derivative and cross derivative terms of ξ in (E.42), and for any β, σ in the admissible choice set,

$$\begin{aligned} \beta^2 \sigma^2 m^2 a_1^2 + 2\beta \sigma m^2 a_1 a_2 + m^2 a_2^2 &\geq \underline{\beta}^2 \underline{\sigma}_L^2 \underline{m}^2 a_1^2 + 2\underline{\beta} \underline{\sigma}_L \underline{m} a_1 a_2 + \underline{m}^2 a_2^2 \\ &\geq C(a_1^2 + 2a_1 a_2 + a_2^2) \\ &= C \|a\|_2^2, \end{aligned}$$

where the constant $C := \min\{\underline{\beta}^2 \underline{\sigma}_L^2 \underline{m}^2, 2\underline{\beta} \underline{\sigma}_L \underline{m}, \underline{m}^2\} > 0$ and $\|\cdot\|_2$ is the standard \mathbb{R}^2 Euclidean norm. In particular, this shows PDE (E.42) is *uniformly elliptic* (see Remark E.15) with Dirichlet boundary data. Thus, standard classical existence and uniqueness results are available (see Evans (1983), Fleming and Soner (2006), and Evans (2010)). Hence, a unique C^2 solution ξ on O to (E.42) exists. But from the standard uniqueness results of viscosity solution to (E.42), this implies we have $v = \xi$ on O . From the arbitrariness of O , this proves that v is C^2 smooth on Γ . \square

Remark E.15. We can now see the significance of the IR constraint in Definition 6.1(1b). If in contrast, we do not have the requirement that $M_t \geq \underline{m}$, so that the state space in question is $(w, m) \in (0, \infty) \times (0, \infty)$ rather than $(0, \infty) \times (\underline{m}, \infty)$, then there does not exist a strictly positive constant C for which $\beta^1 \sigma^2 m^1 a_1^2 + 2\beta \sigma m^2 a_1 a_2 + m^2 a_2^2 \geq C \|a\|_2^2$ can hold, in which case the PDE is known as being *degenerate*. The essential problem is that when $m = 0$, the nature of the state variable dynamics significantly changes (i.e. from being fully stochastic to fully deterministic).

See also Figure 3 for an illustration of the state space Γ . And also see Figure 4 for an illustration, for each fixed $m \in \Gamma_M$, the value function $w \mapsto v(w, m)$.

Remark E.16. The set \mathcal{D} of (E.38) is the *payment condition* and \mathcal{C} is the *continuation region* (i.e. no payment condition).

Remark E.17. Proposition E.14 shows that in the continuation region \mathcal{C} , the value function v is C^2 smooth in both (w, m) . Together with the concavity of the value function in the w -direction from Proposition E.1, it implies that in (E.40), when we optimize over the volatility σ choice and the sensitivity β choice, we can use the usual first order conditions to uniquely characterize them. Thus, Proposition E.1 and Proposition E.14 show that the discussions in Section 8 are on meaningful grounds.

E.5 Free boundary

We first introduce the *free (moving) boundary*,

$$\partial^* := \{(\bar{W}(m), m) : m \in \Gamma_M\}, \quad (\text{E.43})$$

where \bar{W} is the map from Γ_M to Γ_W , defined by,

$$\bar{W}(m) := \sup \{w \in \Gamma_W : v_w(w, m) = -1\}, \quad m \in \Gamma_M. \quad (\text{E.44})$$

Through a rather technical and elaborate argument similar to Soner and Shreve (1989), one can show that \bar{W} is finite and indeed twice continuously differentiable. We omit the proof here, but the argument should follow, in spirit and actuality, from Soner and Shreve (1989).

Then we can have a further regularity upgrade of our earlier results. Note that for each $m \in \Gamma_M$, we can partition $\Gamma_W = (R, \infty) = (R, \bar{W}(m)) \cup [\bar{W}(m), \infty)$. Moreover, note that the set $(R, \bar{W}(m)) \times \Gamma_M = \mathcal{C}$, by construction. Define the function V on Γ as follows. For each $m \in \Gamma_M$, define,

$$V(w, m) \quad (\text{E.45})$$

$$:= \begin{cases} \frac{1}{r_1} \max_{\sigma} \sup_{\beta} [(\mathcal{L}_{e_H} v)(w, m; \sigma, \beta) + \kappa(e_H, \sigma)], & w \in (R, \bar{W}(m)), \\ \bar{W}(m) - w + \frac{1}{r_1} \left[-r_0 \bar{W}(m) + \phi_{\sigma} \left(\frac{\sigma_H}{\sigma_L} - 1 \right) + \frac{1}{2} m^2 v_{mm}(\bar{W}(m), m) + \kappa(e_H, \sigma_H) \right], & w \in [\bar{W}(m), \infty). \end{cases} \quad (\text{E.46})$$

In particular, we have simply taken the value function v , and extracted out the dynamics in the continuation region \mathcal{C} as in Proposition E.14(E.40), and then on the payment condition region, linearly extrapolated the value at the slope -1 . By the smoothness of \bar{W} , we have that V is also a viscosity solution to the HJB PDE (7.3). By uniqueness of viscosity solutions, this implies that $V \equiv v$ on Γ , but we note that V is C^2 in the w -direction on Γ_W .

In particular, when we evaluate $w = \bar{W}(m)$ for any $m \in \Gamma_M$ from (E.46), we obtain,

$$V(\bar{W}(m), m) = \frac{1}{r_1} \left[-r_0 \bar{W}(m) + \phi_{\sigma} \left(\frac{\sigma_H}{\sigma_L} - 1 \right) + \frac{1}{2} m^2 v_{mm}(\bar{W}(m), m) + \kappa(e_H, \sigma_H) \right], \quad m \in \Gamma_M. \quad (\text{E.47})$$

But recalling the terminal condition $V = v = L$ on $\partial\Gamma$, and in particular if we take $m \rightarrow \underline{m}$ in (E.47), we have,

$$L = \frac{1}{r_1} \left[-r_0 \bar{W}(\underline{m}) + \phi_{\sigma} \left(\frac{\sigma_H}{\sigma_L} - 1 \right) + \frac{1}{2} \underline{m}^2 v_{mm}(\bar{W}(\underline{m}), \underline{m}) + \kappa(e_H, \sigma_H) \right]. \quad (\text{E.48})$$

But if we view the (moving) free boundary $m \mapsto \bar{W}(m)$ as the object of interest, then (E.47) identifies the nonlinear ODE

References

- ADMATI, A. R. AND P. PFLEIDERER (1997): “Does It All Add Up? Benchmarks and the Compensation of Active Portfolio Managers,” *The Journal of Business*, 70, 323–350.
- ADRIAN, T. AND M. M. WESTERFIELD (2009): “Disagreement and Learning in a Dynamic Contracting Model,” *Review of Financial Studies*, 22, 3873–3906.
- AÏT-SAHALIA, Y. AND J. JACOD (2012): “Analyzing the Spectrum of Asset Returns: Jump and Volatility Components in High Frequency Data,” *Journal of Economic Literature*, 50, 1007–1050.
- ALBANESE, C. AND S. LAWI (2005): “Laplace Transforms for Integrals of Markov Processes,” *Markov Processes Related Fields*, 11, 677–724.
- ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND H. EBENS (2001): “The distribution of realized stock return volatility,” *Journal of Financial Economics*, 61, 43–76.
- ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND P. LABYS (2003): “Modeling and forecasting realized volatility,” *Econometrica*, 71, 579–625.
- ANG, A. (2012): “Asset Owners and Delegated Managers,” Tech. rep.
- BARNDORFF-NIELSEN, O. E. AND N. SHEPHARD (2002a): “Econometric analysis of realized volatility and its use in estimating stochastic volatility models,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64, 253–280.
- (2002b): “Estimating quadratic variation using realized variance,” *Journal of Applied Econometrics*, 17, 457–477.
- (2004): “Power and Bipower Variation with Stochastic Volatility and Jumps,” *Journal of Financial Econometrics*, 2, 1–37.
- BARRAS, L., O. SCAILLET, AND R. WERMERS (2010): “False discoveries in mutual fund performance: Measuring luck in estimated alphas,” *Journal of Finance*, LXV, 179–216.
- BASAK, S., A. PAVLOVA, AND A. SHAPIRO (2007): “Optimal Asset Allocation and Risk Shifting in Money Management,” *Review of Financial Studies*, 20, 1583–1621.
- BIAIS, B., T. MARIOTTI, G. PLANTIN, AND J.-C. ROCHET (2007): “Dynamic Security Design: Convergence to Continuous Time and Asset Pricing Implications,” *The Review of Economic Studies*, 74, 345–390.
- BIAIS, B., T. MARIOTTI, AND J.-C. ROCHET (2011): “Dynamic Financial Contracting,” in *10th World Congress of the Econometric Society*.
- BOLLEN, N. P. B. AND R. E. WHALEY (2009): “Hedge Fund Risk Dynamics: Implications for Performance Appraisal,” *Journal of Finance*, LXIV, 985–1035.
- BUDHIRAJA, A. AND K. ROSS (2008): “Optimal stopping and free boundary characterizations for some Brownian control problems,” *The Annals of Applied Probability*, 18, 2367–2391.
- BURASCHI, A., R. KOSOWSKI, AND W. SRITRAKUL (2013): “Incentives and Endogenous Risk Taking: A Structural View on Hedge Fund Alphas,” *Journal of Finance*.
- CADENILLAS, A., J. CVITANIĆ, AND F. ZAPATERO (2004): “Leverage decision and manager compensation with choice of effort and volatility,” *Journal of Financial Economics*, 73, 71–92.
- (2007): “Optimal risk-sharing with effort and project choice,” *Journal of Economic Theory*, 133, 403–440.
- CARPENTER, J. N. (2000): “Does Option Compensation Increase Managerial Risk Appetite?” *Journal of Finance*, LV, 2311–2331.
- CHEVALIER, J. AND G. ELLISON (1997): “Risk Taking by Mutual Funds as a Response to Incentives Risk,” *Journal of Political Economy*, 105, 1167–1200.
- CUOCO, D. AND R. KANIEL (2011): “Equilibrium prices in the presence of delegated portfolio management,” *Journal of Financial Economics*, 101, 264–296.
- CVITANIĆ, J., D. POSSAMAÏ, AND N. TOUZI (2014): “Moral Hazard in Dynamic Risk Management,” *Working Paper*.

- CVITANIĆ, J. AND J. ZHANG (2012): *Contract Theory in Continuous-Time Models*, Springer.
- DAROLLES, S. AND C. GOURIEROUX (2014a): “The Effects of Management and Provision Accounts on Hedge Fund Returns — Part I: The High Water Mark Scheme,” in *Modeling Dependence in Econometrics*, ed. by V.-N. Huynh, V. Kreinovich, and S. Sriboonchitta, Cham: Springer International Publishing, vol. 251 of *Advances in Intelligent Systems and Computing*, 23–45.
- (2014b): “The Effects of Management and Provision Accounts on Hedge Fund Returns — Part II: The Loss Carry Forward Scheme,” in *Modeling Dependence in Econometrics*, ed. by V.-N. Huynh, V. Kreinovich, and S. Sriboonchitta, Cham: Springer International Publishing, vol. 251 of *Advances in Intelligent Systems and Computing*, 47–62.
- DEMARZO, P. M. AND M. FISHMAN (2007): “Optimal Long-Term Financial Contracting,” *Review of Financial Studies*, 20, 2079–2128.
- DEMARZO, P. M., M. J. FISHMAN, Z. HE, AND N. WANG (2012): “Dynamic Agency and the q Theory of Investment,” *Journal of Finance*, LXVII.
- DEMARZO, P. M., D. LIVDAN, AND A. TCHISTYI (2013): “Risking Other People’s Money: Gambling, Limited Liability, and Optimal Incentives,” *Working Paper*, 1–38.
- DEMARZO, P. M. AND Y. SANNIKOV (2006): “Optimal Security Design and Dynamic Capital Structure in a Continuous Time Agency Model,” *Journal of Finance*, LXI, 2681–2724.
- DUFFIE, D. (2001): *Dynamic Asset Pricing Theory*, Princeton University Press, 3 ed.
- DURRETT, R. (2010): *Probability: Theory and Examples*, Cambridge University Press, 4th ed.
- DYBVIK, P. H., H. K. FARNSWORTH, AND J. N. CARPENTER (2009): “Portfolio Performance and Agency,” *Review of Financial Studies*, 23, 1–23.
- EVANS, L. C. (1983): “Classical Solutions of The Hamilton-Jacobi-Bellman Equation for Uniformly Elliptic Operators,” *Transactions of the American Mathematical Society*, 275, 245–255.
- (2010): *Partial Differential Equations*, American Mathematical Society, 2nd ed.
- FAMA, E. F. AND K. R. FRENCH (2010): “Luck versus Skill in the Cross-Section of Mutual Fund Returns,” *Journal of Finance*, LXV, 1915–1947.
- FLEMING, W. H. AND H. M. SONER (2006): *Controlled Markov Processes and Viscosity Solutions (Stochastic Modelling and Applied Probability)*, Springer.
- FOLLAND, G. B. (1999): *Real Analysis: Modern Techniques and Their Applications*, Wiley, second ed ed.
- FOSTER, D. P. AND H. P. YOUNG (2010): “Gaming Performance Fees by Portfolio Managers,” *Quarterly Journal of Economics*, 125, 1435–1458.
- FRAZZINI, A. AND L. H. PEDERSEN (2014): “Betting against beta,” *Journal of Financial Economics*, 111, 1–25.
- GIAMBONA, E. AND J. GOLEC (2009): “Mutual fund volatility timing and management fees,” *Journal of Banking & Finance*, 33, 589–599.
- GIAT, Y. AND A. SUBRAMANIAN (2013): “Dynamic contracting under imperfect public information and asymmetric beliefs,” *Journal of Economic Dynamics and Control*, 37, 2833–2861.
- GROCHULSKI, B. AND Y. ZHANG (2011): “Optimal risk sharing and borrowing constraints in a continuous-time model with limited commitment,” *Journal of Economic Theory*, 146, 2356–2388.
- HE, Z. (2009): “Optimal Executive Compensation when Firm Size Follows Geometric Brownian Motion,” *Review of Financial Studies*, 22, 859–892.
- (2011): “A model of dynamic compensation and capital structure,” *Journal of Financial Economics*, 100, 351–366.
- (2012): “Dynamic Compensation Contracts with Private Savings,” *Review of Financial Studies*, 25, 1494–1549.
- HE, Z., B. WEI, AND J. YU (2014): “Optimal Long-term Contracting with Learning,” *Working Paper*.

- HESTON, S. (1993): “A closed-form solution for options with stochastic volatility with applications to bond and currency options,” *Review of Financial Studies*, 6, 327–343.
- HIRSHLEIFER, D. AND Y. SUH (1992): “Risk, Managerial Effort, and Project Choice,” *Journal of Financial Intermediation*, 2, 308–345.
- HODDER, J. E. AND J. C. JACKWERTH (2009): “Incentive Contracts and Hedge Fund Management,” *Journal of Financial and Quantitative Analysis*, 42, 811.
- HOFFMANN, F. AND S. PFEIL (2010): “Reward for Luck in a Dynamic Agency Model,” *Review of Financial Studies*, 23, 3329–3345.
- (2013): “Delegated Investment in a Dynamic Agency Model,” *Working Paper*, August.
- HOLMSTRÖM, B. AND P. MILGROM (1987): “Aggregation and linearity in the provision of intertemporal incentives,” *Econometrica*, 55, 303–328.
- HUANG, J., C. SIALM, AND H. ZHANG (2011): “Risk Shifting and Mutual Fund Performance,” *Review of Financial Studies*, 24, 2575–2616.
- HUNT, P. AND J. KENNEDY (2004): *Financial Derivatives in Theory and Practice*, John Wiley & Sons, Ltd., revised ed.
- KARATZAS, I. AND S. E. SHREVE (1991): *Brownian Motion and Stochastic Calculus*, Springer, 2nd ed.
- KESSLER, M., A. LINDNER, AND M. SORENSEN (2012): *Statistical Methods for Stochastic Differential Equations (Monographs on Statistics & Applied Probability)*, Chapman & Hall / CRC.
- KIM, K. (2004): “Moment Generating Function of the Reciprocal of an Integral of Geometric Brownian Motion,” *Proceedings of the American Mathematical Society*, 132, 2753–2759.
- LAMBERT, R. A. (1986): “Executive effort and Selection of Risky Projects,” *The RAND Journal of Economics*, 17, 77–88.
- LELAND, H. E. (1998): “Agency Costs, Risk Management, and Capital Structure,” *Journal of Finance*, LIII, 1213–1243.
- LY VATH, V., H. PHAM, AND S. VILLENEUVE (2008): “A mixed singular/switching control problem for a dividend policy with reversible technology investment,” *The Annals of Applied Probability*, 18, 1164–1200.
- MIAO, J. AND A. RIVERA (2013): “Robust Contracts in Continuous Time,” *Working Paper*, April.
- MIAO, J. AND Y. ZHANG (2013): “A Duality Approach to Continuous-Time Contracting Problems with Limited Commitment,” *Working Paper*, September.
- Ø KSENDAL, B. (2010): *Stochastic Differential Equations: An Introduction with Applications (Universitext)*, Springer, 6th ed.
- OU-YANG, H. (2003): “Optimal Contracts in a Continuous-Time Delegated Portfolio Management Problem,” *Review of Financial Studies*, 16, 173–208.
- PROTTER, P. E. (2005): *Stochastic Integration and Differential Equations*, Springer, 2nd ed.
- RAUH, J. D. (2008): “Risk Shifting versus Risk Management: Investment Policy in Corporate Pension Plans,” *Review of Financial Studies*, 22, 2687–2733.
- REVUZ, D. AND M. YOR (2005): *Continuous Martingales and Brownian Motion*, Springer.
- ROCKAFELLAR, R. T. (1970): *Convex Analysis (Princeton Landmarks in Mathematics and Physics)*, Princeton University Press.
- SANNIKOV, Y. (2007): “Games with Imperfectly Observable Actions in Continuous Time,” *Econometrica*, 75, 1285–1329.
- (2008): “A Continuous-Time Version of the Principal-Agent Problem,” *Review of Economic Studies*, 75, 957–984.
- (2012a): “Contracts: The Theory of Dynamic Principal-Agent Relationships and the Continuous-Time Approach,” in *10th World Congress of the Econometric Society*.
- (2012b): “Moral Hazard and Long-Run Incentives,” *Working Paper*, December, 1–36.

- (2013): “Dynamic Security Design and Corporate Financing,” in *Handbook of the Economics of Finance*, ed. by G. M. Constantinides, M. Harris, and R. Stulz, Elsevier B.V., vol. 2, 71–122.
- SCHAETTLER, 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.
- SONER, H. M. AND S. E. SHREVE (1989): “Regularity of the value function for a two-dimensional singular stochastic control problem,” *SIAM Journal on Control and Optimization*, 27, 876–907.
- STOUGHTON, N. M. (1993): “Moral Hazard and the Portfolio Management Problem,” *Journal of Finance*, 48, 2009–2028.
- STRACCA, L. (2006): “Delegated Portfolio Management: A Survey of the Theoretical Literature,” *Journal of Economic Surveys*, 20, 823–848.
- SUNG, J. (1995): “Linearity with Project Selection and Controllable Diffusion Rate in Continuous-Time,” *The RAND Journal of Economics*, 26, 720–743.
- (2004): “Optimal Contracts Under Adverse Selection and Moral Hazard: A Continuous-Time Approach,” *Review of Financial Studies*, 18, 1021–1073.
- SZYDLOWSKI, M. (2012): “Incentives, Project Choice and Dynamic Multitasking,” *Working Paper*, 1–51.
- VAN BINSBERGEN, J. H., M. W. BRANDT, AND R. S. J. KOIJEN (2008): “Optimal Decentralized Investment Management,” *Journal of Finance*, 63, 1849–1895.
- WERMERS, R. (2011): “Performance Measurement of Mutual Funds, Hedge Funds, and Institutional Accounts,” *Annual Review of Financial Economics*, 3, 537–574.
- WILLIAMS, N. (2011): “Persistent Private Information,” *Econometrica*, 79, 1233–1275.
- WRIGHT, E. M. (1959): “Solution Of The Equation $\$ze^z = a\$$,” *Bulletin of the American Mathematical Society*, 65, 37–112.
- YONG, J. AND X. Y. ZHOU (1999): *Stochastic Controls: Hamiltonian Systems and HJB Equations (Stochastic Modelling and Applied Probability)*, Springer.
- YOR, M. (1992): “On Some Exponential Functionals of Brownian Motion,” *Advances in Applied Probability*, 24, 509–531.
- ZHU, J. Y. (2013): “Optimal Contracts with Shirking,” *The Review of Economic Studies*, 80, 812–839.
- ZUCKERMAN, G. (2013): “Yen Bets Don’t Add Up for a Fund Giant,” .