Fund flows and risk-taking*

Konstantin W. Milbradt†
Princeton University

December 20, 2008

Abstract

Hedge funds face performance related withdrawals to their assets under management. As hedge fund managers' compensation is closely linked to the assets under management, they optimally take outflows into account in their investment decisions. I solve a continuous time investment model where fund flows are connected to the relative distance of the current assets under management to their high-water marks – also known as the current drawdown – to show that: (i) managers become more conservative the closer they are to the high-water mark; (ii) the possibility of stronger future outflows makes the manager more conservative close to the high-water mark; (iii) this conservatism is reversed as current outflows become larger - the stronger current outflows, the more risk the manager will take on even though the risk-return tradeoff stays constant; (iv) waiving fees can be optimal for the manager when far below the benchmark.

Keywords: Withdrawals; Dynamic Reference Point; Fund Flows; Drawdowns

*I am indebted to my adviser, Markus Brunnermeier, and to Wei Xiong and Hyun Shin for valuable advice and support. I thank Dante Amengual, Ing-Haw Cheng, Zhiguo He, Martin Oehmke, and Vinayak Tripathi. All remaining errors are my own. For most recent version, please check `http://www.princeton.edu/~milbradt`

†Princeton University, Department of Economics, Bendheim Center for Finance, 26 Prospect Avenue, Princeton, NJ 08540-5296, e-mail: milbradt@princeton.edu

1 Introduction

How does the risk of capital outflows impact a hedge fund's dynamic risk taking? This paper provides a theoretical model in which capital outflows are determined by the drawdown level, the distance of assets under management from their high-water mark. Such behavior embeds path-dependencies in withdrawals that can lead to a downward spiral in assets under management – investors withdraw capital, the current drawdown increases, leading to further withdrawals. Hedge fund manager's are commonly paid via incentive contracts that link a manager's compensation to current performance and assets under management. Thus, path-dependencies in fund flows translate into path-dependencies in a hedge fund's investment strategy, even though the fund's investment opportunities remain unchanged.

The fund-manager has a standard incentive contract with a fixed fee of assets under management and upside compensation if assets under management rise above their high-water mark. Investors start withdrawing money as soon as the fund's assets under management drop below the historic high, resulting in path-dependent withdrawals. The drawdown level, the relative distance of current assets to their high-water marks, is the dynamic reference point for withdrawals – outflows increase for higher drawdown levels. Such withdrawals result in a feedback loop that tilts the manager's 'personal' investment opportunity set, i.e. the investment opportunity set corresponding to his personal compensation. I derive the optimal ordinary differential equation describing the manager's investment decisions that gives rise to the following results: (i) the manager becomes more conservative the closer the fund's assets are to the high-water mark; (ii) furthermore, the possibility of stronger future outflows makes the manager more conservative close to the high-water mark; (iii) this conservatism is reversed as current outflows become larger - the stronger current outflows, the more risk the manager will take on even though the risk-return tradeoff stays constant; (iv) it can be optimal for the manager to waive the fixed investment fee even when there is no prospect of attracting any new capital via capital inflows.

The model is set in a continuous-time stochastic market in which a risk-neutral manager provides access to an investment opportunity to outside investors . Three main assumptions drive the model. First, the manager compensation follows the standard 2/20 incentive contract that has a fixed fee component and performance bonuses, where the bonus is connected to achieving new highs for assets under management. The contract therefore provides a convex payoff structure. Second, outside investors use historical highwater marks of assets under management of the fund to assess the manager's performance. The central assumption here is that investors start withdrawing funds at an increasing rate the further assets under management have fallen below historic highs, i.e. as a function of the current drawdown level. This introduces path dependencies in the form of possible downward spirals into the model. Third, the problem is modeled as stationary or infinite horizon, removing any time dependence. At each point in time, the manager can expect, as long as the fund is still in operation, a positive amount of time remaining on the job. This option value will keep risk-taking from becoming ill-defined.

The manager has access to two investment instruments, a risky asset with a constant investment opportunity set and a risk-free bond. But although the investment opportunity set of the fund as a whole is constant, the manager's personal, i.e. compensation related, investment opportunity set is not. As the manager derives his compensation from the assets under management, withdrawals will directly affect his compensation base, leading to a time-varying drift on his 'personal', i.e. compensation related, investment opportunity set. Our assumptions on withdrawals make this time-variation path-dependent. Consequently, the manager has to optimally account for his actions affecting his investment opportunity set.

When the manager is close to meeting the performance criterion and withdrawals are at a minimum, his risk-taking is at a minimum. The intuition is that there is a fundamental asymmetry in the impact of added risk: added risk leads to a higher drift, increasing the possibility of upside compensation but also pushing put the reference point. Added risk, however, also leads to a higher possibility of a negative downward spiral on assets under management through performance related withdrawals. Trading off these effects, the manager will behave more conservatively the closer the fund is to meeting the benchmark performance. Additionally, when fund A faces uniformly higher outflows than fund B, fund A will behave more risk-averse than fund B when close to the performance criterion – missteps are much more costly to A than to B, as they result in a steeper adverse shift in the investment opportunity set.

As performance deteriorates, withdrawals increase, further pulling down assets under management, leading to even higher outflows. These feedback effects lead to an increasing base downward drift on the fund's assets, reducing the future compensation of the manager. To counteract this effect, the manager will increase risk-taking of the fund, something akin to a gradual doubling up strategy – by risking more, the manager might be able to 'stop the bleeding' for some states of the world and get back to more benign regions of his personal investment space, instead of accepting a failing fund almost surely. Consequently, the manager gambles for resurrection in very bad states of the world, but behaves conservatively in goods states of the world.

In the summer of 2007, Goldman Sachs offered to waive the fixed investment fee on a few of its quantitative hedge funds to stem mounting withdrawals. Although the contract terms for the manager are fixed, the fund foregoing compensation is unilaterally implementable. Within my model, I can show that there can be cases when it is optimal for the manager to waive the fixed investment fee even without the prospect of fund inflows. The fixed fee is usually small enough to not diminish the drift of the asset significantly on its own. However, the feedback effects introduced by investors' withdrawals can lead to a dynamic magnification of such small flows. For certain cases, such magnification effects can be strong enough to outweigh the loss of forgone compensation to the manager.

The recent paper by Panageas and Westerfield 2006 is closest to my model. In their setup, the agent only has upside compensation but no fixed management fee while withdrawals are absent. The focus of the paper is the interaction between this upside compensation with the manager's risk-taking. They derive in closed form that the agent takes constant risk, irrespective of how far away the fund's assets are from the high water mark. In contrast, my paper concentrates on the feedback effect of withdrawals on the manager's risk-taking decisions and the resulting path dependencies introduced in the manager's 'personal' investment opportunity set. Another paper that looks at upside compensation and high-water marks is Goetzmann et al 2003. The authors examine, without optimal risk-taking, the value of the different claims to such a contract and the value of the company's assets when such a high-water mark contract is in place. Using an option-based approach, they derive these values in closed form.

One of the first papers to examine a dynamic asset allocation model under path dependencies is Grossman and Zhou 1993. In the paper, a manager maximizes the fund's long term growth rate, but is faced with an absolute non-zero lower bound in fund value. This lower bound is modeled as a function of the maximum process of the fund's value. The authors find that the manager becomes increasingly risk-averse the closer the fund's value is to this lower bound. At the maximum, however, the manager still takes positive risk.

Related papers that use a martingale approach with a fixed time horizon T are Hugonnier and Kaniel 2007, Carpenter 2000 and Basak et al 2007 amongst others. Hugonnier and Kaniel 2007 present a general equilibrium model in which log utility investors have to use the fund to achieve exposure to the risky asset. This results in fund flows depending on the current position of the fund, but no feedback effects or path-dependencies. The authors find that there is flow hedging on part of the manager, who get paid via a fixed fee, even though the investors for which the manager invests are myopic. In general, the martingale approach relies on linear wealth dynamics, therefore excluding the possibility of feedback effects on the fund's assets[1]. Basak et al 2007 solve a fixed horizon time T model with possible out- and inflow effects at the final time. They find that there can be risk-shifting behavior depending on in what part of the flow region the investor is. Carpenter 2000 similar finds that with an option payoff, the manager will take large risks when the option is out of the money.

Although there is little data on fund flows with regard to hedge funds, there are empirical studies of fund flows in mutual funds. Chevalier and Ellison 1997 provide evidence that past

[1]The martingale approach, as it does not rely on value functions, can handle much more general utility functions. The downside is that, unless more general BSDEs results are used, the formulation does not allow non-linear wealth dynamics and the finiteness of the time-horizon usually introduces a time-dependent component into the model. This makes the martingale approach less suitable for modeling feedback dynamics.

performance influences current fund flows – high past returns lead to fund-inflows, while low past returns lead to fund outflows. Similarly, Sirri and Tufano 1998 find that investors invest disproportionally more in funds with recent superior returns, again leading to a fund-flow relationship based on path-dependent reference points.

The paper is structured as follows. Section 2 sets up the model and discusses the main assumptions. The derivation of the optimality equation is presented in Section 3. Section 4 presents the main proposition of the paper, the solution in terms of an ODE. Numerical solutions to the ODE are presented. The case of when a firm will waive the fixed investment fee is presented in section 5. Other extensions are discussed in section 6. Section 7 concludes.

2 Model setup

The model is set in continuous-time with $t \in [0, \infty)$ on a suitable defined probability space $(\Omega, P, \mathcal{F})$. I will examine the strategic decisions of a manager of a hedge fund or mutual fund. Outside investors use the fund to achieve exposure to investment opportunities they cannot access individually. The fund manager, on the other hand, is unable to access the general financial market with wealth derived from his fund employment, and will therefore use his control of the fund's investment position to optimize his compensation. Additionally, for technical reasons it is assumed that the fund manager has no capital of his own.[2]

Performance evaluation. Investors use the assets under management and past maximal size of assets under management as a basis to evaluate the fund's performance. Investors prefer funds that have continuously grown over time, i.e. funds that have repeatedly set new highs for the size of assets under management, rather than funds that have been very large in the past but only manage a fraction of the assets today. Let W_t denote the value of assets under management of the fund, and let M_t be the maximum process of assets under management, i.e.

$$M_t \equiv \max_{0 \leq s \leq t} W_s$$

By the definition of the maximum process, we have $W_t \in [0, M_t]$ at each t, where we assumed that W_t cannot fall below 0 (as $W_t = 0$ is an absorbing state). The process M_t offers a tractable way to model the path dependencies discussed above by introducing a dynamic reference point. Additionally, M_t can play the dual role of governing possible performance bonuses or fund inflows

[2]For a risk-neutral agent that does not face any bankruptcy constraints for fund value above zero, any own capital base would result in infinite risk-taking as there is an excess return to be earned on the risky asset. We therefore assume a wealth-less manager.

via its differential dM_t. To model external benchmarks, M_t can follow a deterministic evolution away from the boundary, i.e. on $\{W < M\}$, as discussed in the appendix.

Preferences & Contract structure. A risk-neutral manager, without any own capital, runs the investment fund and is compensated via some variant of the widely known 2 and 20 structure: he gets a fixed rate β of assets under management (the '2'), plus a proportional performance bonus on beating the performance benchmark γ (the '20'). In terms of our variables, the agent receives a flow-payoff of βW per unit of time, plus possible instantaneous bonuses of γdM.

We can now define the agent's expected future value $V(W, M)$ as

$$V_t (W_t, M_t) \equiv \mathbb{E}_t \int_t^{\infty} e^{-\rho(s-t)} (\beta W ds + \gamma dM) \tag{1}$$

where $\rho > r$ is the subjective discount rate applied by the agent and $\mathbb{E}_t$ is the conditional expectation w.r.t. to the filtration $\mathcal{F}_t$. This rate is bigger than the market discount rate as the manager faces a positive probability of exogenous liquidation of the fund within each time interval unrelated to his performance. I assume that this time of death takes the form of a Poisson process with a constant intensity that contributes to ρ.

It is important to point out that we take the contract structure as given. The contract is not an optimal contract. First, we observe 2/20 contracts being used in reality. Optimal contracts, such as presented in Sannikov 2007, DeMarzo and Sannikov 2006, He 2007 and other papers in the recent literature on dynamic optimal contracts, often take more complex forms. Second, the situation in this model is not that of a single principal, but a large number of principals, here the outside investors, and a single manager. With multiple principals, the optimal contracting problem becomes much more complex, with few theoretical results.

To summarize, the agent has a consumption process that he cannot control directly, but only indirectly through the state variable W_t. The manager will thus allocate the capital of the fund in such a way as to achieve an optimal allocation w.r.t. V_t, not necessarily w.r.t. W_t. There is an agency conflict that originates from the structure of the contract.

Withdrawals. The second main assumption is that investors withdraw money when the fund underperforms. More precisely, investors start withdrawing money when assets under management fall below the historic high M_t. This drop can happen for two reasons. First, the fund is having a string of bad returns carrying assets under management away from their historic high. Second, some investors have already withdrawn some money, shrinking the assets under management. This second avenue creates a feedback loop, in that withdrawals lower the assets under management, which leads to even more withdrawals. To have this feedback loop that is observed in practice, we need to assume that the high-water mark is not adjusted for outflows. If it were adjusted for outflows, there would be no downward spiral of withdrawals feeding further withdrawals, and the only remaining effect would be that of a time-varying subjective discount rate.

Investors use withdrawals for three reasons. First, as already mentioned, most common management contracts (including those assumed in this paper) are too simple to be optimal. Withdrawals then are an added instrument in incentivizing and disciplining management. Second, investors infer management skills from the returns process. Bad returns signal bad management. On one hand, investors might attempt to rationally learn skills of the management, but the learning is misspecified - every fund has the same skills. On the other hand, investors might behave in a behavioral way when shunning bad performance. An assumption that yields tractability here is that withdrawals influence the returns process via liquidations etc. in such a way that it can be summarized by W_t and M_t. Third, investors might have adaptive expectations for the performance benchmark. This behavioral bias then introduces withdrawals as investors become disappointed with the fund's performance.

I formulate an outflow process that depends on the distance of the current assets under management from the maximum level M_t. The distance metric used here by the outside investors is $\frac{M_t}{W_t}$, i.e. relatively how much below historic highs are the fund's assets under management. Additionally, I assume that outflows are proportional to assets under management. Lastly, when the fund is at its maximum value it does not face any outflows. Putting these assumptions together, outflows can be described by a function of the form $h\left(\frac{W}{M}\right)W$ with $h(1) = 0$ and $h'\left(\frac{W}{M}\right) \leq 0$. Inflows could be easily incorporate into the model either via $h(1) < 0$ or as handled in the extensions section via an inflow at the high-water mark. As long as inflows are bounded and moderate, neither specification will change the subsequent findings of this paper.

Bankruptcy. For technical reasons, I assume that the fund is liquidated completely when assets under management sink to a level $W = \varepsilon M$, with $\varepsilon \in (0, 1)$. As assets under management reach new highs, the funds operation expands and costs grow. Due to downward rigidity in scaling operations, the fund will have to declare bankruptcy at a level $\varepsilon M > 0$ to cover possible bankruptcy costs. The underlying assumption here is that bankruptcy costs scale with the maximum level of assets under management.

Investment opportunity. The fund manager has access to an investment opportunity that the investors cannot access individually. The price process of this investment opportunity (alternative called the risky asset) follows a geometric Brownian Motion process with constant mean

μ and instantaneous variance σ

$$\frac{dP}{P} = \mu dt + \sigma dZ$$

I assume here that the investment opportunity's mean rate of return is higher than the risk-free rate, i.e. $\mu > r$. The access to the investment opportunity and any compensation derived from it cannot be sold on the open market, so that the manager is unable to diversify away his compensation risk. The specification of μ and σ can be understood as broadly parameterizing the skill of the manager in finding good investment opportunities. The outside investor uses the fund to gain access to this investment opportunity, but due to the incentive problems involved has to accept agency costs in doing so.

Assets under management. The manager allocates the assets under management in the exclusive investment opportunity available to him and in the risk-free asset. As it is costly to employ the manager, the investors do not ask him to invest in opportunities directly available to them. Let π denote the proportion of assets under management in the risky asset, with $\pi \in L^2$ for technical reasons. We will sometimes refer to π as leverage and to W as wealth. Imposing self-financing, assets under management have to follow

$$dW = W\left[\pi(\mu - r) + r - \beta - h\left(\frac{W}{M}\right)\right]dt + W\pi\sigma dZ - \gamma dM \tag{2}$$

The wealth equation of the classical Merton model is here supplemented with two non-standard terms, the withdrawal function $h(\cdot)$ and the instantaneous performance related payoff to setting a new high in assets under management γdM. When wealth is away from its maximum, i.e. for $W \in (0, M)$, dW generates a second-order operator $\mathcal{L}$,

$$\mathcal{L}g = W\left[\pi(\mu - r) + r - \beta - h\left(\frac{W}{M}\right)\right]g_W + \frac{1}{2}W^2\pi^2\sigma^2 g_{WW} \tag{3}$$

where we used $dM = 0$ for $W \in (0, M)$ and $g(\cdot)$ is any twice-differentiable function of W.

Path dependencies here are generated by the outflow function $h(\cdot)$: although μ and σ are constant, the investment opportunity set for assets under management, from which the manager' compensation derives, changes according to how much the manager is underperforming the benchmark. These shifts in the 'personal' investment opportunity set are what influences the risk-taking of the manager, as he has to dynamically take into account possible trickle down effects of fund flows when when making investment decisions. In the traditional ICAPM language, the manager has intertemporal hedging demand against changes in the investment opportunity set, only that here the changes in the set are influence by the manager's actions.

3 Optimization

The manager's objective is to optimally allocate the fund's resources to maximize his future utility. As already noted, the resulting investment strategy may consequently not be in the best interest of the shareholders. They accept this agency loss due to the fact that the manager allows them to gain access to a new return process. Under the optimal policy, by simple dynamic programming arguments, the total value I_t of running the fund to the manger has to be a martingale. I_t is defined as

$$I_t\left(W_t, M_t\right) \equiv \int_0^t e^{-\rho s}\left(\beta W ds + \gamma dM\right) + e^{-\rho t} V\left(W_t, M_t\right) \tag{4}$$

where the first term, i.e. the time and stochastic integrals, are the flow payoffs to the manager that have already occurred up to time t. The second term summarizes the expected future payoffs via the recursive definition.

Apply Ito's lemma to I_t and multiplying by through by $e^{\rho t}$, we can derive the differential equation describing the dynamics of the value function V:[3]

$$e^{\rho t} dI = \mathcal{L}V dt + V_W W \pi \sigma dZ - V_W \gamma dM + V_M dM_t + \beta W dt + \gamma dM - \rho V dt \tag{5}$$

Taking conditional expectations $\mathbb{E}_t\left[\cdot\right]$, and for the moment ignoring the technical issue of the square integrability of V_W in the term $V_W W \pi \sigma dZ$, we have the stochastic integral being equal to zero. Under the optimal policy, I_t must have zero drift, so that we are left with

$$\max_{\pi}\left[\mathcal{L}V + \beta W - \rho V\right] dt + \left[-V_W \gamma + V_M + \gamma\right] dM = 0 \tag{6}$$

Let us now interpret this equation.

When assets under management W_t are close to meeting the performance benchmark M_t, the manager faces a tradeoff when deciding how much risk to take. On one hand, taking on more risk leads to a higher drift as $\mu - r > 0$, and therefore to a possible performance bonus on top of increasing the asset base. On the other hand, taking on more risk also has downsides. Realizing a new high in the asset base is not purely beneficial, in that the performance benchmark the manager is measured against is pushed up. Also, taking on more risk can lead to a larger loss, which will then set the agent on a path of withdrawals and more negative drift on the asset base.

Technically, as the maximum process can only increase at the point $W = M$, dM is equal to zero on the whole set where W is away from its maximum, i.e. $\{W < M\}$. When $dM > 0$,

[3] A technical note will be of interest here: dM_t is an increasing process and thus for any (semi-)martingale X we have $\langle X, M\rangle = 0$, i.e. there are no second order terms w.r.t. M_t in the Ito expansion. Therfore, the presence of M_t only results in a first order derivative V_M.

optimality requires the second term of (6) to be zero, so that $\gamma(1-V_W)+V_M=0$. We are thus led to the following boundary condition at $W=M$:

$$\gamma\left(1-V_W\left(M,M\right)\right)+V_M\left(M,M\right)=0 \tag{7}$$

A natural second boundary condition holds when the fund enters bankruptcy, which we assumed to happen at the point $W=\varepsilon M$. At this point, the manager cannot expect any future benefits from continued employment. Consequently, close to the bankruptcy point the manager faces the following tradeoff. On the one hand, taking little risk will slowly push him toward the bankruptcy point as the withdrawals introduce a negative drift on the asset base. Along the way, he will still enjoy a steady stream of flow benefits βW. On the other hand, taking larger risks will lead to something akin to doubling up: either the fund hits bankruptcy even faster, or the additional risk pays off, and the investment opportunity set shifts in the managers favor as performance increases and withdrawals weaken. Thus, although the expected bankruptcy time comes closer, expected future payoffs do not necessarily diminish. We will show in the subsequent section that this doubling up results in a controlled increase in exposure that still allows a solution to the problem.

Technically, at the bankruptcy point εM, the value function must converge to zero, i.e.

$$\lim_{W\to\varepsilon M} V\left(W,M\right)=0 \tag{8}$$

In the intermediate stages $W\in(\varepsilon M,M)$, the manager is essentially facing a deterministically changing investment opportunity set, as M is locally fixed. Consider the behavior of $V\left(W,M\right)$ on $\{W<M\}$. Recall that $dM=0$ on $\{W<M\}$, so that only the first term of (6) is left. This gives rise to the following second-order PDE

$$\max_{\pi}\left\{\mathcal{L}V+\beta W-\bar{r}V\right\}=0 \tag{9}$$

Although only derivatives w.r.t. W are present, this is a PDE as M is still present as the second variable - although M is fixed, the agent still accounts for possible shifts in M that become more like as W increases.

To summarize, the properties of the maximum process allow us to ignore M locally on $\{W<M\}$ as it behaves like a constant, which leads to a classical HJB reminiscent of the original Merton problem. Only at $\{W=M\}$ does the maximum process affect the differential of dI_t directly, which results in a boundary condition. The value function, however, accounts for the impact of M even away from the boundary $W=M$.

4 Solution

In this section, we reduce the equation from a PDE to an ODE by an appropriate change of variable. This will allow us to make more precise statements about the manager's behavior, as summarized in the main proposition in subsection 4.4. We can derive the optimal risk-taking at the maximum point in closed form. Furthermore, we are able to derive some technical results on the concavity of the value function and its implications. Lastly, with the ODE in hand, we are now in a position to easily derive numerical solutions to the dynamic programming problem.

4.1 Change of variables

As the agent is risk-neutral, there is no wealth effect, and our economic intuition tells us that the only things that matter in the current formulation are the distance of W to M that defines the current investment opportunity set and the overall scale of the problem. Recall that the measure of distance is $\frac{W}{M}$. Thus, the same level of $\frac{W}{M}$ should lead to the same portfolio behavior irrespective of the level of M. Conjecture the value function to be of the form $V(W, M) = f(u)M$, where $u \equiv \frac{W}{M} \in [0, 1]$, i.e. the scale of the problem is measured by M. In subsequent discussions, we will refer to $f(u)$ as the value function unless otherwise stated. The derivatives of V can be written as $V_W = f'(u)$, $V_{WW} = f''(u)\frac{1}{M}$ and $V_M = f(u) - uf'(u)$. Plugging these into our above PDE (9) and cancelling out M_t, we get the following ODE for the function $f(u)$ on $u \in (0, 1)$:

$$\rho f(u) = \max_{\pi} \left\{ [\pi(\mu - r) + r - \beta - h(u)] uf'(u) + \frac{1}{2}\pi^2\sigma^2 u^2 f''(u) \right\} + \beta u \tag{10}$$

The function $h(u)$ is not directly controlled by π.

4.2 ODE

For $f'' < 0$, we have a concave value function, and the optimal risk-taking is well defined as

$$\pi^*(u) = -\frac{\mu - r}{\sigma^2}\frac{f'(u)}{f''(u)u} \tag{11}$$

Of course, there is still scope for risk-taking to explode as $u \to \varepsilon$, something we will interpret in this model as gambling for resurrection. This should not be confused with the $f'' = 0$ case, were the gambling is immediate as the value function becomes ill-defined. This can be easily seen by the fact that $(\mu - r)uf'(u) > 0$ by assumption on f' and $\mu - r$, so that the manager will immediately pick the highest possible value for π.

After plugging π^* from (11) into the HJB equation (10), we are left with the following non-linear ODE describing the function $f(u)$ on $u \in (\varepsilon, 1)$:

$$\rho f(u) = -\phi \frac{f'(u)^2}{f''(u)} + [r - \beta - h(u)] f'(u) u + \beta u \tag{12}$$

where we used the shorthand $\phi = \frac{1}{2}\left(\frac{\mu - r}{\sigma}\right)^2 > 0$.

Recall that $V_M = f(u) - u f'(u)$, so that the boundary condition at meeting the performance criterion $W = M$, i.e. $u = 1$, can be written as

$$\gamma + f(1) = (1 + \gamma) f'(1) \tag{13}$$

Using this condition and the function (11), under the assumption $f'' < 0$, we can solve for the risk allocation at $u = 1$, which is

$$\pi^*(1) = \frac{2}{\mu - r}\left\{\rho(\gamma + 1) - (r - \beta) - (\rho\gamma + \beta)\frac{1 + \gamma}{f(1) + \gamma}\right\}$$

The second boundary condition at $W = \varepsilon M$ is easily translated. For any $\varepsilon > 0$, the boundary condition becomes

$$f(\varepsilon) = 0$$

Note that $\varepsilon = 0$ leads to a singularity in the ODE equation.[4]

4.3 Concavity of the value function

We can now establish the concavity of the value function in the cases it exists.

Concavity at $u = 1$. Suppose $f'' \geq 0$ at $u = 1$. We will prove by contradiction that $f'' < 0$. For sake of the argument, assume that the manager faces an maximum leverage constraint $\bar{\pi}$. In the actual problem, we have $\bar{\pi} = \infty$, so that taking large enough $\bar{\pi}$ approximate the actual problem. If the value function is not strictly concave, i.e. $f'' \geq 0$, then the manager will pick

[4] At $u = 0$, if $f(0) = 0$ and $|f''(0)| < \infty$, then $f'(0) = 0$. We differentiate (12) w.r.t u.

$$\bar{r} f' = -\phi \frac{2 f' (f'')^2 - (f')^2 f'''}{(f'')^2} + \theta\left(f'' u + f'\right) + \theta' f' u + \beta$$

If $u = f(0) = f'(0) = 0$ holds, we get a contradiction $0 \neq \beta$. Thus, $f'' = \{-\infty, 0, \infty\}$ and the equation exhibits a singularity at $u = 0$. Therefore, we assumed a bankruptcy point $\varepsilon > 0$, albeit close to 0.

the maximum leverage $\pi^* = \bar{\pi}$. The HJB then becomes

$$\begin{aligned} \rho f(1) &= [\bar{\pi}(\mu - r) + r - \beta] \frac{f(1) + \gamma}{1 + \gamma} + \frac{1}{2}\bar{\pi}^2 \sigma^2 u^2 f''(1) + \beta \\ \Longleftrightarrow \frac{f(1)}{1+\gamma}[\rho(1+\gamma) - \{\bar{\pi}(\mu - r) + r - \beta\}] &= [\bar{\pi}(\mu - r) + r - \beta] \frac{\gamma}{1+\gamma} + \frac{1}{2}\bar{\pi}^2 \sigma^2 u^2 f''(1) + \beta \end{aligned}$$

where we used $h(1) = 0$. As $\mu - r > 0$ and $f(1) > 0$, for a high enough $\bar{\pi}$ the LHS will become negative. However, the RHS is positive for $f''(1) \geq 0$ and high enough $\bar{\pi}$, a contradiction. We conclude that $f''(1) < 0$ for large enough $\bar{\pi}$ – at $u = 1$ the value function has to be concave.

The restriction $f''(1) < 0$ in turn imposes a condition on $f(1)$ as

$$f''(1) = \frac{-\phi(f(1) + \gamma)^2 / (1 + \gamma)}{f(1)\{\rho(1+\gamma) - r + \beta\} - (r\gamma + \beta)} < 0$$

implies $f(1) > 1 - \frac{(\gamma-1)(\bar{r}-r)}{(r\gamma+\beta)}$.

Concavity at $u \in (\varepsilon, 1)$ Suppose now there exists a point $\hat{u}$ where the function becomes convex while moving backwards from $u = 1$, i.e. $f''(\hat{u}) = 0$ at some $\hat{u} \in (\varepsilon, 1)$. By $f''(u) < 0$ on $u \in (\hat{u}, 1]$, we have $f'(\hat{u}) > f'(1) > 0$. Recall that at $f''(\hat{u}) = 0$, the agent takes the maximum allowed risk-position $\pi^* = \bar{\pi}$. Thus, we can write the HJB as

$$\rho f(\hat{u}) - [\bar{\pi}(\mu - r) + r - \beta - h(\hat{u})] \hat{u} f'(\hat{u}) = \beta$$

Assuming $f(\hat{u})$ and $f'(\hat{u})$ are bounded away from both zero and infinity, there exists a large enough $\bar{\pi}$ to make the LHS negative. As the RHS is always non-negative as $\beta \geq 0$, we have a contradiction. The point $\hat{u}$ s.t. $f''(\hat{u}) = 0$ does not exist on $(\varepsilon, 1)$. We conclude that the value function is concave on $[\varepsilon, 1]$. Under the assumption that we are facing a classical solution, i.e. $f'' < 0$ for all $u \in (\varepsilon, 1)$, we have thus that $f'(\varepsilon) > f'(1)$.

Concavity at $u = \varepsilon$**.** At ε we have $f(\varepsilon) = 0$. We can now write $f''(\varepsilon)$ as

$$f''(\varepsilon) = \frac{\phi f'(\varepsilon)^2}{\varepsilon[(r - \beta - h(\varepsilon)) f'(\varepsilon) + \beta]}$$

so that the condition

$$(r - \beta - h(\varepsilon)) f' + \beta = (r - h(\varepsilon)) f' + \beta (1 - f') < 0$$

implies $f''(\varepsilon) < 0$. This condition holds for $f'(\varepsilon) > 1$ and $h(\varepsilon) > r$ as long as both f' and h remain bounded. Otherwise, if $\lim_{u\to\varepsilon} h(u) = \infty$, then the concavity of the value functions vanished at ε, i.e. $\lim_{u\to\varepsilon} f''(u) = 0$. From the concavity of f up to point ε, we know that $f'(\varepsilon) > f'(1) > 0$. By assumption on $h(\cdot)$, $h(\varepsilon) > r$ holds.

We conclude that the function f is concave on $(\varepsilon, 1)$ if it exists.

4.4 Main proposition

Let us summarize the results of this section in a proposition:

Proposition 1 *The manager's value function $V(W, M) = f(u) M$ is described by the following ODE for f*

$$\rho f(u) = -\phi \frac{f'(u)^2}{f''(u)} + [r - \beta - h(u)] f'(u) u + \beta u$$

and the boundary conditions

$$\gamma + f(1) = (1+\gamma) f'(1)$$

and

$$f(\varepsilon) = 0$$

The manager will take positive risk at $u = 1$ according to

$$\pi^*(1) = \frac{2}{\mu - r} \left\{ \rho(\gamma+1) - (r-\beta) - (\rho\gamma+\beta) \frac{1+\gamma}{f(1)+\gamma} \right\}$$

When the value function exists, and is bounded away in its zeroth and first derivative from 0 and ∞ for any $u > \varepsilon$, it is concave on $(\varepsilon, 1]$.

An immediate implication is that for any given set of parameter choices, any further constraints that are imposed on the agents portfolio strategies lead lower risk-taking at $u = 1$. Consider two possible scenarios denoted by A, i.e. $_A$, and B, i.e. $_B$, respectively. Suppose that the outflow for A is everywhere at least as strong as for B, i.e. $h_A(u) \geq h_B(u)$ for all u – A is facing more aggressive investors an is in a more constrained scenario. Then, by optimality, we must have $f_B(1) \geq f_A(1) \geq 0$, and consequently $\pi^*_B(1) > \pi^*_A(1)$ - the less constrained manager will take larger risks at the maximum point than the more constrained manager. The intuition is the following. The more constrained manager still takes some risk, but will tread more carefully as there is a large area of the state space where the tradeoff between risk and return becomes significantly worse. Essentially, a misstep is more costly for the more constrained manager as he is facing a more distorted investment opportunity set away from $u = 1$. Dynamically accounting for this, he will behave more cautiously around $u = 1$, where the investment opportunity set is equal to the less constrained one's.

Without any bounds on π, the value can become infinite for $\beta > 0$ when the withdrawal function $h(\cdot)$ does not increase fast enough as u decreases – the agent can increase risk-taking to have enough probability mass on a perpetual positive wealth drift case, yielding infinite utility. We need a strong enough withdrawal function $h(u)$ on assets under management to shift enough probability mass into the negative drift part of the u space (where $h(u)$ is large) to get finiteness. Another way to ensure finiteness is to have a strong enough hitting time, i.e. $f(\varepsilon) = 0$ for some $\varepsilon > 0$. Although any hitting time τ_ε will be $a.s.$ finite, the integral of future utility, $\int_t^{\tau_\varepsilon} (\cdot)\, ds + (\cdot)\, dM$, might still not converge for too low levels of ε.

We can derive the closed form solution of Panageas and Westerfield 2008 by setting $h(\cdot) = \varepsilon = \beta = 0$, so that the manager is only compensated via a proportion of profits, γ. He neither receives any proportion of asset under management, nor does he face possible outflows. This makes the underlying ODE autonomous, and we can now solve (12) in closed form. Denote the solution in this case by f_0, which takes the form

$$f_0(u) = A_0 * u^\eta$$

where η solves the fundamental quadratic $r\eta^2 - \eta(r + \phi + \rho) + \rho = 0$ with $\eta_1 \in (0, 1)$ and $\eta_2 > 1$. As (12) is non-linear, only one η will apply. As we require concavity, we pick the η that lies in $(0, 1)$. From (13), we pin down $A_0 = \frac{\gamma}{\eta(1+\gamma)-1}$. Thus, $\pi_0(u) = -\frac{1}{\eta-1}\frac{\mu-r}{\sigma^2} > 0$, which is a constant. Agents behave as if they were CRRA risk averse investors in the classic Merton model. A solution, i.e. $\eta \in (0, 1)$, exists only if $\mu - r > 0$ and if $A_0 > 0$, which holds if and only if $\eta(1 + \gamma) > 1$. Thus, we need a large enough γ to get a meaningful solution.[5]

To solve the ODE numerically, we use the following parameter values: the mean of the risky asset return process is $\mu = .1$ and the instantaneous volatility is $\sigma = .2$. The market discount rate is $r = .04$, so that $\phi = .045$. The subjective discount rate of the manager is $\rho = .1$. The contract is literally a 2 and 20 contract, that is $\beta = .02$ and $\gamma = .2$. Bankruptcy occurs at a point $\varepsilon = .1$, i.e. when assets under management have fallen to 10% of their historic high. The withdrawal function $h(u)$ is parameterized by

$$h(u) = \kappa \frac{1 - u}{u - \varepsilon}$$

i.e. withdrawals are accelerating as u decreases. I use $\kappa = 1$ as the benchmark case. We note that, assuming $h(\cdot) = \varepsilon = \beta = 0$, the parameters do not lie within the range that allows for a closed form solution, as $\eta = .625$, so that we would require a large updside payment $\gamma > .6$ to

[5]The closed form solution is the basis for the pertubation method closed-form approximations that is presented in a web-appendix. This method complements the numerical solutions to the ODE provided in the paper.

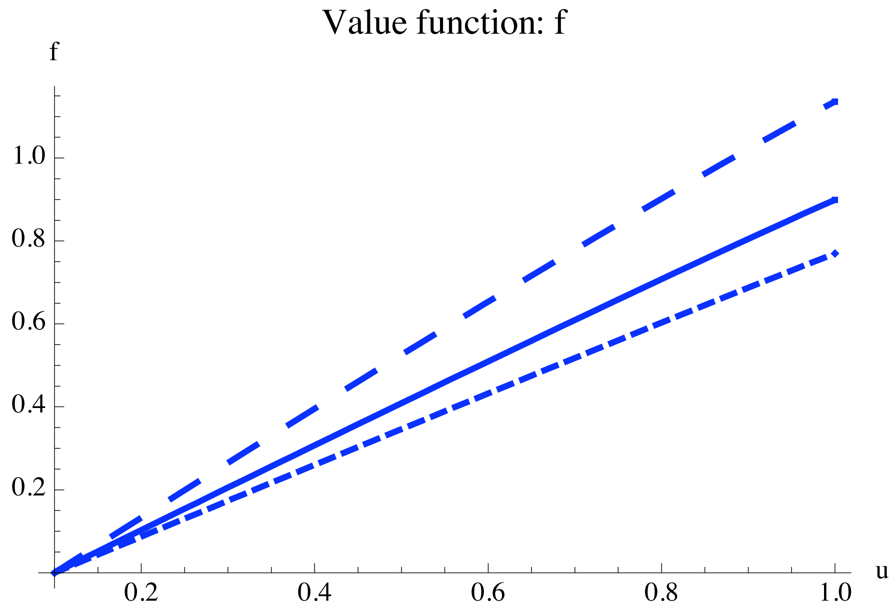

Figure 1: **Value function** $f(u)$ for different withdrawals function parameteriziations: $\kappa = 1$ (solid line), $\kappa = 2$ (dashed line), $\kappa = .5$ (widely dashed line)

get a positive value function, i.e. $A_0 > 0$.

Figures 1 and 2 provide numerical solutions to the value function f and the optimal risk taking π^* respectively. We see from figure 1 that the value functions increase as the withdrawals become less extreme. This translates into more conservative investment strategies as we can see from figure 2. The conservatism close to meeting the performance point for the more constrained (i.e. facing stronger withdrawals) manager is reversed as assets drop further below the reference point – the manager facing the strongest outflows will take more risks at higher levels of u for a large section of the state space.

5 Waiving fees

In the summer of 2007, Goldman Sachs reportedly offered to waive the management fee β for some of its funds under duress. This was generally understood as a move by the investment bank to stem the outflows from the funds in question and attract new capital. This section examines if there can be purely economic reasons outside of raising new capital to waive fees.

Suppose that $W_t \in (\varepsilon M, M)$. Over the interval dt, the management fee reduces assets under management by $\beta W dt$, but leads to a payment of $\beta W dt$ to the manager. If the agent were to waive the fee, W would remain at its previous level. The agent will thus waive the fee if and

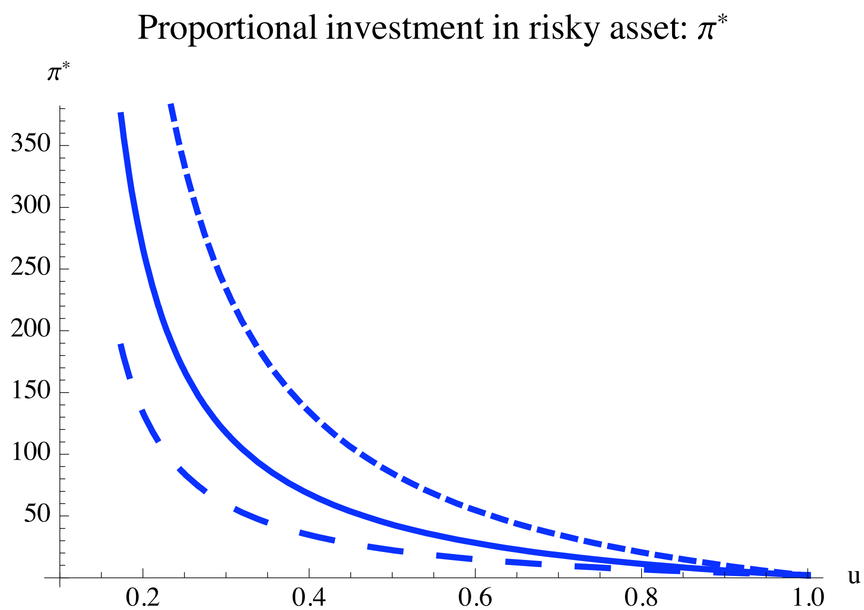

Figure 2: **Proportional investment** $\pi^*(u)$ in the risky asset for different withdrawals function parameteriziations: $\kappa = 1$ (solid line), $\kappa = 2$ (dashed line), $\kappa = .5$ (widely dashed line)

only if

$$\begin{aligned} V(W - \beta W dt, M) + \beta W dt &\leq V(W, M) \\ \iff 1 &\leq \frac{V(W, M) - V(W - \beta W dt, M)}{\beta W dt} = V_W(W, M) \end{aligned}$$

where we used the definition of the derivative in the second line. Applying the change of variables, the manager will waive the fee if and only if

$$1 \leq f'(u)$$

The rule for waiving the fee then is simple: whenever $f'(u)$ lies above the line of $u = 1$, the manager will drop the fee and face an alternative ODE with $\beta = 0$ locally. The intuition for this is the following: By waving the fixed fee, the agent can reduce the instantaneous outflow by a small amount. On its own, this would not make much of a difference, but dynamically the fee causes feedback effects that might be large enough to outweigh foregoing current compensation. This is because dynamically, even small flows have a large impact on the possible future investment opportunity sets. As we are only examining concave f, we know that f' is monotonically decreasing. This means that if a fee is ever waived, it will happen when the company is far away from its maximum.

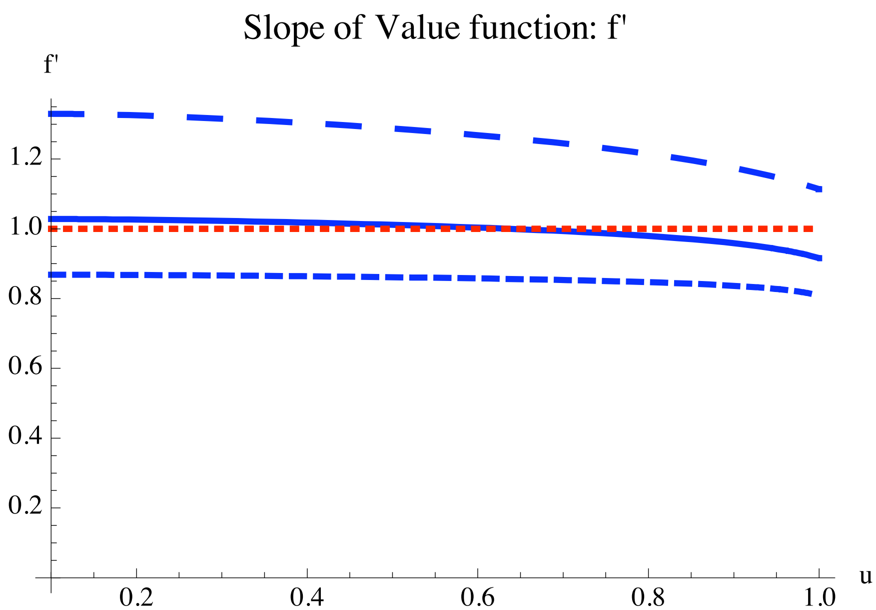

Figure 3: **Slope of value function** $f'(u)$ for different withdrawals function parameteriziations: $\kappa = 1$ (solid line), $\kappa = 2$ (dashed line), $\kappa = .5$ (widely dashed line). The dotted line is the cutoff for waiving the fixed fee β.

Figure 3 shows three different parameterizations of outflows. We see that for our benchmark case $\kappa = 1$, the manager would like to suspend the fixed fee β when assets under management drop below 65% of their historical high, the point where the solid line crosses the dotted line. With $\kappa = 2$, outflows are so strong that the small fixed fee does not impact the future personal investment opportunity set enough to warrant giving up on this sure source of income. In this case, the manager will never waive his fee. This result is reversed for $\kappa = .5$. In this case, the outflows function is sensitive enough to the fixed investment fee that the manager will find it beneficial to always waive the fixed component to his compensation contract and only collect upside payments.

6 Other Extensions

Fund inflows. Investors seek successful funds the same way they shun unsuccessful ones. We can model such fund inflows as taking place when the fund has some period of good returns, i.e. $\delta dM > 0$. Of course, this term does not directly influence the compensation of the agent (it is still γdM), but only influences the agent's value function via changing W. As (4) remains the same, only the boundary condition at $W = M$, i.e. $u = 1$, changes:

$$\gamma + f(1) = (1 + \gamma - \delta) f'(1)$$

Many successful funds are actually closed to the investment public, e.g. the Renaissance flagship fund Magellan. This stems from possible adverse market impact once the fund grows too large for a given set of investment strategies. We will ignore such impacts of the fund's size on the investment opportunity set. For technical reasons, we assume that $\delta < \gamma$.

Our numerical examination showed that δ only influences risk-taking of the fund near $u = 1$. As u moves further away, the influence of δ on risk-taking diminishes and the differences in the risk-taking behavior between $\delta = 0$ and $\delta > 0$ vanish quickly.

Starting a hedge fund. Anecdotal evidence has it that younger funds take more risk. Such behavior can be easily modeled within our framework by simply setting the initial benchmark $M_0 > W_0$ - the fund is immediately subject to withdrawal pressure for non-performance. Once again, such a 'deep out of the money' initial position fundamentally affects the funds risk-taking. Only very risky strategies will ever have a chance of netting the manager a generous payday. If young funds are on average further away from their benchmarks than mature funds, then on average younger funds will take more risks.

Moving benchmark. Investors might have more stringent expectations in that they expect the fund to beat a positive benchmark rate of return. For a moving benchmark, we follow Goetzmann et al 2003 and redefine M_t: M_t still increases when W_t reaches M_t, but also increases deterministically according to some target rate g on $\{W < M\}$. Thus, on $\{W < M\}$, we have

$$dM = gMdt$$

where g is the contractual return rate performance is measured against.

This changes the HJB 6 to a fully fledged PDE, i.e. there is a derivative w.r.t. dM that does not vanish on $\{W < M\}$:

$$\rho V = W\left[\pi(\mu - r) + r - \beta - h(u)\right] V_W + gMV_M + \frac{1}{2} W^2 \pi^2 \sigma^2 V_{WW} + \beta W$$

However, the model is still reducible to an ODE by the previously introduced change of variables $u \equiv W/M$. Recalling that $V_M(W, M) = f(u) - uf'(u)$, we can reduce to the following ODE

$$(\rho - g) f(u) = \max_{\pi} \left\{ \left[\pi (\mu - r) + (r - g) - \beta - h(u)\right] u f'(u) + \frac{1}{2}\pi^2 \sigma^2 u^2 f''(u) \right\} + \beta u \quad (14)$$

Thus, the effect is similar to a uniform downward shift in the subjective and market discount rates ρ and r respectively, while holding the excess return $\mu - r$ constant.

7 Conclusion

I examined the implications of path-dependent fund flows on a fund manager's investment decisions. The fund-manager had a standard incentive contract with a fixed fee of assets under management and upside compensation if assets under management rise above their high-water mark. Investors started withdrawing money as soon as the fund's assets under management dropped below the historic high, resulting in path-dependent withdrawals. The drawdown level, the relative distance of current assets to their high-water marks, supplied the dynamic reference point for withdrawals, with increasing outflows for higher drawdown levels. Such withdrawals resulted in a feedback loop that tilted the manager's 'personal' investment opportunity set, i.e. the investment opportunity set corresponding to his personal compensation. I derived the optimal ordinary differential equation describing the manager's investment decisions that gave rise to the following results: (i) the manager becomes more conservative the closer the fund's assets are to the high-water mark; (ii) furthermore, the possibility of stronger future outflows makes the manager more conservative close to the high-water mark; (iii) this conservatism is reversed as the fund becomes more drawdown - the stronger current outflows, the more risk the manager will take on even though the risk-return tradeoff stays constant; (iv) it can be optimal for the manager to waive the fixed investment fee even when there is no prospect of attracting any new capital via capital inflows.

References

[1] Chevalier, J., G. Ellison, 1997, "Risk taking by mutual funds as a response to incentives," Journal of Political Economy 105, 1167-1200

[2] Grossman, S.J., Z. Zhou, 1993, "Optimal strategies for controlling drawdowns," Mathematical Finance 3, 241-276

[3] Panageas, S., M.M. Westerfield, 2006, "High-water marks: high risk appetites? Convex compensation, long horizons, and portfolio choice," Working Paper

[4] Goetzmann, W.N., J.E. Ingersoll, S.A. Ross, 2003, "High-water marks and hedge fund management contracts," Journal of Finance 58, 1685-1717

[5] Sannikov, Y., 2007, "A continuous-time version of the principal-agent problem," Working paper

[6] DeMarzo, P.M., Y. Sannikov, 2006, "Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model," Journal of Finance 61, 2681-2724

[7] He, Z., 2007, "Optimal executive compensation when cash flows follow a geometric Brownian motion," Working paper

[8] Heinricher, A.C., R.H. Stockbridge, 1991, "Optimal control of the running max," SIAM Journal of Control and Optimization 29, 936-953

[9] Carpenter, J.N., 2000, "Does option compensation increase managerial risk appetite?" Journal of Finance 55, 6, 2311-2331

[10] Hugonnier, J., R. Kaniel, 2007, "Mutual fund portfolio choice in the presence of dynamic flows," Working paper

[11] Basak, S., A. Pavlova, A. Shapiro, 2007, "Optimal asset allocation and risk shifting in money management," Review of Financial Studies forthcoming

[12] Fung, W., D.A. Hsieh, 1997, "Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds," Review of Financial Studies 10, 275–302

[13] Protter, P., 2004, "Stochastic integration and differential equations," 2nd edition, Springer Verlag

[14] Karatzas, I., S.E. Shreve, 1991, "Brownian motion and stochastic calculus", 2nd edition, Springer Verlag

Appendices

A Verification Proof

This section will state the verification proof for the proposition in the text.

Proof.

We need to show that any alternative strategy π does worse than the optimal strategy π^*. We will follow standard forms of verification theorems, e.g. Scheinkman and Xiong 2003, with particular similarity to Panageas and Westerfield 2006. From our numerical estimations, we can immediately see that the optimal f' is everywhere bounded. Let us assume for the moment that the manager exits the market when M hits a level $\bar{M}$. This will allow us to define an appropriate sequence of stopping times so that we have bounded variables W and M. Define the stopping time $\tau = \tau_\varepsilon \wedge \tau_{\bar{M}}$, where we know that τ is $a.s.$ finite. Then for any time $t < \tau$, integrating up dI_t from (5) from 0 to τ, we have

$$I_\tau - I_0 = \int_0^\tau e^{-\rho t} \left[\mathcal{L}V - \rho V + \beta W\right] dt + \int_0^\tau e^{-\rho t} V_W W \pi \sigma dZ + \int_0^\tau e^{-\rho t} \left[\gamma + V_M - V_W \gamma\right] dM$$

From the boundary condition (7), we have the third term on the LHS equal to zero. Furthermore, $\mathcal{L}V - \rho V + \beta W$ is concave in π by assumption $f'' < 0$, and by the HJB we have $\mathcal{L}V - \rho V + \beta W \leq 0$ for all permissible strategies $\pi \in L^2$, with π^* yielding equality. Finally, we know by $0 < W_t \leq M_t \leq \bar{M}$ that W_t is bounded for all $t \leq \tau$. For arbitrary strategies π we know that the stochastic integral is bounded below by $-I_0$, as the value function is everywhere non-negative and payouts are positive

$$\begin{aligned} \int_0^\tau e^{-\rho t} V_W W \pi \sigma dZ &= I_\tau - I_0 - \int_0^\tau e^{-\rho t} \left[\mathcal{L}V - \rho V + \beta W\right] dt \\ &\geq I_\tau - I_0 \geq I_0 \end{aligned}$$

As the stochastic integral w.r.t. to dZ is a local martingale, and as it is bounded from below by $-I_0$, we conclude that it is a supermartingale. Only when π^* is followed does $I_\tau \leq I_0 = V_0$ hold with equality, as f' bounded implies the stochastic integral is a true martingale, and $\mathcal{L}V - \rho V + \beta W = 0$ by our optimization. Taking $\bar{M} \to \infty$ concludes the proof. ■