Fake Alphas, Tail Risk and Reputation Traps

Marco Di Maggio*
Columbia Business School
mdimaggio@columbia.edu

June 6, 2015

Abstract

This paper develops a model of active asset management where a fraction of managers have skill and invest alongside unskilled managers who can generate active returns at a disutility. Because of agency frictions, star funds exploit their status by extracting higher rents from investors and by exposing them to tail risk, while poor performers may end up in a reputation trap, limiting their ability to attract investment. These effects exacerbate fluctuations, especially in times of high-volatility. Moreover, there exists a feedback effect between the managers' reputation and the compensation for selling disaster insurance, which exacerbates agency frictions.

*I thank Daron Acemoglu, Abhijit Banerjee, Jonathan Berk, Hank Bessembinder (discussant), Jules H. van Binsbergen, Amil Dasgupta, Peter DeMarzo, Fernando Duarte (discussant), Mehmet Ekmekci, Vincent Glode, Brett Green, Zhiguo He, Andrew Karolyi, Arvind Krishnamurthy, Qingmin Liu, Andrey Malenko, Gustavo Manso, Alan Moreira, Martin Oehmke, Marco Ottaviani, Andrea Prat, Steve Ross, Yuliy Sannikov, Ravindra Sastry, Antoinette Schoar, Alp Simsek, Jeremy Stein, Alireza Tahbaz-Salehi and seminar participants at MIT, Columbia University, the 2015 FIRS Conference in Reykjavik, the 3rd ITAM Finance Conference, the 2014 FIRS Conference in Québec City, the 2014 FMA Annual Meeting and the 9th Csef-Igier Symposium on Economics and Institutions for helpful comments and suggestions. Byeong-Je An provided outstanding research assistance. Usual disclaimer applies.

1 Introduction

In 2008 the global asset management industry managed a total of $90 trillion, through pension funds, mutual funds and hedge funds. This tallies with the substantial increase in the institutional ownership of corporate equity, from 7.2% in 1950 to 78% in 2007 (French (2008)).[1] Their sheer size and the ability to engage in sophisticated trading strategies through leverage, derivatives and short positions makes institutional investors key players in the financial markets. Moreover, their performance is monitored constantly by capital providers (investors and lenders), which means fund managers must be concerned with perceptions of their ability or reputation.

However, the effectiveness of market discipline in shaping institutional investors' conduct has been seriously challenged by the recent financial crisis. The collapse of financial markets' discipline has been cited as a key element of crises by the Federal Reserve chairman, Ben Bernanke: "Market discipline has in some cases broken down."[2] What is more, institutional investors can amplify aggregate shocks by taking excess risk (Huang et al. (2011)), riding bubbles (Brunnermeier and Nagel (2004)) or hoarding liquidity when it is most needed (Ben-David et al. (2012)).[3] Is it possible that these distortions are due to managers' desire to make and keep a good reputation? How does market discipline interact with aggregate uncertainty to affect fund managers' strategies? In this paper, we show that market discipline may in fact generate perverse incentives and induce more risk-taking and higher agency costs, especially in times of high-volatility.

We begin with a model of intermediation with two classes of agents, households and investment specialists (fund managers), and analyze the households' and specialist's problems jointly. Specialists have the expertise to actively invest in financial markets, while households can only invest in a benchmark portfolio. The specialist invests in the risky asset on the households' behalf and is rewarded with a fixed fee and a performance fee. We can think of the specialist as the manager of a financial intermediary that raises funds from investors or as a fund manager who can actively manipulate the returns on his fund's portfolio. However, this intermediation is subject to moral hazard, in that specialists can generate returns by exploiting their stocks-picking ability or by timing the

[1]The Investment Company Institute (2009) reports that assets of equity mutual funds total $3.8 trillion at the end of 2008, and that 87% of those assets are under active management, as opposed to being index funds.

[2]Statement by Chairman Ben S. Bernanke on December 18, 2007.

[3]In a good number of instances financial institutions deliberately misled investors. Two among many cases are those of Bear Stearns High-Grade Structured Credit Fund, which misled investors into a sinking fund, and fund manager David Einhorn, whose short selling notoriously helped to bring down Lehman Brothers. (*New York Times, June 20, 2008*)

market, but they have short-sighted incentives to mislead the households by manipulating returns and riskiness.

We first capture the specialist's stock-picking ability, defined as the possibility of generating active returns at a disutility as in Glode (2011) and Buffa et al. (2014), to show how reputational concerns aimed at relaxing agency problems impact on specialists' incentives. The model envisages "skilled" specialists, who have an advantage in generating persistent positive returns (*true alphas*), and "unskilled" ones who incur a cost to find the most profitable investment opportunities (*fake alphas*). A specialist's reputation is defined as the households' belief that the specialist is "skilled". Specialists' reputations change over time as households observe the returns on investment. Unskilled specialists want to pool themselves with the skilled in order to attract more investment and earn higher commissions. The fear of reputation loss, therefore, leads unskilled specialists to identify the investment strategies that will generate excess returns over the risk-free asset.

The first main result of the model is its rich equilibrium dynamics. One can identify three regions, depending on the level of the specialist's reputation: a *reputation-exploitation* region, in which good-reputation specialists extract the rents associated with their "star status" by maximizing their short-term payoff; a *reputation-building* region, where specialists with intermediate reputations try to improve them by delivering higher returns, but at a rate that is decreasing in their perceived reputation; and the *reputation trap* region for specialists with poor reputations, where households' confidence in their ability is so low that they refuse to entrust any more capital, which eliminates their chances of improving their reputation.

These results relate our paper to the seminal framework introduced by Berk and Green (2004). However, while Berk and Green (2004) show that performance may not be persistent if the investors allocate their funds to the best-performing funds and the latter have lower returns in managing larger asset volumes, in our model, fund managers are not characterized by decreasing returns to scale in performing their investment strategies. Rather their reputation concerns *endogenously* generate similar implications. In fact, as a manager improves his reputation, he acquires more and more assets and his incentive to exploit his reputation also increases.

In a crisis, the agency costs are exacerbated by the high-reputation specialists' incentives to act myopically. We show, in fact, that as market volatility increases unskilled specialists become more likely to get caught in a reputation trap, while unskilled specialists are more likely to be lucky and attract households' funds. Intuitively, in a more volatile market it is harder for the specialist to affect households' beliefs, which squares with the evidence in Huang et al. (2011). Due to the effect of households' learning processes, high-reputation specialists also have more incentive to exploit their reputation, because they expect their actions not to affect households' investment decisions. This carries two implications: higher turnover among intermediaries during crisis than during boom, if they are forced to exit the market when they lose their households' support; and second, a shrinking of the region in which the specialist exerts higher effort, so that precisely at the time when households would most need the specialist's guidance, the specialist has less incentive to work for higher returns.

We then extend the model to capture aggregate uncertainty about the state of the economy and the possibility that specialists can time the market. The model can capture market-timing ability by assuming that the economy is subject to disaster risk - the relevance of which has been brought out by the recent financial crisis - and the intermediary can determine investors' exposure to it. For instance, Kelly and Jiang (2012) show that a significant part of hedge fund returns can be viewed as compensation for selling disaster insurance. A rationale for this behavior is that specialists seek to enhance their reputation by capturing a positive premium for exposing their households to such risk. We analyze this possibility in a general equilibrium framework where the reward from selling disaster insurance, namely the premium, is endogenously determined by risk-averse market-makers.

Specifically, we show that even when the specialist has the possibility of hedging disaster risk, he may elect not to do so. We consider two different cases: the investors might withdraw their capital from all the funds if the tail risk materializes i.e. a "flight to safety" episode; alternatively, the investors might reward the managers who are able to hedge the tail risk and outperform the market, i.e. investors "bet on the mavericks". The intuition is that in the event of a flight to safety the specialist will lose households' support independently of his behavior (which happens when the tail event is expected to be particularly severe), which increases his incentive to misbehave before the occurrence of such events. In fact, by exposing the households to this tail risk the specialist captures a positive premium, and he might induce returns-chasing investors to update their beliefs about his type and entrust more capital to the fund. This enables him to improve his reputation, which means that agency costs increase as the specialist can exert lower effort than in absence of tail risk.

However, fund managers like John Paulson, Greg Lippman or Michael Burry shot to fame and fortune with investment strategies that paid off during the subprime housing market crash and were rewarded with significant inflows after the crisis. We can capture this possibility in reduced form by assuming that if they hedge against the tail risk, fund managers are going to be rewarded with higher expected profits in the future. We show that even in this case, managers might prefer to increase short-term returns by exposing the households to tail risk, because betting on the crash is costly, and it leads to a decline in reputation in the short run. More importantly, we show that there exists a novel feedback loop between the specialist's reputation and the premium. In fact, in equilibrium as the specialist's reputation increases, the market makers expect the specialist to be skilled with higher probability, they fear to be on the wrong side of the market, which induce them to charge a higher premium. This, in turn, increases the specialist's incentive to expose his portfolio to tail risk even more, as a high premium will boost his short-term performance, which further improves his reputation. Thus, prices and incentives are linked through the effect of reputation.

In other words, when the crisis is expected to be severe and persistent, specialists strategically choose to get over-exposed to tail risk due to their reputation concerns in advance of the crisis. Then market discipline, instead of incentivizing the specialist to protect households from fluctuations actually makes them more vulnerable. This also implies that in normal times, it is very hard for households to distinguish between high-skilled specialists and those who sell disaster insurance, because their returns can be similarly high.

The model delivers several novel empirical implications concerning the specialists' behavior over the business cycle. Some of these predictions find support in existing empirical studies, while others provide new testable implications. We describe some of these results here and provide a more exhaustive discussion of the empirical implications of our results in the body of the paper. First, high-volatility phases are amplified by the intermediaries that manage to attract funds even when unskilled, which increases the riskiness of households' portfolios. Second, these intermediaries' track record should significantly affect their portfolio allocation. In particular, a series of positive shocks to returns affect the type and the riskiness of the strategy the specialist will pursue in a non-monotonic fashion. Third, the incentives to use tail risk exposure to enhance reputation change over time, which means that intermediaries' portfolios should become more and more negatively skewed as they expect market turmoil. Finally, as the premium that can be captured by exposing to tail risk increases, unskilled managers will have stronger incentives to sell disaster insurance, which will boost their performance and their reputation, accelerating the downward spiral. These implications can be expected to reinforce policy makers' efforts to regulate the financial industry. In fact, the financial crisis has demonstrated the importance of the role played by institutional investors, but so far the discussion has focused on their explicit incentives, mainly bonuses and stock options, and how these affect risk-taking. Here, we show that a neglected factor is *implicit incentives*, which might make specialists' actions procyclical.

Relation with the literature. To the best of our knowledge, the study of the role played by reputation concerns in conjunction with the possibility of both stock-picking and market-timing ability has been largely unexplored by both the theoretical and the empirical literature. Few exceptions are Kacperczyk et al. (2012), Kacperczyk et al. (2014), Kelly and Jiang (2012), DeMarzo et al. (2012) and Makarov and Plantin (2015). Kacperczyk et al. (2014) finds that the same fund managers that pick stocks well in expansions also time the market well in recessions, and Kacperczyk et al. (2012) propose a model based on rational inattention to explain this evidence. Kelly and Jiang (2012) show that tail risk is a key driver of hedge fund returns in both the time-series and cross-section, while DeMarzo et al. (2012) and Makarov and Plantin (2015) consider optimal incentive contracts that deter fund managers from putting the firm at risk of a low probability "disaster" or from creating "fake alphas". We provide a unifying framework to study the interaction between stock-picking and market-timing ability when no optimal contracts are available and the specialist is concerned about his reputation.

More generally, this paper contributes to the literature on career concerns in financial markets, which has shown that institutional investors's incentives to boost their reputations might lead to several distortions such as: herding behavior (Scharfstein and Stein (1990), Zwiebel (1995), Dasgupta et al. (2011) and Ottaviani and Sorensen (2006)), the prevention of information aggregation (Dasgupta and Prat (2006) and Dasgupta and Prat (2008)), limited ability to pursue arbitrage opportunities (Shleifer and Vishny (1997)), and the amplification of price volatility in financial markets (Guerrieri and Kondor (2012)).[4] An important difference between our work and these papers is that we assume that the fund manager knows his type, which generates asymmetric information between the investors and the manager. Moreover, we jointly characterize the manager's and the investors' problems and we endogenize the returns on the tail-risk strategy. Our results on the feedback effect between the premium generated by selling disaster insurance and the fund manager's incentives relate to a recent paper by Buffa et al. (2014). They study the joint determination of equilibrium asset prices and fund managers' contracts including the extent of benchmarking, and argue that agency frictions exacerbate price distortions and generate a negative relationship between risk and returns. We fix the contract and derive the manager's optimal behavior as a function of the capital invested in the fund and his reputation.

Other studies that model learning about managerial skill include Lynch and Musto

[4]Other related papers on delegated asset management include Cuoco and Kaniel (2011), Glode (2011), Gervais and Strobl (2012) and Kaniel and Kondor (2013).

(2003), Berk and Green (2004), Dangl et al. (2008) and Malliaris and Yan (2010).[5] In a seminal paper, Berk and Green (2004) were able to reconcile the early evidence on fund managers' behavior (Jensen (1968), Carhart (1997), Chevalier and Ellison (1997), and Sirri and Tufano (1998)), showing that performance may not be persistent even when managers have heterogeneous skills, if the investors allocate their funds to the best-performing funds and the latter have lower returns in managing larger asset volumes. In our model, fund managers are not characterized by decreasing returns to scale, but their reputation concerns endogenously generate similar implications: as a manager improves his reputation, he acquires more and more assets and his incentives to exploit his reputation also increase. With a focus similar to ours, Malliaris and Yan (2010) link the skewness of funds' strategies with managers' career concerns. Unlike that paper, here we set out a rich but tractable dynamic model that allows us to characterize the investors' delegation decision and the fund manager's investment strategy simultaneously. Moreover, as we shall see in Section 5, the combination of reputation-building motives and tail risk generates a new set of empirical implications.[6]

Finally, our dynamic model relates to the growing literature that uses continuous time techniques to find optimal contracts in dynamic principal-agent relationships (Sannikov (2008), Biais et al. (2007), He (2009), Zhu (2012)); characterize the equilibrium in games of imperfect public monitoring (Sannikov (2007) and Sannikov and Skrzypacz (2010)); analyze the capital structure of the firm (DeMarzo and Sannikov (2006) and Bolton et al. (2011)); analyze financial frictions in macro general equilibrium models (Brunnermeier and Sannikov (2014), He and Krishnamurthy (2013) and He and Krishnamurthy (2012)) and provide a recursive characterization of equilibria in reputation games (Faingold (2005), Board and Meyer-ter Vehn (2013) and Faingold and Sannikov (2011)).[7] Methodologically, the paper most closely related to the present one is Faingold and Sannikov (2011), which characterizes the conditions that guarantee a unique equilibrium outcome. We complement that analysis by highlighting the role of reputation in shaping the risk taken by fund managers and how it is affected by the endogenous response of the premium captured by selling disaster insurance.

The paper is structured as follows. The next section presents the basic framework, highlighting the key friction between investors and intermediaries. Section 3 shows how market discipline shapes managers' behavior over time. Section 4 introduces aggregate

[5] On the theory side Basak et al. (2007) and Chapman et al. (2007), among others, characterize managers' optimal portfolio choice in the context of a convex flow-performance relationship.

[6] Other related papers include Garcia and Vanden (2009) and Malamud and Petrov (2014) which show the effect of competition and of the convexity of the managers' fee on prices in static settings respectively.

[7] For a survey of this literature see Sannikov (2012).

uncertainty to capture tail risk and to show that reputational concerns can lead managers to amplify aggregate fluctuations by trying to boost short-term performance. To conclude, Section 5 reviews the empirical evidence on the model's testable predictions and discusses extensions, and Section 6 concludes.

2 The Basic Model

Overview. We propose a continuous-time model with two classes of agents, retail investors and specialists. We are going to examine the behavior of the representative specialist. The specialist has the technology or the know-how to actively invest in a risky strategy that the retail investors cannot follow directly, so specialists receive compensation for investing in the risky asset on the investors' behalf. In terms of banking models, we can think of the specialist as the manager of a financial intermediary that procures resources from the investors or, alternatively, as a fund manager whose skill is unknown to the investors. This specialist will then provide the investors access to investment opportunities and securities that might potentially deliver higher returns. The key friction is that this intermediation relationship is subject to moral hazard. That is, the returns of the risky asset depend crucially on the specialist's actions. Whereas the literature has focused mainly on designing the optimal contract to mitigate moral hazard, our model analyzes the effect that reputational concerns have on the specialists' behavior. In this section, we present the basic framework, and in Section 4 we extend it to allow the specialist to attempt to time the market by affecting the investors' exposure to tail risk.

Setup. We consider an economy with a risk-neutral specialist facing a continuum of risk neutral investors in a continuous-time repeated game. At each time $t \in [0, \infty)$, each investor $i \in I \triangleq [0, 1]$ chooses the fraction of his unit wealth $k_t^i \in [\underline{k}, 1]$ to invest with the specialist, while the remainder $1 - k_t^i$ is invested in a benchmark portfolio that pays a return r_s. For simplicity and without loss of generality, we assume that the benchmark portfolio is riskless. The lower bound $\underline{k} \geq 0$ can be interpreted as the fraction of funds invested that are not costless to withdraw.[8] The specialist invests in a strategy that is subject to idiosyncratic diffusion risk, and whose per dollar returns in excess of the riskless rate are

$$dR_t = a_t dt + \sigma dZ_t \tag{1}$$

where $a_t \in [0, \alpha]$ is an unobserved action that the specialist takes at each point in time,

[8] If $\underline{k} = 0$, then the analysis would be unaffected, with the only difference being that the specialists with no funds to manage would leave the market.

while $Z_t \equiv \{Z_t, \mathcal{F}_t; 0 \leq t < \infty\}$ is a standard Brownian motion on a complete probability space $(\Omega, \mathcal{F}, \mathcal{P})$, and σ is a positive constant capturing the volatility of returns. Then, $(R_t)_{t\geq 0}$ is a noisy public signal whose evolution depends on the choice of effort a_t. The assumption that only the drift of dR_t depends on the specialist's action corresponds to the standard constant support assumption in discrete-time repeated games, and there is no loss of generality in assuming that the specialist's action affects the drift of the strategy returns (1) linearly.[9] The specialist cannot directly influence the volatility of returns σ, as in this case his action would become immediately observable.

We assume that the investors are anonymous: at each time t the public information includes the aggregate distribution of the investors' actions $\bar{k}_t$, which captures the assets under management at time t, but not any individual investor's action. This represents the fact that investors are dispersed and cannot write an explicit contract with the specialist, as is usually the case in the retail asset management industry. Therefore, the specialist must rely on his reputation to induce investors to place capital with him.

Contracts and Payoffs. Contracts can be described by two parameters: a management fee f and a performance fee γ. We assume that both are fixed characteristics of the fund. The management fee is restricted to be non-negative and is per-dollar-managed. The performance fee is symmetric: for example, if at instant t the specialist has a return R_t, his compensation will be $\gamma (R_t - f) + f$ per dollar managed, while the investors would receive $(1-\gamma)(R_t - f)$. The performance fee contract can be thought of as an option of dt maturity on the performance and continuation value of the relationship with investors. If the specialist decides to exercise the option, he "receives" the (potentially negative) performance payment and continues the relationship; if he chooses to rescind the contract the last transfer payment is not made and the specialist exits the market.[10]

Investors have identical preferences, and by investing a fraction k_t^i of his capital with the specialist, each investor i receives the following expected flow payoff at time t:

$$u\left(a_t, k_t^i\right) = k_t^i (1-\gamma)(a_t - f) + \left(1 - k_t^i\right) r_s.$$

At each point in time, the investor has one unit of capital to invest and decides what fraction to invest in the fund and what fraction to allocate to the benchmark portfolio, anticipating that the returns to the fund depend on the specialist's effort choice a_t.

Given the aggregate investment strategy $\bar{k}$, the specialist obtains the following flow

[9]By Girsanov's theorem the probability measures over the paths of two diffusion processes with the same volatility but different bounded drifts are equivalent, i.e. they have the same zero-probability events.

[10]This symmetry in fees is crucial for tractability as noted by Moreira (2013).

payoff

$$\pi\left(a_t, \bar{k}_t\right) = \bar{k}_t\left(\gamma\left(a_t - f\right) + f\right) - a_t b \bar{k}_t - L$$

where parameter $b \in (\gamma, 1)$ captures the degree of misalignment between the two agents, and $\bar{k}_t$ is the aggregate investment with the specialist. That is, there exists a *conflict of interest* between the specialist and the investors, in that a specialist might suggest certain financial products instead of others in order to gain additional fees if that product is sold. As noted in the introduction, this possibility gained particular attention recently during the financial crisis, with subprime mortgages sold to investors that were instead eligible for the prime market.[11] Parameter b might also capture the private benefits to the specialist from diverting the funds into other privately optimal strategies.

Action a_t can be interpreted in various ways. First, it can be seen as a reduced form for the effort exerted by the specialist in picking stocks, i.e. acquiring information on the companies to invest in if (1) is a value strategy. Second, we can think of a_t as the probability with which the specialist suggests the best products given his costumers' preferences; in this case, a lower a_t means that the specialist is more prone to sell the products on which he earns higher commissions, no matter what his investors' preferences are.[12] Finally, the model is isomorphic to a setup in which the specialist can divert cash flow as assumed in the literature on dynamic contracting (see DeMarzo and Sannikov (2006) and Biais et al. (2007)).

While a specialist has more incentive to mislead his customers as the capital invested increases (as captured by the term $a_t b \bar{k}_t$), he might incur additional costs that are unrelated to the capital managed, which we capture with the parameter L. This cost might capture, for example, the running cost of the company, or the cost of hiring a quantitative research team.

Information Structure. While the investors' payoff is common knowledge, there is uncertainty about the type of the specialist θ. At time $t = 0$ investors believe that with probability $p_0 \in (0, 1)$ the specialist is a skilled type ($\theta = S$), who always chooses action $a_t = \alpha > r_s$, and that with probability $1 - p_0$ the specialist is an unskilled type ($\theta = U$), who maximizes the expected value of his profits discounted at rate r. This is meant to capture the idea that investors are unaware of the specialist's incentives, that is, they do not

[11] In a comment on the causes of the subprime mortgage crisis, *The Economist* observes that "many customers appear to have been encouraged to take out loans by brokers more bothered about their fees than their clients' ability to repay their debts" ("The trouble with the housing market," March 24, 2007, 11).

[12] More precisely, as in the credence good literature (see among others Pesendorfer and Wolinsky (2003) and Bolton et al. (2007)) his action a_t is the effort spent in matching investors' preferences with financial products. The signal dR_t could be the returns of the products suggested or the utilities derived by the investors and the parameter σ can capture the transparency of the market.

know whether there may be products to which the specialist would prefer to steer his customers. The skilled type can be thought of as the specialist who has no conflict of interest $b = 0$, who then chooses α, or as the specialist who possesses superior skill and can persistently generate a true alpha. Moreover, these incentives are persistent, which means that the specialist has the opportunity to build a reputation for honest management. Thus, the skilled specialist can be thought as a commitment type (Kreps and Wilson (1982)) while the unskilled is a strategic agent who tries to boost his reputation by imitating him.

Although fund managers and investors observe the fund's returns $(dR_s, s \leq t)$ before choosing a_t and k_t, for the purpose of their respective decision problems, the history of fund returns can be summarized in the specialist's reputation p_t.

Strategies and Equilibrium. A *public strategy* of the specialist is a random process $(a_t)_{t\geq 0}$ with values in $[0, \alpha]$ and progressively measurable with respect to $(\mathcal{F}_t)_{t\geq 0}$. Similarly, a public strategy of investor i is a progressively measurable process $\left(k_t^i\right)_{t\geq 0}$ taking values in $[\underline{k}, 1]$. In the repeated game, the investors formulate a belief about the specialist's type following their observations of $(R_t)_{t\geq 0}$. The belief process is a progressively measurable process $(p_t)_{t\geq 0}$ taking values in $[0, 1]$, where p_t denotes the probability that the investors assign at time t to the specialist being the skilled type.

Definition 1 *A public sequential equilibrium consists of a public strategy profile $(a_t)_{t\geq 0}$ of the unskilled specialist, a public strategy $\left(k_t^i\right)_{t\geq 0}$ for each investor i and a belief process $(p_t)_{t\geq 0}$ such that at all time $t \geq 0$ and after all public histories,*

1. *The strategy of the specialist maximizes his expected payoff*

$$\mathbb{E}_t \left[\int_0^\infty re^{-rs}\pi \left(a_s, \bar{k}_s\right) ds - L \mid \theta = U \right],$$

2. *The strategy of each investor i maximizes his expected payoff*

$$p_t \mathbb{E}_t \left[\int_0^\infty re^{-rs} u \left(\alpha, k_s^i\right) ds \mid \theta = S \right] + (1 - p_t)\, \mathbb{E}_t \left[\int_0^\infty re^{-rs} u \left(a_s, k_s^i\right) ds \mid \theta = U \right]$$

3. *Beliefs $(p_t)_{t\geq 0}$ are determined by Bayes' rule given the common prior p_0.*

A strategy profile satisfying condition (1) and (2) is called *sequentially rational*. A belief process satisfying condition (3) is called *consistent*. We are going to solve for an equilibrium Markovian in p.

We can simplify the above definition in two ways. First, since the investors have identical preferences, we work with the aggregate strategy $\left(\bar{k}_t\right)_{t\geq 0}$ rather than the individual strategies $\left(k_t^i\right)_{t\geq 0}$. Second, since the behavior of any individual investor is not observed by any other player and cannot influence the evolution of the public signal, the investors' strategies must be myopically optimal. Thus, we say a tuple $\left(a_t, \bar{k}_t, p_t\right)_{t\geq 0}$ is a public sequential equilibrium if, for all $t \geq 0$ and after all public histories, conditions (1) and (3) are satisfied, and if

$$k \in \arg\max_{k' \in [\underline{k}, 1]} p_t u\left(\alpha, k'\right) + \left(1 - p_t\right) u\left(a_t, k'\right) \ \forall k \in \operatorname{supp}(\bar{k}_t).$$

Finally, note that for both pure- and mixed-strategy equilibria, the restriction to public strategies is without loss of generality. For pure strategies, it is redundant to condition a player's current action on his private history, as every private strategy is outcome-equivalent to a public strategy. For mixed strategies, the restriction to public strategies is without loss of generality in repeated games with signals that have a product structure, as in the repeated games that we consider. To form a belief about his opponent's private histories in a game with product structure, a player can ignore his own past actions as they do not influence the signal about his opponent's actions.

3 Analysis

In this section we develop a recursive characterization of public sequential equilibria, which we then use throughout the paper. Lemma 1 and Lemma 2 are intermediate steps that characterize the evolution of investors' beliefs and the specialist's Hamilton-Jacobi-Bellman (HJB) equation. Proposition 1 shows a few basic properties of the specialist's equilibrium value function and sets the stage for the equilibrium characterization presented in Proposition 2. The main result of this section is set forth in Proposition 2, which describes the equilibrium behavior of specialist and investors and the resulting dynamics. The next section is devoted to the comparative statics with respect to the main parameters of the model.

A natural state variable in our model is the investors' belief p about the specialist's type. We start by characterizing the stochastic evolution of the investors' posterior beliefs on and off- the equilibrium path in the following lemma.

Lemma 1 (Belief Consistency) *Given prior belief $p_o \in [0,1]$, a belief process $\left(p_t\right)_{t\geq 0}$ is consistent with a public strategy profile $\left(\hat{a}_t, k_t\right)_{t\geq 0}$ if and only if*

$$dp_t = \left[\chi\left(\alpha, \hat{a}, p_t\right)\left(a_t - \bar{a}\left(p_t\right)\right)/\sigma\right] dt + \chi\left(\alpha, \hat{a}, p_t\right) dZ_t^p \tag{2}$$

where for each $(\hat{a}, p) \in [0, \alpha] \times [0, 1]$,

$$\chi(\alpha, \hat{a}, p) \triangleq p(1-p)\sigma^{-1}(\alpha - \hat{a}_t)$$

$$\bar{a}(p) \triangleq p\alpha + (1-p)\hat{a}_t$$

$$\sigma dZ_t^p = a_t dt + \sigma dZ_t - (p_t\alpha + (1-p_t)\hat{a}_t) dt$$

Note that in the statement of Lemma 1, $(\hat{a}_t)_{t\geq 0}$ is the strategy that the investor *thinks* the unskilled specialist is following. Thus, when the specialist deviates from his equilibrium strategy, the deviation affects only the drift of $(R_t)_{t\geq 0}$, but not the other terms in equation (2). In fact, unexpected changes in the observation process cannot raise volatility, since they are unobserved.

Equation (2) highlights the two separate forces that drive the updating. The drift term $\chi(\alpha, \hat{a}, p)(a - \bar{a}(p))/\sigma dt$ takes into account the possibility that the specialist may deviate from the expected choice of effort. In expectation, this term is zero, i.e. $\mathbb{E}(a) = \bar{a}(p)$, but this is useful in the computation of the optimal action for the unskilled advisor. In contrast, the diffusion term in (2) captures the influence of the observed signal on the evolution of beliefs. Z_t^p being a Brownian motion, this part of the updating is completely unpredictable. Intuitively, this expresses the fact that the current belief already incorporates everything that is known, so any change must come as a surprise. The representation

$$\sigma dZ_t^p = a_t dt + \sigma dZ_t - (p_t\alpha + (1-p_t)\hat{a}_t) dt$$

confirms this, showing that the change in beliefs depends on the difference between the realized signal, $a_t dt + \sigma dZ_t$, and the expected signal $\bar{a}(p) dt$.

The coefficient $\chi(\alpha, \hat{a}, p)$ of equation (2) is the volatility of beliefs: it reflects the speed with which the investor learns about the specialist's type. The lower the noise level σ, or the greater the difference between the drifts produced by the two types $\alpha - \hat{a}_t$, the more informative the signal and more pronounced the change in beliefs when the signal is observed. This is only relevant, of course, when the investors are not certain of the current type. For $p = 0$ or 1, the investors rule out any possibility of learning from the signal, i.e. the diffusion term vanishes no matter what action is taken. Once investors are certain of the specialist's type, their beliefs become insensitive to his performance.

We turn to the analysis of the specialist's problem. Since we are looking for an equilibrium that is Markovian in the investors' belief, we derive his HJB equation $V(p)$ as a function of the belief p. The specialist's problem is to find a policy function $a(p_t)$ that solves the following problem

$$V(p) \triangleq \max_{a:[0,1]\to[0,\alpha]} V(a,p)$$

subject to the stochastic evolution of beliefs p derived in Lemma 1 and $a_t \triangleq a(p_t)$. The following lemma derives the specialist's value function as a function of investors' beliefs p.[13]

Lemma 2 (HJB Equation) *Given the investor's beliefs p, investor's strategy k and the expected action â, the specialist's HJB equation is given by*

$$\begin{aligned} rV(p) &= \max_{a\in[0,\alpha]} r\pi\left(a,\bar{k}\right) + \frac{\chi\left(\alpha,\hat{a}_t,p_t\right)}{\sigma}\left[a-\bar{a}(p)\right]V'(p) \\ &\quad -\frac{\chi\left(\alpha,\hat{a}_t,p_t\right)^2}{(1-p)}V'(p) + \frac{\chi\left(\alpha,\hat{a}_t,p_t\right)^2}{2}V''(p) \end{aligned} \tag{3}$$

The specialist's value function has an intuitive interpretation. When choosing his effort, the specialist maximizes the sum of his flow payoff and his continuation value, considering how his decision a affects what the market learns about his type and the resulting impact on his reputation. That is, he weighs the short-run benefits of lowering the optimal action a against the change in his value function $V'(p)$ due to the market's updating about his type θ. Specifically, exerting greater effort enables the specialist to capture higher incentive fees γ but prevents him from capturing the broker fees b. This describes his flow payoff. What prevents the specialist from fully exploiting his position, however, is that lower a with the consequent poor performance, adversely affects his reputation p, which in equilibrium results in lower future investment k.

To show how these different forces play out in equilibrium, we now turn to the conditions that characterize sequential rationality. The investors' problem is straightforward; in fact, since they are anonymous, they cannot do better than maximize their short-sighted payoff:

$$k \in \arg\max_{k'\in[\underline{k},1]} \left[k'(1-\gamma)\left(\bar{a}(p_t)-f\right)+\left(1-k'\right)r_s\right]. \tag{4}$$

Recall that $\bar{a}(p) = p_t\alpha + (1-p_t)\hat{a}_t$ is the expected return on the investment as a function of the expected effort level $\hat{a}$.

The specialist's problem is considerably more complex, because in optimizing he takes into account the effect that his action at time t will have on the investors' beliefs, which

[13] We assume $L < \frac{f(1-\gamma)}{r} = V(1)$. If this parametric assumption is not satisfied, then the specialist never participates to the market.

changes his continuation value. A public strategy $(a_t)_{t\geq 0}$ is sequentially optimal if for all $t \geq 0$ and after all public histories

$$a_t \in \arg\max_{a' \in [0,\alpha]} \bar{k}_t (\gamma - b)\, a' + \frac{\chi(\alpha, \hat{a}_t, p_t)}{r\sigma} V'(p)\, a'.$$

This condition characterizes the specialist's optimal choice of effort as a function of the belief p and of the reputational sensitivity $\Lambda \triangleq \frac{\chi(\alpha,\hat{a}_t,p_t)}{r\sigma} V'(p)$. The latter measures the importance of his continuation value for the specialist. The higher the reputational value Λ, the greater the effort he exerts. Intuitively, the specialist will make greater effort when the benefit $V'(p)$ of doing so is greater, i.e. when his future payoff is more sensitive; when the market is more transparent, i.e. low σ and when he is more patient, i.e. lower r. The incentives to imitate the skilled specialist are weaker when the investors are more strongly convinced about the specialist's type, namely when p is close to zero or one as $\Lambda \to 0$ in these cases.

Before characterizing the equilibrium, we show the following result:

Proposition 1 (Value function) *There exists a unique bounded value function $V(p)$ - increasing in the investors' belief p and decreasing in the volatility σ and in the conflict of interest b - that solves the specialist's HJB equation.*

This result ensures the existence of a solution to the HJB equation and highlights the effects of the main parameters of the model on the unskilled specialist's equilibrium payoff. The first thing to notice is that this is a model in which reputation is good, i.e. allows the specialist to commit (at least imperfectly) to the efficient action and generate a positive alpha.[14] This explains why his equilibrium payoff is increasing in the reputation p: intuitively, better reputation means he attracts more funds and at the same time, irrespective of his effort choice, the fund is more attractive to investors.

It is interesting to consider how the equilibrium payoff of the specialist is affected by the volatility of returns. Figure 1 shows that greater volatility, or equivalently lower market transparency, reduces the specialist's equilibrium payoff. The main channel is the way in which volatility σ affects investors' learning. Higher σ means that it is harder for the specialist to improve his reputation, because positive returns will be interpreted as the result of luck rather than effort.

Finally, since b is the main parameter gauging moral hazard, we study how an increase in the severity of moral hazard might affect the specialist's ability to gain profits. Proposition 1 shows that a sharper conflict of interest with the investors reduces the specialist's

[14] For an example of a model of "bad" reputation see Ely and Valimaki (2003).

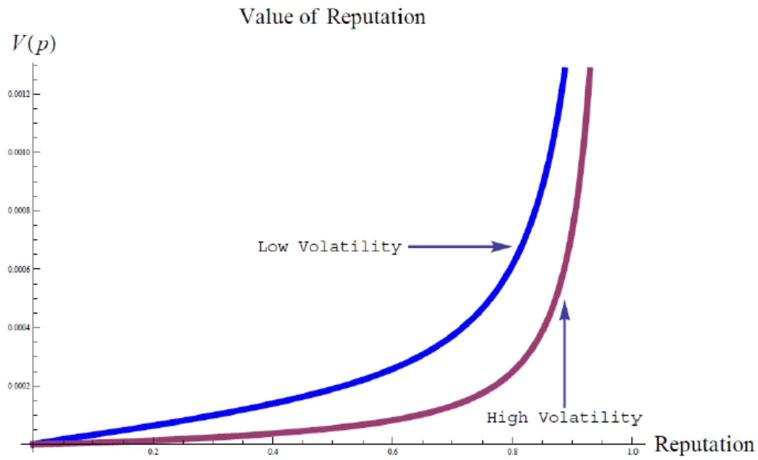

Figure 1: The specialist's value function as a function of his reputation.

equilibrium payoff, because it becomes more costly for him to take the long-term optimal action. This affects the investors' optimal investment policy k adversely.

Now we can characterize the equilibrium behavior of the specialist and the investors. Since effort is costly for the specialist and the investors' payoff is linear in expected effort $\bar{a}$, the specialist chooses the lowest value of a that makes the investors choose a positive level of capital k. Call this value $a^*(p)$. This choice of effort can then be substituted into the first order condition for the specialist in order to specify the level of capital $k^*(p)$ that is incentive-compatible with the specialist's exerting effort $a^*(p)$. The strategy profile $(a^*(p), k^*(p))$ can then be substituted into the specialist's HJB equation (3) in order to get a second-order differential equation in p.

We can now state the main result of this section:

Proposition 2 (Equilibrium) *There exists a unique public sequential equilibrium characterized by two cutoff values, $\underline{p}$ and $\bar{p}$ for the specialist's reputation, such that*

$$\text{For } p < \underline{p},\ (a^*, k^*) = (0, \underline{k});$$
$$\text{For } p < \bar{p},\ (a^*, k^*) = \left(\frac{r_s + f(1-\gamma)}{(1-\gamma)(1-p)} - \frac{p}{(1-p)}\alpha, \min\left\{\frac{\Lambda}{(b-\gamma)}, 1\right\}\right);$$
$$\text{For } p > \bar{p},\ (a^*, k^*) = (0, 1).$$

The dynamics of the effort choice in equilibrium is shown in Figure 2. First, when investors are not going to respond to good performance, the specialist's continuation value becomes insensitive to his performance, in turn reducing his incentive to exert effort. Lower expected effort leads investors to allocate a smaller fraction of their capital to the

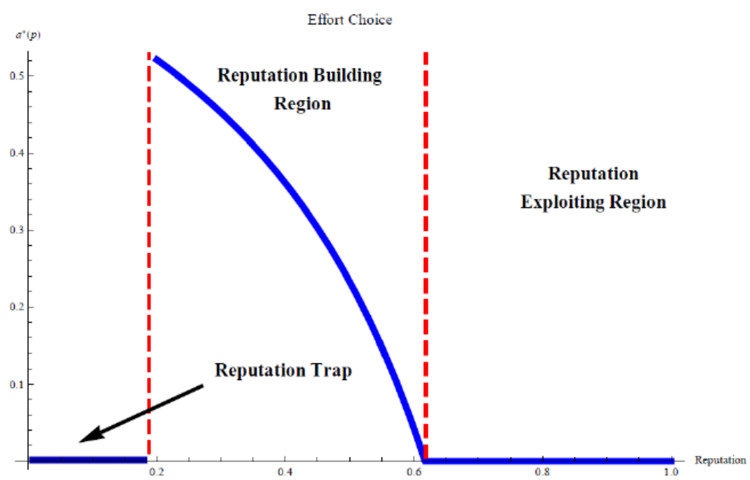

Figure 2: The equilibrium strategy of the specialist as a function of his reputation.

fund, up to the point where they do not find it worthwhile to invest at all. That is, there exists a *reputation trap:* the specialist does not exert any effort because the cost is greater than the gain to his continuation value.

Notice that even if in expectation the specialist is not going to escape the trap – because the investors' beliefs have a negative drift, which means that in the long-run the specialist's reputation converges to zero – after a sufficiently long series of good performance the specialist's reputation can get back above the threshold. That is, since performance is noisy and subject to shocks, reputation can improve enough for investors to find it optimal to invest again.

Second, for intermediate reputation values there exists a *reputation-building* region, where the specialist exerts positive effort enabling him to mimic (imperfectly) the choice of the skilled specialist. Greater effort increases expected returns and incentivizes investors to allocate capital to the fund. This is where we expect most of the funds to be. Good performance affects reputation positively, because investors are still uncertain about the specialist's type, which means that their choice k will be sensitive to the observed returns. This in turn incentivizes the specialist to exert effort over time, which boosts their returns as shown in Figure 2.

Finally, when reputation is sufficiently good, a positive level of effort becomes unsustainable in equilibrium, because the investors are willing to invest even if they expect zero effort from the unskilled specialist. Intuitively, when reputation is high enough, the specialist starts to behave more short-sightedly, extracting more surplus from the relationship with the investors. That is, there exists a *reputation-exploitation* region. This region corresponds to the case of a fund that has already gained the "star status" that enables

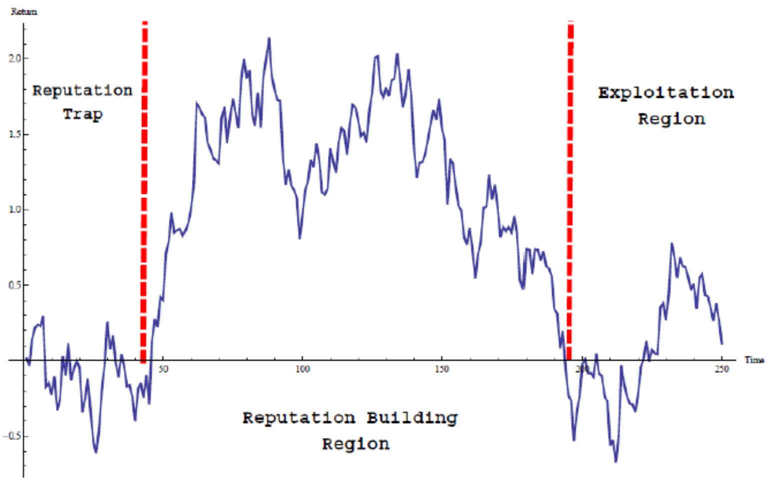

Figure 3: Sample path of returns generated by the strategic specialist in the three different regions.

the specialist to attract more investors thanks to past performance, while current performance becomes less relevant to the investors' choice. That is, the specialist is rewarded with capital inflows, even if he does not outperform the relevant benchmarks as shown in the rightmost region in Figure 2. One might well interpret the equilibrium in this region in the light of some recent scandals. Fund managers like Bernie Madoff or financial advisors like those at Washington Mutual or Goldman Sachs, had the opportunity to deceive their clients, because they were able to gain the trust of the market. Figure 3 shows an example of the corresponding returns generated by the unskilled specialist as a function of his reputation.

Interestingly, while Berk and Green (2004) show that performance may not be persistent if the investors allocate their funds to the best-performing funds and the latter have lower returns in managing larger asset volumes, in our model, fund managers are not characterized by decreasing returns to scale. Rather their reputation concerns *endogenously* generate implications that are similar to Berk and Green (2004). In fact, as a manager improves his reputation, he acquires more and more assets and his incentives to exploit his reputation choosing a lower a also increases, meaning that performance will not persist over time. The effect of size on performance has been recently analyzed empirically by Reuter and Zitzewitz (2010) and Pastor et al. (2013). Reuter and Zitzewitz (2010) exploit the fact that small differences in mutual fund returns can cause discrete changes in Morningstar ratings that, in turn, generate discrete differences in mutual fund size to identify the causal impact of fund size on performance. Their estimates do not support the hypothesis that fund size erodes fund returns, which is consistent with the non-monotone relationship between reputation and returns of our model. Similarly, Pastor et al. (2013) do not find decreasing returns at the fund level. Interestingly, they show, instead, that a fund's performance typically declines over its lifetime, which is consistent with the prediction of Proposition 2 and Figure 2.

A similar rich dynamics to the one presented in Figure 2 is provided in Liu (2011). He proposes a model where customers must pay to observe the firm's past behavior and the equilibrium structure features accumulation, consumption, and restoration of reputation. Limited record-keeping is shown by Liu and Skrzypacz (2013) to give rise to "reputation bubbles" in which short-run players drive up the reputation bubble by giving more and more trust to the opportunistic long-run player under perfect knowledge of his type, because they understand that the opportunistic player has incentives to build up his reputation in order to exploit it even more in the future.[15] In our setting there is no limited record-keeping, nevertheless, owing to the noise in the returns dR_t, the specialist can still exploit his reputation without being found out immediately. The role of reputation in our model is similar to the role played by trust in a recent work by Gennaioli et al. (2015). They propose a model in which trust in the fund manager reduces an investor's perception of the riskiness of a given investment, and gives managers the ability to charge fees. In contrast, higher reputation in our setting results from better past performance and it allows the manager to shirk more without inducing outflows. As in the reputation trap we identify in Proposition 2, Bar-Isaac (2003) shows that a sufficiently long run of bad luck could induce a seller to stop selling, because he cannot convince the buyers that his products are high quality. In our setting, after a long enough sequence of negative shocks σZ_t driving his reputation down, the specialist can end up in a reputation trap.

3.1 Equilibrium Properties

The main point of this section is to further investigate the properties of this equilibrium with a series of results on the equilibrium dynamics and comparative statics results for the optimal action a^* and for the investors' investment strategy k^*. One of the main advantages of having a continuous time framework is the greater flexibility in characterizing the effect of the parameters of the model on the equilibrium outcome.

One striking feature of the asset management industry is the extreme persistence and

[15] Another interesting dynamics is proposed by Phelan (2006): a model in which the long-run player – say the government – switches types over time, showing that governments that betray the public trust do so erratically, that public trust is regained only slowly after a betrayal, and that governments with recent betrayals betray with higher probability than other governments.

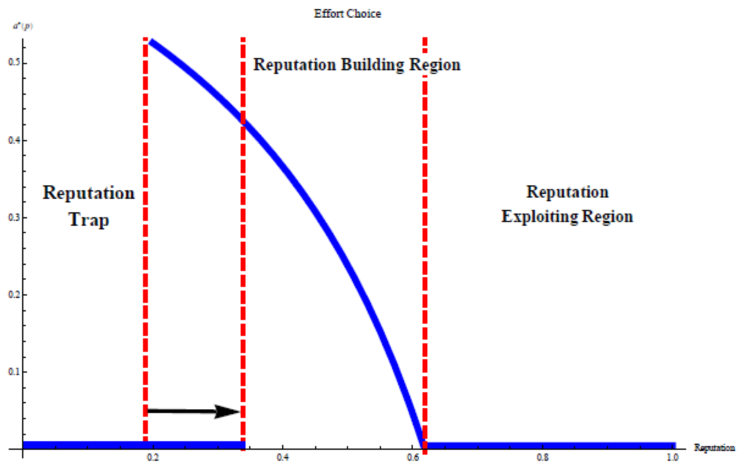

Figure 4: The equilibrium strategy of the investors as a function of the specialist's reputation.

very low variation in the charging fees.[16] Then, it is important to understand how a potential change in the fees might shape the specialist's incentives, and how these formal incentives interact with his objective to build a reputation. In particular, we are interested in understanding how the contracting features affect the equilibrium effort choice. The following proposition shows the answer:

Proposition 3 (Optimal action) *The specialist's action is non-monotone in the market's belief p, and increasing with the fees f and γ and with the investor's outside option r_s.*

Note that, while f, γ and r_s are exogenous parameters, the market belief p is endogenously determined in equilibrium. The specialist's action is first zero for low reputation levels, then it becomes positive but decreasing in p, and then it is zero again for high enough level of reputation. In other words, for $p > \underline{p}$, the specialist's incentives to deceive investors are increasing in the reputation he has built. The effect of reputation p captures the idea that at a high level of reputation the specialist has a greater incentive to deceive investors by exploiting their trust. But this effect is mitigated by the investors' ability to choose the safe investment r_s, and by the performance fee γ. Intuitively, the fees f and γ measure how much the specialist cares about the returns, hence in a sense how closely his interests are aligned with those of his investors. The effect of the investors' outside option is driven by the specialist's need to exert greater effort to compensate the investors for the lost opportunity r_s.

[16]For instance, Deuskar et al. (2011) show that during the period from 2000 to 2009 only 8% of all hedge funds changed fees at least once.

We turn next to the effect of market characteristics on the equilibrium dynamics of Figure 2. In particular, how does the volatility of returns σ affect the possibility of ending up in a reputation trap?

Proposition 4 (Reputation trap and business cycle) *The threshold $\underline{p}$ is increasing in the cost L, in the specialist's patience r, and in the market volatility σ.*

Proposition 4 shows that the region in which the reputation trap can occur widens as the volatility of returns (or, equivalently, the transparency of the market) decreases, as shown in Figure 4. This is because the specialist expects his action to be less effective in influencing investors' beliefs (i.e. his effort is less productive). One implication is that during a crisis, when the volatility of returns is particularly high (Schwert (2011)), the specialist is more likely to be trapped in the $\left[0, \underline{p}\right]$ region and be unable to attract any more funds from investors. This in turn carries two implications. First, there is greater turnover among specialists during bust than during boom if they are forced to exit the market when they lose the support of their investors. Second, the region in which the specialist exerts more effort $\left[\underline{p}, \overline{p}\right]$ shrinks (as $\overline{p}$ stays the same), which means that just when investors need the specialist's guidance to get better returns, the specialist himself has less incentive to produce higher returns. This connection between the business cycle and the specialist's incentives to exploit his reputation recalls Bar-Isaac and Shapiro (2013): in relation to credit rating agencies, they show that the value of reputation depends on economic fundamentals subject to cyclical variations and that the quality of the ratings is countercyclical, because in boom periods, the outside options of current and prospective employees improve substantially, making it harder and costlier for the agency to retain high quality analysts.[17] Our mechanism, by contrast, depends on the investors' attempt to learn the specialist's type. Intuitively, Proposition 4 shows that an increase in the specialist's patience r corresponds to a decrease in his effort level, because he attaches greater weight to the cost of effort at time t and his reputation sensitivity Λ is reduced.

Finally, we can characterize investors' reaction to information about the specialist's performance in the following proposition:

Proposition 5 (Fund flows) *The optimal investment k is increasing in the specialist's reputation p, and is decreasing in σ and b.*

[17] That credit rating agency reputation is ineffective during booms is shown by Ashcraft et al. (2010), with evidence that the boom in the issuance of mortgage-backed securities from 2005 to mid-2007 lowered the quality of ratings. Griffin and Tang (2012) demonstrate that the agencies generally made positive adjustments to their models' predictions of credit quality and that these adjustments were correlated with subsequent downgrades.

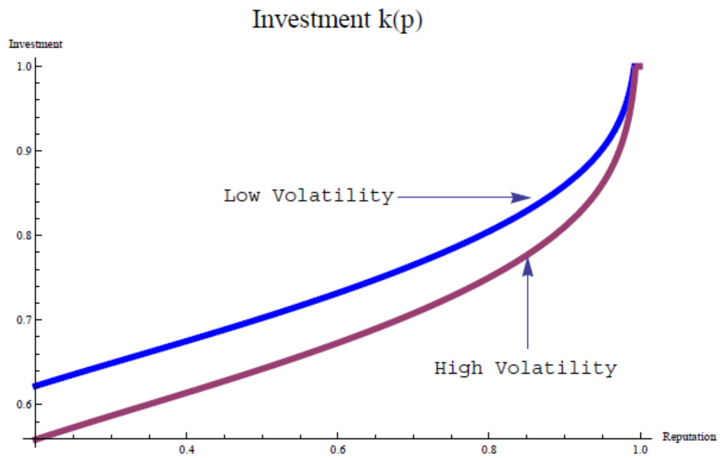

Figure 5: The equilibrium strategy of the investors as a function of the specialist's reputation.

Intuitively, these comparative statics results show that even if the investors are risk-neutral, they are less willing to invest in a volatile environment, i.e. when the volatility of returns σ is high, because it is harder for them to discipline the specialist, which in turn decreases the expected returns of investing with him (see Figure 5). But as the specialist's reputation increases, so does investors' expected payoff, which makes them more willing to allocate wealth to the fund (higher k). Moreover, as the conflict of interest b intensifies, the investors expect the unskilled specialist to have more incentives to exploit his reputation, and accordingly they invest less. The two parameters σ and b can be interpreted as two different aspects of moral hazard: b captures the strength of the incentive for the specialist to manipulate his returns, σ the investors' ability to monitor the specialist.

4 Market Timing and Tail Risk

In the previous section the only source of risk postulated is the strategy risk σZ_t. However, the financial crisis has demonstrated the importance of tail risk. The crisis has also raised the problem of why financial intermediaries that have the expertise and knowledge to detect tail risk nevertheless got so deeply exposed to it. For instance, Kelly and Jiang (2012) show that a significant component of hedge fund returns can be viewed as the compensation for selling disaster insurance. We suggest a rationale for this behavior, namely the expectation that when these extremely rare events do happen they affect specialists' continuation value, which in turn alters their incentives and their trading strategies in the first place. Moreover, in our model the premium on these disaster insurance assets will be endogenously determined, and it will be a function of the specialist's reputation, which provides implications on the feedback loop between asset prices, the specialist's reputation and his incentives to invest in these assets.

Specifically, we assume that there is a disaster insurance asset in zero net supply, let us call it a CDS for simplicity of exposition, which pays a premium ρ in each instant to the seller. Upon the realization of the disaster, a holder of CDS receive $-\xi$ with $\xi < 0$. Thus, the cash flow of CDS is

$$dC_t = \rho dt + \xi dN_t,$$

where dN_t is a Poisson process with an intensity rate φ, that is, the disaster is realized at random time τ_φ and its realization is public knowledge.

We assume that the specialist chooses the exposure $e_t \in [-1, 1]$ to this jump-event risk, or equivalently his holdings of the CDS, being a price-taker with respect to the premium ρ.[18] Thus, the return equation becomes

$$\begin{aligned} dR_t &= a_t dt + \sigma dZ_t + e_t dC_t \\ &= (a_t + e_t \rho)\, dt + \sigma dZ_t + e_t \xi dN_t, \end{aligned} \tag{5}$$

where $e_t\rho$ is the premium component that the specialist earns for holding market-wide tail risk. For tractability, we assume that the skilled specialist always chooses $e_t = -1$, that is, he perfectly hedges this risk as he buys the CDS.[19] We introduce a realistic friction between investors and the specialist: investors cannot know whether fund managers have piled up tail risk in their portfolio, that is, they cannot tell whether the returns are due to the specialist's skill or to the strategic exposure of the investors' funds to these extra risks.[20] This friction allows the unskilled specialist to use the exposure to tail risk as an additional strategy to boost his returns and then his reputation.

Finally, to endogenously determine the premium ρ, we assume that there is a risk-averse market maker who solves the following problem:

$$\max_{c,\bar{e}} E\left[\int_0^\infty e^{-\beta t} u(c_t) dt\right],$$

[18]The analysis of a specialist who is able to directly affect the premium quoted by the market maker is significantly more cumbersome and it is not evident that will add significant insights.

[19]This corresponds to the efficient action for large ξ and φ.

[20]In line with this assumption, Kacperczyk et al. (2008) show that despite extensive disclosure requirements, mutual fund investors do not observe all actions of fund managers. Sato (2014) studies how opacity in financial markets affects investor behavior, asset prices, and welfare, and show that an "opacity price premium" arises as opaque assets trade at a premium over transparent ones despite identical payoffs.

subject to the budget constraint

$$dW_t = (rW_t + \bar{e}_t\rho - c_t)\, dt + \bar{e}_t\xi dN_t,$$

where c is his consumption rate and $\bar{e}$ captures the quantity of CDS in his portfolio. To make the problem interesting, we also assume that the market maker faces a demand o_t of CDS shares from noisy traders, where o is a random variable drawn from a distribution $G(\cdot)$. Thus, market clearing condition is given by the following:

$$e_t + \bar{e}_t = o_t.$$

In a similar spirit to Glosten and Milgrom (1985), a market maker only observes $o_t - e_t$, which means that he is not able to distinguish between an order from a skilled specialist and one from an unskilled one. We can define $J(W)$ the market maker's value function as

$$J(W) = pJ_s(W) + (1-p)J_u(W),$$

where $J_s(W)$ and $J_u(W)$ are the value functions when the specialist is known to be skilled and unskilled, respectively, and p is the specialist's reputation.

Interpretation. The jump event can capture, for example, the possibility of a fall in house prices with (5) being the return on real estate investment. Another moment when negative tail risk materialized was the Russian default of August 1998. The assumption that investors cannot tell whether the high returns stem from exposure to tail risk or from the alpha-generating skill of the specialist may represent the complexity of the financial instruments employed. For instance, before the crisis investment managers would buy AAA rated tranche of CDOs to get a return of 50 to 60 basis points more than a similar AAA rated corporate bond (i.e. which is captured by ρ).[21] That "excess" return was compensating investors for the "tail" risk the CDO would default, which at the time was perceived small, but certainly not zero (as captured by $\varphi > 0$).[22] Similarly, Acharya et al. (2009) argue that commercial and investment banks had set up a way to sell deep out-of-the-money options through an intricate structure of ABCP guarantees.

Equilibrium Premium. We start the analysis by solving for the premium ρ as a function of the specialist's reputation and his demand for the asset. The Hamilton-Jacobi-Bellman

[21] CDOs are pools of loans sliced into tranches and sold to investors based on the credit quality of the underlying securities.

[22] "Bankers' pay is deeply flawed" by Raghuram Rajan on *Financial Times*, 01/08/2008.

equation for the market maker is

$$rJ(W) = \max_{c,\bar{e}} u(c) + (rW + \bar{e}\rho - c)\, J(W)' + \varphi\, (J(W + \bar{e}\xi) - J(W))\,,$$

where the first term captures the flow utility, the second term describes the change in the continuation utility in response to changes in wealth, whereas the last term captures the change in the value function once the event is realized. We can take first order conditions with respect to c and $\bar{e}$:

$$\begin{aligned} u'(c) &= pJ_s(W)' + (1-p)J_u(W)' \\ \rho &= -\xi\varphi \frac{pJ_s\,(W + \bar{e}\xi)' + (1-p)J_u\,(W + \bar{e}\xi)'}{pJ_s(W)' + (1-p)J_u(W)'}. \end{aligned}$$

To obtain closed form solutions, we assume that the market maker has exponential utility with risk-aversion parameter δ. Then, we can show the following result:

Proposition 6 (Equilibrium Premium) *The equilibrium premium is*

$$\rho \approx -\xi\varphi \left(1 - \delta r \xi \frac{o + p - e(1-p)}{1 - \delta r \bar{b}}\right),$$

which converges to the expected payment $-\xi\varphi$ when the market maker is risk-neutral, i.e. $\delta \to 0$. Moreover, ρ is increasing in the specialist's reputation p (for given e) and in the noise traders' demand o.

Intuitively, since the skilled specialist always chooses $e_t = -1$, when the market maker becomes more convinced that the specialist is skilled, namely when p increases, he quotes a higher premium ρ. Moreover, more risk averse market makers, and those facing higher noise traders' demand quote a higher premium as well to protect themselves against the jump risk.

We can now analyze the optimal strategy of the unskilled specialist, but in order to do so, we need to analyze the equilibrium in the case of realization of the jump event. There are two possible cases. First, if the returns turn out to be negatively skewed, and investors incur a loss, they might minimize their exposure to risky assets and invest all their resources into safe securities. This happens, for example, when the negative shock is large and persistent, i.e. large ξ. We label this possibility "flight to safety". Second, investors might reward the specialists who are able to outperform the market during a crisis by granting them higher capital. We label this case "betting on mavericks".

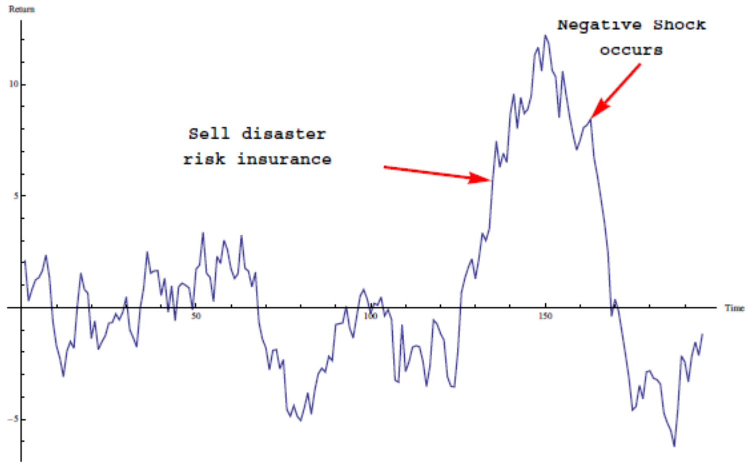

Figure 6: Sample path of returns generated by the strategic specialist when he can increase his exposure to tail risk.

4.1 Flight to Safety

We first analyze the case in which the realization of the tail risk triggers a sharp reaction on the part of investors, say a "flight to safety", captured in the model by assuming that investors set $k = \underline{k}$, i.e. during these periods investors invest all their resources into safe securities or in passive index funds (see, among others, Caballero and Kurlat (2008), Vayanos (2004), and Beber et al. (2009)).

Now we can analyze the dynamics prior to the realization of the jump event. The possibility of producing extra returns by taking tail risk enables the specialist to enhance his reputation if $e_t > 0$. That is, there exists a substitutability between the effort a_t in picking the right investment strategy and the exposure to tail risk e_t. This leads to the following result:

Proposition 7 (Tail risk exposure) *If there is a "flight to safety", the unskilled specialist exposes his investors to tail risk, i.e.* $e_t = 1$ *and the premium is* $\rho = -\xi\varphi\left(1 - \delta r \xi \frac{2p+o-1}{1-\delta r \bar{b}}\right)$*. Moreover, for any* $t < \tau_\varphi$ *his effort choice is lower than when there is no disaster risk.*

Proposition 7 provides two results. First, even when the specialist can hedge disaster risks, he may opt not to do so. The intuition is that if there is a negative event the specialist will lose investors' support regardless of what he does, which increases his incentive to act poorly prior to the occurrence of such events. So it is optimal for him to expose the investors to this risk. Second, the possibility of the jump event ζ allows the specialist to improve his reputation by capturing the premium ρ due to tail risk. Furthermore, this premium is itself increasing in p, which leads the unskilled specialist to attract higher flows by generating higher short-term returns. This means that as investors cannot distinguish between the different sources of returns, agency cost increases as the specialist exerts less effort than in absence of tail risk.

Intuitively, when a crisis is expected to be severe and persistent, so that investors decide to invest exclusively in safe assets, reputation concerns induce specialists to choose strategically to get over-exposed to tail risk. Thus market discipline, instead of incentivizing the specialist to protect investors from fluctuations, actually aggravates the investors' vulnerability. Figure 6 provides a numerical example of a potential path for the returns generated by the unskilled specialist. It shows that the returns generated by a unskilled specialist increase during normal times but fall sharply when the tail event materializes. This further implies that it is very hard for investors to distinguish between truly-skilled specialists and those who are merely selling disaster insurance, because their returns might be similarly high in normal times.

4.2 Betting on Mavericks

The recent crisis shows that while most investors and intermediaries lost money when the subprime market collapsed in late 2006, a handful of hedge-fund managers actually made a fortune by betting against a housing bubble that few, at the time, believed was real. The most notable examples are John Paulson and Greg Lippman. Paulson's firm made $15 billion in 2007, earning the founder $4 billion and the respect of his peers. Greg Lippman was able to rake in $100 million in a single week in February of 2007 by betting against the ABX subprime index, which tracks the demand for credit default swaps. The investors rewarded these successful specialists by allocating them a higher amount of capital from then on.[23]

It is important to capture this possibility as it might lead the specialists to hedge against the tail risk, rather than sell disaster insurance. We can capture this possibility by assuming that with probability $1 - \eta(e)$, where $\eta(e) = 0$ for $e < 0$ and $\eta'(e), \eta''(e) > 0$, the specialist is able to capture an extra expected profit $\overline{V}$. This captures in reduced form the possibility for the specialist to leverage the boosted reputation in the aftermath of the downturn, or the opportunity to have access to a new set of investors willing to allocate capital with the specialist. For instance, specialists that outperform during severe downturns might have access to institutional investors, that are going to reward their ability to time the market.

[23]This effect seems to persist, in fact, investors did not leave Paulson's fund, even when in 2011 he made losing trades in Bank of America, Citigroup and Sino-Forest Corporation, which caused his flagship fund, Paulson Advantage Fund, to record a negative 40% return in 2011.

Formally, in this case the specialist's value function becomes

$$(r+\varphi)\,V(p) = \max_{a,e} r\pi\left(a,e,\bar{k}\right) + \frac{\chi(\alpha,\hat{a}_t,p_t)}{\sigma}\left[(a+e)-\bar{a}(p)\right]V'(p) \tag{6}$$

$$+\varphi\left(1-\eta(e)\right)\overline{V} - \frac{\chi(\alpha,\hat{a}_t,p_t)^2}{(1-p)}V'(p) + \frac{\chi(\alpha,\hat{a}_t,p_t)^2}{2}V''(p) \tag{7}$$

Comparing (6) with (3), we can see that there are three extra terms. First, the flow payoff takes into account the possibility for the specialist to gain the premium $e\rho$. Second, by manipulating the investment strategy and exposing it to tail risk, the specialist can affect the drift of the investors' beliefs, as captured by the second term. Finally, at arrival rate φ the specialist will obtain the extra return $\overline{V}$ with probability $(1-\eta(e))$ when the tail risk is realized.

Taking the first order condition with respect to e, we obtain the following:

$$\underbrace{r\gamma\bar{k}\rho}_{\text{Myopic incentive}} + \underbrace{\frac{\chi(\alpha,\hat{a}_t,p_t)}{\sigma}V'(p)}_{\text{Gains in reputation at } t<\tau_\varphi} = \underbrace{\varphi\eta'(e)\,\overline{V}}_{\text{Expected Gain at } t>\tau_\varphi}. \tag{8}$$

The first term depicts the specialist's incentive to increase e to capture the excess returns thanks to the positive premium ρ. By increasing the short-term returns, the specialist is able to improve his reputation, which will benefit him in the future as he will be able to gain higher assets under management, and higher fees, which is captured in the second term. However, increasing the investors' exposure to tail risk has a cost as well, that is, the foregone possibility to attain the premium $\overline{V}$.

We can now address the following question: what is the effect of the specialist's reputation on his incentives to pile up tail risk? In order to answer this question, we complete the analysis by considering the interplay between the reputation, the premium and the specialist's choice e_t. The analysis is significantly more complex in this case when e_t is interior as the specialist's reputation has both a direct effect on the premium ρ and an indirect one through the specialist's choice e. However, we can show the following result:

Proposition 8 *If the expected size of disaster, $-\varphi\xi$, and the noisy demands for CDS are sufficiently large, the unskilled specialist has a higher incentive to sell disaster insurance as his reputation increases, i.e. $\frac{de}{dp} > 0$.*

Proposition 8 presents the main result of this section. Intuitively, as the specialist's reputation increases, the market maker fears more and more to be on the wrong side of the market, and thus start charging a higher premium for the insurance against the tail risk. However, by doing so he encourages the unskilled specialist to expose the fund to tail risk in order to capture the higher premium. This holds exactly when the investors would benefit the most from hedging against this risk, i.e. when the expected size of the disaster is large, and when the market maker faces noisier demand for the CDS, which makes him less able to price the CDS correctly. In sum, the specialist's desire to boost his reputation increases his incentive to expose the investors to tail risk, moreover, as he captures the premium and improves his performance, through a general equilibrium effect the premium rises, which makes "selling disaster insurance" an even more attractive strategy.

5 Discussion

5.1 Empirical Implications

We have proposed a framework for analyzing the behavior of specialists who manage investors' funds and are compensated with a fixed fraction of the asset under management and a fraction of the realized returns, but use strategies that are driven mainly by implicit incentives. The model offers a rich set of testable predictions.

History and Performance. First, in contrast to the existing literature, our model predicts that the performance history of fund managers has a significant impact on their future behavior. In particular, a long sequence of positive returns should be associated with lower returns in the future, while a long enough series of negative shocks can lead the fund out of the market, with more shocks needed for the more established funds. This can be tested by looking at how past performance affects portfolio holdings and excess returns over time.

Reputation and Volatility. Second, the strategies adopted by fund managers depend crucially on their reputation and on market volatility. This implies a testable cross-sectional variation across managers with different histories. In particular, managers with poorer reputations should increase the riskiness of their portfolio in hopes of improving their status by realizing higher returns. At the same time, fund managers with better reputations tend to underperform in periods of high volatility but nevertheless suffer less disinvestment, due to slow investor learning.

We can relate the results implied by Proposition 4 to some recent empirical evidence on inflows in the asset management industry. For instance, Huang et al. (2011) analyze mutual fund data from 1993 through 2006, finding that the sensitivity of flows to past performance is weaker for funds with more volatile past returns and that older funds have weaker sensitivity than younger ones. These two empirical findings suggest that investors update their beliefs about the fund managers over time, and are reluctant to invest with a manager whose returns are so volatile that they can be presumed to be due to luck rather than skill; and that investors tend to downgrade the reputation of already established managers hit by negative shocks less sharply.[24]

Reputation and Tail Risk. Third, the portfolio returns of fund managers with poor reputations should become more and more skewed as the probability of a crisis or of a flight to safety increases. That is, "reaching for yields" behavior is more likely to be observed during a bull market, as the probability of lower returns in the future is higher, or when it is more difficult to produce returns in excess of the risk-free rate, as when the rates on treasuries bills are high. Kelly and Jiang (2012) show that a significant part of hedge fund returns can be viewed as compensation for selling disaster insurance and, consistent with the predictions of the model, they show that young managers with limited performance history tend to generate higher returns by exposing their investors to tail risk. There should also be a positive correlation between the premium paid for selling disaster insurance and the accumulation of tail risk in the managers' portfolios.

Flight to safety and Volatility. Another implication is provided by the solution to the investors' portfolio choice problem depicted in Figure 5: in periods of high volatility, investors' funds tend to flow towards safer assets, i.e. k_t is lower. Intuitively, an increase in volatility induce investors to shield themselves from the agency conflict with the fund manager by diverting their resources towards the benchmark portfolio r_s. This endogenous flight to safety is confirmed empirically by several recent studies about the financial crisis such as Mitchell and Pulvino (2012). They note that, around the Lehman Brothers collapse, investors shifted assets out of risky securities, such as the investment with fund managers, into U.S. Treasuries, as captured in the model by the benchmark safe portfolio r_s, which caused the 28-day U.S. Treasury bill yield to decrease from 1.4% on September 12, 2008 to 0.1% on September 17, 2008.[25]

Size and Returns. Finally, the size of the fund has a non-monotone effect on the manager's incentives. This follows from the existence of a reputation trap and a reputation exploitation region, in which the managers' incentives are the lowest, while in the reputation building region their incentives to behave in the interest of households are the strongest; and the increasing relationship between the capital allocated by the households to the funds and the managers' reputation. As predicted by Proposition 2, Pastor et al.

[24] Similar evidence, corroborating these implications, is provided by Franzoni and Schmalz (2013).

[25] These movements in the short-term rates during the Lehman period are outside the distribution of the daily change in rates during the previous seven years.

(2013) show that a fund's performance typically declines over its lifetime.

5.2 Discussion of Assumptions

The model abstracts from a number of important features of reality. In what follows we discuss how the basic environment can accommodate these features and how they would modify the results.

Management fees. In the baseline version we take the contract between the specialist and the investors as given for several reasons. First, the performance fee aligns specialists' and investors' interests except for the former's reputation concerns. In reality, asset management contracts usually give the manager a fixed fraction of the asset under management plus a performance fee, and there is evidence of persistence in these contracts.[26] For instance, Deuskar et al. (2011) analyze the hedge fund industry and show that during the period from 2000 to 2009, a very small fraction, around 8%, of all hedge funds changed fees at least once. Second, we can endogenize the fees by allowing competition among intermediaries or by assuming that investors can bargain with managers over the fees at discrete intervals. This would not affect the main results. Finally, the literature has extensively examined the role of contracts and, more generally, of explicit incentives for managers.[27] It might be interesting in the spirit of DeMarzo and Sannikov (2006) to analyze the optimal contract when investors can commit at $t = 0$. In this case, there would be no shirking in equilibrium, in that, the specialist would always have the incentive to exert high effort. In our view, while the analysis of the optimal contract can generate a number of insights on executives' compensation and how the policy maker should intervene, at the same time, the way in which market discipline shapes the managers' behavior remains an important issue. An important topic for future research is to extend the current setting to study the interactions between implicit and explicit incentives.

Monitoring. The model has assumed that investors observe the fund's returns continuously which enables them to update their beliefs about the specialist's type. We could change the environment by making the returns visible only at discrete intervals. This might capture the inability of investors to monitor the specialist's decisions, either because monitoring is costly or because investors may be rationally inattentive. At the same time, monitoring difficulties could represent the complexity of the financial instruments that a specialist uses. This would not change the main implications of the model except

[26]For studies on the role of more complex contracting features, such as high-water marks, see the papers by Goetzmann et al. (2003) and Panageas and Westerfield (2009) among others.

[27]For instance, Stoughton (1993) investigates the significance of nonlinear contracts on the incentive for portfolio managers to collect information.

that suboptimal strategies by the specialist could go undetected for a longer time, making reputation more persistent.

Competition. We could allow investors to search for the best-performing specialist at a cost. This would endogenize the investors' outside opportunity and also avoid a situation in which only the specialist with the best past performance remains in the market. In this modified framework, the managers' exploitation region would depend crucially on the presence of other managers with better reputation and on investors' search costs. This would strengthen the result that during a crisis, when more managers underperform, agency costs rise as the managers with good reputation exploit investors' perception of their ability to avoid exerting effort, given the lack of strong competition.

6 Conclusion

In this paper, we have provided a novel model of intermediation with two classes of agents, households and investment specialists, e.g. fund managers. Specialists have the expertise to actively invest in risky assets, which households cannot purchase directly. However, this intermediation relationship is subject to moral hazard, in that specialists can generate returns by exploiting their stocks-picking ability or by timing the market, but they have short-sighted incentives to mislead the households by manipulating both the returns and the riskiness of their investment strategy.

The model generates a rich equilibrium dynamics as specialists' reputation has a non-monotone effect on their incentives to take advantage of the investors' trust. We have also shown that, although the presence of implicit incentives motivate the specialist to generate excess returns, it might also lead to several distortions. For instance, somewhat paradoxically, we show that in a crisis the agency costs are exacerbated by the high-reputation specialists' incentives to act myopically.

We have then extended the model to capture the possibility that the economy is subject to tail risk. In this setting, we have shown that by exposing the households to this tail risk the specialist captures a positive premium, and he might induce returns-chasing investors to update their beliefs about his type and entrust more capital to the fund. We have also shown that when the premium on disaster insurance is endogenously determined, it increases the incentives of the specialists with higher reputation to reach for yield.

The model is able to generate several novel empirical implications concerning the specialists' behavior, which can be tested employing data on fund managers' holdings.

References

Acharya, V. V., T. Cooley, M. Richardson, and I. Walter (2009). Manufacturing Tail Risk: A Perspective on the Financial Crisis of 2007–2009. *Foundations and Trends in Finance 4*(4), 247–325.

Ashcraft, A. B., P. Goldsmith-Pinkham, and J. I. Vickery (2010). MBS ratings and the mortgage credit boom. *FRB of New York Staff Report* (449).

Bar-Isaac, H. (2003). Reputation and Survival: learning in a dynamic signalling model. *The Review of Economic Studies 70*(2), 231–251.

Bar-Isaac, H. and J. Shapiro (2013). Ratings quality over the business cycle. *Journal of Financial Economics 108*(1), 62–78.

Basak, S., A. Pavlova, and A. Shapiro (2007). Optimal asset allocation and risk shifting in money management. *Review of Financial Studies 20*(5), 1583–1621.

Beber, A., M. W. Brandt, and K. A. Kavajecz (2009). Flight-to-quality or flight-to-liquidity? Evidence from the euro-area bond market. *Review of Financial Studies 22*(3), 925–957.

Ben-David, I., F. Franzoni, and R. Moussawi (2012). Hedge fund stock trading in the financial crisis of 2007–2009. *Review of Financial Studies 25*(1), 1–54.

Berk, J. and R. Green (2004). Mutual Fund Flows and Performance in Rational Markets. *Journal of Political Economy 112*(6).

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*(2), 345–390.

Board, S. and M. Meyer-ter Vehn (2013). Reputation for quality. *Econometrica 81*(6), 2381–2462.

Bolton, P., H. Chen, and N. Wang (2011). A unified theory of Tobin's q, corporate investment, financing, and risk management. *The Journal of Finance 66*(5), 1545–1578.

Bolton, P., X. Freixas, and J. Shapiro (2007). Conflicts of interest, information provision, and competition in the financial services industry. *Journal of Financial Economics 85*(2), 297–330.

Bolton, P. and C. Harris (1999). Strategic experimentation. *Econometrica 67*(2), 349–374.

Brunnermeier, M. and S. Nagel (2004). Hedge funds and the technology bubble. *The Journal of Finance 59*(5), 2013–2040.

Brunnermeier, M. K. and Y. Sannikov (2014). A macroeconomic model with a financial sector. *The American Economic Review 104*(2), 379–421.

Buffa, A. M., D. Vayanos, and P. Woolley (2014). Asset management contracts and equilibrium prices. Technical report, National Bureau of Economic Research.

Caballero, R. J. and P. D. Kurlat (2008). Flight to quality and bailouts: Policy remarks and a literature review.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of finance 52*(1), 57–82.

Chapman, D. A., R. Evans, and Z. Xu (2007). Career Concerns and the Active Fund Managers Problem. Technical report, Working Paper Boston College.

Chevalier, J. and G. Ellison (1997). Risk Taking by Mutual Funds as a Response to Incentives. *Journal of Political Economy 105*(6), 1167–1200.

Cuoco, D. and R. Kaniel (2011). Equilibrium prices in the presence of delegated portfolio management. *Journal of Financial Economics 101*(2), 264–296.

Dangl, T., Y. Wu, and J. Zechner (2008). Market discipline and internal governance in the mutual fund industry. *Review of Financial Studies 21*(5), 2307–2343.

Dasgupta, A. and A. Prat (2006). Financial equilibrium with career concerns. *Theoretical Economics 1*(1), 67–93.

Dasgupta, A. and A. Prat (2008). Information aggregation in financial markets with career concerns. *Journal of Economic Theory 143*(1), 83–113.

Dasgupta, A., A. Prat, and M. Verardo (2011). Institutional Trade Persistence and Long-Term Equity Returns. *The Journal of Finance 66*(2), 635–653.

De Coster, C. and P. Habets (2006). *Two-Point Boundary Value Problems: Lower and Upper Solutions: Lower and Upper Solutions*, Volume 205. Elsevier.

DeMarzo, P. and Y. Sannikov (2006). Optimal Security Design and Dynamic Capital Structure in a Continuous-Time Agency Model. *The Journal of Finance 61*(6), 2681–2724.

DeMarzo, P. M., D. Livdan, and A. Tchistyi (2012). Risking Other People's Money: Gambling, Limited Liability, and Optimal Incentives. *University of California, Berkeley*.

Deuskar, P., Z. Wang, Y. Wu, and Q. Nguyen (2011). The dynamics of hedge fund fees. *Available at SSRN 1659275*.

Ely, J. C. and J. Valimaki (2003). Bad reputation. *The Quarterly Journal of Economics 118*(3), 785–814.

Faingold, E. (2005). Building a reputation under frequent decisions. Technical report, mimeo.

Faingold, E. and Y. Sannikov (2011). Reputation in Continuous-Time Games. *Econometrica 79*(3), 773–876.

Franzoni, F. A. and M. C. Schmalz (2013). Fund Flows in Rational Markets. *SSRN Electronic Journal*.

French, K. R. (2008). Presidential address: The cost of active investing. *The Journal of Finance 63*(4), 1537–1573.

Garcia, D. and J. M. Vanden (2009). Information acquisition and mutual funds. *Journal of Economic Theory 144*(5), 1965–1995.

Garleanu, N., L. H. Pedersen, and A. M. Poteshman (2009). Demand-based option pricing. *Review of Financial Studies 22*(10), 4259–4299.

Gennaioli, N., A. Shleifer, and R. Vishny (2015). Money doctors. *The Journal of Finance 70*(1), 91–114.

Gervais, S. and G. Strobl (2012). The Industrial Organization of Money Management. *Available at SSRN 2023884*.

Glode, V. (2011). Why mutual funds underperform. *Journal of Financial Economics 99*(3), 546–559.

Glosten, L. and P. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics 14*(1), 71–100.

Goetzmann, W. N., J. E. Ingersoll, and S. A. Ross (2003). High-Water Marks and Hedge Fund Management Contracts. *The Journal of Finance 58*(4), 1685–1718.

Griffin, J. M. and D. Y. Tang (2012). Did subjectivity play a role in CDO credit ratings? *The Journal of Finance 67*(4), 1293–1328.

Guerrieri, V. and P. Kondor (2012). Fund Managers, Career Concerns, and Asset Price Volatility. *American Economic Review 102*(5), 1986–2017.

He, Z. (2009). Optimal executive compensation when firm size follows geometric brownian motion. *Review of Financial Studies 22*(2), 859–892.

He, Z. and A. Krishnamurthy (2012). A model of capital and crises. *The Review of Economic Studies 79*(2), 735–777.

He, Z. and A. Krishnamurthy (2013). Intermediary Asset Pricing. *The American Economic Review 103*(2), 732–70.

Huang, J., C. Sialm, and H. Zhang (2011). Risk shifting and mutual fund performance. *Review of Financial Studies 24*(8), 2575–2616.

Huang, J., K. Wei, and H. Yan (2011). Investor Learning and Mutual Fund Flows.

Jensen, M. C. (1968). The performance of mutual funds in the period 1945–1964. *The Journal of Finance 23*(2), 389–416.

Kacperczyk, M., C. Sialm, and L. Zheng (2008). Unobserved actions of mutual funds. *Review of Financial Studies 21*(6), 2379–2416.

Kacperczyk, M., S. Van Nieuwerburgh, and L. Veldkamp (2012). Rational Attention Allocation over the Business Cycle. *Available at SSRN 1411367*.

Kacperczyk, M., S. Van Nieuwerburgh, and L. Veldkamp (2014). Time-Varying Fund Manager Skill. *The Journal of Finance 69*(4), 1.

Kaniel, R. and P. Kondor (2013). The delegated Lucas tree. *Review of Financial Studies 26*(4), 929–984.

Keller, G. and S. Rady (1999). Optimal experimentation in a changing environment. *The Review of Economic Studiesc 66*(3), 475–507.

Kelly, B. and H. Jiang (2012). Tail risk and hedge fund returns. *Available at SSRN*.

Kreps, D. M. and R. Wilson (1982). Reputation and imperfect information. *Journal of economic theory 27*(2), 253–279.

Liu, Q. (2011). Information acquisition and reputation dynamics. *The Review of Economic Studies 78*(4), 1400–1425.

Liu, Q. and A. Skrzypacz (2013). Limited Records and Reputation Bubbles. *Journal of Economic Theory*.

Lynch, A. W. and D. K. Musto (2003). How investors interpret past fund returns. *The Journal of Finance 58*(5), 2033–2058.

Makarov, I. and G. Plantin (2015). Rewarding trading skills without inducing gambling. *The Journal of Finance 70*, 925–962.

Malamud, S. and E. Petrov (2014). Portfolio Delegation and Market E fficiency. *Swiss Finance Institute Research Paper* (14-09).

Malliaris, S. and H. Yan (2010). Reputation concerns and slow-moving capital. *Yale International Center for Finance working paper*.

Mitchell, M. and T. Pulvino (2012). Arbitrage crashes and the speed of capital. *Journal of Financial Economics 104*(3), 469–490.

Moreira, A. (2013). A Reputation Based Model of Limited Arbitrage.

Ottaviani, M. and P. Sorensen (2006). Reputational cheap talk. *The Rand Journal of Economics 37*(1), 155–175.

Panageas, S. and M. M. Westerfield (2009). High-Water Marks: High Risk Appetites? Convex Compensation, Long Horizons, and Portfolio Choice. *The Journal of Finance 64*(1), 1–36.

Pastor, L., R. F. Stambaugh, and L. A. Taylor (2013). Scale and Skill in Active Management. *Available at SSRN 2318788*.

Pesendorfer, W. and A. Wolinsky (2003). Second opinions and price competition: Inefficiency in the market for expert advice. *Review of Economic Studies 70*(2), 417–437.

Phelan, C. (2006). Public trust and government betrayal. *Journal of Economic Theory 130*(1), 27–43.

Reuter, J. and E. Zitzewitz (2010). How much does size erode mutual fund performance? A regression discontinuity approach. Technical report, National Bureau of Economic Research.

Sannikov, Y. (2007). Games with imperfectly observable actions in continuous time. *Econometrica 75*(5), 1285–1329.

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

Sannikov, Y. (2012). Dynamic Security Design and Corporate Financing. *Handbook of Economics and Finance*.

Sannikov, Y. and A. Skrzypacz (2010). The role of information in repeated games with frequent actions. *Econometrica 78*(3), 847–882.

Sato, Y. (2014). Opacity in Financial Markets. *Review of Financial Studies 27*(12), 3502–3546.

Scharfstein, D. and J. Stein (1990). Herd behavior and investment. *The American Economic Review 80*(3), 465–479.

Schwert, G. W. (2011). Stock volatility during the recent financial crisis. *European Financial Management 17*(5), 789–805.

Shleifer, A. and R. Vishny (1997). The Limits of Arbitrage. *The Journal of Finance 52*(1), 35–55.

Sirri, E. R. and P. Tufano (1998). Costly search and mutual fund flows. *The journal of finance 53*(5), 1589–1622.

Stoughton, N. M. (1993). Moral hazard and the portfolio management problem. *The Journal of Finance 48*(5), 2009–2028.

Vayanos, D. (2004). Flight to quality, flight to liquidity, and the pricing of risk. Technical report, National Bureau of Economic Research.

Zhu, J. Y. (2012). Optimal contracts with shirking. *The Review of Economic Studies*, rds038.

Zwiebel, J. (1995). Corporate conservatism and relative compensation. *Journal of Political Economy*, 1–25.

7 Appendix

Proof of Lemma 1

We could derive the stochastic evolution of investors' beliefs in two different ways. The first involve applying Girsanov's theorem and is the one used by Faingold and Sannikov (2011), while the second one follows directly from Bayes' rule and is adapted from Bolton and Harris (1999). We are going to follow the second approach, which allows for off-equilibrium beliefs, and can be directly employed to derive the Bellman equation for the specialist.

Consider the random change in belief from p_t over $[t, t+dt]$. Let dR_t be the signal change observed by the investors over this time interval. By Bayes' rule we have:

$$dp_t \equiv p_{t+dt} - p_t = \frac{p_t\left(1-p_t\right)\left(f_S\left(dR\right) - f_U\left(dR\right)\right)}{p_t f_S\left(dR\right) + \left(1-p_t\right) f_U\left(dR\right)} \tag{9}$$

If the actual action is a, while the investor expects $\hat{a}\left(p\right)$, then the probability density $f_U\left(dR\right)$ of increment dR is a normal random variable with mean $\hat{a}\left(p\right)dt$ and variance $\sigma^2 dt$. We exploit the second order Taylor series approximation $e^z \approx 1 + z + z^2/2$:

$$f_U\left(dR\right) \propto e^{-\left(dR - \hat{a}dt\right)^2/2\sigma^2 dt} = e^{\left(\hat{a}dR - \hat{a}^2 dt/2\right)/\sigma^2} \approx 1 + (\hat{a}dR - \frac{\hat{a}^2}{2}dt)/\sigma^2 + \frac{1}{2}(\hat{a}dR - \frac{\hat{a}^2}{2}dt)^2/\sigma^4$$

The realized signal dR depends on the actual action a taken by the specialist: $dR = adt + \sigma dZ$. Substituting this and cancelling all terms above order dt, we find:

$$f_U\left(adt + \sigma dZ\right) \approx 1 + \hat{a}\left(adt + \sigma dZ\right)/\sigma^2$$

Now we can substitute back into equation (9) to get:

$$\begin{aligned} dp &\approx \frac{p\left(1-p\right)\left(\alpha - \hat{a}\right)\left(adt + \sigma dZ\right)/\sigma^2}{1 + \bar{a}\left(p\right)\left(adt + \sigma dZ\right)/\sigma^2} \\ &\approx p\left(1-p\right)\left(\alpha - \hat{a}\right)\left(adt + \sigma dZ\right)/\sigma^2 \left[1 - \bar{a}\left(p\right)\left(adt + \sigma dZ\right)/\sigma^2\right] \\ &\approx \left[p\left(1-p\right)\left(\alpha - \hat{a}\right)\left(a - \bar{a}\left(p\right)\right)/\sigma^2\right] dt + \left[p\left(1-p\right)\left(\alpha - \hat{a}\right)/\sigma\right] dZ \end{aligned}$$

Then, we obtain the equation in the statement of the proposition.

Proof of Lemma 2

First, notice that the specialist has private information about his type θ. Then, we first need to derive the stochastic evolution of beliefs from the specialist's point of view. We can rewrite the Brownian motion Z^p as follows

$$\begin{aligned} \sigma dZ^p &= dR_t - (p\alpha + (1-p)\hat{a})\,dt \\ &= dR_t - adt - p(\alpha - \hat{a})\,dt \\ &= dZ^S - p(\alpha - \hat{a})\,dt \end{aligned}$$

Substituting in (2) we get

$$dp = \left[\frac{\chi(\alpha,\hat{a},p)(a - \bar{a}(p))}{\sigma} - \frac{|\chi(\alpha,\hat{a},p)|^2}{(1-p)}\right]dt + \chi(\alpha,\hat{a},p)\,dZ^s$$

Notice that conditional on the specialist being the unskilled type, the posterior on the skilled type must be a supermartingale, because the specialist expects the beliefs about the skilled type to decrease over time.

We can now derive the value function. From the Principle of Optimality, we know that $V(p)$ satisfies

$$V(p) = \max_{a\in[0,\alpha]}\left\{r(k(\gamma(a-f)+f) - abk)\,dt + e^{-rdt}\mathbb{E}_p[V(p+dp)]\right\}$$

where

$$\mathbb{E}_p[V(p+dp)] = V(p) + V'(p)\,\mathbb{E}_p[dp] + \frac{1}{2}V''(p)\,\mathbb{E}_p[dp]^2$$

follows from Ito's Lemma. Substituting in the expression for the evolution of beliefs dp we find

$$\mathbb{E}_p[V(p+dp)] = V(p) + V'(p)\left[\frac{\chi(\alpha,\hat{a},p)(a-\bar{a}(p))}{\sigma} - \frac{|\chi(\alpha,\hat{a},p)|^2}{(1-p)}\right]dt + V''(p)\frac{|\chi(\alpha,\hat{a},p)|^2}{2}dt$$

Finally, using the approximation $e^{-rdt} = 1 - rdt$ we get

$$V(p) = \max_{a\in[0,\alpha]}\left\{\begin{array}{c} (k(\gamma(a-f)+f) - abk)\,dt \\ +(1-rdt)\left[V(p) + V'(p)\left[\frac{\chi(\alpha,\hat{a},p)(a-\bar{a}(p))}{\sigma} - \frac{|\chi(\alpha,\hat{a},p)|^2}{(1-p)}\right]dt + V''(p)\frac{|\chi(\alpha,\hat{a},p)|^2}{2}dt\right]\end{array}\right\}$$

rearranging and eliminating terms of order $(dt)^2$, we get the expression for the value function displayed in the proposition.

Proof of Proposition 1

We start with some preliminaries about the specialist's value function. To simplify notation, let us assume that $\underline{k} = 0$. A triple $(a_t, k_t, p_t)_t$ constitutes an equilibrium if $(a_t, k_t)_{t\geq 0}$

belongs to the image of the function $\Gamma\left(p_t,\Lambda\right):[0,1]\times\mathbb{R}\rightrightarrows[0,\alpha]\times[0,1]:$

$$\Gamma\left(p,\Lambda\right)\triangleq\left\{(a,k):\begin{array}{l}\bar{k}\in\arg\max_{\bar{k}'\in[0,1]}\bar{k}'\left(1-\gamma\right)\left(\bar{a}\left(p\right)-f\right)+\left(1-\bar{k}'\right)r_s\\ a\in\arg\max_{a'\in[0,\alpha]}a'\bar{k}\left(\gamma-b\right)+\Lambda a'\end{array}\right\}$$

Fix (p,Λ) and consider the following correspondence

$$\Sigma\left(a,k\right)\triangleq\left\{\begin{array}{l}k\in\arg\max_{k'\in[0,1]}\left[k'\left(1-\gamma\right)\left(\bar{a}\left(p\right)-f\right)+\left(1-k'\right)r_s\right]\\ a\in\arg\max_{a'\in[0,\alpha]}a'k\left(\gamma-b\right)+\Lambda a'\end{array}\right\}$$

Then an action profile (a_t,k_t) belongs to $\Gamma\left(p,\Lambda\right)$ if and only if it is a fixed point of $\Sigma\left(a,k\right)$. Fix $(a,k)\in[0,\alpha]\times[0,1]$ and notice that the assumptions on the agents' flow payoff and on the drift of the diffusion process R_t imply that $\pi\left(a,k\right)$ and $u\left(a,k\right)$ are weakly concave and hence by *Brouwer's fixed point* theorem $\Sigma\left(a,k\right)$ has a fixed point. Then, the correspondence $\Gamma\left(p,\Lambda\right)$ is non-empty, which shows existence.

The proof of the existence of a solution to the ODE follows the techniques introduced by Keller and Rady (1999) and Faingold and Sannikov (2011). We divide the proofs in three steps. First we show existence. We then show that the solution must be unique and finally we show that the unique solution must be an increasing function.

The proof of Proposition 1 relies on standard results from the theory of boundary-value problems for second-order equations (see for example De Coster and Habets (2006)). We now review the part of that theory that is relevant for our existence result.

Given a continuous function $H:[a,b]\times\mathbb{R}^2\to\mathbb{R}$ and real numbers c and d, consider the following boundary value problem:

$$\begin{aligned} V''\left(x\right) &= H\left(x,V\left(x\right),V'\left(x\right)\right),\ x\in[a,b] \\ V\left(a\right) &= c, V\left(b\right)=d. \end{aligned} \tag{10}$$

Given real numbers α and β, we are interested in sufficient conditions for the previous problem to admit a $\mathcal{C}^2-$solution $V:[a,b]\to\mathbb{R}$ with $\alpha\leq V\left(x\right)\leq\beta$ for all $x\in[a,b]$. One sufficient condition is called *Nagumo condition*, which posits the existence of a positive continuous function $\psi:[0,\infty)\to\mathbb{R}$ satisfying

$$\int_0^\infty\frac{vdv}{\psi\left(v\right)}=\infty$$

and

$$\left|H\left(x,v,v'\right)\right|\leq\psi\left(\left|v'\right|\right),\ \forall\left(x,v,v'\right)\in[a,b]\times[\alpha,\beta]\times\mathbb{R}.$$

To prove existence we use the following result, which follows from Theorems II.3.1 and I.4.4 in de Coster and Habets (2006):

Lemma 3 *Suppose that* $\alpha \leq c \leq \beta, \alpha \leq d \leq \beta$ *and that* $H : [a,b] \times \mathbb{R}^2 \to \mathbb{R}$ *satisfies the Nagumo condition relative to* α *and* β*. Then:*

(a) the boundary value problem (10) admits a solution satisfying $\alpha \leq V(x) \leq \beta$ *for all* $x \in [a,b]$*;*

(b) there is a constant $R > 0$ *such that every* C^2-*function* $V : [a,b] \to \mathbb{R}$ *that satisfies* $\alpha \leq V(x) \leq \beta$ *for all* $x \in [a,b]$ *and solves*

$$V''(x) = H\left(x, V(x), V'(x)\right), \; x \in [a,b],$$

satisfies $|V'(x)| \leq R$ *for all* $x \in [a,b]$*.*

Step 1. Existence. Since the right hand side of the ODE blows up at $p = 0$ and $p = 1$, our strategy of proof is to construct the solution as the limit of a sequence of solutions on expanding closed subintervals of $(0,1)$. Indeed, let $H : (0,1) \times \mathbb{R}^2 \to \mathbb{R}$ denote the right-hand side of the ODE and for each $n \in \mathbb{N}$ consider the boundary value problem:

$$\begin{aligned} V''(p) &= H\left(p, V(p), V'(p)\right), \; p \in [1/n, 1-1/n] \\ V(1/n) &= \underline{g}, V(1-1/n) = \bar{g}. \end{aligned} \quad (11)$$

There exists a constant $K_n > 0$ such that

$$|H(p,v,v')| \leq K_n\left(1 + |v'|^2\right), \; \forall \left(p,v,v'\right) \in [1/n, 1-1/n] \times \left[\underline{g}, \bar{g}\right] \times \mathbb{R}.$$

Since $\int_0^\infty K_n^{-1}\left(1+v^2\right)^{-1} v dv = \infty$, for each $n \in \mathbb{N}$ the boundary value problem above satisfies the hypothesis of the Lemma relative to $\alpha = \underline{g}$ and $\beta = \bar{g}$. Therefore, for each $n \in \mathbb{N}$ there exists a $\mathcal{C}^2-$function $V_n : [1/n, 1-1/n] \to \mathbb{R}$ which solves the ODE on $[1/n, 1-1/n]$ and satisfies $\underline{g} \leq V_n \leq \bar{g}$. Since for $m \geq n$ the restriction of V_m to $[1/n, 1-1/n]$ also solves the ODE on $[1/n, 1-1/n]$, by the quadratic growth condition above the first and the second derivatives of V_m are uniformly bounded for $m \geq n$, and hence the sequence $(V_m, V_m')_{m \geq n}$ is bounded and equicontinuous over the domain $[1/n, 1-1/n]$. By the Arzela'-Ascoli Theorem, for every $n \in \mathbb{N}$ there exists a subsequence of $(V_m, V_m')_{m \geq n}$ which converges uniformly on $[1/n, 1-1/n]$. Then, using a diagonalization argument, we can find a subsequence of $(V_n)_{n \in \mathbb{N}}$, denoted $(V_{n_k})_{k \in \mathbb{N}}$, which converges pointwise to a continuously differentiable function $V : (0,1) \to \left[\underline{g}, \bar{g}\right]$ such that on every closed subinterval of $(0,1)$ the convergence takes place in $\mathcal{C}^1$.

Finally, V must solve the ODE on $(0,1)$, since $V''_{n_k}(p) = H\left(p, V_{n_k}(p), V'_{n_k}(p)\right)$ converges to $H(p, V(p), V'(p))$ uniformly on every closed subinterval of $(0,1)$, by the continuity of H and the uniform convergence $\left(V_{n_k}, V'_{n_k}\right) \to (V, V')$ on closed subintervals of $(0,1)$. We also need to impose that

$$\lim_{p \to \underline{p}^+} V(p) = \lim_{p \to \underline{p}^-} V(p)$$

and

$$\lim_{p \to \underline{p}^+} V'(p) = \lim_{p \to \underline{p}^-} V'(p).$$

Step 2. Uniqueness. Suppose V and U are two bounded solutions. Assuming that $U(p) > V(p)$ for some $p \in (0,1)$, let $p_0 \in (0,1)$ be the point where the difference $U - V$ is maximized. Thus we have $U(p_0) - V(p_0) > 0$ and $U'(p_0) - V'(p_0) = 0$. But then, the difference $U(p) - V(p)$ must be strictly increasing for $p > p_0$, a contradiction. In fact, suppose that $U(p_0) \leq V(p_0)$ and $U'(p_0) \leq V'(p_0)$. If $U'(p) \leq V'(p)$ for all $p > p_0$ then we must also have $U(p) < V(p)$ on that range. Otherwise, let

$$p_1 \triangleq \inf\left\{p \in [p_0, 1) : U'(p) > V'(p)\right\}$$

then $U'(p_1) = V'(p_1)$ by continuity, and $U(p_1) < V(p_1)$ since $U(p_0) \leq V(p_0)$ and $U'(p) < V'(p)$ on $[p_0, p_1)$. By the optimality equation, it follows that $U''(p_1) - V''(p_1) < 0$, therefore $U'(p_1 - \varepsilon) > V'(p_1 - \varepsilon)$ for sufficiently small $\varepsilon > 0$, and this contradicts the definition of p_1.

Step 3. Monotonicity. We want to show that the unskilled specialist's equilibrium payoff is weakly increasing in the investors' belief p. Suppose V is not weakly increasing on $[0,1]$. Take a maximal subinterval $[p_0, p_1]$ on which V is strictly decreasing. Recall that $V(0) = 0$ and $V(1) = \frac{f(1-\gamma)}{r}$, then since $V(0) < V(1)$ it follows that $[p_0, p_1] \neq [0,1]$. Take $p_1 < 1$. Since p_1 is a local minimum, $V'(p_1) = 0$. Also, $V(p_1) \geq \pi_{|p_1}(a^*, k^*)$ otherwise $V''(p_1) < 0$. Hence

$$V(p_0) > V(p_1) \geq \pi_{|p_1}(a^*, k^*) \geq \pi_{|p=0}(a^*, k^*) = V(0).$$

Then $p_0 > 0$. Therefore $V'(p_0) = 0$ and $V''(p_0) > 0$ (because $\pi_{|p_1}(a^*, k^*) \geq \pi_{|p_0}(a^*, k^*)$), and so p_0 is a strict local minimum, a contradiction.

Step 4. Comparative statics. We can now show that the unskilled specialist's value function is decreasing in the volatility σ, in the discount rate r and in the conflict of interest b. We are going to focus on the effect of the volatility σ, as the proof for the comparative static with respect to r and b follows the same steps.

First, we can reformulate the Hamilton–Jacobi–Bellman equation as

$$V'' = G\left(p, V, V'\right)$$

with boundary conditions $V(0) = 0$ and $V(1) = \frac{1+f(1-\gamma)}{r}$.

Then, the function V_L (V_H) is called a *subsolution* (*supersolution*) of the restated problem if $V_L'' \geq G(p, V_L, V_L')$ ($V_H'' \leq G(p, V_H, V_H')$). One of the properties of the sub and supersolution (see Berfeld and Lakshmikantham, 1974) is that if $V_L(p)$ and $V_H(p)$ are sub- and supersolution of V, and $V_H(p) > V_L(p)$ then

$$V_L(p) \leq V(p) \leq V_H(p).$$

Let us assume that $\sigma_2 > \sigma_1$ and suppose by contradiction that for some p, $V_{\sigma_2}(p) > V_{\sigma_1}(p)$. Then $V_{\sigma_2}(p) - V_{\sigma_1}(p)$ must attain a local maximum. At the maximum point we then have

$$V_{\sigma_2}''(p) - V_{\sigma_1}''(p) \leq 0.$$

Hence, the formulation of the HJB equation implies

$$G\left(p, V_{\sigma_1}, V_{\sigma_1}'\right) \geq G\left(p, V_{\sigma_2}, V_{\sigma_2}'\right),$$

which contradicts the assumption that $\sigma_2 > \sigma_1$ and $V_{\sigma_2}(p) > V_{\sigma_1}(p)$. Then, the specialist's value function is decreasing in the volatility of the returns dR_t.

Proof of Proposition 2

To simplify notation, let us assume that $\underline{k} = 0$. Let us start by showing that $(a^*, k^*) = (0, 0)$ cannot be an equilibrium for every p. If $\hat{a} = 0$ and $p < \bar{p}$ then the best response is to set $k = 0$, but if $k = 0$ the FOC for the specialist becomes $\frac{p(1-p)V'(p)}{r\sigma^2} > 0$ which means that the specialist has an incentive to increase his effort a^*, which is a contradiction.

Similarly, $(a^*, k^*) = (\alpha, 1)$ is not an equilibrium for any p. Suppose that $\hat{a} = \alpha$, then $k = 1$, however the FOC for the specialist becomes $(\gamma - b) < 0$ which means that the specialist has an incentive to decrease his optimal choice a^*, which is a contradiction.

It is straightforward to see that $(a^*, k^*) = (\alpha, 0)$ cannot be an equilibrium either, because if $\hat{a} = \alpha$ the investors' best response is to set $k = 1$.

Now we can find the threshold for the specialist's reputation that identify the exploitation region. Suppose that $\hat{a} = 0$. The investor's best response is $k = 1$ if and only if $p > \bar{p}$, where $\bar{p}$ is defined as

$$(1-\gamma)(\bar{p}\alpha - f) - r_s = 0$$

For $p < \bar{p}$, we have $k^* = 0$.

For reputation values $p \in \left(\underline{p}, \bar{p}\right)$, the first order conditions hold with equality. The investor's FOC defines the specialist's effort choice, while the specialist's FOC defines the optimal investment, as function of p.

We can now find the interval of p supporting $(0, 0)$ as an equilibrium. The lower threshold $\underline{p}$ is defined by the following condition

$$V\left(\underline{p}\right) = L \tag{12}$$

We know that $V(0) = 0 < L$, which means that given the monotonicity of the specialist's value function $\underline{p} > 0$. In order to make sure that $\underline{p} < 1$ we need to impose the following parametric restriction:

$$V(1) = \frac{f(1-\gamma)}{r} > L$$

that is, the payoff that the specialist can capture when his maximum reputation is achieved needs to be higher than the cost of running the fund L. Otherwise, the specialist would never participate in this market.

Proof of Proposition 3

Given the closed form solution for the effort choice given by:

$$a^* = \frac{r_s + f(1-\gamma)}{(1-\gamma)(1-p)} - \frac{p}{(1-p)}\alpha,$$

we can just differentiate it to immediately obtain

$$\frac{da}{df} > 0, \frac{da}{dr_s} > 0$$

and

$$\frac{da}{d\gamma} = \frac{(f + r_s)(1-p)}{(1-\gamma)^2(1-p)^2} > 0$$

We can then apply Cramer's rule to the set of FOCs to show that the optimal effort choice is decreasing in the reputation value p :

$$\frac{da}{dp} = \frac{\begin{vmatrix} u_{kk} & -u_{kp} \\ \pi_{ak} & -\pi_{ap} \end{vmatrix}}{\begin{vmatrix} u_{kk} & u_{ka} \\ \pi_{ak} & \pi_{aa} \end{vmatrix}} = \frac{(1-\gamma)(\alpha-\hat{a})(\gamma-b)}{(b-\gamma)(1-\gamma)(1-p)} < 0$$

Proof of Proposition 4

The threshold $\underline{p}$ is given by the condition

$$V\left(\underline{p}\right) = L$$

Since the left hand side is increasing in p, an increase in L leads to an increase in the threshold $\underline{p}$. Moreover, we have shown in Proposition 1 that the specialist's value function is decreasing in σ, r and b, which means that an increase in the value of the threshold $\underline{p}$ is needed in order to keep satisfying condition (12).

Proof of Proposition 5

The optimal investment function $k(p)$ is increasing in p because the investors' expected payoff is supermodular in k and p. That is, we have that

$$\frac{\partial \mathbb{E}u\left(\bar{a}(p), k\right)}{\partial k \partial p} = (1-\gamma)\,\alpha > 0.$$

The comparative statics result with respect to σ and b follow from the inspection of the FOC defining the optimal $k^*(p)$.

Proof of Proposition 6

The proof is similar to the one proposed by Garleanu et al. (2009). We start with the case in which the unskilled specialist chooses $e_t = 1$. The HJB for the market maker $J(W)$ is

$$rJ(W) = \max_{c,\bar{e}} u(c) + (rW + \bar{e}\rho - c)\,J(W)' + \varphi\left(J(W+\bar{e}\xi) - J(W)\right),$$

and the FOC with respect to c and $\bar{e}$ are given by the following:

$$\begin{aligned} u'(c) &= pJ_s(W)' + (1-p)J_u(W)', \\ \rho &= -\xi\varphi \frac{pJ_s(W+\bar{e}\xi)' + (1-p)J_u(W+\bar{e}\xi)'}{pJ_s(W)' + (1-p)J_u(W)'}. \end{aligned} \tag{13}$$

We can now conjecture that

$$\begin{aligned} J_s(W) &= -\frac{1}{\delta r}e^{-\delta r(W+b_s)} \\ & \quad \textit{and} \\ J_u(W) &= -\frac{1}{\delta r}e^{-\delta r(W+b_u)}, \end{aligned}$$

where b_s and b_u are some constant, with $b_s > b_u$, since the value function in high demand of CDS is higher. We can now impose a market clearing condition for each case. Then, we have

$$\begin{aligned} \frac{J_s(W+(o+1)\xi)'}{e^{-\delta rW}} &\backsimeq 1 - \delta r\xi(o+1) - \delta r b_s \\ \frac{J_s(W)'}{e^{-\delta rW}} &\backsimeq 1 - \delta r b_s \\ \frac{J_u(W+(o-1)\xi)'}{e^{-\delta rW}} &\backsimeq 1 - \delta r\xi(o-1) - \delta r b_u \\ \frac{J_u(W)'}{e^{-\delta rW}} &\backsimeq 1 - \delta r b_u. \end{aligned} \tag{14}$$

Thus, the equilibrium premium is approximately given by

$$\begin{aligned} \rho &= -\xi\varphi \frac{p(1-\delta r\xi(o+1) - \delta r b_s) + (1-p)(1-\delta r\xi(o-1) - \delta r b_u))}{p(1-\delta r b_s) + (1-p)(1-\delta r b_u)} \\ &= -\xi\varphi\left(1 - \delta r\xi\frac{2p+o-1}{1-\delta r\bar{b}}\right) \end{aligned}$$

where $\bar{b} = pb_s + (1-p)b_u$.

We can also solve b_s and b_u by plugging the optimal consumption rule and equilibrium premium in HJB:

$$\begin{aligned} (o+1)\delta r\xi^2\varphi\frac{2p+o-1}{1-\delta r\bar{b}} &= (o+1)\xi\varphi + r\tilde{b} + \frac{\varphi}{\delta r}\left(e^{-\delta r(o+1)\xi} - 1\right), \\ (o-1)\delta r\xi^2\varphi\frac{2p+o-1}{1-\delta r\bar{b}} &= (o-1)\xi\varphi + r\tilde{b} + \frac{\varphi}{\delta r}\left(e^{-\delta r(o-1)\xi} - 1\right), \end{aligned}$$

where $\tilde{b}$ satisfies the following equation:

$$e^{-\delta r\left(W+\tilde{b}\right)} = pe^{-\delta r(W+b_s)} + (1-p)e^{-\delta r(W+b_u)}.$$

For the case in which the unskilled specialist chooses an interior solution, the optimal e solves the following

$$r\gamma\bar{k}\rho + \frac{\chi\left(\alpha,\hat{a},p\right)}{\sigma}V'(p) = \varphi\eta'(e)\bar{V}. \tag{15}$$

From equation (13), we know that the equilibrium premium is equal to

$$\rho = -\xi\varphi\frac{pJ_s\left(W+\bar{e}\xi\right)' + (1-p)J_u\left(W+\bar{e}\xi\right)'}{pJ_s(W)' + (1-p)J_u(W)'}.$$

We can then impose market clearing conditions:

$$\begin{aligned}
J_s\left(W+\bar{e}\xi\right)' &= J_s\left(W+(o+1)\xi\right)' = \exp\left(-\delta r(W+(o+1)\xi+b_s)\right) \\
J_u\left(W+\bar{e}\xi\right)' &= J_u\left(W+(o-e)\xi\right)' = \exp\left(-\delta r(W+(o-e)\xi+b_u)\right)
\end{aligned}$$

To find a closed form solution, we can use an approximation as equation (14). Thus, the equilibrium premium is

$$\begin{aligned}
\rho &= -\xi\varphi\frac{p\left(1-\delta r\xi(o+1)-\delta rb_s\right) + (1-p)\left(1-\delta r\xi(o-e)-\delta rb_u\right)}{p(1-\delta rb_s) + (1-p)\left(1-\delta rb_u\right)} \\
&= -\xi\varphi\left(1-\delta r\xi\frac{o+p-e(1-p)}{1-\delta r\bar{b}}\right).
\end{aligned} \tag{16}$$

Proof of Proposition 7

The HJB in the case of flight to safety is given by the following:

$$\begin{aligned}
\left(r+\varphi\right)V\left(p\right) &= \max_{a,e} r\pi\left(a,e,\bar{k}\right) + \frac{\chi\left(\alpha,\hat{a}_t,p_t\right)}{\sigma}\left[(a+e)-\bar{a}\left(p\right)\right]V'\left(p\right) \\
&\quad -\frac{\chi\left(\alpha,\hat{a}_t,p_t\right)^2}{(1-p)}V'\left(p\right) + \frac{\chi\left(\alpha,\hat{a}_t,p_t\right)^2}{2}V''\left(p\right).
\end{aligned} \tag{17}$$

Hence, the FOC with respect to e is given by

$$r\gamma\bar{k}\rho + \frac{\chi\left(\alpha,\hat{a}_t,p_t\right)}{\sigma}V'\left(p\right) > 0$$

since $V'(p) > 0$. This means that the optimal exposure e is a corner solution $e = 1$. Given this result, the expected returns of the investors are increased by an amount equal to the premium ρ collected by the specialist. This means that he needs to provide a lower a^* to make the investors willing to delegate a higher amount of capital to the specialist.

Proof of Proposition 8

Define the following function:

$$F(p,e) = r\gamma\bar{k}\rho + \frac{\chi(\alpha,\hat{a},p)}{\sigma}V'(p) - \varphi\eta'(e)\bar{V}$$

where the premium is given by

$$\rho = -\xi\varphi\left(1 - \delta r\xi\frac{o + p - e(1-p)}{1-\delta r\bar{b}}\right)$$

Then, we can totally differentiate the above equation to get $\frac{de}{dp}$:

$$\frac{de}{dp} = -\frac{F_p}{F_e},$$

where

$$\begin{aligned}
F_p &= r\gamma\bar{k}\frac{d\rho}{dp} + \frac{\chi'}{\sigma}V' + \frac{\chi}{\sigma}V'', && (18)\\
F_e &= r\gamma\bar{k}\frac{d\rho}{de} - \varphi\eta''\bar{V} < 0, \\
\frac{d\rho}{dp} &= \delta r\varphi\xi^2\frac{1 + e + \delta r\left(o(b_s - b_u) - eb_s - b_u\right)}{\left(1-\delta r\bar{b}\right)^2}, && (19)\\
\frac{d\rho}{de} &= -\frac{\delta r\varphi\xi^2(1-p)}{1-\delta r\bar{b}} < 0.
\end{aligned}$$

Thus, the sign of de/dp depends on the sign of F_p. The first term of F_p depends on $d\rho/dp$. The second term of F_p depends on p. When p is close to zero, the second term is positive, and when p is close to one, the second term is negative. The third term of F_p is assumed to be negative given that specialist's value function is concave. If the first term is positive and sufficiently large, we can say that F_p is positive, and thus the optimal e is increasing in the specialist's reputation. From the equations in 18, we have that F_p is positive when the expected size of disaster, $-\varphi\xi$, is large and the noisy demands for CDS is large.