

Rewarding Trading Skills without Inducing Gambling

IGOR MAKAROV and GUILLAUME PLANTIN*

ABSTRACT

This paper develops a model of active asset management in which fund managers may forgo alpha-generating strategies, preferring instead to make negative-alpha trades that enable them to temporarily manipulate investors' perceptions of their skills. We show that such trades are optimally generated by taking on hidden tail risk, and are more likely to occur when fund managers are impatient and when their trading skills are scalable, and generate a high profit per unit of risk. We propose long-term contracts that deter this behavior by dynamically adjusting the dates on which the manager is compensated in response to her cumulative performance.

The last 30 years have witnessed two important developments in financial markets. First, financial innovation has made it possible to slice and combine a large variety of risks by trading a rich set of financial instruments. Second, management of large amounts of capital has been delegated to entities such as hedge funds and bank proprietary desks that are not subject to significant trading restrictions and not required to disclose publicly the details of their positions.

The amount of capital available to such entities crucially depends on investors' perceptions of their "alpha," that is, their ability to generate excess returns above the level of fair compensation for risk. Combined with the relative opaqueness of these entities and their vast risk-taking opportunities, this creates pervasive incentives for money managers. In particular, fund managers running out of alpha-generating strategies may find it tempting to pretend otherwise, taking risky positions with zero or even negative alpha that may temporarily improve their perceived reputation in the event of favorable outcomes. Strategies that generate frequent small positive excess returns that

*Makarov is with the London School of Economics and Plantin is with the Toulouse School of Economics. We are most grateful for helpful comments made by the Editor (Cam Harvey) and two anonymous referees, as well as Gilles Chemla, Zhiguo He, Christopher Hennessy, Roni Kisin, Peter Kondor, Semyon Malamud, Gustavo Manso, Andrea Prat, Jan Schneider, Lars Stole, Ilya Strebulaev, and seminar participants at Chicago Booth, Imperial College, Kellogg, LBS, Pompeu Fabra, Rochester, Queen Mary, University of London, University of Lugano, University of Tokyo, Vienna University, American Finance Association 2011 Denver meetings, HKUST Finance Symposium, Second Paris Hedge Fund Conference, Third Annual Paul Woolley Centre Conference, Society for Economic Dynamics 2010 Montreal Meeting, Second Theory Workshop on Corporate Finance and Financial Markets, New York, and UBC Summer Conference 2010. Plantin benefitted from a European Research Council Starting Grant (No. 263673 – RIFIFI).

DOI: 10.1111/jofi.12257

are offset by very rare and large losses seem especially well suited to disguising luck as skill. As Rajan (2008) puts it, “How can untalented investment managers justify their pay? Unfortunately, all too often it is by creating fake alpha – appearing to create excess returns but actually taking on hidden tail risk.”¹ Consistent with this view, Jiang and Kelly (2012) show that a significant number of hedge funds are indeed exposed to tail risk.

Creating fake alpha by taking on hidden tail risk does not appear to be limited to the hedge fund industry. For example, in its 2008 report to shareholders on the causes of its subprime losses, UBS concludes that “The UBS compensation and incentivization structure did not effectively differentiate between the creation of alpha versus the creation of return based on a low cost of funding.” More generally, Acharya et al. (2010) argue that the manufacture of tail risk through deliberate retention of senior tranches on poor collateral by U.S. banks was an important ingredient of the 2008 banking crisis.

The perverse incentives to enter into (at best) zero-alpha gambles are associated with a number of costs. First, they defeat the purpose of delegated asset management, which is meant to achieve superior returns by optimally combining “brains and resources.” Second, they lead to a misallocation of capital. The manufacture of tail risk also has far-reaching consequences for overall financial stability when gambling institutions are systemically important.

In this paper, we develop a new framework for the study of these risk-taking incentives. We study situations in which managers find it optimal to fake their alpha and propose a new class of contracts that eliminate incentives to employ such strategies.

Our model builds upon the frictionless benchmark of Berk and Green (2004), who study career concerns in delegated fund management. In their model, a fund manager and investors discover the manager’s alpha-generating skill by observing her realized returns. The excess returns that a manager is expected to generate increase with her skill, but decrease in the number of funds under her management. Competitive investors supply funds to the manager until they earn a zero net (after fees) expected return. At the beginning of each period, the manager sets fees that enable her to reach the optimal fund size, and extracts the entirety of the surplus that she generates. Learning and competition among investors imply that both fund flows and managerial compensation strongly depend on the manager’s record.

We extend this model by allowing the manager to secretly enter into zero-alpha trades with the sole purpose of manipulating investors’ perceptions of her skill. In what follows, we refer to this opportunistic behavior as inefficient risk shifting or gambling. In contrast to earlier papers on risk shifting, we propose a general setting in which the fund manager can secretly choose to take on positions with arbitrary payoff distributions. This captures the large set of trading opportunities available to modern managers, and is therefore an important case to consider.

¹ Raghuram Rajan, “Bankers’ pay is deeply flawed,” *Financial Times*, January 9, 2008.

We first study the impact of this friction in the case in which the manager and investors sign only short-term contracts. Three factors conducive to inefficient risk shifting emerge from our analysis. The first is the size of the alpha per unit of risk that can be generated by a skilled manager. If it is large, the history of returns has a large impact on investors' beliefs about the manager's ability to generate future excess returns. The second factor is the scalability of trading skill, that is, the sensitivity of expected excess returns to fund size. If trading skill is scalable, a good reputation translates into a large future fund size and in turn large future profits. Third, because the manager can manipulate her reputation only temporarily, it is more valuable for her to do so when she is more impatient. These three factors determine the convexity of future expected gains as a function of realized returns, and thus affect inefficient risk-shifting incentives. In particular, the model predicts that "fallen star" managers (those who show high initial potential but realize disappointing returns) are particularly prone to gambling. For a calibration consistent with Berk and Green (2004), we find that their efficient equilibrium with short-term contracts breaks down in the sense that any equilibrium must involve some degree of risk shifting.

We are able to fully characterize such equilibria with risk shifting in a simplified version of the model where the manager maximizes both her expected current return and the expected reputation that results from it. Interestingly, even though we impose no restriction on the risk profiles available to the manager, we show that she optimally manufactures hidden tail risk. In other words, she sells disaster insurance, adding some noisy payoff to the collected premium so that investors cannot discover the exact nature of the trade.

We next consider long-term contracts. Here we follow two distinct lines of inquiry. First, we consider a contract popular in the hedge fund industry, namely, a contract in which the manager's profits are given by a performance fee above a high-water mark. Similar to Panageas and Westerfield (2009), we find that, without new inflows or outflows triggered by realized performance, the performance fee based on a high-water mark does not induce inefficient risk shifting. In the presence of fund flows, however, we show that the high-water mark contract does not generally solve the risk-shifting problem.

Second, we consider an optimal contract that fully eliminates risk-shifting incentives. This contract is designed to discriminate between skill and luck, exploiting the fact that the impact of gambling on investors' beliefs vanishes in the long run, when true skill is eventually revealed. The contract defers payments to the manager at dates that vary depending on her cumulative performance. The promised payment also evolves such that it always at least matches the manager's outside options, so she does not renegotiate the contract. As we detail in Section II, this mechanism is highly reminiscent of recent proposals for bankers' compensation reforms issued by both public authorities and the industry itself. All such proposals consist of a bonus deferral together with a clawback mechanism that revises the initial promised payment with the benefit of hindsight. We offer theoretical foundations for these proposals. More importantly, we qualify them, suggesting that it is important to adjust

the timing of compensation dynamically, further postponing it upon observing a poor track record. In our model, committing to pay a bonus at a fixed date may generate inefficient gambling in general, even if this date is remote and even if the bonus is adjusted via a clawback provision. The dynamic revision of the payment date turns out to be important, and yet it is absent from the current suite of proposed reforms.

To our knowledge, this paper is the first to derive a formal connection between fund managers' career concerns and alpha-faking through hidden tail risk in a fully rational environment. More precisely, we bridge two strands of literature, namely, that on risk shifting and that on career concerns. The risk-shifting friction was first introduced by Jensen and Meckling (1976) as a source of value destruction within overly leveraged firms. Arguably, this friction is particularly relevant in the context of sophisticated financial institutions that can swiftly alter their risk profiles. Accordingly, a large asset pricing literature considers the impact of nonconcave objective functions on the risk-shifting incentives of fund managers who have access to dynamically complete markets. Contributions include those of Basak, Pavlova, and Shapiro (2007), Carpenter (2000), and Ross (2004). In line with these studies, we seek to identify the risk-taking strategies that optimally respond to nonconcave objectives. Different from prior research, however, nonconcavities in the manager's objective are not assumed in our model. Rather, they arise endogenously from reputation concerns in a truly dynamic environment. Further, we extend the literature by identifying optimal contracts that eliminate the costs of this friction.

Acharya, Pagano, and Volpin (2012) also develop a model in which career concerns may lead managers to destroy value, but by a mechanism other than manipulating a payoff distribution. As in our model, limited commitment prevents managers from receiving insurance against the risk that their reputation deteriorates. The authors assume that learning about managerial skill can take place only if managers run the same project sufficiently long. Thus, risk-averse managers may prefer to churn across projects to prevent learning. While such churning shields managers from reputation risk, it inefficiently slows down the identification of good managers. Malliaris and Yan (2012) consider a two-period model in which a manager may be tempted to take on tail risk to manipulate her expected reputation. Their setup is related to the static version of our model, which we solve for equilibria with risk shifting. The main difference is that we do not impose binary payoffs, as they do.

Our paper is also related to Goetzmann et al. (2007), who study manipulation-proof measures of managerial performance. They show that to be manipulation-proof a measure should take the form of a concave utility function averaged over the return history. We further show that if the fund manager has a nonconcave continuation utility she can engage in inefficient risk shifting, and that the optimal contracts aim to concavify the manager's objective.

Finally, our study relates to two recent extensions of the Berk and Green (2004) model. First, Berk and Stanton (2007) apply the Berk and Green (2004) setup to closed-end funds. In this case, learning affects the net asset value of

the fund and not its size, which is fixed by construction. Berk and Stanton (2007) show that the impact of learning explains several characteristics of the closed-end fund discount, and that the behavior of this discount crucially depends on the nature of the compensation contract. Second, Dangl, Wu, and Zechner (2008) study an extension of Berk and Green’s (2004) approach in which a management company can fire a manager if her performance is not good enough. They restrict their analysis to short-term compensation contracts and solve for the optimal firing rule.

The remainder of our paper is organized as follows. Section I develops our baseline model of career concerns, and studies the impact of the risk-shifting friction when the manager and investors sign only short-term contracts. Section II studies long-term contracts. Section III concludes. Technical proofs are relegated to the Appendix.

I. Career Concerns and Inefficient Risk Shifting

In Section I.A, we introduce and solve a frictionless model of career concerns in delegated asset management that closely follows the approach used by Berk and Green (2004). In Section I.B, we introduce a risk-shifting friction to this benchmark approach, and study its impact on the equilibrium in the presence of short-term contracts. In Section I.C, we fully characterize equilibria with risk shifting in a simplified version of our baseline model. In Section I.D, we discuss the costs of risk-shifting strategies.

A. Frictionless Benchmark: The Berk and Green model

Time is discrete and is indexed by $\{n\Delta t\}$, where $n \in \mathbb{N}$ and $\Delta t > 0$. There is a single consumption good that serves as the numéraire. Agents are of two types: a manager and investors. Agents live forever, are risk neutral, and discount future consumption at the instantaneous rate $r > 0$.

The manager is protected by limited liability and thus is unable to have negative consumption. Investors receive a large endowment of the consumption good at each date $n\Delta t$; the manager does not. The manager has exclusive access to an investment technology. If the manager invests q_t consumption units at date t using her technology, she generates $q_{t+\Delta t}$ units at date $t + \Delta t$ such that

$$q_{t+\Delta t} = q_t e^{\left(r+a\theta-c(q_t)-\frac{\sigma^2}{2}\right)\Delta t+\sigma(B_{t+\Delta t}-B_t)}, \tag{1}$$

where $(B_t)_{t \geq 0}$ is a standard Wiener process, and a and σ are strictly positive numbers.

The parameter $\theta \in \{0, 1\}$ measures the manager’s skill, which is unobservable by both the manager and investors. All other parameters are common knowledge. The parameter a is the alpha that a skilled manager can generate with her first dollar. As in Berk and Green (2004), the cost $c(q_t)$ captures the fact that many arbitrage opportunities or informational rents in financial

markets are not perfectly scalable. The function c is increasing, taking the form

$$c(q) = \beta q^{\frac{1}{\alpha-1}}, \quad (2)$$

where $\alpha \geq 1$ and $\beta > 0$.

All agents share a common date 0 prior that the manager is endowed with high skill, that is, $\theta = 1$, with probability $\pi_0 \in (0, 1)$. Except for the manager's skill, each action and realized return of the manager are publicly observable at each date $n\Delta t$. Thus, information is symmetric across agents.

Let π_t denote the probability that the agents assign to the possibility that the manager is skilled at a given date t . We refer to π_t as the manager's perceived skill. Given $\pi_t = \pi$, the net expected surplus created over $[t, t + \Delta t]$ if the manager invests q units is

$$q\left(\pi e^{(a-c(q))\Delta t} + (1-\pi)e^{-c(q)\Delta t} - 1\right),$$

which in the limit as $\Delta t \rightarrow 0$ becomes $q(a\pi - c(q))$. Thus, the manager creates the maximal net expected surplus over $[t, t + \Delta t]$ if she invests $q(\pi)$ such that

$$q(\pi) = \arg \max_q \left(\pi e^{(a-c(q))\Delta t} + (1-\pi)e^{-c(q)\Delta t} - 1 \right).$$

We refer to such $q(\pi)$ as the optimal fund size. Denoting $v(\pi)\Delta t$ as the maximal expected profit over $[t, t + \Delta t]$ that corresponds to this optimal fund size, we have the following result.

LEMMA 1:

$$\lim_{\Delta t \rightarrow 0} q(\pi) = \left(\frac{\alpha - 1}{\alpha\beta} a\pi \right)^{\alpha-1}, \quad (3)$$

$$\lim_{\Delta t \rightarrow 0} v(\pi) = \beta^{1-\alpha} \frac{(\alpha - 1)^{\alpha-1}}{\alpha^\alpha} (a\pi)^\alpha. \quad (4)$$

PROOF OF LEMMA 1: See the Appendix. □

Using a power specification (2) for the cost function $c(q)$, the expected instantaneous surplus $v(\pi)$ is proportional to π^α as Δt becomes small. The parameter α captures the scalability of trading skill. As α increases, the manager's skill π becomes more scalable and therefore her expected profit becomes more sensitive to her reputation, in other words, becomes more convex in π . In the hedge fund universe, global macro strategies are typically quite scalable, while strategies based on shareholder activism may be more difficult to spread over increasing amounts of capital.

Berk and Green (2004) specify a linear cost function c corresponding to $\alpha = 2$. They show that this specification matches quantitatively well the empirically observed relationship between realized returns of mutual funds and fund flows. In the remainder of the paper, we generally restrict the analysis to the somewhat simpler limiting case in which $\alpha = 1$. In this case, the cost is zero for $q < 1$

and infinite for $q > 1$. Thus, the manager can scale up the fund at no cost up to an upper bound that is (without loss of generality) normalized to one. The optimal fund size is thus always one and $v(\pi)$ is linear in π . Thus, in this case perceived skill affects the expected rate of return of the fund but the optimal fund size remains constant.

To realize any surplus, investors must delegate their funds to the manager and agree with her on the profit-sharing rule. We assume that investors are competitive and can fully commit to a contract. However, at the end of any period, the fund manager is free to walk away from a contract and sign a new one with competing investors. More precisely, at the end of each period $[t, t + \Delta t]$, after returns are realized and all contractual transfers for the period are made, the fund manager is free to terminate a contract without financial obligation to its investors and enter a new one with new investors starting at period $[t + \Delta t, t + 2\Delta t]$. In other words, commitment is one-sided. This is a common assumption in labor economics. We find it all the more plausible in the financial services industry, where fund managers can swiftly move between jobs and financial centers because their activity entails few specific investments. Limited cross-border enforcement precludes covenants that would make such moves costly.

In this section, we study incentives to take risk when investors and the manager cannot enter into long-term contracts, but rather interact in a spot labor market at each date. Our goal is to determine whether incentives created by market forces alone can discipline managers who are concerned about their reputation in the labor market. We assume, like Berk and Green (2004), that the manager only enters into one-period contracts with investors.

ASSUMPTION 1: *At each date t , the manager offers investors a one-period asset management contract.*

Assumption 1 imposes a restriction only on the horizon of contracts: the parties cannot contract at date t on actions or transfers beyond date $t + \Delta t$. Parties are free to write any one-period contracts, however, subject to the limited liability constraint of the manager. We next show that this restriction to short-term contracts does not lead to misallocation of capital in the absence of frictions.

LEMMA 2: *Under Assumption 1, the manager adopts the optimal fund size $q(\pi)$ at each date and extracts the maximal expected surplus $v(\pi)\Delta t$.*

PROOF OF LEMMA 2: One only needs to exhibit a particular contractual arrangement that enables the manager to raise $q(\pi)$ in each period and receive expected compensation $v(\pi)\Delta t$ over $[t, t + \Delta t]$. In the absence of any frictions, there are many different arrangements in which this can be achieved. For example, the manager can simply ask investors at the beginning of each period to pay her a salary $v(\pi)\Delta t$, collect funds $q(\pi)$ from them, invest, and leave them the date $(t + \Delta t)$ proceeds. As assumed by Berk and Green (2004), she may alternatively quote a fee $f\Delta t$ at the beginning of each period. The fee is the fraction of the date $(t + \Delta t)$ assets under management (before any new inflows or outflows of funds) that accrues to the manager. If a manager with perceived skills π quotes

a fee $f \Delta t$, competitive investors will supply funds as long as their net expected rate of return is equal to r . Thus, their fund supply $q(f)$ solves

$$(1 - f \Delta t) \left(\pi e^{(r+a-c(q(f)))\Delta t} + (1 - \pi)e^{(r-c(q(f)))\Delta t} \right) = e^{r\Delta t}.$$

Therefore, as $\Delta t \rightarrow 0$, $q(f)$ solves

$$\pi a = c(q(f)) + f. \tag{5}$$

The manager maximizes her expected profits by choosing f such that

$$f = \arg \max_f f \Delta t \times q(f) = \arg \max_f (\pi a - c(q(f))) \Delta t \times q(f),$$

which implies that the manager collects the maximal expected surplus $v(\pi) \Delta t$. \square

This baseline model is essentially identical to that of Berk and Green (2004). The main modeling difference is that the distribution of skills is binomial in our setup while it is Gaussian in theirs. With our specification, the model is stationary in perceived skill π and is therefore more tractable. In particular, this specification yields a tractable formulation of the manager’s total expected profit from date 0 onwards. While this is not particularly useful in the frictionless environment of Berk and Green (2004), it turns out to be instrumental when we introduce asymmetric information between the manager and investors.

Let $(\pi_{n\Delta t})_{n \geq 0}$ denote the process that describes the manager’s perceived skill at each date. From Lemma 2, the manager’s continuation utility is

$$V(\pi, \Delta t) = E \left[\sum_{n=0}^{\infty} e^{-rn\Delta t} v(\pi_{n\Delta t}) \Delta t \mid \pi_0 = \pi \right],$$

where $v(\pi)$ is given by (4). The following proposition shows that this continuation utility converges to a simple limit when Δt becomes small.

PROPOSITION 1: *Let*

$$V(\pi) = \lim_{\Delta t \rightarrow 0} V(\pi, \Delta t).$$

We have

$$V(\pi) = \int_0^1 G(\pi, x) v(x) dx, \tag{6}$$

where

$$G(\pi, x) = \frac{2\sigma^2}{\psi a^2 x^2 (1-x)^2} \begin{cases} g(1-\pi)g(x) & \text{if } 0 \leq x \leq \pi \\ g(\pi)g(1-x) & \text{if } \pi \leq x \leq 1 \end{cases}, \tag{7}$$

and

$$\psi = \sqrt{1 + \frac{8r\sigma^2}{a^2}}, \tag{8}$$

$$g(u) = u^{\frac{1}{2} + \frac{1}{2}\psi} (1 - u)^{\frac{1}{2} - \frac{1}{2}\psi}.$$

Convergence of $V(\pi, \Delta t)$ to $V(\pi)$ when $\Delta t \rightarrow 0$ is uniform over $\pi \in (0, 1)$.

PROOF OF PROPOSITION 1: See the Appendix. □

In the remainder of the paper, all results are established for Δt sufficiently small to approximate the manager’s continuation utility with its continuous-time limit (6).

Expression (6) can be interpreted as an expectation over the instantaneous surplus $v(x)$ weighted by a discount factor $G(\pi, x)$. The factor $G(\pi, x)$ has an intuitive interpretation: it measures the discounted frequency of the future dates at which the manager will have a perceived ability x given that she starts out with perceived skill π . Notice that the factor G depends only on a/σ , which governs the speed with which agents learn about the manager’s skill, and on the discount rate r . The cost parameters α and β affect only the instantaneous profit v from (4). This is because the cost is known and thus filtered out by the agents when inferring skills from realized returns. It is easy to verify that

$$\int_0^1 G(\pi, x) x dx = \frac{\pi}{r}. \tag{9}$$

Thus, in the case $\alpha = 1$, the continuation utility $V(\pi)$ is proportional to π .

We now depart from this frictionless benchmark set by Berk and Green (2004) and introduce informational asymmetry between the manager and investors. We posit that the manager may secretly enter into zero-alpha trades in order to manipulate her perceived skill instead of efficiently investing in the alpha-generating technology described in (1). We show that the spot labor market interactions considered thus far may generate perverse incentives in the presence of this friction.

B. Risk Shifting

The Berk and Green (2004) model assumes a frictionless interaction between the manager and investors. While this is a very useful benchmark, in reality several informational frictions may affect the delegation of asset management to sophisticated entities such as hedge funds and investment banks. As the following quote from the Financial Stability Forum (FSF) “Principles for Sound Compensation Practices” suggests, the impossibility of perfectly monitoring shifts in risk exposures within such institutions can be particularly problematic:

In principle, if risk management and control systems were strong and highly effective, the risk-taking incentives provided by compensation

systems would not matter because risk would stay within the firm's appetite. In practice, all risk management and control systems have limitations and, as the current crisis has shown, they can fail to properly control risks. The incentives provided by compensation can be extremely powerful. Without attention to the risk implications of the compensation system, risk management and control systems can be overwhelmed, evaded, or captured by risk-takers.

To study such incentives for surreptitiously shifting risk exposures, we introduce the following friction to the baseline model.

ASSUMPTION 2: At each date, the manager can secretly invest all or part of her funds in an alternative (risk-shifting) technology whose returns are perfectly scalable and independent of the returns on the technology described in (1). This technology enables her to generate a one-period gross return with any arbitrary distribution over $[0, \infty)$ with mean $e^{r'\Delta t}$, where $r' \leq r$. Investors observe returns realized at the reporting and contracting dates $n\Delta t$, at which the manager's position is marked-to-market.

We now briefly comment on our modeling choice. We interpret the alpha-generating technology (1) available to the manager as the investment strategy that she and the investors and/or her supervisors agree upon. In practice this strategy may vary over time, and involve shifts in asset selection, asset allocation, or overall risk exposure, but such shifts are agreed between the parties. For simplicity, we abstract from changes in strategy and, following Berk and Green (2004), model the agreed trading strategy as a simple production function (1).

The risk-shifting technology that the manager may secretly use reflects the positions that she can conceal from the investors and/or her supervisors by evading the various control and risk management systems put in place to monitor her. To bring our results into sharper focus, we assume that these trades are detrimental to investors in that they do not generate a positive risk-adjusted expected excess return. Moreover, we allow for the possibility that concealing these trades from investors comes at a cost $r - r'$. The assumption that these trades are independent of the alpha-generating technology ensures that they cannot be used for arbitrage purposes. Because such trades yield less in terms of expectation than the manager's alpha-generating technology (1), she will not invest her own funds in them. In the presence of career concerns, however, the manager may be tempted to use the risk-shifting technology because it provides her an opportunity to manipulate investors' beliefs about her reputation.

Given the large set of trading opportunities available to sophisticated managers, we consider a general setting in which the fund manager can secretly choose any arbitrary payoff distribution. One of the questions we are interested in relates to understanding which trades work best in manipulating investors' beliefs. For expositional simplicity, we assume perfect scalability of these trades. It is a straightforward matter to extend our analysis to the case in which large trades are more difficult to hide than small trades.

The manager in our model may be literally interpreted as an individual trader, or alternatively may be interpreted as a desk or division that collectively decides to breach its mandate or exceed its risk limits. In the remainder of the paper, we refer to the secret use of this inefficient technology by the manager, as opposed to the transparent risk taking induced by the alpha-generating technology (1), as risk shifting or gambling.

We now study whether the friction introduced in Assumption 2 affects the outcomes in the Berk and Green (2004) environment for short-term contracts (Assumption 1). The first type of opportunistic behavior could occur if the manager secretly invests a portion of her funds at the risk-free rate r' rather than using her alpha-generating technology. This represents the case in which a manager performs well from January to November, say, and then stops investing actively in order to lock in a profit. We have the following result.

PROPOSITION 2: *If*

$$r' \leq r - \sigma^2/2, \tag{10}$$

then the manager does not secretly invest at the risk-free rate.

PROOF OF PROPOSITION 2: See the Appendix. □

The remainder of the paper focuses on situations in which the manager uses the risk-shifting technology to increase rather than reduce risk. We believe that this is a more important question because such excessive risk taking could contribute to financial instability, as suggested by the recent financial crisis. In what follows we therefore assume that condition (10) holds.

We now characterize the manager’s incentives to gamble secretly (and inefficiently). Suppose that the economy is in an equilibrium in which investors believe that the manager always invests in the alpha-generating technology (1). We assess whether the manager could be tempted to deviate and enter into a one-shot gamble.

Suppose that the manager gambles during her first trading round, realizes return R , and does not gamble further. Also suppose that investors believe that the manager has never gambled. Let $\pi_{n\Delta t}$ be the manager’s (correct) belief about her skill at date $n\Delta t$ and let $\pi_{n\Delta t}^R$ be investors’ (incorrect) perception of her skills at date $n\Delta t$. We have the following result.

PROPOSITION 3:

$$\pi_t^R \equiv \lim_{\Delta t \rightarrow 0, n\Delta t \rightarrow t} \pi_{n\Delta t}^R = \frac{\pi_t R \frac{\alpha}{\sigma^2}}{1 - \pi_t + \pi_t R \frac{\alpha}{\sigma^2}}. \tag{11}$$

PROOF OF PROPOSITION 3: See the Appendix. □

REMARK 1: *In principle, a realization of $R = 0$ should perfectly reveal to investors that the manager gambled because her alpha-generating technology delivers strictly positive returns with a probability of one. We assume instead that π_t^R is continuous at zero, so that investors infer $\pi = 0$ from observing $R = 0$. Equivalently, we could assume that traders who are caught gambling are excluded from the market.*

The continuation utility of the manager after the return R is realized is

$$W(\pi, R, \Delta t) = E \left[\sum_{n=0}^{\infty} e^{-rn\Delta t} v(\pi_{n\Delta t}^R) \Delta t \mid \pi_0 = \pi \right].$$

For Δt small, (6) and (11) imply that

$$W(\pi, R) \equiv \lim_{\Delta t \rightarrow 0} W(\pi, R, \Delta t) = \int_0^1 G(\pi, x) v \left(\frac{xR^{\frac{\alpha}{\sigma^2}}}{1 - x + xR^{\frac{\alpha}{\sigma^2}}} \right) dx. \tag{12}$$

Notice that as $\Delta t \rightarrow 0$ the mean of the gamble, $e^{r'\Delta t}$, converges to one irrespective of r' . The manager therefore chooses a unit mean gamble whose distribution maximizes her expectation over $W(\pi, R)$. Formally, denoting by M the set of Borelian probability measures over $[0, +\infty)$, the manager solves

$$\begin{aligned} \max_{\mu \in M} \int_0^{\infty} W(\pi, R) d\mu(R) \\ \text{s.t.} \quad \int_0^{\infty} R d\mu(R) = 1. \end{aligned} \tag{13}$$

In the Internet Appendix,² we show that a generic solution to (13) is attained using a binary gamble—or a measure that consists of two atoms, one above one and the other below one. If these two atoms coincide at one, the manager does not find it worthwhile to gamble and an equilibrium without inefficient risk shifting can be sustained. If they do not, then the optimal binary gamble is not degenerate, which means that such an equilibrium without risk shifting does not exist.

Obviously, a sufficient condition that makes gambling undesirable is that $W(\pi, R)$ is concave in R . More generally, the optimal gamble is determined by the convexity properties of $W(\pi, R)$, which in turn depend on the parameters α, σ, α , and r . We have the following result.

PROPOSITION 4: *Suppose that Assumptions 1 and 2 hold.*

(i) *If*

$$\frac{\alpha}{\sigma^2} \leq \frac{1}{\alpha}, \tag{14}$$

then there exists an equilibrium in which the manager extracts the entire expected surplus and does not engage in risk shifting.

(ii) *If*

$$\frac{\alpha}{\sigma^2} > \frac{1}{\alpha} \quad \text{and} \quad r > \frac{\alpha(\alpha - 1)\alpha^2}{2\sigma^2}, \tag{15}$$

then such an equilibrium does not exist.

² The Internet Appendix is available in the online version of the article on the *Journal of Finance* website.

PROOF OF PROPOSITION 4: See the Appendix. □

To gain a better understanding of the results in Proposition 4, suppose that the manager tries to “pick up nickels in front of a steamroller,” in other words, she gambles and realizes an instantaneous return of $1 + \varepsilon$ with probability $1/(1 + \varepsilon)$, where ε is small, or loses everything. Then, from (11), in the case of success, her new reputation is

$$\pi^R = \frac{\pi (1 + \varepsilon)^{\frac{\alpha}{\sigma^2}}}{1 - \pi + \pi(1 + \varepsilon)^{\frac{\alpha}{\sigma^2}}},$$

which for π and ε small enough is approximately

$$\pi^R \simeq \pi_0 \left(1 + \frac{\alpha}{\sigma^2} \varepsilon \right). \tag{16}$$

Further, Lemma 6 in the Appendix shows that, if r is not too small, then the continuation utility of the manager $V(\pi)$ behaves as π^α as $\pi \rightarrow 0$. Therefore, the manager’s net expected gain from the gamble for π and ε small enough is proportional to

$$\frac{1}{1 + \varepsilon} \left(\pi^R \right)^\alpha - \pi^\alpha \simeq \pi^\alpha \left(\frac{\alpha\alpha}{\sigma^2} - 1 \right) \varepsilon. \tag{17}$$

Thus, whether there is risk shifting depends on whether the ratio $\alpha\alpha/\sigma^2$ is greater or less than one.

Proposition 4 shows that risk shifting is particularly appealing when three conditions are met. First, managerial skill generates a high alpha per unit of risk (α/σ^2 is large). One can see from (16) that, in this case, realized returns have a large impact on investors’ beliefs. Second, the manager’s skill is sufficiently scalable (α is large). In this case, positive news about her skill translates into a large increase in expected future fund size, and thus into large future expected profits. Finally, the manager should be sufficiently impatient (r is large). If the manager is patient, she only cares about the long run in which she ends up with the reputation that she deserves regardless of her earlier attempts to gamble. In this regard, it is worth noting that condition $r > \frac{\alpha(\alpha-1)\alpha^2}{2\sigma^2}$ in (15) is only a sufficient condition for risk shifting. Numerical analysis shows that the manager is tempted to gamble under much milder conditions on r when $\alpha/\sigma^2 > 1/\alpha$ is satisfied.

To assess whether the risk-shifting friction is likely to be important in practice, we consider a calibration consistent with that of Berk and Green (2004). We set $\alpha = 2$, $a = 5\%$, $\sigma = 25\%$, and $r = 5\%$. Simple calculations show that conditions (15) are satisfied. More generally, condition $\alpha/\sigma^2 > 1/\alpha$ is very likely to be satisfied in practice as it holds whenever the Sharpe ratio of a portfolio strategy is larger than its volatility ($a/\sigma > \sigma$), which is true for almost all investment strategies. In sum, we find that risk shifting matters in this model for parameter values that are empirically plausible.³

³ Our results are also broadly supported by empirical findings of Brown, Harlow, and Starks (1996) and Chevalier and Ellison (1997), who show that mutual fund managers tend to take on more risk following disappointing performance.

C. Equilibria with Risk Shifting

Proposition 4 establishes that, under the plausible conditions (15), an equilibrium without gambling does not exist. This raises two questions: under these conditions, which gambles emerge in equilibrium, and how does gambling in equilibrium affect learning and the distribution of realized returns?

To answer these questions, we consider a reduced-form version of our model that allows us to abstract from complex signaling issues that would arise in an infinite horizon dynamic model, but nevertheless yields interesting insights into equilibrium gambling strategies.⁴ Specifically, we study the following static model. We assume that the manager makes only one investment decision: she can invest one unit of capital using one of two technologies. First, she can use an alpha-generating technology, which produces a gross return of

$$e^{a\theta - \sigma^2/2 + \sigma\xi}, \quad \xi \sim N(0, 1), \quad (18)$$

where again $\theta \in \{0; 1\}$ is the unknown ability of the manager. Alternatively, she can invest her funds in the risk-shifting technology that enables her to generate a gross return with any arbitrary distribution over $[0, +\infty)$ with unit mean.

Investors and the manager share the common prior belief π_0 that the manager is skilled ($\theta = 1$). Upon observing the return R realized by the manager, but without knowing if she gambled or used the alpha-generating technology, investors update their beliefs about the manager's skill. As before, we denote by $\pi_1^R = \text{Prob}\{\theta = 1|R\}$ investors' posterior belief that $\theta = 1$, and refer to it as the manager's reputation.

We assume that the manager invests to maximize the sum of her expected current return R and a reduced-form continuation utility that is proportional to her reputation π_1^R :

$$\max E\left(\gamma R + \pi_1^R\right), \quad \gamma > 0.$$

Assuming a continuation utility that is linear in reputation is commonplace in the literature on career concerns (see, for example, Dasgupta and Prat (2008), Holmstrom and Ricart I Costa (1986), and Scharfstein and Stein (1990)). It is also consistent with what the manager obtains in our dynamic model when the fund size is fixed ($\alpha = 1$).

Lemma 3, similar to Proposition 4, shows that, if $a > \sigma^2$, then an equilibrium without gambling does not exist.

LEMMA 3: *If $a > \sigma^2$ and π_0 is sufficiently low, then there exists no equilibrium without gambling.*

⁴In an equilibrium with gambling, unlike in an equilibrium without gambling, the manager and investors possess different information about the manager's ability. While the manager knows whether she gambled, investors can only imperfectly infer this from observed returns. This creates room for signaling, whereby a truly skilled manager would like to credibly distinguish herself from a lucky gambler. Solving for such equilibria with asymmetric information is a very difficult problem and is beyond the scope of this paper.

PROOF OF LEMMA 3: See the Appendix. □

We now assume that the conditions in Lemma 3 are satisfied, and characterize an equilibrium with gambling. As before, we assume that the manager who is caught gambling is excluded from the market and attains zero utility. Thus, in any equilibrium with gambling, the manager cannot gamble with a probability of one. Notice that if the manager decides to gamble she should choose a distribution for her returns that has no atoms other than at zero. Suppose otherwise that investors believe a particular return $\hat{R} > 0$ occurs with a strictly positive probability. Because the distribution of returns of the efficient technology has no atoms, then upon observing \hat{R} investors would conclude that the manager gambles and penalize her with zero payoff. In this case, the manager would be better off setting \hat{R} to zero in the first place and using the realized surplus to improve the distribution of returns.

The forgoing observations imply that the investment strategy of the manager can be summarized as follows. The manager invests in her alpha-generating technology with probability $(1 - q)$ and gambles with probability $q \in (0, 1)$. In the latter case, she chooses a gamble that pays off zero with probability $x \geq 0$ and, conditional on not yielding zero, admits an atomless density φ over $[0, \infty)$. The following lemma further characterizes this strategy.

LEMMA 4: *The gamble yields zero with a positive probability ($x > 0$). The density φ is single-peaked and has the support $[z_1, z_2]$, where $1 < z_1 < z_2$.*

PROOF OF LEMMA 4: See the Appendix. □

Lemma 4 shows that equilibria with risk shifting involve trading strategies that take on hidden tail risk. With some probability, the manager sells disaster insurance. She uses the insurance premium to generate a noisy excess return if the disaster does not occur. Upon observing such fake excess returns, investors still revise their views about the manager’s skill upwards because the manager also uses her alpha-generating technology with some probability.

The proof of Lemma 4 provides an analytical expression for φ . Here we consider a numerical example. We assume the following parameter values: $\sigma = 10\%$, $a = 2\sigma^2 = 2\%$, $\pi_0 = 40\%$, and $\gamma = 10\%$.

First, we check that, with these parameter values, an equilibrium without gambling does not exist. Suppose, by contradiction, that such an equilibrium does exist. In this case, after observing a return R , investors form a posterior belief about the manager’s skill:

$$\pi_1^R = \frac{\pi_0 \chi R^{a/\sigma^2}}{\pi_0 \chi R^{a/\sigma^2} + 1 - \pi_0}, \quad \chi = e^{-a(a-\sigma^2)/2\sigma^2}.$$

Figure 1, Panel (A), depicts π_1^R . It is increasing in R and initially convex and then concave.

If the manager does not gamble and invests in the efficient technology, her expected future reputation coincides with the current reputation, and is π_0 . The Internet Appendix offers a simple procedure to check whether the manager




Figure 1. Taking on hidden tail risk. Figure 1 illustrates an example considered in Section I.C. Parameter values are as follows: $\sigma = 10\%$, $a = 2\sigma^2 = 2\%$, $\pi_0 = 40\%$, and $\gamma = 10\%$. In Panel (A), the solid line shows a reputation of the manager at date 1 as a function of the realized gross return R when investors believe that the manager does not gamble. The optimal gamble delivers either zero or a gross return of 1.23 with a probability of 80%. In Panel (B), the solid line displays a reputation of the manager in the equilibrium in which investors rationally take into account that the manager gambles. In equilibrium, the manager gambles with probability $q = 5.6\%$. When she gambles, the manager loses everything with probability $x = 8.8\%$ and obtains a gross return between 1.02 and 1.49. Panel (C) shows the gamble’s density.

can enhance her expected reputation by resorting to gambling. This procedure consists of finding the straight line that has the smallest value at one among all the straight lines that are above the graph of π_1^R . The manager manipulates her reputation if and only if this straight line takes a value strictly larger than π_0 at $R = 1$. An optimal gamble has its support included in the set where this line coincides with π_1^R .

Panel (A) shows that, in our example, there is an optimal gamble that delivers either zero or a gross return of 1.23 with a probability of 80%. The expected reputation from such a gamble is 40.6% > 40%. The gains in expected reputation from gambling come at the cost of lower expected returns. If the manager gambles, the expected return is only one, while it is $\pi_0 e^a + (1 - \pi_0) = 1.008$ if she uses the alpha-generating technology. Because the cost is less than the reputation gain, an equilibrium without gambling is not sustainable.

Next, we solve for the manager’s equilibrium strategy with randomized gambling. Figure 1, Panel (B), shows the equilibrium posterior π_1^R . In contrast to that in Panel (A), the posterior now has a linear portion over [1.02, 1.49]. In equilibrium, the manager gambles with probability $q = 5.6\%$. When she gambles, the manager loses everything with probability $x = 8.8\%$ and obtains an excess return between 2% and 49% with a density shown in Figure 1, Panel (C). The range of these returns [1.02, 1.49] corresponds to the linear part of π_1^R in Figure 1, Panel (B). The unconditional probability that the manager goes bust is $q \times x = 0.5\%$. Thus, the manager trades an excess return with probability 99.5% with the risk of losing everything.

While binary gambles are optimal deviations in an equilibrium in which investors believe that the manager does not gamble, the equilibrium gambles are not binary because this would be detected by investors. The equilibrium posterior with gambling is flatter than the one without gambling in the region [1.02, 1.49], where the returns from gambling are realized. As a result, learning about managerial skill is slower in the equilibrium with gambling.

D. What Are the Costs of Risk Shifting?

The results of Sections I.B and I.C suggest that there are three types of costs associated with risk shifting induced by career concerns in delegated asset management. First, realized returns are riskier in equilibria with gambling. While in our setting investors are risk-neutral, it would be straightforward to extend our model to one in which investors are risk-averse. In this case, additional risk would be costly to them, unless one assumes that gambles are purely idiosyncratic and investors can diversify them away. From a financial stability perspective (which is beyond the scope of our model), the most worrying aspect of equilibrium risk shifting is that it thickens the left tail of returns, since the manager finances small and frequent positive excess returns with rare but devastating losses. When such incentives prevail within institutions that have legal access to public safety nets, or that are too big or too systemic to fail, the induced cost for taxpayers of the occurrence of such fat-tailed returns can be very high.

Second, the manager may give up pursuing alpha-generating strategies and invest instead in fairly priced portfolios, which have a risk profile better suited to the manipulation of her reputation. As a result, markets may not be as efficient as they would be if the manager tried to correct any mispricing.

Finally, gambling in equilibrium slows down the discovery of managerial skill because the manager applies her skill less often. This leads to a less efficient allocation of capital.

II. Long-Term Contracts and Risk Shifting

In the previous section, we show that risk shifting reduces the gains from matching “brains and resources” through delegated asset management, and generates significant tail risk for financial institutions. We also show that risk shifting is likely to occur in equilibrium if only one-period contracts are used.

This section introduces the possibility of long-term contracting. We perform two exercises. First, we consider the risk-taking incentives induced by performance fees with high-water marks, which are commonly used in the hedge fund industry. We show that this contract does not eliminate the risk-shifting incentives created by the flow-performance relationship. Second, we solve for optimal contracts that eliminate risk shifting even when conditions (15) are satisfied.

For simplicity, this section restricts the analysis to the case in which the alpha-generating technology is such that $\alpha = 1$ in (2). In this case, reputation affects the manager’s expected alpha but not the optimal fund size, which is constant and equal to one.

A. Performance Fees with High-Water Marks

The typical fee structure in the hedge fund industry usually includes a management fee and a performance fee. Management fees are a fixed fraction of

the fund's net asset value, ranging from 1% to 4% per annum, and are meant to cover the fund's operating costs. The performance fee is a fraction of the fund's profits over a given year, ranging from 15% to 50%. The most widespread combination is a 2% management fee and a 20% performance fee, the so-called "2-20" contract. The performance fee typically includes a high-water mark (HWM) provision, meaning that the fee applies only to the profits in excess of the previous fund maximum. The goal of this section is to study whether performance fees are conducive to risk shifting in our model.

Because the terms of the contract are usually fixed, in general the manager does not extract the maximum expected surplus but leaves some surplus to investors. This would not be possible with the one-sided commitment assumed thus far because the manager would enter into a new contract with higher fees if her reputation improved. Such renegotiation by hedge fund managers is not just a theoretical possibility. In practice, managers do adjust their fees in response to their performance. There is also ample anecdotal evidence that fund managers tend to close funds that have underperformed in order to reset the HWM.⁵ Nevertheless, to provide a meaningful analysis, we only assume in this subsection that both investors and the manager agree on the fees and can commit not to renegotiate them.

To understand the risk-taking incentives created by the performance fee, we first consider a situation in which there are no new inflows after the manager and investors sign the contract at date 0. Specifically, we assume that both the manager and investors can commit to the following contract: (i) at date 0, investors make an initial unit investment in the fund; (ii) the manager receives a performance fee k with an HWM equal to one and returns the fraction of the fund value in excess of one that she does not receive as a performance fee to investors; and (iii) there are no new inflows to the fund.

Under the above contract, if the manager uses the alpha-generating technology, the fund size $(q_t)_{t \geq 0}$ evolves according to

$$dq_t = q_t[(r + a\theta)dt + \sigma dB_t] - dH_t, \quad (19)$$

$$H_t = \int_0^t \mathbf{1}_{\{q_s=1\}} dq_s. \quad (20)$$

H_t is the fund value in excess of one that is redeemed to investors after the performance fee is paid. It is a nondecreasing adapted process that increases only when q_t exceeds the optimal fund size equal to one. At any date t , the continuation utility of the manager is

$$W_t = E_t \left[\int_t^\infty e^{-rs} kdH_{t+s} \right]. \quad (21)$$

⁵ See, for example, Gregory Zuckerman, "Andor Haunted by a Bad Bet," *Wall Street Journal*, July 15, 2004; Amey Stone, "Hedge Funds: Fees Down? Close Shop," *Bloomberg Business Week*, August 7, 2005; or "Hedge-Fund Closures: Quitting While They're Behind," *The Economist*, February 18, 2012.

The next proposition computes this utility W_t and shows that the manager never finds it worthwhile to gamble in this case.

PROPOSITION 5: *Suppose that the manager commits to the contract described in (i), (ii), and (iii) above. Then she never finds gambling worthwhile and her continuation utility is*

$$W_t = k(\pi_t \rho^{-1} q_t^\rho + (1 - \pi_t) q_t), \tag{22}$$

$$\rho = \frac{-(r + \alpha - \sigma^2/2) + \sqrt{(r + \alpha - \sigma^2/2)^2 + 2r\sigma^2}}{\sigma^2} < 1.$$

PROOF OF PROPOSITION 5: See the Appendix. □

Proposition 5 shows that, without new inflows, a performance fee with an HWM is not sufficiently convex to trigger gambling. A similar result is established by Panageas and Westerfield (2009). Although our results are closely related, our setting differs from that of Panageas and Westerfield (2009) along several interesting dimensions.

Panageas and Westerfield (2009) also consider a risk-neutral fund manager who is compensated according to (21). The manager cannot secretly gamble but she can take positions in a risk-free asset and in a risky asset that carries an exogenous constant excess return. Panageas and Westerfield (2009) find that, despite risk neutrality, the manager optimally maintains a finite and constant leverage. One way to understand their result is to notice that the performance fee with an HWM can be viewed as a continuum of call options with varying maturities. An increase in risk raises the value of the options with the closest maturity but also increases the possibility that more remote options will become far out-of-the money in the case of adverse realizations. The resulting trade-off yields a constant leverage that depends on both the manager’s discount rate and the excess return on the risky asset. When the latter goes to zero, Panageas and Westerfield (2009) show that the optimal leverage goes to infinity. This result crucially depends, however, on the assumption that the manager’s discount rate is strictly greater than the risk-free rate.

If the manager had a discount rate greater than r in our setting, she would find it optimal to gamble as her reputation deteriorates. In our model, as the manager’s reputation declines, the perceived expected return on the alpha-generating technology goes to zero. Therefore, at some point the manager would choose to forfeit investing in a technology that provides a small excess return but limited risk in favor of gambling with unrestricted risk-taking possibilities. This does not occur here only because the manager’s discount rate is equal to the risk-free rate. In this case, the manager never trades, even in the hope of gaining small expected excess returns in exchange for more risk.

Proposition 5 may suggest that the HWM contract does not lead to gambling. This conclusion, however, does not apply if fund inflows are considered. The case of $\alpha = 1$ is especially illustrative and simple. Provided that the reputation of the manager is high enough to cover her performance fees, investors will

always be willing to supply funds until the assets under management reach their optimal size equal to one. This drastically reduces the penalty that accrues from realizing negative returns and therefore from gambling.

To demonstrate this, suppose that the manager has an option to reopen her fund to new investors once. Let q^* be the fund size when she decides to reopen her fund. In this case, the manager raises $1 - q^*$ units of fresh capital with an HWM set to one. We assume that the new investors are the first to withdraw their capital if the total fund size exceeds the optimal fund size of one.⁶

Let q_t^o and q_t^n be the respective stakes in the fund of the old and new investors after the fund is reopened and before the stake of the old investors reaches one. Then, if the manager uses her alpha-generating technology, q_t^o and q_t^n evolve according to

$$dq_t^o = q_t^o [(r + a\theta)dt + \sigma dB_t], \tag{23}$$

$$d(q_t^o + q_t^n) = (q_t^o + q_t^n) [(r + a\theta)dt + \sigma dB_t] - dH_t, \tag{24}$$

$$H_t = \int_0^t \mathbf{1}_{(q_t^o + q_t^n = 1)} d(q_t^o + q_t^n). \tag{25}$$

When the stake of the old investors reaches one, they become the sole investors in the fund again, and their capital evolves according to (19) and (20). We denote this moment as τ_1 . Let τ_{q^*} denote the time when the funds of the original investors reach a level q^* and the manager reopens the fund to new inflows. The manager's expected fees from new investors are

$$W^n(q^*) = E_t \left[\int_{\tau_{q^*}}^{\tau_1} e^{-rs} k q_t^n dH_{t+s} \right], \tag{26}$$

where H_t is given by (25). The present value of these fees is

$$Ee^{-r\tau_{q^*}} W^n(q^*).$$

While $W^n(q^*)$ decreases in q^* and reaches maximum at $q^* = 0$, $Ee^{-r\tau_{q^*}}$ increases in q^* . Thus, in setting q^* that solves

$$\sup_{q^*} Ee^{-r\tau_{q^*}} W^n(q^*), \tag{27}$$

the manager faces a trade-off between the amount of fees she can generate from new investors and how quickly she can receive these fees.

Gambling affects the distribution of τ_{q^*} . In particular, if the manager takes on hidden tail risk and realizes a low return, the manager can exercise her option to raise new funds sooner. This greatly reduces the cost of gambling and therefore provides strong incentives to engage in risk shifting. Of course,

⁶ This assumption makes gambling less likely because it decreases the manager's expected fees from opening the fund to new investors, and thus makes gambling more costly.

the option to raise new funds is only of value to managers whose reputation is good enough that they are still able to attract new flows after a loss. Therefore, incentives to take on hidden tail risk are highest for managers with a good reputation. However, this result depends on the assumption made in this section that the manager does not renegotiate her contract as her reputation improves. In the next subsection, we show that, when the manager is paid the maximum surplus she can generate, it is a manager with a low reputation who has the highest incentives to gamble. This is because a manager who is perceived to be highly skilled has much to lose if her reputation becomes tarnished. Under the “2-20” contract, however, a good manager does not generally extract the maximum surplus, which makes gambling less costly. We thus have the following proposition.

PROPOSITION 6: *Suppose that the manager commits to the contract described in (i) and (ii). In addition, suppose that the manager has an option to reopen her fund to new investors who supply funds as long as they expect to break even. If $r > \sigma^2/10$, then there exists $\bar{\pi}$ such that for all $\pi \geq \bar{\pi}$, the manager with a reputation π gambles.*

PROOF OF PROPOSITION 6: See the Appendix. □

Proposition 6 shows that an option to receive more inflows, even if it occurs only once, is sufficient to create gambling incentives. This suggests that typical performance fees may be problematic in practice given that hedge funds are typically open-ended and do have the option to increase their capital.

B. Optimal Contract with Contingent Bonus Deferral

Under the strong, and in our view unrealistic, assumption that the manager can commit not to renegotiate a contract after her reputation has improved, very simple contracts can eliminate risk shifting. For example, fully insuring the manager by guaranteeing a fixed wage equal to $a\pi_0$ per period would eliminate any incentives to gamble. This is not feasible in the more realistic environment with one-sided commitment considered here—the manager can walk away from such an insurance contract as soon as her reputation improves.

When commitment is one-sided, contracts must be structured in such a way that the continuation utility of the manager at any date is at least as large as her outside options given her current reputation. This implies that a contract for a given initial π_0 cannot be determined in isolation. Instead, all contracts for all initial skill levels depend on each other through the channel of managerial outside options. In this section, we present contracts that eliminate risk shifting by exploiting the fact that the impact of gambling on investor beliefs vanishes in the long run.

Before constructing these contracts, we illustrate the way they work using the following simple contract. Suppose that investors and the manager sign the following contract at date 0. The investors commit to invest one unit in the

fund at each date between 0 and T . The manager does not receive any payment until date T , at which she receives a single payment w_T such that

$$w_T = \frac{a\pi_T}{r}(e^{rT} - 1). \tag{28}$$

Suppose, for now only, that after date T the manager no longer works. The date 0 present value of the payment is

$$E_0[e^{-rT} w_T] = \frac{a\pi_0}{r}(1 - e^{-rT}),$$

which is exactly the present value of the excess returns that investors expect to receive between 0 and T if the manager uses her alpha-generating technology throughout. The contract resembles a deferred bonus with clawback provisions, whereby the manager receives a single terminal payment that is contingent on her entire track record over the period $[0, T]$, as summarized by π_T .

Consider the manager's incentives to gamble at date 0 given such a contract. As in (13), if the manager decides to gamble, she chooses a distribution μ that solves

$$\begin{aligned} \Pi(\pi, T) \equiv \max_{\mu} \int_0^{\infty} E_0 \left[\frac{\pi_T R^{\frac{a}{\sigma^2}}}{1 - \pi_T + \pi_T R^{\frac{a}{\sigma^2}}} \mid \pi_0 = \pi \right] d\mu(R) \\ \text{s. t. } \int_0^{\infty} R d\mu(R) = 1. \end{aligned} \tag{29}$$

The manager will not be tempted to gamble at date 0 if and only if

$$\Pi(\pi_0, T) = \pi_0. \tag{30}$$

The following lemma establishes important properties of $\Pi(\pi, T)$.

LEMMA 5: For all $\pi \in (0, 1)$, the function $\Pi(\pi, T)$ decreases with respect to T and is equal to π for all T sufficiently large. Furthermore, there exists $\underline{\pi}$ such that, for all $\pi \geq \underline{\pi}$, $\Pi(\pi, 0) = \pi$.

PROOF OF LEMMA 5: See the Appendix. □

Because $\Pi(\pi, T)$ decreases with respect to T , date 0 gambling becomes less appealing as the payment date becomes more remote. This formalizes the insight that the impact of gambling on investors' beliefs diminishes over time. Lemma 5 implies in particular that we can define for all π

$$\tau(\pi) \equiv \inf \{ t \geq 0 : \Pi(\pi, t) = \pi \}. \tag{31}$$

If $T \geq \tau(\pi_0)$, then the manager does not gamble at date 0. As the reputation of the manager improves, incentives to gamble decrease. Lemma 5 shows that a manager with a sufficiently high reputation does not gamble even if the payment is made at date 0.

In light of the forgoing, we are now ready to define an optimal contract. There are two matters that the fixed-date single-payment contract outlined

above does not address. First, while setting the payment date at $\tau(\pi_0)$ deters date 0 gambling, there is no guarantee that the manager will not be tempted to gamble as the payment date approaches. Second, it is necessary to check that, if the manager enters into a new single-payment contract after the current one expires, this does not affect her incentives to gamble within the current contract.

We construct the optimal contract as follows. At date 0, the investors commit to pay $w_{T_0} = \frac{a\pi_{T_0}}{r}(e^{rT_0} - 1)$ at a random date T_0 that is defined as follows. Initially, T_0 is equal to $\tau(\pi_0)$. Then, for all $n \geq 1$, if $T_{(n-1)\Delta t} = 1$ the payment is made and the contract ends. Otherwise, the date is revised as $T_{n\Delta t} = \max\{T_{(n-1)\Delta t} - 1; \tau(\pi_{n\Delta t})\}$. After the payment of this current contract is made, a new similar contract is initiated at date T with initial reputation π_T . We have the following proposition.

PROPOSITION 7: *If the manager is compensated according to a sequence of single-payment contracts such that a contract that starts at date t promises the payment $\frac{a\pi_T}{r}(e^{rT} - 1)$ at a random date T as defined as above, she never gambles and extracts the maximum expected surplus.*

PROOF OF PROPOSITION 7: Investors can commit to this contract because, by construction, investors break even ex ante. We therefore only need to prove two remaining results: first, the manager’s continuation utility is higher than her outside option at any date, and second, the manager never finds it optimal to gamble.

With respect to the first point, we define the random payment date of the current contract as T . At date t , the manager expects to receive

$$E_t \left[e^{-r(T-t)} \left(w_T + \frac{a\pi_T}{r} \right) \right] = \frac{e^{rt}a\pi_t}{r} \geq \frac{a\pi_t}{r}.$$

Turning to the second point, we observe that, by construction, the random payment date T is such that the manager has no incentive to gamble in order to increase the expected payment from the current contract. It may still be the case that the manager finds it worthwhile to gamble to increase her expected payoff on the subsequent contracts that follow the terminal payment of the current one. Notice that, from Lemma 5, the case that is most conducive to gambling is that in which all the payments from subsequent contracts are made once and for all at date T . Even in this case, the expected payment is $\frac{a\pi_T}{r}$, which is proportional to $w_T = \frac{a\pi_T}{r}(e^{rT} - 1)$ and thus not conducive to gambling.

Notice that this contract can be implemented in practice with share grants to the manager instead of cash bonuses provided that (i) the manager is sufficiently senior or important within the firm that her decisions actually affect the share price, and (ii) the vesting schedule is stochastic, depending on the manager’s entire track record as described above.

It is interesting to compare this contract with the one that emerges from models of dynamic moral hazard such as, for example, DeMarzo and Sannikov

(2006) or He (2009). In these models, an agent may secretly divert cash flows instead of reporting her entire output to the principal. The optimal contract in this case consists, as in ours, of adjusting the present value of the future payments promised to the agent as her track record evolves.

However, the timing conditions of the future payments are rather different in the two settings. In our contract it is crucial that the investors can commit to not accelerating the timing of payments upon observing stellar performance. In moral hazard models, a stellar track record results in accelerated payments being made to the agent.

This difference arises because the contracts are driven by different economic considerations. In our model, a stream of fixed payments would completely eliminate any risk-shifting incentives. However, such fixed payments are not viable because investors are forced to adjust the manager's continuation value as her reputation improves. If not, a good track record would lead the manager to repudiate her current labor contract and sign a new one that reflects her new improved reputation. The necessity of adjusting the continuation value creates risk-shifting incentives that can be addressed by the investors' commitment to deferring payments.

In contrast, in moral hazard models, an adjustment of continuation values is meant to provide the agent with incentives to report the highest possible output instead of diverting some of it. Thus, the sooner the agent reports high income, the better.

C. Current Financial Reforms and Contingent Bonus Deferral Contract

There is a widely shared view that inappropriate compensation schemes within the financial services industry contributed to the financial crisis that erupted in 2008. According to "Principles for Sound Compensation Practices" published by the FSF and summarizing the outcome of multiple surveys of financial institutions, "over 80 percent of market participants believe that compensation practices played a role in promoting the accumulation of risks that led to the current crisis."

In the face of these compensation issues, public authorities around the world have issued guidelines for compensation reforms. These guidelines prominently feature the deferral of bonuses and the introduction of clawback mechanisms. For example, in the United States, "Guidance on Sound Incentive Compensation Policies," jointly issued by several authorities,⁷ lists four methods of making compensation more sensitive to risk, including "deferral of payment" with explicit mention of clawbacks, "longer performance periods," and "reduced sensitivity to short-term performance." The FSF similarly advises that "One way to align time horizons is to place a portion, and in some cases up to the entirety, of any given year's bonus grant, both cash and equity, into the equivalent of an escrow account. All or part of the grant is reversed if the firm as a whole

⁷ The authorities include Office of the Comptroller of the Currency, Federal Reserve, Federal Deposit Insurance Corporation, Treasury, and Office of Thrift Supervision.

performs poorly, or if the exposures the employee caused the firm to assume in the year for which the bonus was granted perform poorly (a clawback).”

The industry is gradually beginning to follow this guidance. The International Institute of Finance, surveying the practices of 37 financial institutions representing more than half of global banking activity, concludes that⁸ “The industry has begun to take steps to strengthen the link between delivery of deferred compensation and the continued performance of the individual. Over 40% of the firms surveyed include performance-based criteria in their deferred compensation schemes, although in a majority of cases this takes simply the form of a penalty for gross misconduct or large-scale unexpected losses. A number of firms have developed more sophisticated approaches that incorporate a final payout multiplier that adjusts compensation up or down based on current year or historical performance.”

Interestingly, the stochastic payment dates that we introduce in this section strongly resemble this introduction of bonus deferral and clawbacks. While our optimal contract follows the spirit of these reforms, we believe that it also raises some important caveats. Our model suggests that setting the bonus payment date once and for all is in general not optimal, even if the terminal bonus is adjusted for subsequent performance. Additionally, it is important to commit to adjusting the payment date as events unfold. For example, our analysis shows that a three-year deferral might induce excessive risk taking at the end of year 2 if the manager has not performed well and the bonus is due in one year.

III. Conclusion

Financial innovation has come under severe criticism after the crisis that erupted in 2007. Several papers document that a number of structured products appeared to be aimed mostly at exploiting investors’ weaknesses or ignorance.⁹ In this paper, we suggest that, by enlarging the set of financial instruments available to fund managers, financial innovation may exacerbate agency costs even when investors are sophisticated, provided that investors cannot perfectly monitor the positions of the fund managers.

We introduce a novel framework to study this agency problem between managers and investors. In this framework, a manager’s compensation depends on investors’ perception of the manager’s ability to generate excess returns above a fair compensation for risk (α). The manager can temporarily distort the perception of her α -generating skill by trading a rich menu of financial instruments.

⁸ In “Compensation in Financial Services, Industry Progress and the Agenda for Change,” March 2009. The International Institute of Finance is the leading global association of financial institutions.

⁹ For example, Coval, Jurek, and Stafford (2009a) and Coval, Jurek, and Stafford (2009b) argue that senior CDO tranches aimed at exploiting the misperception of correlation risk by rating agencies, and that of systematic exposure by investors. Henderson and Pearson (2011) consider a class of structured equity products offered at prices that are hard to reconcile with purchasers’ rationality.

The model delivers two main implications. First, in the absence of careful contracting, this friction may lead managers to take on hidden tail risk in order to distort their perceived skill temporarily. Second, while current compensation reforms based on the deferral and clawback of bonuses appear to be an appropriate tool for addressing this issue, in their current form these reforms miss the point that the payment dates must be dynamically adjusted as the cumulative performance of the manager evolves.

In general, there are ways to address risk shifting in delegated asset management other than through the compensation design. Possible solutions include increased transparency, or restrictions on the set of instruments that managers can trade. In this paper, we focus on solving the risk-shifting problem using compensation design alone. This makes the problem more challenging. Future research could combine optimal contracting with these additional means of addressing the risk-shifting problem. On the other hand, we assume that positions are always valued at a fair market price. Instruments that are more difficult to value such as illiquid securities or exotic derivative contracts are likely to provide fund managers additional risk-shifting incentives if trading losses can be concealed for some time. Future research could further explore these channels for engaging in risk transformation.

Initial submission: December 17, 2011; Final version received: December 18, 2014
 Editor: Campbell Harvey

Appendix: Proofs

PROOF OF LEMMA 1: As $\Delta t \rightarrow 0$, the optimal fund size and maximal expected surplus solve

$$\lim_{\Delta t \rightarrow 0} q(\pi) = \arg \max_q q(\pi a - \beta q^{\frac{1}{\alpha-1}}),$$

$$\lim_{\Delta t \rightarrow 0} v(\pi) = \max_q q(\pi a - \beta q^{\frac{1}{\alpha-1}}).$$

Direct computations show that

$$\lim_{\Delta t \rightarrow 0} q(\pi) = \left(\frac{\alpha - 1}{\alpha \beta} a \pi \right)^{\alpha-1},$$

$$\lim_{\Delta t \rightarrow 0} v(\pi) = \beta^{1-\alpha} \frac{(\alpha - 1)^{\alpha-1}}{\alpha^\alpha} (a \pi)^\alpha. \quad \square$$

PROOF OF PROPOSITION 1: We first show point-wise convergence. That is, we establish (6) for a fixed $\pi_0 = \pi$. By Bayes's theorem, $\pi_{n\Delta t}$, the perceived skill at date $n\Delta t$ satisfies

$$\pi_{n\Delta t} = \frac{\pi_0 \varphi_{n\Delta t}}{1 - \pi_0 + \pi_0 \varphi_{n\Delta t}}, \tag{A1}$$

where

$$\varphi_{n\Delta t} = \exp \left\{ \frac{a}{\sigma^2} \left(a \left(\theta - \frac{1}{2} \right) n\Delta t + \sigma B_{n\Delta t} \right) \right\} \tag{A2}$$

is the likelihood ratio process. Let us introduce the continuous-time process $(\pi_t)_{t \geq 0}$ that obeys

$$d\pi_t = \frac{a}{\sigma} \pi_t (1 - \pi_t) d\bar{B}_t, \quad \pi_0 = \pi,$$

where $\bar{B}_t = \frac{1}{\sigma}(\theta at + \sigma B_t - a \int_0^t \pi_s ds)$. Then $(\bar{B}_t)_{t \geq 0}$ is a standard Wiener process under the agents' filtration (see Liptser and Shiryaev (1978)). Further, as $\Delta t \rightarrow 0$ and $n\Delta t \rightarrow t$, $\pi_{n\Delta t} \rightarrow \pi_t$ a.s. (see Liptser and Shiryaev (1978)). Hence, $V(\pi)$ can be written as

$$\begin{aligned} V(\pi) &= E_0 \int_0^\infty e^{-rt} v(\pi_t) dt, \\ \text{s.t. } d\pi_t &= \frac{a}{\sigma} \pi_t (1 - \pi_t) d\bar{B}_t, \quad \pi_0 = \pi. \end{aligned} \tag{A3}$$

By the Feynman–Kac formula, the function V solves the following linear second-order differential equation:

$$\frac{a^2}{2\sigma^2} \pi^2 (1 - \pi)^2 V''(\pi) - rV(\pi) + v(\pi) = 0. \tag{A4}$$

From (A3) it follows that

$$V(0) = v(0)/r, \quad V(1) = v(1)/r. \tag{A5}$$

The corresponding homogeneous equation

$$\frac{a^2}{2\sigma^2} \pi^2 (1 - \pi)^2 V''(\pi) - rV(\pi) = 0 \tag{A6}$$

has two regular singular points at zero and one. All solutions of the homogeneous equation are linear combinations of the two independent solutions

$$\begin{aligned} g(\pi) &= (1 - \pi)^{\frac{1}{2} + \frac{1}{2}\psi} \pi^{\frac{1}{2} - \frac{1}{2}\psi}, \quad \psi = \sqrt{1 + 8r\sigma^2/a^2}, \\ h(\pi) &= g(1 - \pi). \end{aligned}$$

From here, formulas (6) and (7) are standard results in the theory of nonhomogeneous differential equations. The function G is the Dirichlet–Green function for the differential operator associated with the homogeneous differential equation (see, for example, Driver (2003)).

We now show that $V(\pi, \Delta t)$ converges to $V(\pi)$ uniformly in π as $\Delta t \rightarrow 0$. We have

$$V(\pi) - V(\pi, \Delta t) = E \left[\sum_{n=0}^{\infty} \int_{n\Delta t}^{(n+1)\Delta t} (e^{-rt} v(\pi_t) - e^{-rn\Delta t} v(\pi_{n\Delta t})) dt \mid \pi_0 = \pi \right].$$

Thus, it is enough to show that $\forall \varepsilon > 0, \exists \bar{\Delta t}$ such that $\forall \Delta t < \bar{\Delta t}$ and $\forall \pi \in [0, 1]$

$$\sup_{s \leq \Delta t} \sup_{\pi \in [0,1]} |E(v(\pi_s) - v(\pi))| < \varepsilon. \tag{A7}$$

By a change of variables, (A7) can be written as

$$\sup_{s \leq \Delta t} \sup_{\pi \in [0,1]} |E(\widehat{v}(\pi, B_s, s) - \widehat{v}(\pi, 0, 0))| < \varepsilon,$$

where

$$\widehat{v}(\pi, x, t) = v \left(\frac{\pi \exp \left\{ \frac{a}{\sigma^2} \left(a \left(\theta - \frac{1}{2} \right) t + \sigma x \right) \right\}}{1 - \pi + \pi \exp \left\{ \frac{a}{\sigma^2} \left(a \left(\theta - \frac{1}{2} \right) t + \sigma x \right) \right\}} \right). \tag{A8}$$

Because v is uniformly continuous over $[0, 1]$, it is enough to show that $\forall \pi \in [0, 1]$ and $\forall \varepsilon > 0, \exists \bar{\Delta t}$ such that $\forall \Delta t < \bar{\Delta t}$

$$\sup_{s \leq \Delta t} |E(\widehat{v}(\pi, B_s, s) - \widehat{v}(\pi, 0, 0))| < \varepsilon. \tag{A9}$$

This follows from the weak convergence of the measures induced by B_s to the measure concentrated at zero as $s \rightarrow 0$. □

PROOF OF PROPOSITION 2: Suppose that investors believe that the manager invests in her private technology. In this case, if the manager does invest in her private technology, then π_t evolves according to (A3). If, on the other hand, she invests in the risk-free asset, π_t evolves according to

$$d\pi = \frac{a}{\sigma} \pi(1 - \pi) \left(r' - r + \frac{\sigma^2}{2} - (\pi a - c(q_t)) \right) dt.$$

Suppose that at time t the manager allocates percentage x_t of her funds to her alpha-generating technology and invests the rest in the risk-free asset. Then her continuation utility is

$$V(\pi, x) = E_0 \int_0^{\infty} e^{-rt} v(\pi_t) dt, \tag{A10}$$

s.t. $d\pi_t = \frac{a}{\sigma} \pi_t(1 - \pi_t) \left((1 - x_t) \left(r' - r + \frac{\sigma^2}{2} - (\pi a - c(q_t)) \right) + x_t d\bar{B}_t \right), \pi_0 = \pi.$

The optimal investment policy x_t that maximizes (A10) satisfies the HJB equation:

$$\sup_{x \in [0,1]} x^2 \psi_2^2 V'' + (1-x)\psi_1 V' - rV + v = 0, \quad (\text{A11})$$

where

$$\begin{aligned} \psi_1 &= \frac{\alpha}{\sigma} \pi_t (1 - \pi_t) \left(r' - r + \frac{\sigma^2}{2} - (\pi\alpha - c(q_t)) \right) < 0, \\ \psi_2 &= \frac{\alpha}{\sigma} \pi_t (1 - \pi_t) > 0. \end{aligned}$$

If $x_t \equiv 1$, then π_t is a martingale and by Jensen's inequality ($v(\pi) = \pi^\alpha$)

$$Ev(\pi_t) \geq v(\pi_0).$$

Therefore, at the optimal investment policy x_t , $rV(\pi, x) \geq v(\pi)$. Thus, (A11) implies that the optimal policy is indeed $x_t \equiv 1$. \square

PROOF OF PROPOSITION 3: Suppose that the manager gambles and realizes return R over $[0, \Delta t]$, and from then on invests in her alpha-generating technology. Let $(\pi_{n\Delta t}^R)_{n \in \mathbb{N}}$ denote the process under the *manager's filtration* of her skill as perceived by investors who believe instead that she has invested in her storage technology at date 0. These investors believe that

$$R = e^{\left(r + \theta\alpha - c(q_0) - \frac{\sigma^2}{2}\right)\Delta t + \sigma B_{\Delta t}}.$$

From (A1) and (A2) it follows that

$$\begin{aligned} \pi_{\Delta t}^R &= \frac{\pi_0 R^{\frac{\alpha}{\sigma^2}} e^{\frac{\alpha}{\sigma^2} \left(\frac{\sigma^2}{2} + c(q_0) - r - \frac{\alpha}{2}\right)\Delta t}}{1 - \pi_0 + \pi_0 R^{\frac{\alpha}{\sigma^2}} e^{\frac{\alpha}{\sigma^2} \left(\frac{\sigma^2}{2} + c(q_0) - r - \frac{\alpha}{2}\right)\Delta t}}, \\ \forall n \geq 0, \quad \pi_{(n+1)\Delta t}^R &= \frac{\pi_{\Delta t}^R \frac{\varphi_{(n+1)\Delta t}}{\varphi_{\Delta t}}}{1 - \pi_{\Delta t}^R + \pi_{\Delta t}^R \frac{\varphi_{(n+1)\Delta t}}{\varphi_{\Delta t}}}. \end{aligned} \quad (\text{A12})$$

As $\Delta t \rightarrow 0$,

$$\lim_{\Delta t \rightarrow 0} \pi_{\Delta t}^R = \frac{\pi_0 R^{\frac{\alpha}{\sigma^2}}}{1 - \pi_0 + \pi_0 R^{\frac{\alpha}{\sigma^2}}},$$

and

$$\lim_{\Delta t \rightarrow 0} \frac{\varphi_{(n+1)\Delta t}}{\varphi_{\Delta t}} = \frac{\pi_t}{1 - \pi_t} \frac{1 - \pi_0}{\pi_0}.$$

Therefore,

$$\lim_{\Delta t \rightarrow 0, n\Delta t \rightarrow t} \pi_{n\Delta t}^R = \frac{\pi_t R^{\frac{\alpha}{\sigma^2}}}{1 - \pi_t + \pi_t R^{\frac{\alpha}{\sigma^2}}}. \quad (\text{A13})$$

\square

PROOF OF PROPOSITION 4: *Part 1.* We first show that, if $\sigma^2 > \alpha a$, then the manager does not engage in risk shifting. Proposition 2 demonstrates that the manager will not invest in the alternative technology at the risk-free rate provided that $r' < r - \sigma^2/2$. We now show that she will not invest in any risky gamble as well.

Suppose the manager believes that she is skilled with probability π . At the same time, suppose that investors believe that the manager is skilled with probability π' and that the manager never engages in risk shifting. We show that the manager has no incentives to deviate by taking a one-shot risky gamble in this case.

Suppose the manager takes a gamble and realizes return R . Let $W(\pi, \pi', R, \Delta t)$ be the expected utility of the manager conditional on realizing a first-period return of R . Similar to the proof of Proposition 3, one can show that investors' perception of the manager's skill π_t^R is

$$\pi_t^R = \frac{\pi_t R^{\frac{\alpha}{\sigma^2}}}{(1 - \pi_t) \frac{(1-\pi)\pi}{(1-\pi)\pi'} + \pi_t R^{\frac{\alpha}{\sigma^2}}}. \tag{A14}$$

Proposition 1 implies that

$$\lim_{\Delta t \rightarrow 0} W(\pi, \pi', R, \Delta t) = \int_0^1 G(\pi, x) v \left(\frac{x R^{\frac{\alpha}{\sigma^2}}}{(1-x) \frac{(1-\pi)\pi}{(1-\pi)\pi'} + x R^{\frac{\alpha}{\sigma^2}}} \right) dx. \tag{A15}$$

Differentiating twice with respect to R shows that this function is concave in R when $\sigma^2 \geq \alpha a$. Hence, the manager has no incentives to take a one-shot risky gamble in this case. Because this holds for arbitrary heterogeneous priors π, π' , this implies that multi-period deviations cannot be desirable by backward induction.

Part 2. We now show that, if $\sigma^2 < \alpha a$ and $r\sigma^2 > \frac{\alpha^2}{2}(\alpha - 1)$, then for $\pi_0 = \pi$ small enough there exists a one-period gamble that makes the manager better off. Let $R = (1 - \rho)^{-\sigma^2/\alpha}$, $\rho \in [0, 1)$. Consider the following gamble:

$$\begin{cases} R & \text{Prob. } 1/R \\ 0 & \text{Prob. } 1 - 1/R. \end{cases}$$

From (A15) the expected net gain from the above one-period gamble over perpetual investment in the efficient storage technology is

$$\int_0^1 G(\pi, x) x^\alpha u(x, \rho) dx, \tag{A16}$$

where

$$u(x, \rho) = \frac{(1 - \rho)^{\sigma^2/\alpha}}{(1 - \rho(1 - x))^\alpha} - 1. \tag{A17}$$

Since $\sigma^2 < \alpha a$, there exists \bar{x} and some $\hat{\rho} \in (0, 1)$ such that for all $x \in [0, \bar{x}]$, $\frac{(1-\hat{\rho})^{\sigma^2/a}}{(1-\hat{\rho}(1-x))^\alpha} > 1 + \varepsilon$ for some $\varepsilon > 0$ and therefore $u(x, \hat{\rho}) > \varepsilon > 0$. Thus, for π small enough

$$\int_0^1 G(\pi, x) x^\alpha u(x, \hat{\rho}) dx > \int_\pi^1 G(\pi, x) x^\alpha u(x, \hat{\rho}) dx.$$

Using (7) we have

$$\int_\pi^1 G(\pi, x) x^\alpha u(x, \hat{\rho}) dx = \frac{2\sigma^2}{\psi \alpha^2} g(\pi) \int_\pi^1 x^{\alpha-\frac{3}{2}-\frac{1}{2}\psi} (1-x)^{-\frac{3}{2}+\frac{1}{2}\psi} u(x, \hat{\rho}) dx. \tag{A18}$$

Notice that $r\sigma^2 > \frac{\alpha^2}{2} \alpha (\alpha - 1)$ implies that $\psi > 2\alpha - 1$. Therefore, the integral

$$\int_\pi^1 x^{\alpha-\frac{3}{2}-\frac{1}{2}\psi} (1-x)^{-\frac{3}{2}+\frac{1}{2}\psi} u(x, \hat{\rho}) dx$$

diverges as $\pi \rightarrow 0$. In this case, the sign is determined by the sign of $u(\cdot, \hat{\rho})$ in the neighborhood of zero, which is positive. Thus, the net gain from the gamble is positive. \square

PROOF OF LEMMA 3: Suppose that investors believe that the manager does not gamble. Then their posterior probability about her ability upon observing R is

$$\pi_1(R) = \frac{\pi_0 \chi R^{\alpha/\sigma^2}}{\pi_0 \chi R^{\alpha/\sigma^2} + 1 - \pi_0}, \quad \chi = e^{-a(\alpha-\sigma^2)/2\sigma^2}.$$

If the manager does not gamble, she expects to get

$$\gamma(\pi_0 e^a + 1 - \pi_0) + \pi_0. \tag{A19}$$

Consider π_0 such that

$$4\pi_0(1 - \pi_0) < \chi(1 + \gamma(e^a - 1))^{-2}. \tag{A20}$$

Suppose that the manager deviates and enters into a gamble that generates gross return $\sqrt{\frac{1-\pi_0}{\pi_0 \chi}}$ with probability $1/\sqrt{\frac{1-\pi_0}{\pi_0 \chi}}$ and zero otherwise. Direct computations show that, in this case, the manager expects to get

$$\gamma + \frac{1}{2} \sqrt{\frac{\pi_0 \chi}{1 - \pi_0}}. \tag{A21}$$

Comparing (A21) with (A19) one can see that the manager obtains a higher utility if she gambles, provided that (A20) holds. \square

PROOF OF LEMMA 4: We solve for an equilibrium in which the manager invests in her private technology with probability $1 - q$ and gambles otherwise. If the manager gambles, she receives zero with probability x and a return between

z and $z + dz$ with probability $(1 - x)\varphi(z)dz$. For the ease of computations, we write $\varphi(z)$ as

$$\varphi(z) = \frac{1}{z\sqrt{2\pi\sigma^2}} e^{-\ln^2(z)/2\sigma^2} g(z), \quad (\text{A22})$$

where $g(z) \geq 0$. In such an equilibrium, upon observing a realized return z , investors' posterior is¹⁰

$$\pi_1(z) = \pi_0 \frac{(1 - q)\chi z^{a/\sigma^2} + q(1 - x)g(z)}{(1 - q)(\pi_0 \chi z^{a/\sigma^2} + 1 - \pi_0) + q(1 - x)g(z)}, \quad \chi = e^{-a(a - \sigma^2)/2\sigma^2}. \quad (\text{A23})$$

Let $E_\mu(E_0)$ be the expectation operator if the manager gambles (invests in her alpha-generating technology). When the manager gambles she takes investors' posterior π_1 in (A23) as given and chooses a gamble that solves

$$\begin{aligned} & \sup E_\mu \pi_1(z) \\ & \text{s. t. } E_\mu z = 1. \end{aligned}$$

We show in the Internet Appendix that the solution to this problem coincides with the solution to its dual problem:

$$\begin{aligned} & \inf_{(A, B) \in \mathbb{R}^2} A + B \\ & \text{s. t. } \forall z \geq 0, A + Bz \geq \pi_1(z), \end{aligned} \quad (\text{A24})$$

where A and B are some real numbers.

The dual problem admits a simple and practical interpretation: it minimizes the value at one of a straight line that is above the graph of $\pi_1(z)$. The optimal gamble then has its support included in the set where this line coincides with $\pi_1(z)$. We also show in the Internet Appendix that this solution can be realized with a binary gamble. In this particular case, a binary gamble cannot be the equilibrium solution, however, because it would be detected. Thus, it must be that $\pi_1(z)$ has linear portions that coincide with the minimal straight line $A + Bz$.

Given that $\pi_1(z)$ is first convex and then concave, it is easy to see that there must be two real numbers z_1 and z_2 such that $\pi_1(z) = Bz$ for $z \in 0 \cup [z_1, z_2]$ and $\pi_1(z) < Bz$ otherwise. Figure 1, Panel (B), illustrates this for a particular numerical example. The interval $[z_1, z_2]$ is the support of the gamble when it does not yield zero.

We are now ready to solve for the manager's gambling strategy. We first determine B . Direct computations show

$$qE_\mu \pi_1(z) + (1 - q)E_0 \pi_1(z) = \pi_0(1 - qx). \quad (\text{A25})$$

¹⁰ We continue to assume that $\pi_1(z)$ is continuous at zero.

Because the manager randomizes between gambling and investing in her alpha-generating technology, it must be that she is indifferent between gambling or investing efficiently:

$$E_\mu(\gamma z + \pi_1(z)) = E_0(\gamma z + \pi_1(z)),$$

or

$$E_\mu \pi_1(z) = E_0 \pi_1(z) + \gamma \pi_0 (e^\alpha - 1). \tag{A26}$$

Expressions (A25) and (A26) together imply that

$$B = E_\mu \pi_1(z) = (1 - qx)\pi_0 + \gamma(1 - q)\pi_0(e^\alpha - 1). \tag{A27}$$

Since $\pi_1(z) = Bz$ for $z \in 0 \cup [z_1, z_2]$, we can solve for $g(z)$ from (A23):

$$g(z) = \begin{cases} \frac{(1-q)\pi_0\chi z^{a/\sigma^2} (Bz + B(1-\pi_0)(\pi_0\chi)^{-1}z^{1-a/\sigma^2} - 1)}{q(1-x)(\pi_0 - Bz)} & \text{if } z \in [z_1, z_2], \\ 0 & \text{otherwise.} \end{cases} \tag{A28}$$

Continuity of $\pi_1(z)$ implies that it must be the case that

$$\frac{\pi_0\chi z_i^{a/\sigma^2}}{\pi_0\chi z_i^{a/\sigma^2} + 1 - \pi_0} = Bz_i, \quad i = 1, 2,$$

or

$$1 = Bz_i + B(1 - \pi_0)(\pi_0\chi)^{-1}z_i^{1-a/\sigma^2}, \quad i = 1, 2. \tag{A29}$$

Notice that, because $a > \sigma^2$, equation (A29) can have at most two solutions.

We are left with two free parameters: q and x . The parameters are determined in equilibrium from requiring that (i) $\varphi(z)$ be a density and (ii) the gamble have expected return equal to one:

$$\int_0^\infty \varphi(z) dz = 1, \tag{A30}$$

$$(1 - x) \int_0^\infty z\varphi(z) dz = 1. \tag{A31}$$

□

PROOF OF PROPOSITION 5: First, suppose that the manager is skilled so that the excess return on the alpha-generating technology is a . Whenever $q_t < 1$, the manager’s continuation utility is a martingale and therefore solves

$$0 = -rW + W_q q_t (r + a) + \frac{1}{2} W_{qq} \sigma^2 q_t^2. \tag{A32}$$

□

The ODE (A32) has a general solution of the form

$$W(q_t) = C_+ q_t^{\rho^+} + C_- q_t^{\rho^-},$$

where the constants ρ^+ and ρ^- solve the quadratic equation

$$-r + \rho \left(r + a - \frac{1}{2}\sigma^2 \right) + \frac{1}{2}\sigma^2\rho^2 = 0. \tag{A33}$$

Solving (A33), we have

$$\rho^\pm = \frac{-(r + a - \frac{1}{2}\sigma^2) \pm \sqrt{(r + a - \frac{1}{2}\sigma^2)^2 + 2r\sigma^2}}{\sigma^2}. \tag{A34}$$

There are two boundary conditions (at $q = 0$ and at $q = 1$):

$$\begin{aligned} \lim_{q_t \rightarrow 0} W(q_t) &= 0, \\ W'(1) &= k. \end{aligned}$$

Thus,

$$W(q_t) = k\rho^{-1}q_t^\rho, \quad \rho = \frac{-(r + a - \frac{1}{2}\sigma^2) + \sqrt{(r + a - \frac{1}{2}\sigma^2)^2 + 2r\sigma^2}}{\sigma^2}. \tag{A35}$$

If the manager is unskilled, then $a = 0$ and $\rho = 1$. Thus, if the manager is skilled with probability π_t , then her continuation utility is

$$W(q_t) = k(\pi_t\rho^{-1}q_t^\rho + (1 - \pi_t)q_t). \tag{A36}$$

PROOF OF PROPOSITION 6: First, we compute the expected surplus that the manager expects to get from new investors if she reopens her fund when its size is q^* . As in Proposition 5, we first consider the case in which the manager is skilled with probability one. Let $x_t = \ln(q_t^\rho/q^*)$. Then

$$dx_t = \left(r + a - \frac{1}{2}\sigma^2 \right) dt + \sigma dB_t, \quad x_0 = 0. \tag{A37}$$

Let m_t be the running maximum of x_t , $m_t = \max_{0 \leq s \leq t} x_s$. The manager is paid whenever $x_t = m_t$. Whenever $x_t < m_t$, the manager's continuation utility follows a martingale and therefore solves

$$0 = -rW^n + W_x^n \left(r + a - \frac{1}{2}\sigma^2 \right) + \frac{1}{2}W_{xx}^n\sigma^2. \tag{A38}$$

The ODE (A38) has a general solution of the form

$$W^n(x_t, m_t) = f(m_t)e^{\rho^+x_t} + g(m_t)e^{\rho^-x_t},$$

where the constants ρ^+ and ρ^- solve the quadratic equation (A33), and therefore are as in (A34). The functions $f(m_t)$ and $g(m_t)$ are arbitrary functions of m_t . Because $\lim_{x_t \rightarrow -\infty} W^n(x_t) = 0$ it must be the case that $g(m_t) \equiv 0$. Without loss of generality, the solution is

$$W^n(x_t, m_t) = e^{\rho(x_t - m_t)} f(m_t), \quad \rho = \rho^+.$$

The boundary condition at $x_t = m_t$ is

$$f'(m_t) - \rho f(m_t) + k(1 - q^* e^{m_t}) = 0. \tag{A39}$$

A general solution to (A39) is

$$f(m_t) = k \left(\frac{q e^{m_t}}{1 - \rho} + \frac{1}{\rho} \right) + C e^{\rho m_t}, \tag{A40}$$

where C is a constant. When q_t^o reaches one for the first time, the stake of the new investors in the fund becomes zero. This implies that $f(-\ln(q^*)) = 0$, and therefore

$$W^n(x_t^o, h_t) = k \left(\frac{x_t^o}{h_t} \right)^\rho \left(\frac{1 - h_t^\rho}{\rho} + \frac{h_t - h_t^\rho}{1 - \rho} \right), \tag{A41}$$

where $h_t = q^* e^{m_t}$. Thus,

$$W^n(q^*) \equiv W^n(q^*, q^*) = k \left(\frac{1 - (q^*)^\rho}{\rho} + \frac{q^* - (q^*)^\rho}{1 - \rho} \right). \tag{A42}$$

Cox and Miller (1965) show that

$$E e^{-r v_{q^*}} = \frac{\rho - \rho^-}{\rho (q^*)^{\rho^-} - \rho^- (q^*)^\rho}. \tag{A43}$$

Thus, the present value of the opportunity to reopen the fund at date 0 if the manager uses only her alpha-generating technology is

$$V = \sup_{q^*} k \frac{\rho - \rho^-}{\rho (q^*)^{\rho^-} - \rho^- (q^*)^\rho} \left(\frac{1 - (q^*)^\rho}{\rho} + \frac{q^* - (q^*)^\rho}{1 - \rho} \right). \tag{A44}$$

Suppose at time 0 the manager takes a gamble that delivers a gross return $q^* < 1$ with probability p and a gross return $1 + p(1 - q^*)/(1 - p)$ with probability $(1 - p)$. Suppose also that, if a return q^* is realized, then the manager reopens the fund. Then the expected continuation utility of the manager is

$$k \left[p \left(\frac{q^*}{\rho} + W^n(q^*) \right) + (1 - p) \left(p(1 - q^*)/(1 - p) + \frac{1}{\rho} + V \right) \right]. \tag{A45}$$

The manager will gamble if and only if the above utility is greater than her expected utility if she does not gamble. Thus, she will gamble if and only if

$$\sup_{q^*} W^n(q^*) - \frac{(1 - q^*)(1 - \rho)}{\rho} > V.$$

Direct computations show that

$$\sup_{q^*} W^n(q^*) - \frac{(1 - q^*)(1 - \rho)}{\rho} = 1.$$

The supremum is achieved at $q^* = 0$. So the skilled manager will gamble if $V < 1$. By looking at ρ and ρ^- one can see that they depend only on r/σ^2 and a/σ^2 . Direct but tedious calculations show that $V < 1$ as long as $r/\sigma^2 > 0.1$. \square

PROOF OF LEMMA 5: The integrand in (29) is initially convex and then concave in R . We show in the Internet Appendix that this implies that the optimal gamble is a simple binary gamble such that, for some $R > 1$, the manager earns R with probability $1/R$ and loses everything otherwise. Therefore,

$$\Pi(\pi, t) \equiv \sup_{R \geq 1} R^{\lambda-1} E_0[h(\pi_t, R) | \pi_0 = \pi],$$

where

$$h(\pi, R) = \frac{\pi}{1 - \pi + \pi R^\lambda}, \quad \lambda = \frac{a}{\sigma^2}. \tag{A46}$$

For any $R \geq 1$, $h(\pi, R)$ is a concave function of π , which implies that $h(\pi_t, R)$ is a supermartingale. Therefore, $\Pi(\pi, t)$ is a decreasing function of t . As $t \rightarrow \infty$, $E_0[h(\pi_t, R) | \pi_0 = \pi] \rightarrow \pi R^{-\lambda}$. Therefore, for large t , $\Pi(\pi, t) \leq \pi R^{-1}$, and the optimal choice of R is one, that is not to gamble. Direct computations show that, if $\pi > \frac{\lambda-1}{\lambda}$, then

$$\pi \leq \Pi(\pi, t) \leq \sup_{R \geq 1} R^{\lambda-1} h(\pi, R) = \pi. \tag{A47}$$

\square

LEMMA 6: Suppose conditions (15) hold. Then there exists a finite limit

$$0 < \lim_{\pi \rightarrow 0} V(\pi) \pi^{-\alpha} < \infty. \tag{A47}$$

PROOF OF LEMMA 6: We have

$$\begin{aligned} V(\pi) &= \int_0^1 G(\pi, x) x^\alpha dx \\ &= \frac{2\sigma^2}{a^2 \psi} \left[(1 - \pi)^{\frac{1}{2}(1+\psi)} \pi^{\frac{1}{2}(1-\psi)} \int_0^\pi (1 - x)^{-\frac{1}{2}(\psi+3)} x^{\alpha+\frac{1}{2}(\psi-3)} dx \right. \\ &\quad \left. + (1 - \pi)^{\frac{1}{2}(1-\psi)} \pi^{\frac{1}{2}(1+\psi)} \int_\pi^1 (1 - x)^{\frac{1}{2}(\psi-3)} x^{\alpha-\frac{1}{2}(3+\psi)} dx \right], \end{aligned}$$

where ψ is defined in (8). Further, $r\sigma^2 > \frac{a^2}{2} \alpha (\alpha - 1)$ implies that $\psi > 2\alpha - 1$. Therefore,

$$\exists \lim_{\pi \rightarrow 0} (1 - \pi)^{\frac{1}{2}(1-\psi)} \pi^{\frac{1}{2}(1+\psi)} \int_\pi^1 (1 - x)^{\frac{1}{2}(\psi-3)} x^{\alpha-\frac{1}{2}(3+\psi)} dx \times \pi^{-\alpha} < \infty,$$

and

$$\exists \lim_{\pi \rightarrow 0} (1 - \pi)^{\frac{1}{2}(1+\psi)} \pi^{\frac{1}{2}(1-\psi)} \int_0^\pi (1 - x)^{-\frac{1}{2}(\psi+3)} x^{\alpha+\frac{1}{2}(\psi-3)} dx \times \pi^{-\alpha} < \infty. \tag{A47}$$

\square

REFERENCES

- Acharya, Viral, Thomas Cooley, Matthew Richardson, and Ingo Waler, 2010, Manufacturing tail risk: A perspective on the financial crisis of 2007-09, *Foundations and Trends in Finance* 4, 247–325.
- Acharya, Viral, Marco Pagano, and Paolo Volpin, 2013, Seeking alpha: Excess risk taking and competition for managerial talent, Working paper, New York University.
- Basak, Suleyman, Anna Pavlova, and Alex Shapiro, 2007, Optimal asset allocation and risk shifting in money management, *Review of Financial Studies* 20, 1583–1621.
- Berk, Jonathan, and Richard Green, 2004, Mutual fund flows and performance in rational markets, *Journal of Political Economy* 112, 1269–1295.
- Berk, Jonathan, and Richard Stanton, 2007, Managerial ability, compensation, and the closed-end fund discount, *Journal of Finance* 62, 529–556.
- Brown, Keith, W. Van Harlow, and Laura Starks, 1996, Of tournaments and temptations: An analysis of managerial incentives in the mutual fund industry, *Journal of Finance* 51, 85–110.
- Carpenter, Jennifer, 2000, Does option compensation increase managerial risk appetite?, *Journal of Finance* 55, 2311–2331.
- Chevalier, Judith, and Glenn Ellison, 1997, Risk taking by mutual funds as a response to incentives, *Journal of Political Economy*, 105, 1167–1200.
- Coval, Joshua, Jakub Jurek, and Erik Stafford, 2009a, The economics of structured finance, *Journal of Economic Perspectives* 23, 3–25.
- Coval, Joshua, Jakub Jurek, and Erik Stafford, 2009b, Economic catastrophe bonds, *American Economic Review* 99, 628–666.
- Cox, David, and Michael Miller, 1965, *The Theory of Stochastic Processes* (Chapman and Hall, London).
- Dangl, Thomas, Youchang Wu, and Josef Zechner, 2008, Market discipline and internal governance in the mutual fund industry, *Review of Financial Studies* 21, 2307–2343.
- Dasgupta, Amil, and Andrea Prat, 2008, Information aggregation in financial markets with career concerns, *Journal of Economic Theory* 143, 83–113.
- DeMarzo, Peter, and Yuliy Sannikov, 2006, Optimal security design and dynamic capital structure in a continuous-time agency model, *Journal of Finance* 61, 2681–2724.
- Driver, Bruce, 2003, Analysis tools with applications, Lecture notes, University of California San Diego.
- Goetzmann, William, Jonathan Ingersoll Matthew Spiegel, and Ivo Welch, 2007, Portfolio performance manipulation and manipulation-proof performance measures, *Review of Financial Studies* 20, 1503–1546.
- He, Zhiguo, 2009, Optimal executive compensation when firm size follows geometric Brownian motion, *Review of Financial Studies* 22, 859–892.
- Henderson, Brian, and Neil Pearson, 2011, The dark side of financial innovation: A case study of the pricing of a retail financial product, *Journal of Financial Economics* 100, 227–247.
- Holmstrom, Bengt, and Joan Ricart I Costa, 1986, Managerial incentives and capital management, *Quarterly Journal of Economics* 101, 835–860.
- Jensen, Michael, and William Meckling, 1976, Theory of the firm: Managerial behavior, agency costs and ownership structure, *Journal of Financial Economics* 3, 305–360.
- Jiang, Hao, and Bryan Kelly, 2012, Tail risk and hedge fund returns, *Review of Financial Studies* 27, 2841–2871.
- Liptser, Robert, and Albert Shiryaev, 1978, *Statistics of Random Processes* (Springer-Verlag, Berlin and Heidelberg).
- Malliaris, Steven, and Hongjun Yan, 2012, Reputation concerns and slow-moving capital, Working paper, Yale University.
- Panageas, Stavros, and Mark Westerfield, 2009, High-water marks: High risk appetites? Convex compensation, long horizons, and portfolio choice, *Journal of Finance* 64, 1–36.

Ross, Stephen, 2004, Compensation, incentives, and the duality of risk aversion and riskiness, *Journal of Finance* 59, 207–225.

Scharfstein, David, and Jeremy Stein, 1990, Herd behavior and investment, *American Economic Review* 80, 465–479.

Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Appendix S1: Internet Appendix.